Efficient Matching in the School Choice Problem*

Philip J. Reny

Department of Economics
University of Chicago

Abstract

Stable matchings in school choice needn't be Pareto efficient and can leave thousands of students worse off than necessary. Call a matching μ *priority-neutral* if no matching can make any student whose priority is violated by μ better off without violating the priority of some student who is made worse off. Call a matching *priority-efficient* if it is priority-neutral and Pareto efficient. We show that there is a unique priority-efficient matching and that it dominates every priority-neutral matching and every stable matching. Moreover, truth-telling is a maxmin optimal strategy for every student in the mechanism that selects the priority-efficient matching.

Keywords: school choice, stable matchings, fair matchings, Pareto efficient matchings, priority-efficiency, priority-neutrality, truth-telling, maxmin optimality.

1 Introduction

Many U.S. cities (including New York City, Boston, Seattle, Cambridge, Charlotte, Denver, Minneapolis, and Columbus) have some form of school choice that allows families to choose a school for their children that is outside the district in which they live. But because there may not be enough seats at any given school to accommodate all students for whom that school is their first choice, school districts often use priority rules together with a lottery to resolve the conflicts that inevitably arise.

For example, it is not unusual for applicants in a school's predefined walk zone who have a sibling already enrolled in that school to have priority over applicants who only have a sibling at the school, and for the latter to have priority over applicants who are only in the school's walk zone. Any conflicts between students in the same priority group are then resolved according to the students' randomly assigned lottery numbers. In effect then, each school

*This version: March 13, 2022. I thank Atila Abdulkadiroğlu, Ben Brooks, Yeon-Koo Che, Aram Grigoryan, Thayer Morrill, Parag Pathak, Motty Perry, Doron Ravid, Joe Root, Al Roth, Tayfun Sönmez, Qianfeng Tang, Leeat Yariv, Utku Unver, Tina Danting-Zhang, and the editor and several referees for helpful comments. Financial support from the National Science Foundation (SES-1724747, SES-2049810) is gratefully acknowledged. Author contact: preny@uchicago.edu

can strictly rank any two students by considering which of them has higher priority, or, if they have the same priority, which of them has a higher-ranked random number. Henceforth, we will refer to this strict ranking as the school's priority order even though, in practice, a school's strict ranking of all the students may be a consequence of combining priorities with randomly assigned numbers.

A matching of students to schools creates a priority violation when some student i prefers another school s to their own school and, either, school s has a vacant seat, or, student i has higher priority at s than some student assigned to s. Such a matching is said to violate i's priority (at school s). Because schools can strictly rank the students, any conflicts between students over a given school can be resolved by that school's priority order. Even so, it is not at all clear whether the priority orders across all of the schools are mutually compatible, i.e., whether there is a matching of students to schools that does not violate any student's priority at any school. When such a matching does exist, it is called *stable*.[1]

Remarkably, stable matchings always do exist, regardless of the schools' priority orders and regardless of the students' preferences over schools. Even more remarkable is that among all of the stable matchings there is one (and only one) that all of the students agree is best. Both of these results are due to Gale and Shapley (1962) (henceforth GS), who call the stable matching that is best for all the students, *student-optimal*.

Unfortunately, as is well-known, the student-optimal stable matching need not be Pareto efficient.[2] In fact, the extent of the inefficiencies can be very large. For example, Kesten (2010) (henceforth K) shows that for any set of schools and seat quotas, there are school priorities, students, and student preferences over schools such that the student-optimal stable matching assigns each student to his or her worst or second-worst school. While this theoretical possibility is indeed an extremely poor outcome, one might wonder whether any significant inefficiencies actually occur in practice.

According to Abdulkadiroğlu et. al. (2009), in a New York City school district in 2006-2007, over 4,000 grade 8 students could have been made better off by reassigning them to a school different than their match in the student-optimal stable matching, without hurting any other students. Thus the extent of the inefficiencies that can arise in the student-optimal stable matching is a matter of real practical importance.[3]

When the student-optimal stable matching is not Pareto efficient, selecting any Pareto

[1] Some authors instead call such a matching *fair*. The two concepts are equivalent.

[2] We treat schools here as objects to be allocated to students and so Pareto efficiency is always with respect to students only. This is in keeping with much of the school choice literature starting with Balinski and Sönmez (1999).

[3] Abdulkadiroğlu, Pathak, and Roth's (2009) main focus is not on Pareto efficiency but on the quite distinct constrained efficiency problem that seeks a stable matching that is undominated among all stable matchings in the presence of priority rules with ties. See also Erdil and Ergin (2008).

efficient matching at all will necessarily create priority violations. Taking the view that priorities must be respected unless a Pareto improvement is possible leads to the following question. Is there a natural way to select a Pareto efficient matching that weakly Pareto dominates the student-optimal stable matching? The answer is not obvious. Indeed, no such matching mechanism can be strategy proof.[4] In particular, Abdulkadiroğlu and Sönmez's (2003) adaptation to the school choice problem of Gale's top trading cycles mechanism,[5] being Pareto efficient and strategy-proof, cannot solve the problem at hand.[6]

One approach to the problem, pioneered by K, asks students not only to submit preferences over schools but also asks them to either "consent" or "not consent" to allowing their priorities to be violated. With this information, K computes a matching using a novel modification of the deferred acceptance algorithm that violates a student's priority only if that student has given their consent. K shows that students are never harmed by giving their consent—and so we assume henceforth that all students do consent in K's mechanism—and that when all students consent, the algorithm produces a matching that is Pareto efficient and dominates the student-optimal stable matching.[7] These properties are clearly important. However, it turns out that other mechanisms that produce distinct matchings have these properties as well and so the fundamental selection problem remains.[8]

Another approach to the problem is to expand the set of matchings by modifying the conditions that define "stable" matchings. A powerful example of this approach is due to Ehlers and Morrill (2020) (henceforth EM) who apply GS's definition of blocking to von Neumann and Morgenstern's (1944) (henceforth vNM) set-valued definition of stability.[9] This leads EM to define a set of matchings to be *legal* if it contains precisely those matchings that are individually rational and do not violate any student's priority at any school to which that student could be assigned by some matching in the set.[10] The self-referential nature of this definition implies that legal sets of matchings are in fact defined as the fixed points of a set-valued map, which is typical of vNM stable sets.[11] EM show that this map has a

[4]See Balinski and Sönmez (1999, Lemma 3), Abdulkadiroğlu et. al. (2009, Theorem 1), and Kesten (2010, Proposition 1).

[5]Gale's mechanism is described in Gale and Shapley (1974).

[6]Hence, the mechanisms studied in Pápai (2000), being Pareto efficient and group strategy-proof, also cannot solve the present problem.

[7]For some experimental evidence on students' consent decisions, see Cerrone et. al. (2021).

[8]Consider, for example, a mechanism that selects the student-optimal stable matching if one or more students do not consent but selects any dominating Pareto efficient matching otherwise. With this mechanism, if a student does not give consent, it is dominant to report their preferences truthfully (Dubins and Freedman 1981). Moreover, if one reports truthfully, there is never any harm in giving consent since one's school assignment can only improve. Hence, reporting preferences truthfully and giving consent dominates any strategy that withholds consent.

[9]See Section 4.

[10]A matching is individually rational if no student's assigned school is worse for them than being unmatched.

[11]GS stable sets can also be defined as fixed points of set-valued maps. The difference, however, is that

unique fixed point—i.e., that there is a unique legal set of matchings—and that there is a unique matching in this legal set that all students agree is best. In addition, they show that this student-optimal legal matching is Pareto efficient, that it Pareto dominates the student-optimal stable matching, and that it in fact coincides with the matching that is computed by K's algorithm.[12]

Because legal sets are defined relative to themselves, legal matchings can be difficult to justify. For example, when—as often happens—the student-optimal legal matching violates a student's priority at some school, there is no simple explanation as to why no legal matching can place the student there. The approach that we take here emphasizes simplicity. While certainly desirable on general principles, simplicity may be especially important for Pareto efficient solutions to school choice problems because practical success very likely hinges upon one's ability to explain (to students, to parents and, if necessary, to judges) why, under the proposed solution, any student whose priority is violated at a school should not be assigned there.

Technically, the route we take may be seen as "dual" to that of EM in that we apply a more stringent definition of GS's blocking concept to GS's definition of stability.[13] But the basic idea is simply to modify the usual definition of Pareto efficiency so as to pay due respect to school priorities.

Say that a matching is *priority-neutral* if it is not possible to make any student whose priority is violated better off without violating the priority of some student who is made worse off. Notice that all stable matchings are priority-neutral because they violate no student's priority.

While stability captures the idea that students have an absolute right to relief from priority violations—no matter the effect on other students—priority-neutrality captures the idea that students have a right to relief from priority violations but cannot violate others' priorities without remedy in order to gain such relief.

Of course our goal is to select a Pareto efficient matching. Consequently, we seek matchings that are both priority-neutral and Pareto efficient, and so let us call any such matching *priority-efficient*. It is not at all clear whether there are many priority-efficient matchings or whether there are none at all.

Our results are as follows. First, there always exists precisely one priority-efficient matching. Moreover, every student weakly prefers this matching to every priority-neutral matching and so to every stable matching as well. Second, the matching that is singled out both by

GS stable sets can be defined more directly, without any self-reference at all.

[12]EM allow school priorities to be described by substitute choice functions as in Blair (1988), whereas K assumes responsive priorities as do we. So whenever EM's results are compared with K's or with our own, we are restricting EM's priorities to be responsive.

[13]See Section 4.

K and by EM is precisely the priority-efficient matching. Third, the set of priority-neutral matchings is a lattice under the partial order defined by the students' preferences, and, every priority-neutral matching is in EM's legal set but matchings in the legal set can fail to be priority-neutral.

We also briefly consider the mechanism, let us call it the *priority-efficient (PE) mechanism*, that selects the priority-efficient matching after students submit their preferences (school priorities and quotas are assumed known to mechanism administrators). The strategic properties of the PE mechanism are the same as the analogous mechanisms of K and EM because all three mechanisms select the priority-efficient matching. Hence, as shown by K there are information environments under which truth-telling is an equilibrium of the PE mechanism,[14] although truth-telling is not a dominant strategy. Our final result states that truth-telling in the PE mechanism is maxmin optimal for every student when they are unsure of the preferences that other students will submit. Thus, sufficiently cautious students participating in the PE mechanism are well-advised to report truthfully.

The remainder of the paper is organized as follows. Section 2 provides an example of a priority-efficient matching. Section 3 contains our notation and results. Section 4 discusses other related literature and the clarifies the sense in which our approach is dual to that of EM. All proofs are in Section 5.

2 An Example

To gain some familiarity with priority-efficient matchings, consider the school choice problem illustrated in Figure 1 involving five students, $i_1, ..., i_5$, and five schools, $s_1, ..., s_5$ each with a quota of one seat.[15] Student preferences are given by the table on the left and school priorities are given by the table on the right. For example, the table on the left indicates that student i_1 ranks school s_2 highest, s_1 second-highest, etc., while the table on the right indicates that school s_2 gives highest priority to student i_3, second-highest priority to student i_5, etc. Dots indicate that the remaining rankings do not matter for the purposes of this example.

The table of student preferences in Figure 1 depicts four matchings. The shaded squares indicate the student-optimal stable matching, while the three other matchings $\underline{\mu}$ (underlined), μ° (circled), and μ^* (asterisked), are Pareto efficient, and two of them, μ° and μ^*, Pareto dominate the student-optimal stable matching.[16]

[14]See K Theorem 2. See also EM Theorem 5, and Reny (2021) Theorem 4.2.

[15]In this example, students prefer any school to being unmatched and so the latter possibility can be ignored.

[16]There are other Pareto efficient matchings, e.g. the matching that assigns students 1,2,...,5 to schools

To see how priority-efficiency works to select a unique Pareto efficient matching, we will be content here to show that among the three Pareto efficient matchings $\underline{\mu}$, μ°, and μ^{*}, only μ^{*} is priority-efficient. To do so, it suffices to check that only μ^{*} is priority-neutral since we already know that it is Pareto efficient.

Consider the matching $\underline{\mu}$. This matching violates student i_2's priority at school s_3 (and also at school s_1), because student i_2 prefers school s_3 to school s_5 where she is assigned, and student i_2 has priority over student i_3 at school s_3 where i_3 is assigned. Consider now the student-optimal stable matching indicated by the green-shaded cells in Figure 1. Let us call this stable matching σ. Student i_2 prefers σ to $\underline{\mu}$ because σ assigns i_2 to school s_1, which he prefers to school s_5 to which he is assigned under $\underline{\mu}$. Moreover, because σ is stable, it violates no student's priority. Consequently, $\underline{\mu}$ is not priority-neutral because the matching σ makes a student, i_2, whose priority is violated by $\underline{\mu}$ better off without violating the priority of any student at all. So $\underline{\mu}$ is not priority-efficient.

Next, consider the matching μ°. This matching too violates student i_2's priority at school s_3 and for the same reason as given in the previous paragraph. Let us compare μ° to the matching μ^{*}. Student i_2 prefers μ^{*} to μ°, and the only student who finds μ^{*} worse than μ° is student i_3, whose priority is not violated by μ^{*}. Consequently, μ° is not priority-neutral because the matching μ^{*} makes student i_2, whose priority is violated by μ°, better off without violating the priority of any student who is made worse off (notice that μ^{*} violates i_5's priority at s_1, s_2, and s_3 but makes i_5 no worse off than μ°). So μ° too is not priority-efficient.

Student Preferences

i_1	i_2	i_3	i_4	i_5
$\underline{s}_2^{*}$	s_3^{*}	$\underline{s}_3$	$\underline{s}_1^{*}$	$\underline{s}_4$
s_1	s_1	s_4^{*}	s_2	s_1
s_3	$\underline{s}_5$	s_2	s_4	s_3
⋮	⋮	⋮	⋮	s_2
				s_5^{*}

School Priorities

s_1	s_2	s_3	s_4	s_5
i_2	i_3	i_1	i_4	⋮
i_1	i_5	i_5	i_3	
i_5	i_4	i_2	i_5	
i_4	i_1	i_3	⋮	
i_3	i_2	i_4		

(each school has one seat)

Figure 1. The student optimal stable matching is shaded; the underlined and circled matchings are each Pareto efficient but not priority-neutral; the matching indicated by asterisks is Pareto efficient and priority-neutral, hence priority-efficient.

2,3,5,1,4, respectively.

Finally, consider the matching μ^*. We wish to show that μ^* is priority-neutral. To see this, notice first that the only student whose priority is violated by μ^* is student i_5 (at schools s_1, s_2, and s_3). To make student i_5 better off, one of the other students would have to be assigned to school s_5, which would make that student worse off. Moreover, no matter which of the other students is assigned to school s_5, that student's priority will be violated because he has top priority at a school that he prefers to s_5. Consequently, it is not possible to change the matching from μ^* so as to make the only student whose priority is violated by μ^*, student i_5, better off without violating the priority of a student who is made worse off. Hence, μ^* is priority-neutral and therefore, being Pareto efficient, it is priority-efficient. Notice also that μ^* Pareto dominates the student-optimal stable matching.

3 Notation and Results

Let I denote the nonempty finite set of students and let S denote the nonempty finite set of schools. The set S contains a distinguished element, $\emptyset$, called the *null school*, which represents being unmatched. Each school $s \in S$ has a finite number of available seats, or quota, $q_s \in \{1, 2, ...\}$, with $q_\emptyset = \#I$, and each $s \in S$ has a strict total order Π_s over the set of students I.[17] Each student $i \in I$ has a strict total order P_i over the set of schools S, and we write sR_it to mean sP_it or $s = t$. We will call Π_s school s's priority order over students, and we will call P_i student i's preferences over schools. All of these elements are fixed throughout the analysis, unless stated otherwise.

A *matching* is any mapping $\mu : I \to S$ such that $\#\mu^{-1}(s) \leq q_s$ for every $s \in S$.

For any two matchings μ and ν, we reduce notation by writing $\mu P_i \nu$ instead of $\mu(i)P_i\nu(i)$ and by writing $\mu R_i \nu$ instead of $\mu(i)R_i\nu(i)$ and similarly for $sP_i\mu$ for $s \in S$, etc. We begin with some standard definitions.

A matching μ *violates student i's priority* if there is $s \in S$ such that $sP_i\mu$ and, either, $i\Pi_sj$ for some $j \in \mu^{-1}(s)$ or $\#\mu^{-1}(s) < q_s$.[18] We then also say that μ violates i's priority at school s.[19]

A matching μ is *stable* if it does not violate any student's priority (at any school).

A matching μ *dominates* a matching ν if $\mu R_i \nu$ for every $i \in I$. Hence, because student preferences are strict, a matching μ Pareto dominates a matching ν if and only if μ dominates ν and $\mu \neq \nu$.

[17] The null school's priority order, $\Pi_\emptyset$, is included for notational convenience only. See footnote ??.

[18] Thus, by convention, every student has priority over every empty seat.

[19] Because $q_\emptyset = \#I$, μ violates i's priority at school $\emptyset$ if and only if $\emptyset P_i \mu$. In particular, whether i's priority is violated by μ is independent of the null school's priority order $\Pi_\emptyset$. Consequently, none of the definitions in this paper depend in any way on $\Pi_\emptyset$.

A matching is a *student-optimal stable matching* if it is stable and it dominates every other stable matching.

GS show that a student-optimal stable matching always exists and that it is unique. We next introduce our main concepts.

Say that a matching μ is *priority-neutral* if no matching ν can make any student whose priority is violated by μ better off unless ν violates the priority of some student who is made worse off.

Say that a matching is a *student-optimal priority-neutral matching* if it is priority-neutral and it dominates every other priority-neutral matching.

Say that a matching μ is *priority-efficient* if it is priority-neutral and Pareto efficient.

We can now state our first result.

Theorem 3.1 *There is a unique priority-efficient matching μ^*. This priority-efficient matching dominates every priority-neutral matching and so μ^* is the unique student-optimal priority-neutral matching. In particular, since every stable matching is priority-neutral, μ^* dominates the student-optimal stable matching. Finally, any matching μ that is not dominated by μ^* violates the priority of some student who strictly prefers μ^* to μ.*

Remark 3.2 *The last sentence in the theorem statement says that the matching μ^* has the property that it is not possible to make any student—whether their priority is violated or not—better off without violating the priority of some student who is made worse off.*[20]

It turns out that the matching that is produced by K's algorithm and the matching that EM identify as their student-optimal legal matching are both equal to the unique priority-efficient matching. Before providing the formal statement of this result, we remind the reader of EM's main definition and results.

Let IR be the set of individually rational matchings.[21] A set of matchings L is defined by EM to be *legal* if $L = \{\mu \in IR :$ for every $\nu \in L$ and for every student i, μ does not violate i's priority at $\nu(i)\}$. EM show that there is a unique legal set and that there is a matching in it that dominates all the others—so they call this matching the student-optimal legal matching. Moreover, they show that the student-optimal legal matching coincides with the matching that is computed by K's algorithm.

We can now state our equivalence result.

Theorem 3.3 *The matching that is computed by K's algorithm and the matching that EM call the student-optimal legal matching are both equal to the unique priority-efficient matching.*

[20] Since any matching satisfying this property must evidently be both priority-neutral and Pareto efficient, it follows by Theorem 3.1 that μ^* is the unique matching with this property.

[21] A matching μ is individually rational if $\mu R_i \emptyset$ for every student i.

Our next result describes the structure of the set of priority-neutral matchings.

Theorem 3.4 *The set of priority-neutral matchings is a lattice with respect to the partial order defined by the dominance relation.*[22] *The largest element of this lattice is the priority-efficient matching and the smallest element is the worst stable matching for all of the students.*

Remark 3.5 *While the least upper bound of any two priority-neutral matchings is always their coordinatewise (i.e., student-by-student) maximum, we do not know whether the greatest lower bound is always their coordinatewise minimum. Nevertheless, for any two priority-neutral matchings there is always a largest priority-neutral matching that is smaller than both.*

EM show that their legal set of matchings is a lattice with the same partial order employed here. Our next result says that the lattice of priority-neutral matchings is a subset of the lattice of matchings in the legal set.[23]

Theorem 3.6 *Every priority-neutral matching is a member of EM's legal set of matchings.*

But there are legal matchings that fail to be priority neutral as the next example shows.

Example 3.7 *In the school choice problem shown in Figure 2,*[24] *the matching μ^* that is asterisked is the unique priority-efficient matching.*[25] *By Theorem 3.3, μ^* is also EM's student-optimal legal matching. We wish to show that the matching $\underline{\mu}$ that is underlined in Figure 2 is a member of EM's legal set, L, but that it is not priority-neutral. Let us first show that $\underline{\mu} \in L$. Suppose, by way of contradiction, that $\underline{\mu}$ is not in L. Then, since $\underline{\mu}$ is individually rational, the definition of a legal set (see just above Theorem 3.3), implies that there must be a matching $\nu \in L$ and a student i whose priority is violated by $\underline{\mu}$ at $\nu(i)$. Since the only student whose priority is violated by $\underline{\mu}$ is student i_5, whose priority is violated at school s_2, we must have $\nu(i_5) = s_2$. But then, consulting Figure 2, $\nu(i_5) = s_2 P_{i_5} \mu^*(i_5)$. Together with $\nu \in L$, this contradicts the student-optimality of μ^* among all elements of L. Hence, $\underline{\mu} \in L$. To see that $\underline{\mu}$ is not priority-neutral, observe that student i_5, whose priority is violated by $\underline{\mu}$ at school s_2, is made better off by the circled matching and no student is made worse off.*

[22] That is, the partial order $\geq$ defined by, $\mu \geq \nu$ if $\mu R_i \nu$ for every $i \in I$.

[23] We do not know whether the lattice of priority-neutral matchings is a sublattice of the lattice consisting of the legal set. Specifically, while the former is a join sublattice of the latter, it is an open question as to whether it is a meet sublattice as well.

[24] In this example students prefer any school $s_1, ..., s_5$ to the null school and so we omit the null school from the figure.

[25] This can be verified by applying the Kesten-Tang-Yu algorithm, which is described in Section 5.

Student Preferences

i_1	i_2	i_3	i_4	i_5
s_2^*	s_3^*	s_3	s_1^*	s_4
s_1	$\underline{s_1}$	$\underline{s_4^*}$	$\underline{s_2}$	s_1
s_5	s_5	s_2	s_4	s_2
$\underline{s_3}$	$\vdots$	$\vdots$	$\vdots$	$\underline{s_5^*}$
s_4				s_3

School Priorities

s_1	s_2	s_3	s_4	s_5
i_2	i_3	i_1	i_4	i_5
i_1	i_5	i_5	i_3	$\vdots$
i_5	i_4	i_2	i_5	
i_4	i_1	i_3	$\vdots$	
i_3	i_2	i_4		

(each school has one seat)

Figure 2. The underlined matching is a member of Ehlers and Morrill's (2020) unique legal set, but it is not priority-neutral.

Finally, let us consider the direct mechanism that selects the priority-efficient matching after students submit their preferences, where mechanism administrators are assumed to know the schools' priorities and quotas. Let us call this mechanism the *priority-efficient (PE) mechanism.*

Recall that for each school $s \in S$, Π_s denotes its priority ranking over students and q_s denotes its quota. Let (Π, q) denote the entire profile of school priorities and quotas. For every student i and for every pair of i's school-preferences P_i and P_i', let $w_i(P_i'|P_i, \Pi, q)$ be the P_i-worst school to which student i can be assigned in any priority-efficient matching when student i submits preferences P_i', when school priorities and quotas are fixed and given by (Π, q), and as the other students' submitted preferences vary over all possible preferences over schools.[26] Thus, we are assuming that each student knows (Π, q), which makes the following result stronger (see the remark just below).

Theorem 3.8 *Let (Π, q) be any profile of school priorities and quotas. For any student i and for any pair of preferences P_i and P_i' for i,*

$$w_i(P_i|P_i, \Pi, q) R_i w_i(P_i'|P_i, \Pi, q).$$

That is, for any fixed profile of school priorities and quotas, truth-telling in the PE mechanism is a maxmin optimal strategy for every student, where the minimum is taken over all possible submitted preferences of other students.

[26] I am grateful for conversations with Ben Brooks that led to the consideration here of maxmin student behavior.

Remark 3.9 *The result holds a fortiori if school priorities or quotas are unknown and students minimize also over them. It seems possible to strengthen the result by reducing the set of others' preferences over which the minimum is taken, though we leave this for future work. See footnote 32.*

4 Discussion

Others who have followed K's algorithmic approach include Che and Tercieux (2019), and Dur et. al. (2019), while others who have followed EM insofar as employing vNM stable sets include Ehlers (2007), and Mauleon et. al. (2011).

For approaches that compare mechanisms by looking at their sets of priority violations, see Abdulkadiroğlu et. al. (2020), Kwon and Shorrer (2020), and Tang and Zhang (2020). From results in the last of these papers, we may conclude that switching from the priority-efficient matching to any other Pareto efficient matching may eliminate some priority violations but always creates at least one new priority violation.

Finally, let us briefly explain the technical sense in which the approach we have taken here can be seen as "dual" to the approach taken by EM. For any two matchings μ and ν, say that ν *GS-blocks* μ if there is a student i whose priority is violated by μ at $\nu(i)$.[27] Say that ν *neutrally-blocks* μ if (a) μ violates some student i's priority, (b) $\nu P_i \mu$, and (c) $\nu R_j \mu$ for every student j whose priority is violated by ν. For any matching μ, say that μ is *GS-stable under GS-blocking* if μ is not GS-blocked by any matching ν, and say that μ is *GS-stable under neutral-blocking* if μ is not neutrally-blocked by any matching ν. For any set of matchings S, say that S is *vNM-stable under GS-blocking* if $S = \{\mu \in M : \mu$ is not GS-blocked by any $\nu \in S\}$, where M is the set of all matchings.

With these definitions, GS's stable matchings are those matchings that are GS-stable under GS-blocking, EM's legal sets are those sets of individually rational matchings that are vNM-stable under GS-blocking, and priority-neutral matchings are those matchings that are GS-stable under neutral-blocking. Thus EM maintain GS's notion of blocking but modify their notion of stability while, "dually," we modify GS's notion of blocking and maintain their notion of stability.

[27] GS-blocking as defined here (called simply "blocking" by EM) is equivalent to the blocking concept that is implicit in GS and made explicit in, e.g. Roth and Sotomayor (1990).

5 Proofs.

For any pair of matchings μ and ν, define the functions $\mu \vee \nu$ and $\mu \wedge \nu$ (the coordinatewise maximum and minimum of μ and ν, respectively), each mapping I into S, as follows. For each $i \in I$, $[\mu \vee \nu](i)$ is the one school, $\mu(i)$ or $\nu(i)$, that student i weakly prefers to the other, and $[\mu \wedge \nu](i)$ is the one school, $\mu(i)$ or $\nu(i)$, that student i prefers weakly less than the other. Note that these functions are well-defined because student preferences are strict and so indifference between $\mu(i)$ and $\nu(i)$ can occur only if $\mu(i) = \nu(i)$. In general neither $\mu \vee \nu$ nor $\mu \wedge \nu$ need be matchings.

Our first lemma provides a modest but useful generalization of the standard lattice property of stable matchings that is basic to much of matching theory. EM (Lemmas 4, 5, and 6) prove a similar result in a more general setting, but restrict attention to matchings that are individually rational, which here would mean restricting attention to matchings that each student weakly prefers to the null school $\emptyset$. Since this restriction is not actually necessary (indeed, the result holds even if the matchings considered are individually irrational for every student) and because in our setting with responsive priorities for schools a very short proof relying on well-known results can be given, we provide the lemma and its proof here.

Lemma 5.1 *Let μ and ν be any two matchings and suppose that for every student i, μ does not violate i's priority at $\nu(i)$ and ν does not violate i's priority at $\mu(i)$. Then, (a) $\mu \vee \nu$ and $\mu \wedge \nu$ are matchings, (b) under each of the four matchings μ, ν, $\mu \vee \nu$, and $\mu \wedge \nu$, the set of students assigned to the null school is the same and, for each school, the number of students assigned to that school is the same, and (c) when the matching switches from μ (or from ν) to $\mu \vee \nu$, each school's students before the switch all have strictly higher priority at that school than every student who transfers in to that school, and, when the matching switches from μ (or from ν) to $\mu \wedge \nu$, each school's students after the switch all have strictly higher priority at that school than every student who transfers out of that school.*

Proof. Let μ and ν satisfy the stated hypotheses, let P denote the profile of all student preferences, and let Π and q denote the profiles of all school priorities and quotas, respectively. For each student i, raise both $\mu(i)$ and $\nu(i)$ to the top of i's ranking, but without changing their relative order (hence $\mu \vee \nu$ and $\mu \wedge \nu$ are unaffected). Call the new profile of student preferences $\hat{P}$. Then, by our hypotheses, both μ and ν are stable matchings in the school choice problem $(\hat{P}, \Pi, q)$. All of the desired conclusions now follow from results for stable matchings in Roth and Sotomayor (1990), i.e., their Corollary 5.32 proves (a), their Theorem 5.12 proves (b), and their Theorems 5.27 and 5.33 prove (c). ■

Say that a school $s \in S$ (note that s can be the null school) is *underdemanded* at a matching μ if no student prefers s to the school to which they are assigned by μ.[28]

Tang and Yu (2014) (henceforth TY) provide a simpler alternative to K's algorithm which they show computes the same matching as K's algorithm. This alternative algorithm, which we will call the *Kesten-Tang-Yu algorithm*, proceeds in rounds and works as follows.[29]

In the first round, run the student-proposing deferred acceptance (DA) algorithm with the entire sets of schools and students, yielding the student-optimal stable matching, μ_1 say. Each school s that is underdemanded at μ_1 (equivalently, each school that did not reject any students during the execution of the DA algorithm),[30] is permanently assigned its students $\mu_1^{-1}(s)$ (which can be the empty set) for the remainder of the algorithm and school s and its students are called *settled*. Remove all settled schools and students and proceed to round two.[31] The second round proceeds exactly as the first but where the DA algorithm is applied only to the "submarket" of unsettled schools and students. This submarket's underdemanded schools at its DA matching are permanently assigned their students, and these students and schools become settled and are removed, altogether resulting in a second-round matching μ_2 (which includes the permanently assigned students and their schools from the first round). These rounds repeat, with each round n producing a matching μ_n, and with smaller sets of unsettled schools and students with each successive round. The algorithm ends after the round, N say, in which all remaining schools and students become settled, thereby defining the matching, μ_N, that is the output of the algorithm. We then call $\mu_1, ..., \mu_N$ the *Kesten-Tang-Yu sequence.*

Following TY, for any matching μ and for any student i, say that student i is *(Pareto) μ-improvable* if there is a matching that dominates μ and that makes student i strictly better off. Say that student i is *(Pareto) μ-unimprovable* if student i is not μ-improvable.

Lemma 5.2 *Suppose that $\mu_1, ..., \mu_N$ is the Kesten-Tang-Yu sequence and that student i is settled in round n. Then student i is μ_n-unimprovable.*

Proof. Throughout the proof, by "round one," "round two," etc., we will always mean the corresponding round of the Kesten-Tang-Yu algorithm.

Let i, n and $\mu_1, ..., \mu_N$ be as given in the statement of the lemma and let ν dominate μ_n. We must show that $\nu(i) = \mu_n(i)$.

[28] Kesten and Kurino (2013) introduced the "underdemanded school" terminology into the school choice literature. The idea can be traced back to at least Gale and Sotomayor's (1985) notion of a man (or woman) with no "admirers" in a marriage market.

[29] We describe here only the special case of the Kesten-Tang-Yu algorithm in which all students consent to allowing their priorities to be violated.

[30] When there are at least as many seats in total across all schools as students, as our convention $q_\emptyset = \#I$ implies, it is well-known that at least one such school always exists. See, e.g. Gale and Sotomayor (1985) for the case of one to one matching.

[31] In contrast to TY, we find it more convenient to remove underdemanded schools at the end of each round rather than at the beginning of each round. This has no effect on the final outcome.

Let J_n and T_n be the sets of students and schools, respectively, that are unsettled at the start of round n (so $J_1 = I$ and $T_1 = S$). Then $i \in J_n$.

Let j be any student in J_n and let s be any school outside T_n. Hence, there is $k < n$ such that s is settled in round k. Therefore, school s is underdemanded at the student-optimal stable matching for round k's submarket of unsettled students and schools. Hence, $\mu_k(j) P_j s$ since student j is included in that submarket (student $j \in J_n$ is not settled until round $n > k$ or later) and since $\mu_k(j)$ is j's school in that submarket's student-optimal stable matching. Hence, $\mu_n(j) P_j s$ since $j \in J_n$ implies that $\mu_n R_j \mu_k$ (by TY Lemma 2). Consequently, the only schools that student j prefers to $\mu_n(j)$ are schools that are unsettled at the start of round n, i.e., schools in T_n.

Since ν dominates μ_n, we have that for any $j \in J_n$, either, $\nu(j) P_j \mu_n(j)$ in which case $\nu(j) \in T_n$ by the conclusion of the previous paragraph, or, $\nu(j) = \mu_n(j)$ in which case $\nu(j) \in T_n$ because, by definition, the round n matching μ_n assigns every student in J_n to a school in T_n. Hence, ν assigns every student in J_n to a school in T_n, which means that the restriction of ν to J_n is a matching for the submarket (J_n, T_n) consisting of students in J_n and schools in T_n. By definition, the restriction of μ_n to J_n is the student-optimal stable matching for the submarket (J_n, T_n). Since, by hypothesis, student i is settled in round n, i's assigned school $\mu_n(i)$ is underdemanded at this student-optimal stable matching in this submarket. Also, since ν dominates μ_n, the restriction of ν to J_n dominates the restriction of μ_n to J_n. Hence, we may apply TY Lemma 1 to the submarket (J_n, T_n) to conclude that $\nu(i) = \mu_n(i)$ as desired. ■

Say that a finite sequence of matchings $\mu_1, \mu_2, ..., \mu_N$ is *feasible* if μ_1 is stable, μ_N is Pareto efficient, and, for each $n > 1$, μ_n dominates μ_{n-1} and μ_n does not violate the priority of any μ_{n-1}-improvable student. Feasible sequences of matchings turn out to play a central role in the study of priority-neutral matchings, and their existence is ensured by the following lemma.

Lemma 5.3 *The Kesten-Tang-Yu sequence of matchings is well-defined and feasible, and (as shown by TY) the last matching in the sequence coincides with the matching that is computed by K's algorithm.*

Proof. By TY Proposition 1, the Kesten-Tang-Yu algorithm is well defined and ends in finitely many rounds, N say. Therefore it produces a finite sequence of matchings, $\mu_1, \mu_2, ..., \mu_N$, where μ_n is the matching produced in the n-th round. By TY Theorem 3, μ_N is the matching that is computed by K's algorithm. So it remains only to show that $\mu_1, \mu_2, ..., \mu_N$ is feasible. Henceforth, by "first round," "second round," etc., we mean the corresponding round of the Kesten-Tang-Yu algorithm.

Observe first that the matching μ_1 that is produced in the first round is the student-optimal stable matching. In particular, μ_1 is stable. Second, μ_N is Pareto efficient by TY Theorem 1. Third, for each $n > 1$, $\mu_{n-1}(j) = \mu_n(j)$ for any student j that is settled before round n, and $\mu_n R_j \mu_{n-1}$ for any student j who is unsettled at the start of round n by TY Lemma 2. Hence, μ_n dominates μ_{n-1}. So it remains only to show that for each $n > 1$, μ_n does not violate the priority of any μ_{n-1}-improvable student.

Suppose that student i is μ_{n-1}-improvable. Then for every $k \leq n-1$, student i is μ_k-improvable because μ_{n-1} dominates μ_k. Hence, by Lemma 5.2, student i is not settled before round n. Therefore student i is included in the submarket consisting of students and schools that are unsettled at the start of round n. Since, by definition, the restriction of μ_n to students in that submarket is the student-optimal stable matching for that submarket, student i's priority is not violated by μ_n at any school that is unsettled at the start of round n. For any other school, i.e., any school s that is settled in some round $k < n$, school s is underdemanded at μ_k restricted to round k's submarket of unsettled students and schools. Hence, $\mu_k(i) P_i s$ since student i is included in that submarket, and so $\mu_n(i) P_i s$ since μ_n dominates μ_k. Therefore, μ_n does not violate i's priority at any school, whether that school is unsettled at the start of round n or not. Since i was arbitrary, we may conclude that μ_n does not violate the priority of any μ_{n-1}-improvable student. ■

Lemma 5.4 *Let $\mu_1, ..., \mu_N$ be any feasible sequence of matchings and let μ be priority-neutral. Then $\mu R_i \mu_N$ holds for every student i whose priority is violated by μ.*

Proof. Let μ be a priority-neutral matching. We must show that $\mu R_i \mu_N$ for every student i whose priority is violated by μ.

We begin by showing that for every $n \in \{1, ..., N\}$,

(1) if μ_n violates any student i's priority, then $\mu_n R_i \mu$, and

(2) if μ violates any student i's priority, then $\mu R_i \mu_n$.

We proceed by induction starting with $n = 1$.

So suppose that $n = 1$. Then (1) holds trivially since μ_1 is stable. To see (2), suppose that μ violates student i's priority. Then $\mu R_i \mu_1$ since otherwise, μ_1 would make i better off without violating the priority of any other student (μ_1 is stable), which would contradict the priority-neutrality of μ. Hence, (2) holds.

Assume as an induction hypothesis that (1) and (2) hold for n. We must show that (1) and (2) hold for $n + 1$.

To see that (1) holds, suppose that μ_{n+1} violates i's priority. We must show that $\mu_{n+1} R_i \mu$. By the induction hypothesis, μ_n and μ satisfy the hypotheses of Lemma 5.1. Consequently, $\mu \vee \mu_n$ is a matching that dominates μ_n. Also, by feasibility μ_{n+1} dominates μ_n. Since μ_{n+1}

violates i's priority, feasibility implies that student i is μ_n-unimprovable. Hence, $[\mu \vee \mu_n](i) = \mu_n(i)$ and $\mu_{n+1}(i) = \mu_n(i)$ (since $\mu \vee \mu_n$ and μ_{n+1} dominate μ_n), from which we obtain $\mu_{n+1}(i) = \mu_n(i) R_i \mu(i)$ and so (1) holds for $n+1$. It remains to show that (2) also holds for $n+1$.

Suppose, by way of contradiction, that (2) fails for $n+1$. That is, suppose that μ violates i's priority and that $\mu_{n+1} P_i \mu$. Since $\mu_{n+1}(i)$ makes i better off than μ, and since μ is priority-neutral, there must be a student j whose priority is violated by μ_{n+1} such that $\mu P_j \mu_{n+1}$. But then (1) would fail for $n+1$, which is a contradiction and completes the induction argument.

Since (2) holds for each $n = 1, ..., N$, we have $\mu R_i \mu_N$ for every i such that μ violates i's priority, as desired. ■

Lemma 5.5 *Let $\mu_1, ..., \mu_N$ be any feasible sequence of matchings and let μ be any matching such that $\mu R_i \mu_N$ holds for every student i whose priority is violated by μ. Then μ_N dominates μ.*

Proof. Let $\mu_1, ..., \mu_N$, and μ satisfy the hypothesis of the Lemma. We will first show by induction on n that,

$$\mu \vee \mu_n \text{ is a matching for every } n \in \{1, ..., N\}. \tag{5.1}$$

For $n = 1$, we must show that $\mu \vee \mu_1$ is a matching. It suffices to show that μ and μ_1 satisfy the hypotheses of Lemma 5.1. For any student i, μ cannot violate i's priority at $\mu_1(i)$. Otherwise, $\mu_1 P_i \mu R_i \mu_N$, where the weak preference follows by the choice of μ_N and μ. But then $\mu_1 P_i \mu_N$ contradicting the fact that, by feasibility, μ_N dominates μ_1. Hence, μ does not violate the priority of any student i at $\mu_1(i)$. Therefore, since μ_1 is stable, μ and μ_1 satisfy the hypotheses of Lemma 5.1 and so $\mu \vee \mu_1$ is a matching.

Next, assume as an induction hypothesis that (5.1) holds for n. We must show that $\mu \vee \mu_{n+1}$ is a matching. So it suffices to show that μ and μ_{n+1} satisfy the hypotheses of Lemma 5.1.

For any student i, μ cannot violate i's priority at $\mu_{n+1}(i)$. Otherwise, $\mu_{n+1} P_i \mu R_i \mu_N$, where the weak preference follows by the choice of μ_N and μ. But then $\mu_{n+1} P_i \mu_N$ contradicting the fact that, by feasibility, μ_N dominates μ_{n+1}. Hence, μ does not violate the priority of any student i at $\mu_{n+1}(i)$.

Suppose next that μ_{n+1} violates i's priority. Then, by feasibility, i is μ_n-unimprovable. Since $\mu \vee \mu_n$ is a matching (induction hypothesis), $\mu \vee \mu_n$ dominates μ_n. Also, by feasibility, μ_{n+1} dominates μ_n. Therefore, since i is μ_n-unimprovable, $[\mu \vee \mu_n](i) = \mu_n(i) = \mu_{n+1}(i)$ and so $\mu_{n+1}(i) = [\mu \vee \mu_n](i) R_i \mu(i)$. So μ_{n+1} does not violate i's priority at $\mu(i)$. Hence, μ and μ_{n+1}

satisfy the hypotheses of Lemma 5.1 and so we may conclude that $\mu \vee \mu_{n+1}$ is a matching. This completes the induction and establishes (5.1).

Setting $n = N$ in (5.1), we may conclude that $\mu \vee \mu_N$ is a matching. Consequently, because, by definition, $\mu \vee \mu_N$ dominates μ_N, and because, by feasibility, μ_N is Pareto efficient, we must have $\mu \vee \mu_N = \mu_N$. But this means that μ_N dominates μ as desired. ■

Lemma 5.6 *If $\mu_1, ..., \mu_N$ is any feasible sequence of matchings then μ_N is the unique priority-efficient matching and it dominates every priority-neutral matching.*

Proof. Let $\mu_1, ..., \mu_N$ be any feasible sequence. We first show that μ_N is priority-efficient. By the feasibility of the sequence, μ_N is Pareto efficient and so it remains to show that μ_N is priority-neutral. So suppose that student i's priority is violated by μ_N and that there is a matching ν such that $\nu P_i \mu_N$. Then μ_N does not dominate ν. Hence, by Lemma 5.5, there must be a student j whose priority is violated by ν such that $\mu_N P_j \nu$. Therefore, μ_N is priority-neutral as desired.

Next, we show that μ_N dominates every priority-neutral matching. But this follows immediately from Lemmas 5.4 and 5.5.

Finally, let us show that μ_N is the unique priority-efficient matching. Let ν be priority-efficient. Then ν is priority-neutral and so by what we have just shown, μ_N dominates ν. However, because ν is priority-efficient, it is Pareto efficient. Hence we must have $\nu = \mu_N$, as desired. ■

Proof of Theorem 3.1. By Lemma 5.3, the Kesten-Tang-Yu sequence of matchings is well-defined and feasible. Hence, by Lemma 5.6, there is a unique priority-efficient matching μ^* and it dominates every priority-neutral matching. In particular, because every stable matching is priority-neutral, μ^* dominates every stable matching, including the student-optimal stable matching. To prove the theorem statement's last sentence, suppose that μ^* does not dominate some matching μ. Then $\mu^* \vee \mu$ is not a matching since, if it were, it would Pareto dominate the Pareto efficient matching μ^*. Hence, by Lemma 5.1 there is a student i such that either (a) μ violates i's priority at $\mu^*(i)$ or (b) μ^* violates i's priority at $\mu(i)$. If (a) holds the desired conclusion follows immediately. If (b) holds then going from μ^* to μ makes student i, whose priority is violated by μ^*, better off. Hence, by the priority-neutrality of μ^*, μ must violate the priority of a student who is made worse off. ■

Proof of Theorem 3.3. That the matching computed by K's algorithm and EM's student-optimal legal matching are the same follows from EM (Theorem 3, Lemma 15, and Remark 3). That the matching computed by K's algorithm is the priority-efficient matching follows from Lemmas 5.3 and 5.6. ■

Lemma 5.7 *Let μ^* be the priority-efficient matching and let μ be any priority-neutral matching. Then $\mu(i) = \mu^*(i)$ for any student i whose priority is violated by μ.*

Proof. Let μ be priority-neutral and suppose that μ violates i's priority. By Theorem 3.1, μ^* dominates μ. So we must have $\mu^*(i) = \mu(i)$ since otherwise μ^* would make student i better off (by strict preferences) and would not make any student worse off, contradicting μ's priority-neutrality. ■

Lemma 5.8 *If μ and ν are priority-neutral, then $\mu \vee \nu$ is priority neutral.*

Proof. Let μ and ν be priority-neutral, let μ^* be the unique priority-efficient matching and let i be any student. By Lemma 5.7, if μ violates i's priority, then $\mu(i) = \mu^*(i)$ and so $\mu(i) = \mu^*(i) R_i \nu(i)$ since μ^* dominates every priority-neutral matching by Theorem 3.1. Hence, μ does not violate i's priority at $\nu(i)$. Similarly ν does not violate i's priority at $\mu(i)$ and so μ and ν satisfy the hypotheses of Lemma 5.1. Therefore by part (a) of that lemma $\mu \vee \nu$ is a matching. It remains to show that $\mu \vee \nu$ is priority-neutral.

Suppose, by way of contradiction, that $\mu \vee \nu$ is not priority neutral. Then there is a student i and a matching w such that (a) $\mu \vee \nu$ violates i's priority, (b) $w P_i (\mu \vee \nu)$, and (c) $w R_j (\mu \vee \nu)$ for every student j whose priority is violated by w.

By (a), there is a school s at which i's priority is violated by $\mu \vee \nu$. Hence, $s P_i (\mu \vee \nu)$ and there is a student j with $[\mu \vee \nu](j) = s$ who has lower priority than i at s. Assume, without loss of generality, that $[\mu \vee \nu](j) = \mu(j)$.

We claim that, (a') μ violates i's priority, (b') $w P_i \mu$, and (c') $w R_j \mu$ for every student j whose priority is violated by w. If this claim is true, it will contradict the priority-neutrality of μ and complete the proof. Since (b') and (c') follow immediately from (b) and (c), respectively, the proof will be complete if we show (a'). But (a') follows because $\mu(j) = s$, $s P_i (\mu \vee \nu) R_i \mu$, and i has higher priority at s than j. ■

Lemma 5.9 *Every priority-neutral matching dominates the worst stable matching for all the students.*

Proof. By Corollary 5.32 in Roth and Sotomayor (1990) there is a stable matching, $\underline{\sigma}$ say, that is the worst stable matching for all the students. Let μ be any priority-neutral matching, and let μ^* be the priority-efficient matching. Then, for every student i whose priority is violated by μ, we have $\mu(i) = \mu^*(i)$ (Lemma 5.7) and so $\mu(i) = \mu^*(i) R_i \underline{\sigma}(i)$ since μ^* dominates every stable matching (Theorem 3.1). In particular, μ does not violate any student i's priority at $\underline{\sigma}(i)$. Also, because $\underline{\sigma}$ is stable, $\underline{\sigma}$ does not violate any student i's priority at $\mu(i)$. Hence μ and $\underline{\sigma}$ satisfy the hypotheses of Lemma 5.1 and so by part (a) of that lemma, $\mu \wedge \underline{\sigma}$ is a matching. We next show that $\mu \wedge \underline{\sigma}$ is stable.

Consider any student i such that $\mu R_i \underline{\sigma}$. When the matching switches from $\underline{\sigma}$ to $\mu \wedge \underline{\sigma}$ student i does not change schools, and by part (c) of Lemma 5.1, the lowest priority student at each school has weakly higher priority after the switch than before. Consequently, because student i's priority is not violated before the switch (because $\underline{\sigma}$ is stable), and because student i's school remains the same, i's priority is not violated after the switch. Hence, $\mu \wedge \underline{\sigma}$ does not violate i's priority.

Next, consider any student i such that $\underline{\sigma} P_i \mu$. Then, by what was shown in the first paragraph, student i's priority is not violated by μ. When the matching switches from μ to $\mu \wedge \underline{\sigma}$, student i does not change schools, and by part (c) of Lemma 5.1, the lowest priority student at each school has weakly higher priority after the switch than before. Consequently, because student i's priority is not violated by μ before the switch, and because student i's school remains the same, i's priority is not violated after the switch. Hence, $\mu \wedge \underline{\sigma}$ does not violate i's priority.

Since $\mu \wedge \underline{\sigma}$ does not violate any student's priority it is stable. Hence, $\underline{\sigma}$, the worst stable matching for all the students, is dominated by $\mu \wedge \underline{\sigma}$. Since, by definition, $\mu \wedge \underline{\sigma}$ is dominated by μ, we conclude by transitivity of the dominance relation that $\underline{\sigma}$ is dominated by μ, as desired. ■

Proof of Theorem 3.4. Let $\geq$ be the partial order over matchings defined by the dominance relation (i.e., $\mu \geq \nu$ iff $\mu R_i \nu$ for all $i \in I$) and let N be the set of priority-neutral matchings. Then $(N, \geq)$ is a partially ordered set. Let μ and ν be any members of N. By Lemma 5.8, their least upper bound in N is the priority-neutral matching $\mu \vee \nu$. We next show that μ and ν have a greatest lower bound among elements of N. By Corollary 5.32 in Roth and Sotomayor (1990) there is a stable matching, $\underline{\sigma}$ say, that is the worst stable matching for all the students. Since every stable matching is priority-neutral, Lemma 5.9 implies that the set $\{w \in N : w \leq \mu \text{ and } w \leq \nu\}$ is nonempty because it contains $\underline{\sigma}$. By what we have just proven about least upper bounds, this set, being finite, has a least upper bound which is therefore the greatest lower bound in N of μ and ν. Hence, $(N, \geq)$ is a lattice with smallest element $\underline{\sigma}$ by Lemma 5.9 and largest element equal to the priority-efficient matching by Theorem 3.1. ■

Proof of Theorem 3.6. Let μ be priority-neutral. Then μ is individually rational because any student can be transferred to the null school, $\emptyset$, without violating any other student's priority (since $q_\emptyset = \#I$). Let L be the unique legal set of matchings (well-defined by EM Theorem 2). By Theorem 3.3, the priority-efficient matching, μ^*, is the student-optimal legal matching. Assume, by way of contradiction, that $\mu \notin L$. Then, since μ is individually rational, by the definition of L there is $\nu \in L$ and $i \in I$ such that i's priority is violated by

μ at $\nu(i)$. Hence, by Lemma 5.7, $\mu(i) = \mu^*(i)$ and so $\mu(i) = \mu^*(i) R_i \nu(i)$ since μ^* dominates every matching in L. But then i's priority is not violated by μ at $\nu(i)$. ■

Proof of Theorem 3.8. Fix the profile (Π, q) of all school priorities and quotas for the remainder of this proof. For any profile P of student preferences and for any student j, let $\mu^*(P)$ be the priority-efficient matching and let $\sigma(P)$ be the student-optimal stable matching in the school-choice problem (P, Π, q), and let $\mu_j^*(P)$ and $\sigma_j(P)$ denote j's assigned school in those matchings.

Consider now any student i, any pair of preferences P_i and P_i' for i, and let $\hat{P}_{-i}$ be any preference profile for the other students that leads to i's P_i-worst priority-efficient matching when he submits preferences P_i. Then,

$$w_i(P_i|P_i, \Pi, q) = \mu_i^*(P_i, \hat{P}_{-i}) R_i \sigma_i(P_i, \hat{P}_{-i}) R_i \sigma_i(P_i', \hat{P}_{-i}), \tag{5.2}$$

where the equality follows by the definition of $\hat{P}_{-i}$, the first weak preference follows because (by Theorem 3.1) the priority-efficient matching dominates (according to the reported preferences) every stable matching , and the second weak preference follows by Dubins and Freedman (1981).

For each student $j \neq i$, modify $\hat{P}_j$ by raising school $\sigma_j(P_i', \hat{P}_{-i})$ to the top of j's ranking,[32] and let $\bar{P}_j$ denote the ranking that results. Then, the matching $\sigma(P_i', \hat{P}_{-i})$ remains stable with respect to the modified profile of student preferences $(P_i', \bar{P}_{-i})$ and there is no better matching for any $j \neq i$. Hence, with respect to $(P_i', \bar{P}_{-i})$, no matching can Pareto improve upon $\sigma(P_i, \hat{P}_{-i})$ because, by stability, student i can be made better off only by being assigned to a school that is at its quota under $\sigma(P_i, \hat{P}_{-i})$ which would require reassigning some student $j \neq i$ to a different, and hence less $\bar{P}_j$-preferred, school. So the matching $\sigma(P_i, \hat{P}_{-i})$ is stable and Pareto efficient with respect to $(P_i', \bar{P}_{-i})$, and hence also the student-optimal stable matching and so, being Pareto efficient, it is priority-efficient by Theorem 3.1. Thus, $\mu^*(P_i', \bar{P}_{-i}) = \sigma(P_i', \hat{P}_{-i})$.

So by transitivity and (5.2), $w_i(P_i|P_i, \Pi, q) R_i \mu_i^*(P_i', \bar{P}_{-i}) R_i w_i(P_i'|P_i, \Pi, q)$, as desired, where the last weak preference follows from the definition of $w_i(P_i'|P_i, \Pi, q)$. ■

References

Abdulkadiroğlu, A., P. A. Pathak, and A. E. Roth. 2009. "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *American Economic Review* 99 (5): 1954-1978.

[32] It would suffice here to raise each $\sigma_j(P_i, \hat{P}_{-i})$ above $\sigma_i(P_i, \hat{P}_{-i})$ if it is not already above it. In fact, it would suffice to modify each $\hat{P}_j$ so that the set of schools ranked below $\sigma_j(P_i, \hat{P}_{-i})$ remains unchanged and so that student i is $\sigma(P_i, \hat{P}_{-i})$-unimprovable.

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

Abdulkadiroğlu, A., Y.-K. Che, P. A. Pathak, A. E. Roth, and O. Tercieux. 2020. "Efficiency, Justified Envy, and Incentives in Priority-Based Matching." *AER: Insights* 2 (4): 425–442.

Alcalde, J., and A. Romero-Medina. 2017. "Fair Student Placement." *Theory and Decision* 83: 293-307.

Balinski, M., and T. Sönmez. 1999. "A Tale of Two Mechanisms: Student Placement." *Journal of Economic Theory* 84: 73-94.

Blair, C. 1988. "The Lattice Structure of the Set of Stable Matchings with Multiple Partners." *Mathematics of Operations Research.* 13 (4): 619-628.

Cerrone, C., Y. Hermstrüwer, and O. Kesten. 2021. "School Choice with Consent: An Experiment," Unpublished.

Che, Y.-K., and Olivier Tercieux. 2019. "Efficiency and Stability in Large Matching Markets" *Journal of Political Economy.* 127 (5): 2301-2342.

Chwe, M. S.-Y. 1994. "Farsighted Coalitional Stability." *Journal of Economic Theory.* 63: 299-325.

Dubins, L. E., and D. A. Freedman. 1981. "Machiavelli and the Gale–Shapley Algorithm." *American Mathematical Monthly* 88 (7): 485–494.

Dur, U., A. A. Gitmez, and Ö. Yilmaz. 2019. "School Choice Under Partial Fairness." *Theoretical Economics.* 14: 1309–1346.

Ehlers, L. 2007. "von Neumann-Morgenstern Stable Sets in Matching Problems." *Journal of Economic Theory* 134: 537-547.

Ehlers, L., and T. Morrill. 2020. "(Il)legal Assignments in School Choice." *Review of Economic Studies.* 87: 1837–1875.

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

Gale, D., and L. S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly.* 69 (1): 9-15.

Gale, D., and M. Sotomayor. 1985. "Some Remarks on the Stable Matching Problem." *Discrete Applied Mathematics* 11: 223-232.

Kesten, O. 2010. "School Choice with Consent" *The Quarterly Journal of Economics.* August 1297-1348.

Kesten, O., and M. Kurino. 2013. "Do Outside Options Matter in School Choice? A New Perspective on the Efficiency vs. Strategy-Proofness Trade-off. Unpublished.

Kwon, H., and R. I. Shorrer. 2018. "Justified-Envy-Minimal Efficient Mechanisms for Priority-Based Matching." Unpublished.

Mauleon, A., V. J. Vannetelbosch, and W. Vergote. 2011. "von Neumann–Morgenstern Farsightedly Stable Sets in Two-Sided Matching." *Theoretical Economics.* 6: 499–521.

Pápai, S. 2000. "Strategy-proof Assignment by Hierarchical Exchange." *Econometrica.* 68 (6): 1403- 1433.

Reny, P. J. 2020. "A Simple Sufficient Condition for a Unique and Student-Efficient Stable Matching in the College Admissions Problem." *Economic Theory Bulletin.* 9: 7-9.

Reny, P. J. 2021. "Efficient Matching in the School Choice Problem." Unpublished.

Roth. A. E., and M. A. Oliveira Sotomayor. 1990. *Two-Sided Matching: A Study in Game Theoretic Modeling and Analysis.* New York: Cambridge University Press.

Shapley, L. and H. Scarf. 1974. "On Cores and Divisibility." *Journal of Mathematical Economics.* 1: 23-37.

Tang, Q., and J. Yu. 2014. "A New Perspective on Kesten's School Choice with Consent Idea." *Journal of Economic Theory.* 154: 543–561.

Tang, Q., and Zhang, Y. 2021. "Weak Stability and Pareto Efficiency in School Choice." *Economic Theory.* 71: 533–552. https://doi.org/10.1007/s00199-020-01255-3

Troyan, P., D. Delacrétaz, and A. Kloosterman. 2020. "Essentially Stable Matchings." *Games and Economic Behavior.* 120: 370-390.

Von Neumann, J,. and O. Morgenstern. 1944. *Theory of Games and Economic Behavior.* Princeton: Princeton University Press.