



Munich Personal RePEc Archive

The structure of priority in the school choice problem

Duddy, Conal

National University of Ireland Galway

24 May 2017

Online at <https://mpra.ub.uni-muenchen.de/81204/>

MPRA Paper No. 81204, posted 08 Sep 2017 06:13 UTC

THE STRUCTURE OF PRIORITY IN THE SCHOOL CHOICE PROBLEM*

Conal Duddy[†]

7th September 2017

Abstract

In a school choice problem each school has a priority ordering over the set of students. These priority orderings depend on criteria such as whether a student lives within walking distance or has a sibling already at the school. I argue that by including just the priority orderings in the problem, and not the criteria themselves, we lose important information. More particularly, the priority orderings fail to capture important aspects of the information from which they are derived when a student may satisfy a given criterion across multiple schools. This loss of information results in mechanisms that discriminate between students in ways that are not easy to justify. I propose an extended formulation of the school choice problem wherein a “priority matrix”, indicating which criteria are satisfied by each student-school pair, replaces the usual profile of priority orderings.

Keywords. School choice; matching; priority ordering; deferred acceptance.

JEL classification. C78, H40, I28.

*I thank Ashley Piggins, Bill Zwicker and Bettina Klaus for many helpful comments.

[†]J.E. Cairnes School of Business and Economics and the Whitaker Institute, National University of Ireland Galway, University Road, Galway, Ireland. E-mail: conal.duddy@gmail.com.

1 Introduction

In many school districts in the United States students are assigned to public schools via a matching mechanism. Districts vary in the particular mechanism that they use. Each mechanism is a solution to a matching problem called the *school choice problem*.

This problem was first formulated as a mechanism design problem by Abdulkadiroğlu and Sönmez (2003). Their article and the literature that followed it have led to many school districts adopting mechanisms based on the Deferred Acceptance algorithm and the Top Trading Cycles algorithm. Indeed, this literature is notable for its direct connection to policy making (see Roth, 2008). Certainly, the formulation of the school choice problem as a problem of mechanism design has had a highly significant and very positive impact, greatly improving how student-school matching is conducted in many districts.

That formulation, which is the subject of this paper, is based on the classic *college admissions problem* of Gale and Shapley (1962). It consists of five items: (i) a set of students, (ii) a set of schools, (iii) a list of school capacity numbers, (iv) a profile of student preference orderings over the schools and (v) a profile of school priority orderings over the students. A school's priority ordering is a ranking of students based on criteria such as having a sibling already at the school or living within walking distance of the school.

In this paper I argue that item (v), the profile of school priority orderings, can fail to capture important aspects of the information from which it is derived. In particular, important information is lost when a student satisfies a priority criterion across multiple schools. This loss of information means that matching mechanisms must treat situations that are substantively different from one another as though they were identical. I show how this can result in unfair matches and how it disqualifies mechanisms that are reasonable.

I define the school choice problem more fully in the next section and I use simple examples in section 3 to show that this formulation can suppress crucial information. I elaborate on those examples in section 4. On the basis of that discussion I propose a new formulation of the problem in section 5. Then I show in

sections 6 and 7 how the concept of justified envy and the Deferred Acceptance algorithm can be adapted to this new version of the school choice problem. I discuss in section 8 the difference between my contribution and some recent literature on the school choice problem. Section 9 concludes the paper.

2 The school choice problem

The standard formulation of the school choice problem consists of five items:

1. a set I of students,
2. a set S of schools,
3. a list of natural numbers, each indicating the capacity of a school,
4. a list of strict preference orderings over S , one for each student, and
5. a list of weak priority orderings over I , one for each school.

The total number of available seats across all of the schools must be at least as great as the number of students, with each school having at least one available seat. This list of five items can be found in, for example, Ergin and Sönmez (2006) and Abdulkadiroğlu *et al.* (2017).

A *matching* assigns each student to a school. The number of students assigned to a school must not exceed the capacity of that school, and each student must be assigned to exactly one school. A method of assigning students to schools may involve some randomization. A *mechanism* is a function that associates a probability distribution over the set of possible matchings to each school choice problem. An *expected matching* is a matrix that gives each student's probability of being matched to each school.

It is important to note here that the schools do not set their priority orderings autonomously. This is a key point of difference between the school choice problem and the older college admissions problem. Priority criteria are chosen by the district school board and these criteria induce priority orderings for all of the schools in the district. One consequence of this is that schools are not considered to be strategic

agents in the school choice problem. By contrast, colleges *are* considered to be strategic agents in the college admissions problem. Similarly, in the case of the school choice problem the usual definition of Pareto efficiency is “one-sided” (only the preferences of students matter) whereas for the college admissions problem it is “two-sided”.

3 Motivating examples

To help motivate a change to the standard formulation of the school choice problem let us consider two simple scenarios. In scenario A there are three students i , j and k and three schools s_1 , s_2 and s_3 . Student i and schools s_1 and s_2 are in the Oak Hill neighbourhood. Students j and k and school s_3 are in Elm Hill. There is just one available place at each school.

All three students agree that school s_1 is excellent and that s_3 is a low-performing school. The students share the same preference ordering over the schools; they all rank s_1 first, s_2 second and s_3 third. This means that the preferences of the students, without any other information, do not provide us with any reason to prefer any particular matching over any other one.

However, when we take priority criteria into account we may find cause to discriminate between possible matchings. The district school board has determined that two priority criteria are applicable. We denote them by p and q . Criterion p can be read as “lives within walking distance” and q as “has a sibling already at the school”. These criteria are satisfied by student-school pairs as indicated in the following *priority matrix*.

$$\begin{array}{ccccc}
 & s_1 & s_2 & s_3 & \\
 i & p, q & p & - & \\
 j & - & - & p & \\
 k & - & - & p &
 \end{array} \tag{1}$$

We see that schools s_1 and s_2 are within walking distance of i while s_3 is within walking distance of j and k . We also see that student i has a sibling who already attends s_1 .

How should we match the three students to the three schools in this scenario? It surely is sensible to match i to s_1 . After all, i has a sibling at that school and lives within walking distance. We can then use a fair coin to decide how to match j and k to s_2 and s_3 . Under this approach, the students face the following expected matching. The entries show each student's probability of being assigned to each school.

$$\begin{array}{rcccc}
 & s_1 & s_2 & s_3 & \\
 i & 1 & 0 & 0 & \\
 j & 0 & 1/2 & 1/2 & \\
 k & 0 & 1/2 & 1/2 &
 \end{array} \tag{2}$$

This is in fact the solution that all three of the standard mechanisms generate for this scenario. Under the Boston mechanism, we assign as many students as possible to their first-choice schools. Where a school is over-subscribed we refer to that school's priority ordering to determine which students are accepted. So the Boston mechanism would give the available place at s_1 to i . The other two standard mechanisms, the Gale-Shapley student-optimal stable mechanism and the top trading cycles mechanism, are slightly more complex. For now it suffices to say that they satisfy a principle called "Mutual Best" (see Morrill, 2013), just as the Boston mechanism does. This principle says that if school s is the top choice of student i , and student i is at the top of the priority ordering for school s , then i should be assigned to s (unless the school cannot accommodate all such "mutual best" students). Thus those two mechanisms agree with the Boston mechanism in this case; student i should be assigned to s_1 while j and k should be assigned s_2 and s_3 according to a lottery.

Students j and k , and their parents, may be unhappy that the place at the most desirable school, s_1 , is given to i with a probability of one and not included in the lottery. But this assignment is entirely defensible. It can be defended on the grounds that i has a sibling at s_1 . So the standard mechanisms have delivered a perfectly reasonable outcome here.

Now let us consider scenario B. This is the same as scenario A but with one feature removed. We now suppose that i does not have a sibling at s_1 . That is, the priority criteria are satisfied as follows.

$$\begin{array}{cccc}
& s_1 & s_2 & s_3 \\
i & p & p & - \\
j & - & - & p \\
k & - & - & p
\end{array} \tag{3}$$

This is a significant change. The s_1 and s_2 columns are now identical. This means that schools s_1 and s_2 differ from each other only in their desirability. They do not differ with regard to the criteria set out by the district school board; s_1 is within walking distance of i but so is s_2 , s_1 is beyond walking distance for j and k but so is s_2 and no student has a sibling at either s_1 or s_2 . How should we match students and schools in this scenario?

3.1 Three proposals for scenario B

Consider the following proposal. First, select one of the students j and k by lottery and assign that student to s_3 . It seems appropriate to assign either j or k to s_3 rather than forcing i to travel beyond walking distance. This means that either j or k will travel to Oak Hill for their schooling. The limited capacity of the sole Elm Hill school makes this unavoidable. Thus one Elm Hill student together with i will attend school in Oak Hill. Let us say that k is the student who is matched to s_3 .

We must now match i and j to s_1 and s_2 . Regardless of how we match them, i will be attending a school within walking distance and j will not. So in choosing between the two possible ways to match i and j to s_1 and s_2 the issue of walking distance is not a discriminant. And both students have the same preference over the schools. I propose, then, that we use a fair coin to decide who will be assigned to s_1 and who to s_2 .

In summary, by using a coin to decide who of the Elm Hill students will attend s_3 and then using a coin again to decide who of the remaining students will attend

s_1 , we have the following expected matching.

$$\begin{array}{rcccc}
 & s_1 & s_2 & s_3 & \\
 i & 1/2 & 1/2 & 0 & \\
 j & 1/4 & 1/4 & 1/2 & \\
 k & 1/4 & 1/4 & 1/2 &
 \end{array} \tag{4}$$

This seems like a reasonable solution for scenario B. However, it is not the only approach we could take. An alternative approach would be the following. We begin by assigning the place at s_1 by fair lottery over all three students. This means that every student has a probability of one third of being matched to s_1 . If student i is not given the place at s_1 then i is matched to s_2 . This ensures that i is not forced to travel to Elm Hill. If the place at s_1 is given to i then the place at s_2 is assigned to j or k according to a fair coin.

It is easy to see that j and k each have a probability of one third of being matched to s_1 and a probability of one half of being matched to s_3 . It follows that they each have a probability of one sixth of being matched to s_2 . Hence, under this second proposal we have the following expected matching.

$$\begin{array}{rcccc}
 & s_1 & s_2 & s_3 & \\
 i & 1/3 & 2/3 & 0 & \\
 j & 1/3 & 1/6 & 1/2 & \\
 k & 1/3 & 1/6 & 1/2 &
 \end{array} \tag{5}$$

A third proposal is the following. Let us assign a distinct number to each student by lottery. The students queue up to choose a school in ascending order of their lottery numbers. Under this simple approach there is a risk that i may be left with the place at s_3 . To avoid this, we add a caveat as follows. We reserve a place at either s_1 or s_2 for i , using a fair coin to choose which one. The other students may not take this reserved place unless i has already been matched to another school. For example, suppose that we reserve the seat at s_2 for i and also that i receives the lowest lottery number. Then i will choose to take the seat at s_1 and the reserved seat at s_2 is released so that the next student in the queue may take it. To take another

example, suppose again that the seat at s_2 is reserved for i but this time student i receives the highest lottery number and j receives the lowest number. Then, first in the queue, j takes the place at s_1 . Though second in the queue, k must take the place at s_3 because s_2 is reserved for i . Finally, i takes the reserved place at s_2 .

Under this approach student i has a probability of two thirds of being matched to s_1 . This is because i is matched to s_1 if i comes first in the lottery (a probability of one third) or if the place at s_1 is reserved for i (a probability of one half). When we subtract the probability of both events occurring (one sixth) we arrive at a probability of two thirds. Thus, under this third proposal the students face the following expected matching.

$$\begin{array}{rcccc}
 & s_1 & s_2 & s_3 & \\
 i & 2/3 & 1/3 & 0 & \\
 j & 1/6 & 1/3 & 1/2 & \\
 k & 1/6 & 1/3 & 1/2 &
 \end{array} \tag{6}$$

I submit that these three proposals are among a number of reasonable solutions that are worth considering in the case of scenario B.

3.2 An unexpected difficulty

Naturally, these proposals would entail treating scenario B differently from scenario A. This seems sensible given that the scenarios are indeed quite different from one another. Yet, when we represent scenarios A and B as school choice problems, that is, using the five items listed in section 2, we find something surprising. We find that the difference between the two scenarios is lost. Both scenarios correspond to exactly the *same* school choice problem. This is because no school's priority ordering changed when we moved from scenario A to B.

The school priority orderings are given by the following three columns. Let us refer to this as a *profile* of priority orderings. Students j and k rank equally in each school's priority ordering.

$$\begin{array}{rcccc}
 & s_1 & s_2 & s_3 & \\
 \hline
 i & i & i & j,k & \\
 j,k & j,k & j,k & i &
 \end{array} \tag{7}$$

Of course, in scenario B we removed i 's sibling from s_1 but this does not change the priority ordering for s_1 . Student i continues to have higher priority for that school.

Since all five of the items that define the school choice problem, as given in section 2, are unchanged across these two scenarios, a matching mechanism receives exactly the same input for both scenarios. An immediate consequence of this is that all mechanisms must generate the same probability distribution for both scenarios. Therefore, if we are to assign the place at s_1 to i in scenario A, on the basis that i has a sibling at that school, then we *must* assign the place at s_1 to i in scenario B too. Indeed, this is what all three of the standard mechanisms do. Similarly, in scenario B whomever of the Elm Hill students travels to Oak Hill for their schooling is automatically assigned to the inferior Oak Hill school, s_2 . This is in spite of the fact that our reason for making s_1 exclusive to i in scenario A is absent in scenario B.

The way in which the school choice problem is defined imposes this cross-scenario restriction. It severely limits the set of mechanisms that we can consider and makes it impossible to treat students fairly in both scenarios. Expected matching (2) is fair in scenario A but quite unfair in scenario B. Expected matchings (4), (5) and (6) are arguably fairer than (2) in scenario B but they would be inappropriate in scenario A. Yet we must choose a single expected matching to fit *both* scenarios. The standard mechanisms generate expected matching (2) for both scenarios. In doing this they impose a degree of unequal access to better schools that appears to be needlessly high in the case of scenario B.

Of course, part of the inequality of expected matching (2) in that scenario can readily be justified. In particular, the fact that i will definitely not be assigned to the least desirable school s_3 can be justified on the grounds that there is sufficient capacity in local Oak Hill schools for i . However, another part of the inequality, the exclusion of j and k from s_1 , arises because crucial information is missing from the school choice problem itself.

4 The structure of priority

The scenarios that we have considered motivate us to reconsider the definition of the school choice problem. The first step in developing a new definition of the problem

is to identify precisely the particular issue that is revealed by those examples and to discuss that issue in more general terms. That is what I do in this section.

In the preceding section we encountered two kinds of priority structure: the priority matrix, as in (1), and the profile of priority orderings, as in (7). A priority matrix indicates which priority criteria are satisfied by each student-school pair. A profile of priority orderings consists of a ranking of students for each school, wherein students are ranked according to the strength of their respective claims to priority for each school. The earlier scenarios A and B are essentially about the relationship between these two structures.

These two kinds of priority structure are, of course, very closely related. If we have a priority matrix and we know the relative importance of the priority criteria then we can determine the relative strength of each student's claim to a place at each school. That is, we can derive a profile of priority orderings. To see why the relative importance of criteria matters, consider the case that one student has a sibling at a particular school but does not live within walking distance while another student does live within walking distance but does not have a sibling at the school. To construct a profile of priority orderings in this case we must know which of these criteria is the more important (or that they are of equal importance). Indeed, since each student-school pair may satisfy multiple criteria, we need a ranking not just of individual criteria but of combinations of criteria.

To make this more formal, let P be a set of priority criteria and let 2^P denote the power set of P . Let f be a mapping from $I \times S$ to 2^P . For each student i in I and each school s in S , $f(i, s)$ is the set of priority criteria that are satisfied by the pair (i, s) . Thus, the mapping f describes the priority matrix (since the ordering of the columns and rows is not important). As we have noted, f by itself is not sufficient to induce a profile of priority orderings. Let \succeq be a weak ordering over 2^P that ranks criteria, and combinations of criteria, by importance. Let us write $iR_s j$ to mean that student i is ranked equal to or above student j in the priority ordering for school s . Then $iR_s j$ if and only if $f(i, s) \succeq f(j, s)$. In this way, f and \succeq together induce a profile of priority orderings.

As we have noted, it is the profile of priority orderings, and not the pair (f, \succeq) , that is primitive in the formulation of the school choice problem. This formula-

tion implies, then, that the profile of priority orderings captures all of the relevant information contained in the pair (f, \succeq) from which it is derived. However, the examples that we considered in section 3 show that this is not the case.

To help clarify this point it is useful to make the argument more specific. If $f(i, s)$ contains p then let us say that student i has a *claim* to a place at school s on the basis of p . The strength of this claim depends on the importance of the criterion that underpins it, p . A student who satisfies multiple priority criteria for a given school has multiple claims to a place at that school. In the following subsection I make a distinction between *connected* and *unconnected* claims. Then I further divide connected claims into two kinds: *conjunctive* and *disjunctive*. I argue that the profile of priority orderings can be taken as primitive only if the claims they represent are *unconnected* or *conjunctive* in nature. However, I argue that, for matters of school choice, connected claims are *disjunctive* in nature and so this approach is not appropriate.

4.1 The nature of claims

The distinction between these kinds of claims is central to my discussion of the school choice problem. However, this distinction is not special to school choice. Indeed, let us briefly step away from the school choice problem and consider the following set of four very simple economic problems.

In each problem there are two individuals i and j . In two of the problems we are tasked with allocating two items of food to these individuals, and we must give one item to each person. In the other two problems we must allocate a place on a program at a professional school to each of i and j , with just one place available at each school. There are two priority criteria involved in the problems and they are labelled p and q . Let criterion p be, “this person is a vegetarian and this food is suitable for vegetarians”, and let q be, “this person has achieved excellent grades in subjects relevant to this program”.

The following grid of four priority matrices describes the four problems.

	Apple pie	Meatloaf	Apple pie	Turnip
i	p	–	p	p
j	–	–	–	–
	Medical	Law	Medical	Business
i	q	–	q	q
j	–	–	–	–

In the upper-left problem individual i satisfies a priority criterion for one item, the apple pie, and thereby obtains a relatively strong claim to that item. This is an example of an unconnected claim. This means that the criterion that i satisfies for this item is not satisfied by i for any other item. Similarly, in the lower-left problem individual i has an unconnected claim to a place at the medical school. By contrast, in each of the two problems on the right, individual i satisfies a single priority criterion over multiple items. In the case of the upper-right problem, individual i has a connected claim that is based on p and that spans the apple pie and the turnip. This is a connected claim because it is underpinned by the same criterion across *both* items. In the lower-right problem we find again that i has a connected claim spanning multiple items, this time on the basis of q .

We have made a distinction between unconnected and connected claims. To make the next distinction, between connected claims that are conjunctive and those that are disjunctive, let us carefully consider each of the four problems. Let us suppose that in the two upper problems both individuals would most like to have the apple pie, and that in the two lower problems both individuals would prefer to attend the medical school. How then should we match the individuals to the items in each case?

Let us first consider the two problems on the left-hand side of the grid, in which i has an unconnected claim to an item. In the case of the upper-left problem it is sensible to give the apple pie to i even though both individuals would like to have that item. We can point to p as our reason for doing so: we should assign the apple pie to i because i is a vegetarian. In the case of the lower-left problem it is sensible to allocate the place at the medical school to i . This time we may point to q as the reason: i is the more deserving of the place on that program. The two left-hand

problems, then, seem to be just superficially different from one another.

The two right-hand problems, in which i has a connected claim spanning two items, appear to differ just superficially too. But this appearance is deceptive. In the case of the upper-right problem it would be nonsensical to point to p as a reason for allocating the apple pie to i . We may well take the view that i is entitled to receive a food item that is suitable for vegetarians but *both* items satisfy that entitlement. Perhaps the best solution here would be to use a fair coin to decide who receives the apple pie.

By contrast, in the case of the lower-right problem it is perfectly sensible to give the place at the medical school to i on the basis of q . Indeed, the same reason applies as in the lower-left problem: i is the more deserving of the place on that program. Individual i 's meriting of a place on the medicine program is not diminished by the fact that i is deemed to be more deserving of a place at the business school too.

When we compare the two problems on the right of the grid we find that they are very different from one another even though they share the same formal structure. This observation motivates a division of connected claims into the two aforementioned categories: conjunctive and disjunctive.

Priority criterion q , that refers to merit, is an example of one that grants a conjunctive claim. In the case of the lower-right problem, criterion q grants i priority for the medicine program and for the business program, with emphasis here on the conjunction. If, say, we were to allocate the business program to individual i , she could appeal to q , together with her preference ordering of the programs, to argue that she should be given the place on the medicine program instead, even though q applies to both programs for her.

Priority criterion p , that refers to a dietary requirement, is an example of one that grants a disjunctive claim. In the case of the upper-right problem, criterion p is satisfied by individual i in respect of both items. However, i cannot sensibly invoke p to argue that she should be the one to receive the apple pie when the other item is a turnip.

4.2 Loss of information

Suppose again that we seek to match a set of individuals to a set of goods. Let i , j and k be the individuals and a , b and c be the goods. Let us suppose that all three individuals have the same preference ordering; they all prefer a to b and b to c . There are two priority criteria r and t . To make the example more concrete let r be, “this person adheres to religious laws and this good is compatible with those laws”, while t is, “this good is fully wheelchair-friendly and this person uses a wheelchair”. Now suppose that we have the following profile of priority orderings.

$$\begin{array}{ccc} a & b & c \\ \hline i & i & i, j, k \\ j, k & j, k & \end{array} \quad (8)$$

Here we see that all three individuals i , j and k have equally strong claims to c , and that i has highest priority for a and b .

In this situation we find that important information is suppressed by the profile of priority orderings. For example, it could be the case that i adheres to religious laws and that both a and b are compatible with those laws while c is not. In other words, (8) is consistent with the following priority matrix.

$$\begin{array}{ccc} a & b & c \\ i & r & r - \\ j & - & - - \\ k & - & - - \end{array} \quad (9)$$

However, (8) is also consistent with a very different situation. Consider the following priority matrix, for example, in which good a is the only one that is compatible with i 's religious beliefs.

$$\begin{array}{ccc} a & b & c \\ i & r & t - \\ j & - & - - \\ k & - & - - \end{array} \quad (10)$$

We cannot determine, then, whether i has a disjunctive claim spanning a and b , as in (9), or unconnected claims to each of a and b , as in (10). Yet it is perfectly reasonable to wish to distinguish between those two possibilities when assigning the goods to the individuals. After all, in the case of (10) we may point to r , together with the fact that i would prefer to have a rather than b , as a reason to assign a to i with a probability of one whereas in the case of (9) we cannot. So we see in this example how a profile of priority orderings can fail to convey important information.

On the other hand, when connected claims are *conjunctive* in nature a profile of priority orderings arguably *does* convey all of the important information. To explain this point, it may be helpful to draw an analogy to the rule of “conjunction elimination” in propositional logic. According to this rule we may infer P and Q from $P \wedge Q$. In a similar way, a conjunctive claim that spans, say, goods a and b , is no different to having an unconnected claim to each of a and b . If we “eliminate” conjunction in this way, then we can simply consider each column of a priority matrix separately. In this case we may find that a profile of priority orderings captures all of the important information about the claims of the individuals.

To clarify this point, let us consider again profile (8) but this time suppose that r and t are criteria that, when satisfied by one person over multiple goods, grant a conjunctive rather than a disjunctive claim to those goods. To make the example more concrete, suppose that r means, “deserving in virtue of good behaviour”, and t means, “deserving in virtue of hard work”. As before, we do not know the underlying priority matrix. It *could* be (9) or (10), as they are both consistent with (8), and there are other possibilities too. But, in this case does it matter which is the underlying priority matrix? When we compare (9) and (10), for example, there is no obvious reason to prefer one matching for (9) and then some other matching for (10). Arguably, since we can eliminate conjunction and effectively apply the criteria to each good separately, all that matters is the relative strength of the claims *within each column* of the priority matrix. Indeed, it follows from this argument that the profile of priority orderings is a suitable priority structure in this case since that is exactly the information the profile conveys.

In summary then, a profile of priority orderings captures the important inform-

ation from its underlying priority matrix when we may regard the claims of individuals as applying to each item separately. In the case of a conjunctive claim that spans multiple items we may “eliminate” the conjunction, and consider each column of a priority matrix separately. But we cannot eliminate disjunction in this way. This is why a profile of priority orderings suppresses important information when individuals have disjunctive claims spanning multiple items.

4.3 Relevance to school choice

Since the school choice problem includes just priority orderings, the standard school choice mechanisms effectively treat all priority criteria as though they grant conjunctive claims to priority. In the context of student-school matching, we can certainly conceive of criteria of this kind. Criteria that are based on deservingness and merit typically belong in this category. A student may be deemed more deserving of a place at a particular school on the basis of grades or good behaviour. And, though it would be very controversial, one could also argue that deservingness derives from financial contributions. In the United States public high schools are funded partly by property taxes. Perhaps a student could be deemed especially deserving of a place at a public school on the basis that his or her parents have paid a large amount in property tax.

However, public school choice programs in the United States do not involve criteria of this kind. Following its Roundtable on Public School Choice, the Office of Educational Research and Improvement (1992) noted that “on principle, all members of the Roundtable do not favor student-based admissions criteria”. Examples of student-based criteria are those based on grades, behaviour and criminal records. Consistent with the views of the Roundtable, the priority criteria that are applied in public school choice programs are about practical issues, such as the cost of transport, and not deservingness. Typical criteria refer to walking distance or the availability of bilingual teaching programs. These are pragmatic criteria that would seem to confer disjunctive claims to school places.

This also explains why my argument applies specifically to the school choice problem and not to the college admissions problem. In the case of college admis-

sions, each college ranks students according to some combination of test scores, grades, interviews and so on. Thus, students do not have disjunctive claims spanning multiple colleges, and so it is sufficient to have a priority structure that consists of a ranking of students for each college.

5 A new definition

In this section I propose a new definition of the school choice problem. We have seen that a profile of priority orderings, item 5 in the original definition, can fail to capture important information from the priority matrix when students have disjunctive claims spanning multiple schools. I propose that we include the priority matrix and the ranking of priority criteria, themselves, in the definition of the problem, replacing the profile of priority orderings that is derivable from them. In other words, I propose that we make f and \succeq primitive notions in the problem. Crucially, this allows us to see where a school's priority ordering comes from. We can see where there is a connected claim and we can identify the criterion that underpins it.

Accordingly, the seven items that constitute an *extended school choice problem* are:

1. a set I of students,
2. a set S of schools,
3. a list of natural numbers, each indicating the capacity of a school,
4. a list of strict preference orderings over S , one for each student,
5. a set P of priority criteria,
6. a mapping f from $I \times S$ to 2^P , the power set of P , and
7. a weak ordering \preceq over 2^P such that $A \subseteq B \rightarrow A \preceq B$.

Returning to the earlier examples, let us observe that scenarios A and B become distinct from one another when they are represented as extended school choice problems. Let p denote one priority criterion, “lives within walking distance”, and let q

denote the other criterion, “has a sibling at this school”. Then in the case of scenario A $f(i, s_1)$ is $\{p, q\}$ whereas in scenario B $f(i, s_1)$ is $\{p\}$. So item 6 is the formal tool that accounts for the difference in this case.

One consequence of this extension of the school choice problem is that it expands the set of mechanisms that we may consider. Of course, this expansion is worthwhile only if it encompasses mechanisms that are reasonable. To demonstrate that this is the case, I provide in section 7 an example of how a standard mechanism can be adapted to the extended school choice problem in such a way that it becomes sensitive to information about priority factors. In this case the mechanism that I adapt is the Gale-Shapley student-optimal stable mechanism, which may be implemented by the celebrated Deferred Acceptance algorithm.

A second consequence is that we may consider normative principles that are precluded by the standard definition of the problem. Indeed, I suspect that important normative aspects of student-school matching have been overlooked because of the way in which the problem has been defined. In the next section I discuss a normative concept that arises naturally once we extend the school choice problem.

Note also that the extended definition opens up a rich set of possible mechanisms and normative principles without removing anything that was possible under the standard definition. Since we can construct a priority ordering for each school we can derive a standard school choice problem from an extended problem. This means that any of the standard mechanisms of student-school matching can be applied to an extended problem.

6 Justified envy

Justified envy, also called “priority violation” by Kesten (2010), is a central concept in much of the normative analysis in the literature on the school choice problem. It is defined as follows.

Justified envy. There is a student i and a school s such that (i) student i prefers school s to the one that i has been assigned to and (ii) a student has been assigned to s despite having lower priority for s than i has.

Take either of the earlier scenarios A or B, for example, and suppose that student j or k is assigned to s_1 while i is assigned to s_2 . This is an instance of justified envy because i prefers s_1 to s_2 and has higher priority for s_1 than the student who has been assigned to that school. However, I argue that this envy is rightly justified in the case of scenario A only.

In scenario A, if student i is turned away from s_1 then she misses out on attending the same school as her sibling. She and her family can justifiably feel that they have been hard done by, and their complaint is clearly relevant to public policy. In the case of scenario B, on the other hand, i 's claim to priority for s_1 is based solely on priority criterion p , that she lives within walking distance of the school. But, I argue, this priority criterion grants her a disjunctive claim that encompasses both s_1 and s_2 . Since she has been assigned to s_2 she cannot sensibly appeal to p to justify her envy. Indeed, if we swap the assignments of the two students we achieve no goal of public policy; the number of students attending a school within walking distance is unchanged.

I propose therefore an alternative concept that I call *strongly justified envy*. To help define this concept I introduce some notation here. Given a student i and a set of priority criteria C , let $S(C, i)$ be the set of all schools s such that C is a subset of $f(i, s)$. In other words, $S(C, i)$ is the set of all schools that meet all of the criteria in C with respect to student i .

Strongly justified envy. There is a student i and a set of priority criteria C such that (i) student i prefers every school in $S(C, i)$ to the one that i has been assigned to and (ii) a student has been assigned to a school in $S(C, i)$ despite having lower priority for that school than i has.

The intuition behind strongly justified envy is the following. A student i cannot sensibly appeal to C to justify his or her envy if it is the case either that i has in fact been assigned to a school in $S(C, i)$ or if there is a school in $S(C, i)$ to which i would *not* like to be moved.

The concepts of justified envy and strongly justified envy coincide in scenario A and diverge in scenario B. In scenario A, if i is assigned to either s_2 or s_3 then i regards the student assigned to s_1 with strongly justified envy. This is because

$S(\{p, q\}, i)$ contains just s_1 in that scenario. In scenario B, on the other hand, the matching exhibits strongly justified envy only in the case that i is assigned to s_3 . In that scenario, $S(\{p, q\}, i)$ is empty and $S(\{p\}, i)$ is $\{s_1, s_2\}$ and so there is no case of strongly justified envy if i is assigned to s_2 .

The concept of strongly justified envy is relevant to the argument I make in this paper because it is an example of a normative concept that is surely relevant to school choice but that is excluded by the standard definition of the school choice problem. In the next section I provide an example of a *mechanism* that, like the concept of strongly justified envy, is sensitive to the additional information contained in an extended school choice problem.

7 Deferred Acceptance

The famous Deferred Acceptance algorithm originates in a seminal article by Gale and Shapley (1962). The algorithm is first applied to the school choice problem by Abdulkadiroğlu and Sönmez (2003). The resulting mechanism is called the Gale-Shapley student-optimal stable mechanism. A key property of the mechanism is that it always generates a matching that is free from justified envy (and for this reason is said to be stable).

The standard algorithm, applied to school choice, works as follows. First, we assign a distinct number to each student in I by lottery. These numbers are used to break ties in the school's priority ordering so that the ordering becomes strict.¹ In each round of the procedure students “propose” to schools. In the first round each student proposes to his or her top-choice school. Each school considers all of its proposers and decides who to hold and who to reject. In subsequent rounds, each rejected student proposes to his or her next most preferred school. If a student is “held” by a school at the end of a round then he or she proposes again to the same school in the following round. This is because the schools release all students at the beginning of each round. The procedure continues until a round ends with every student being held by a school.

¹For analysis of the surprisingly complex issues around tie-breaking in school priority orderings see Erdil and Ergin (2008), Abdulkadiroğlu, Che and Yasuda (2015) and Özek (2016).

In each round of the procedure each school has a set of proposers. This set includes those students who were held by the school at the end of the previous round, and new proposers who were rejected by other schools in the previous round. Let q be the capacity of a particular school. If the number of proposers is less than or equal to q then the school holds all of its proposers in that round. If the number of proposers is greater than q then the school holds the top q proposers according to its priority ordering.

In this description, a school considers its proposers one-by-one and in ascending ordering of their lottery numbers. As a school receives each proposal in turn, it may possibly decide initially to hold a student and then later in the same round decide to reverse this decision in light of a better proposal and reject that student. The school always tentatively holds the first q proposing students, where q is the capacity of the school. Then, as it considers each subsequent proposer, it may choose to hold that student and thereby *displace* a student whom it had been holding. Thus, after the first q proposals in a round, *every* subsequent decision to hold a student means rejecting a student whom the school had previously chosen to hold. Whenever a student is to be displaced, it is always the lowest-priority student held by the school who is rejected. Where multiple students are tied for lowest priority the one who has the highest (unluckiest) lottery number is rejected.

To decide whether to hold or reject each student the school applies the decision tree shown in figure 1. Recall that the school releases all students at the beginning of each round, thereby making all seats available. Hence the answer to the first question, “Do we have a free seat?”, will be “yes” for the first q students and “no” for all subsequent students. However, these first q students may be displaced before the end of the round as the school considers successive proposing students.

To motivate a change in the decision tree shown in figure 1 let us return to our earlier scenario B and apply the Deferred Acceptance algorithm. All three students i , j and k have the same preference ordering. They all rank s_1 first, s_2 second and s_3 third. Thus in the first round of the procedure all three students propose to s_1 . Let us suppose that j has the lowest lottery number and i has the second-lowest lottery number. Then s_1 considers j 's proposal first and tentatively holds j , simply because this is the first proposal. Next s_1 considers i 's proposal and because i has

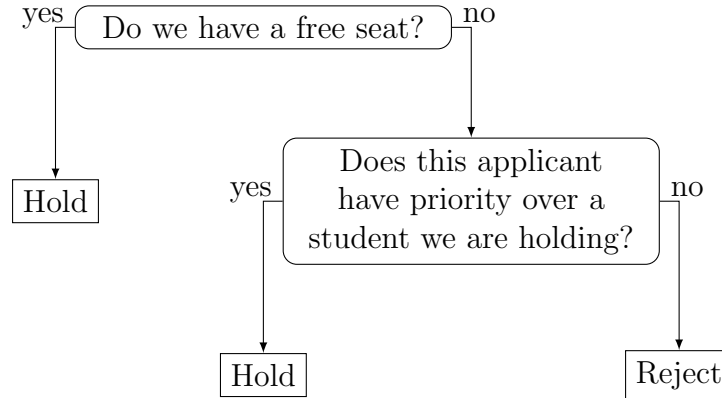


Figure 1: Standard decision tree

higher priority than j has, s_1 decides to hold i , which means that j is displaced. To complete the round, s_1 considers k 's proposal and rejects it. In the second round, j and k move on to s_2 while i proposes again to s_1 . In the end, i is matched to s_1 , j to s_2 and k to s_3 .

Note that during the first round i automatically displaced j on the basis that i lives within walking distance of s_1 and j does not. However, we know that s_2 is also within walking distance of i . Why then must j be immediately displaced from s_1 once i proposes? I argue that in this situation there should be a non-zero probability that s_1 continues to hold j and rejects i on the basis that s_2 can accommodate i just as well as it can accommodate j .

The following definitions are helpful to me in describing a modification of this mechanism.

High priority. A student has *high priority* for a school if he or she has a top- q position in the school's priority ordering after ties are broken according to the student lottery numbers (favouring lower numbers), where q is the capacity of the school.

Exchangeable. A school s' is *exchangeable* for school s with respect to student i if $f(i, s)$ is a subset of $f(i, s')$. That is, (i, s') satisfies every priority criterion that (i, s) satisfies (and perhaps more besides).

Selection number. Separately from the lottery numbers assigned to students to de-

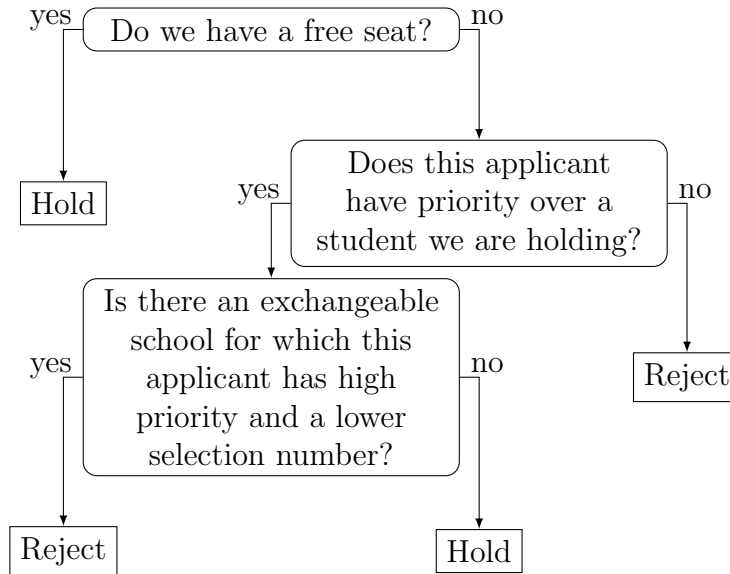


Figure 2: Alternative decision tree

termine the order in which they propose to schools, we randomly assign *selection numbers* to student-school pairs. For each student i we assign a distinct number in $\{1, \dots, |S|\}$ to each school in S by lottery. So a given student has a distinct selection number for each school, but two students may share the same selection number for a given school.

I propose that we replace the decision tree shown in figure 1 with the one shown in figure 2.

The new decision node, the difference between the two mechanisms, is relevant when a school receives a proposal from a student who (i) is not one of the first q to propose to the school in that round and (ii) has higher priority than a student who is tentatively being held by the school. Suppose that a school s receives such a proposal from a student i . Let E be the set of all schools that are exchangeable for s with respect to i . And let H be the (possibly empty) set of all schools at which i has high priority. It is the intersection of these two sets, $E \cap H$, that matters at this new decision node. It contains those schools that meet all the same criteria with i as s does and from which i can never be displaced by a higher priority student. If i has high priority for s then s itself is in $E \cap H$.

If $E \cap H$ is empty or contains only s then s will hold i . Otherwise s decides whether to hold i based on the selection numbers than have been issued to i for the schools in $E \cap H$ and for s . If i has a lower selection number for a school in $E \cap H$ than he or she has for s then s rejects i . If not then s holds i . This means that the decision of a school in regard to a proposing student does not depend on the position of that school in the student's preference ordering, which is a key characteristic of the original algorithm.

To help clarify this second decision tree, let us consider again the case of scenario B in which student j has lottery number one and i has lottery number two. As before, all three students propose to s_1 in the first round. Then s_1 tentatively holds j , the first proposer. Next s_1 considers i 's proposal. Here we must identify the schools s such that (i) school s is exchangeable for s_1 with respect to i and (ii) student i has high priority for s_1 . Only s_1 and s_2 meets these conditions. Thus whether s_1 holds i (displacing j) depends on the selection numbers that i has for schools s_1 and s_2 . If i has a lower selection for s_1 then s_1 holds i , otherwise s_1 rejects i .

This means that before the student lottery numbers and selection numbers are drawn, student i has a probability of two thirds of being matched to s_1 . This is because i is matched to s_1 if i has the lowest lottery number (probability of one third) or if i 's selection number for s_1 is greater than i 's selection number for s_2 (probability of one half) and those two events are independent. Indeed, the expected matching that the students face in scenario B under this mechanism is (6).

Since the probability of i being assigned to s_3 is zero we can be certain that there will be no strongly justified envy in the outcome. In fact, this mechanism generates a matching free from strongly justified envy for *all* extended school choice problems. To see that this is true, when a school rejects a student l it is implied that either (i) the school has already reached capacity, holding students who are equal to or higher than l in its priority ordering, or (ii) there is exchangeable school that is certain to hold l should he or she propose to it. Thus, l can fail to displace a lower-priority student at a school but this can happen only if l is guaranteed a place at an exchangeable school. Hence, strongly justified envy cannot arise.

Let us briefly consider the case of scenario A in which, as before, students i , j and k receive lottery numbers two, one and three respectively. All three students

propose to s_1 in the first round. Then s_1 initially holds j only to reject j when i proposes. Student j is immediately displaced here because school s_2 is not exchangeable for s_1 in scenario A. In the end, i is matched to s_1 , j is matched to s_2 and k is matched to s_3 . We can see that, before the lottery number are drawn, the students face expected matching (2).

This mechanism, then, yields a different expected matching over the two scenarios A and B. This is the crucial difference between this mechanism and the standard ones and it is the point of my presenting this mechanism here. I submit that this is not an unreasonable mechanism. Yet it is excluded by the standard definition of the school choice problem. This concludes my argument for an expansion of the set of school choice mechanisms. In the next section I discuss the relationship of this paper to some recent literature.

8 Literature on diversity and slot-specific priorities

Socio-economic and racial diversity in schools is an important matter of public policy in the United States and many other countries. Much of the recent literature on the school choice problem addresses this issue. In order to facilitate mechanisms that are sensitive to diversity-related concerns, a number of variations of the standard school choice problem have been defined.

Upon a cursory inspection, this paper might seem to be about diversity in schools too. A reader may wonder whether the issue that I have raised here has not already been dealt with in some way in the extensive literature on that topic. However, this paper is not about diversity in schools. The issue that I raise in this paper is fundamentally different from the issues analysed in the existing literature. In this final section I seek to clarify this point.

An exercise that is helpful in separating this paper from the existing literature is to focus on the case that each school can accommodate just one student. The school choice problem is normally a many-to-one matching problem but in this special case it becomes a one-to-one matching problem. Though not realistic, this exercise will help us to understand the fundamental difference between the topic of this paper and the issues that are addressed in the existing literature. I begin by establishing

that, unsurprisingly, the various models of the school choice problem that address the issue of diversity in schools all collapse into the standard model in this case, because diversity is simply not an issue when each school has one student.

Let us turn first to an important variation of the school choice problem called the *controlled school choice problem*. In the controlled problem, school enrolments are subject to exogenously imposed constraints that maintain diversity in schools. These constraints usually take the form of lower or upper limits for students in particular ethnic, racial or socio-economic groups. For analysis of controlled school choice problems see, for example, Kojima (2012), Hafalir, Yenmez and Yildirim (2013) and Ehlers *et al.* (2014). When the number of students at each school is exactly one then quotas and other limits intended to ensure diversity clearly do not apply.

Echenique and Yenmez (2015) introduce a model in which schools regard some student types as being complementary to others, reflecting a preference for diversity in their classrooms. A model of this kind is also studied by Bó (2016). Each school has a choice function that is defined over *sets* of students instead of just having a ranking of individual students. Given a menu of possible sets of students, the choice function identifies the student body that the school would most like to have. This choice may reflect a preference for diversity across different types of student. Of course, in the special (and unrealistic) case that we are considering in this section, wherein the capacity of each school is just one, a ranking of possible student populations is no different from a ranking of individual students.

In another departure from the standard school choice problem, Kominers and Sönmez (2013) consider slot-specific school priorities. Boston is an example of a city that implements slot-specific priorities. In Boston, a “walk zone” priority factor is applicable to just half of the seats (or slots) at each school. Previously, economists had modelled this case by splitting each school into two schools, one sensitive to the walk zone and the other not. Kominers and Sönmez argue convincingly that this approach is unsuitable and they consider a school choice problem with slot-specific priorities. Their work builds on the contribution of Sönmez and Switzer (2013). In the special case that each school has just one slot, however, the idea of slot-specific priorities is not relevant.

Yet, the extended school choice problem does not collapse into the standard version of the problem when each school has one student. Indeed, whether schools have one student or multiple students is not relevant to the topic of this paper. This helps us to see that the contribution of this paper is independent of the existing literature on diversity in schools. By discussing the special case in which there is one student for each school I do not mean to imply that the difference exists *only* in that case. Rather, focusing on the case of one-to-one matching brings the difference into sharp relief. The two topics are quite distinct from one another irrespective of the size of each school's capacity. This explains why the existing literature, though extremely rigorous on the topic of diversity, does not cover the same ground as this paper.

This paper is about an alternative formulation of the school choice problem in which normatively significant information is not lost. By retaining and using this information we may find that we can construct mechanisms that are potentially fairer than the existing mechanisms. Of course, it may be that a greater degree of diversity emerges as a by-product of mechanisms that treat students more fairly. Nevertheless, fairness is a different matter from diversity and it is fairness that we are concerned with here.

9 Conclusion

I have argued that the canonical definition of the school choice problem excludes some methods of student-school matching in a way that seems arbitrary. I proposed an extended definition in section 5. Items 5–7 in that definition are new and they replace the list of school priority orderings in the original definition.

Items 5–7 have heretofore been “behind the scenes” in this literature. They have always implicitly been the items from which school priority orderings are derived. But they do not feature in the standard definition of the problem because it is tacitly accepted in the literature that school priority orderings capture all of the relevant information contained in those items. If this view is correct then it is convenient to simply treat school priority orderings as primitive objects in the matching problem and to discard those antecedent items.

I compared two simple scenarios to argue that, on the contrary, this approach results in the loss of important information. We saw that this loss of information limits the set of possible solutions to the extent that those two very different scenarios must be treated as though they were identical. It is for this reason that I propose the extended school choice problem in which items 5–7 are restored.

This alternative definition of the problem expands the set of mechanisms that we may consider. My view is that this produces the natural solution space for the school choice problem. And let us note that one need not desire to design new mechanisms in order to find this expansion to be worthwhile. For example, existing impossibility/uniqueness theorems that are relevant to school choice may become more conclusive or may be undermined in interesting ways when they are applied to a larger set of mechanisms.

To provide an example of how a mechanism can be made sensitive to this additional information I proposed an adaptation of the Deferred Acceptance algorithm. Crucially, it would be impossible to define this new mechanism under the standard formulation of the school choice problem.

I have also sought to clarify the difference between the issue I address in this paper and the issue of diversity in student-school matching that is the focus of much of the recent literature. The special case in which each school can accommodate just one student helps us to see that this paper is quite separate from that literature both in its motivation and in its proposals.

References

- Abdulkadiroglu, A., Y.-K. Che, P. A. Pathak, A. E. Roth and O. Tercieux. 2017. Minimizing Justified Envy in School Choice: The Design of New Orleans' One-app. NBER Working Paper Series 23265.
- Abdulkadiroglu, A., Y.-K. Che and Y. Yasuda. 2015. Expanding "choice" in school choice. *American Economic Journal: Microeconomics* 7(1): 1–42.
- Abdulkadiroglu, A. and T. Sönmez. 2003. School choice: A mechanism design approach. *American Economic Review* 93(3): 729–747.
- Bó, I. 2016. Fair implementation of diversity in school choice. *Games and Eco-*

- conomic Behavior* 97: 54–63.
- Echenique, F. and M. B. Yenmez. 2015. How to control controlled school choice. *American Economic Review* 105(8): 2679–2694.
- Ehlers, L., I. E. Hafalir, M. B. Yenmez and M. A. Yildirim. 2014. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory* 153: 648–683.
- 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.
- Ergin, H. and T. Sönmez. 2006. Games of school choice under the Boston mechanism. *Journal Public Economic* 90(1): 215–237.
- Gale, D. and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly* 69(1): 9–15.
- Hafalir, I.E., M. B. Yenmez and M. A. Yildirim. 2013. Effective affirmative action in school choice. *Theoretical Economics* 8(2): 325–363.
- Kesten, O. 2010. School choice with consent. *Quarterly Journal of Economics* 125(3): 1297–1348.
- Kesten, O. and M. U. Ünver. 2015. A theory of school-choice lotteries. *Theoretical Economics* 10(2): 543–595.
- Kojima, F. 2012. School choice: Impossibilities for affirmative action. *Games and Economic Behavior* 75(2): 685–693.
- Kominers, S.D. and T. Sönmez. 2016. Matching with slot-specific priorities: Theory. *Theoretical Economics* 11(2): 683–710.
- Morrill, T. 2013. An alternative characterization of top trading cycles. *Economic Theory* 54(1): 181–197.
- Office of Educational Research and Improvement. 1992. *Getting started: How choice can renew your public schools*. Washington, DC: US Government Printing Office.
- Özek, U. 2016. Fairness versus efficiency in public school assignments. *Journal of Public Economic Theory* 19(1): 234–243.
- Roth, A.E. 2008. What have we learned from market design? *Economic Journal* 118(527): 285–310.
- Sönmez, T. and T. B. Switzer. 2013. Matching with (branch-of-choice) contracts at

the United States Military Academy. *Econometrica* 81(2): 451–488.