

ARTICLE

The structure of priority in the school choice problem

Conal Duddy 

J.E. Cairnes School of Business and Economics and the Whitaker Institute, National University of Ireland Galway, Ireland

Email: conal.duddy@gmail.com

(Received 24 May 2017; revised 7 July 2018; accepted 15 July 2018; first published online 24 January 2019)

Abstract

In a school choice problem, each school has a priority ordering over the set of students. These orderings depend on criteria such as whether a student lives within walking distance or has a sibling at the school. A priority ordering provides a ranking of students but nothing more. I argue that this information is sufficient when priority is based on merit but not when priority is based on criteria such as walking distance. I propose an extended formulation of the problem wherein a ‘priority matrix’, indicating which criteria are satisfied by each student-school pair, replaces the usual priority orderings.

Keywords: School choice; matching; priority ordering; justified envy

1. Introduction

In many school districts in the USA 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 problem of mechanism design 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.

The standard formulation of the school choice problem comprises 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. This five-item formulation is based on the classic *college admissions problem* of Gale and Shapley (1962).

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 Section 6 how the concept of justified envy can be adapted to this new version of the school choice problem. I discuss in Section 7 the difference between my contribution and some recent literature on the school choice problem. Section 8 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 $q = (q_s)_{s \in S}$ of natural numbers, each indicating the capacity of a school,
4. a list $P = (P_i)_{i \in I}$ of strict preference orderings over S , one for each student, and
5. a list $\succeq = (\succeq_s)_{s \in S}$ of weak priority orderings over I , one for each school.

This list of five items can be found in, for example, Ergin and Sönmez (2006) and Abdulkadiroğlu *et al.* (2017). 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. A matching μ is a function from I to S , so that $\mu(i)$ is the school that student i is assigned to. Every student is assigned to a school and the number of students assigned to a school s must be no greater than q_s .

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’.

2.1. Deterministic and probabilistic matching

Typically, the priority criteria applied within a school district do not produce *strict* priority orderings for each school. Rather, most priority orderings will have large

groups of students who are ‘tied’ in the ranking. A simple way to break these ties is to use random lottery numbers. Matching procedures are often composed of two phases as follows. In the first phase, lottery numbers are assigned to students and these are used to break ties in school priority orderings and thereby make those orderings strict. In the second phase, a deterministic algorithm is applied to generate a matching.

The drawing of lottery numbers may be regarded as part of the matching mechanism or as an event that is exogenous to the mechanism. Both of these approaches can be found in the literature. When the lottery is regarded as exogenous then the matching mechanism is purely deterministic and may be represented formally by a function that generates a single matching as its output. Such a function is called a *deterministic mechanism*. On the other hand, when randomization is regarded as part of the matching procedure, a mechanism can instead be represented by a function that generates a probability distribution over possible matchings. This kind of function is called a *random mechanism*.

By combining a deterministic mechanism with a lottery phase we induce a probability distribution over matchings. In this way we can associate a random mechanism to any deterministic mechanism.¹ The mechanisms that we discuss in this paper will be random mechanisms that generate probability distributions over possible matchings. An *expected matching* is a matrix that gives each student’s probability of being matched to each school.

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 c_1 and c_2 . Criterion c_1 can be read as, ‘lives within walking distance’, and c_2 as, ‘has a sibling already at the school’. These criteria are satisfied by student-school pairs as indicated in the following *priority matrix*.

¹There are sophisticated random mechanisms in the literature, such as the Fractional Deferred-Acceptance mechanism (Kesten and Ünver 2015), that are not constructed in this way, i.e. in which there is not simply a lottery followed by a deterministic algorithm. For analysis of the surprisingly complex issues around tie-breaking in school priority orderings see Erdil and Ergin (2008), Abdulkadiroğlu *et al.* (2015) and Özek (2016).

$$\begin{array}{cccc}
 & s_1 & s_2 & s_3 \\
 i & c_1, c_2 & c_1 & - \\
 j & - & - & c_1 \\
 k & - & - & c_1
 \end{array} \tag{1}$$

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.

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

In fact, all three of the standard mechanisms agree that student i should be assigned to school s_1 . 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 student-optimal stable mechanism and the top trading cycles mechanism, are slightly more complex. Here 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 students). Thus, those two mechanisms agree with the Boston mechanism in this case; student i should be assigned to s_1 . When we combine any of these deterministic mechanisms with a lottery, as discussed earlier in Section 2.1, we obtain expected matching (2).

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. But this assignment is entirely defensible. It can be defended on the grounds that i has a sibling at s_1 .

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.

$$\begin{array}{cccc}
 & s_1 & s_2 & s_3 \\
 i & c_1 & c_1 & - \\
 j & - & - & c_1 \\
 k & - & - & c_1
 \end{array} \tag{3}$$

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 place at s_2 for i and also that i receives the

lowest lottery number. Then i will choose to take the place at s_1 and the reserved place at s_2 is released so that the next student in the queue may take it. To take another example, suppose again that the place 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.

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

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.

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

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 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 random 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)–(6) are arguably fairer than (2) in scenario B but they would be inappropriate in scenario A. Yet we are forced to choose a single expected matching to fit *both* scenarios.

Of course, part of the inequality of expected matching (2) in the context of scenario B 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 C be a set of priority criteria and let 2^C denote the power set of C . Let f be a mapping from $I \times S$ to 2^C . 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 us use the symbol \leq to denote a weak ordering over 2^C that ranks criteria, and combinations of criteria, by importance. Let us write $i \succeq_s j$ to mean that student i is ranked equal to or above student j in the priority ordering

for school s . Then $i \succeq_s j$ if and only if $f(j, s) \leq f(i, s)$. In this way, f and \leq together induce a profile of priority orderings.

As we have noted, it is the profile of priority orderings, and not the pair (f, \leq) , that is included in the formulation of the school choice problem. This formulation implies, then, that the profile of priority orderings captures all of the relevant information contained in the pair (f, \leq) 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 criterion c then let us say that student i has a *claim* to a place at school s on the basis of c . The strength of this claim depends on the importance of the criterion that underpins it, c . 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 programme 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 c_1 and c_2 . Let criterion c_1 be, ‘this person is a vegetarian and this food is suitable for vegetarians’, and let c_2 be, ‘this person has achieved excellent grades in subjects relevant to this programme’.

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

	Apple pie	Meat loaf		Apple pie	Turnip
i	c_1	—	i	c_1	c_1
j	—	—	j	—	—
	Medical	Law		Medical	Business
i	c_2	—	i	c_2	c_2
j	—	—	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 c_1 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 c_2 .

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 c_1 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 c_2 as the reason: i is the more deserving of the place on that programme. 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 c_1 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 c_2 . Indeed, the same reason applies as in the lower-left problem: i is the more deserving of the place on that programme. Individual i 's meriting of a place on the medicine programme 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 c_2 , that refers to merit, is an example of one that grants a conjunctive claim. In the case of the lower-right problem, criterion c_2 grants i priority for the medicine programme and for the business programme, with emphasis here on the conjunction. If, say, we were to allocate the business programme to individual i , she could appeal to c_2 , together with her preference ordering of the programmes, to argue that she should be given the place on the medicine programme instead, even though c_2 applies to both programmes for her.

Priority criterion c_1 , 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 c_1 is

satisfied by individual i in respect of both items. However, i cannot sensibly invoke c_1 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 x, y and z be the goods. Let us suppose that all three individuals have the same preference ordering; they all prefer x to y and y to z . There are two priority criteria c_R and c_W . To make the example more concrete let c_R be, ‘this person adheres to religious laws and this good is compatible with those laws’, while c_W is, ‘this good is fully wheelchair-friendly and this person uses a wheelchair’. This set of priority criteria $\{c_R, c_W\}$ is determined by some authority that is external to us. That is, we, in the role of match-maker, do not have any say in what the criteria ought to be and we cannot alter the set $\{c_R, c_W\}$. We must respect these criteria only and not any other criteria of our own devising or that are proposed by the individuals.

Now suppose that we have the following profile of priority orderings.

$$\begin{array}{ccc}
 x & y & z \\
 \hline
 i & i & i, j, k \\
 j, k & j, k &
 \end{array} \tag{8}$$

Here we see that all three individuals i, j and k have equally strong claims to z , and that i has highest priority for x and y .

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 x and y are compatible with those laws while z is not. In other words, (8) is consistent with the following priority matrix.

$$\begin{array}{ccc}
 & x & y & z \\
 i & c_R & c_R & - \\
 j & - & - & - \\
 k & - & - & -
 \end{array} \tag{9}$$

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

$$\begin{array}{ccc}
 & x & y & z \\
 i & c_R & c_W & - \\
 j & - & - & - \\
 k & - & - & -
 \end{array} \tag{10}$$

We cannot determine, then, whether i has a disjunctive claim spanning x and y , as in (9), or unconnected claims to each of x and y , 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 c_R , together with the fact that i would prefer to have x rather than y , as a reason to assign x 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 A and B from $A \wedge B$. In a similar way, a conjunctive claim that spans, say, goods x and y , is no different to having an unconnected claim to each of x and y . 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 c_R and c_W 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 c_R means, ‘deserving in virtue of good behaviour’, and c_W 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 information 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 USA 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 programmes in the USA 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 programmes 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 programmes. 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 admissions, 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 \leq 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, an *extended school choice problem* comprises these seven items:

1. a set I of students,
2. a set S of schools,
3. a list $q = (q_s)_{s \in S}$ of natural numbers, each indicating the capacity of a school,
4. a list $P = (P_i)_{i \in I}$ of strict preference orderings over S , one for each student,
5. a set C of priority criteria,
6. a mapping f from $I \times S$ to 2^C , the power set of C , and
7. a weak ordering \leq over 2^C such that $A \subset B \rightarrow A < B$.

Items 1–4 are unchanged from the original formulation. Items 5–7 in the new formulation are, as it were, the basic ingredients for a profile of priority orderings, item 5 in the original formulation. I have argued that by including these ‘ingredients’ instead of a ‘ready-made’ profile of priority orderings we expose important information that is relevant to matching. Moreover, since we may derive a profile of priority orderings from items 5–7, nothing is lost in the alternative formulation.

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. Recall that c_1 denotes one priority criterion, 'lives within walking distance', and c_2 denotes the other criterion, 'has a sibling at this school'. In the case of scenario A, $f(i, s_1)$ is $\{c_1, c_2\}$ whereas in scenario B $f(i, s_1)$ is $\{c_1\}$. 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 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 criterion that emerges once we extend the school choice problem.

6. Justified envy

The extended formulation of the school choice problem gives us cause to revisit familiar normative concepts. In this section we consider the concept of *justified envy*, also called 'priority violation' by Kesten (2010), which is a central concept in the literature on the school choice problem.²

To be precise in our definitions, let us note some of the different kinds of 'object' to which formal definitions may apply in this context. Recall that a school choice problem is a list of five items (I, S, q, P, \succeq) and that an extended school choice problem is a list of seven items (I, S, q, P, C, f, \leq) . In addition to these items there are also matchings, which we denote by μ , and then there are deterministic mechanisms and random mechanisms. Naturally, if we define a property that is applicable to one kind of object, we may find that this helps us to define a property that is applicable to another kind of object. With this in mind, we begin by defining a property that is applicable to *students* relative to a given matching and a given school choice problem.

Take any school choice problem (I, S, q, P, \succeq) and any arbitrary matching μ . Relative to these objects, each student in I may or may not have justified envy. Note that this concept does not refer to mechanisms at all, neither deterministic nor random. However, we will use the concept of justified envy to define properties applicable to both kinds of mechanism later. Justified envy is defined as follows.

Justified envy. Relative to a given school choice problem (I, S, q, P, \succeq) and a given matching μ , a student i has justified envy if there exists a student j such that (i) $i \succ_{\mu(j)} j$ and (ii) $\mu(j)P_i\mu(i)$.

Part (i) requires that i has higher priority than j has at the school that j has been assigned to. Recall that $\mu(j)$ is the school that j is assigned to at matching μ , and so $\succ_{\mu(j)}$ is the strict part of the priority ordering for that school. Part (ii) requires that i prefers j 's school, $\mu(j)$, to the one that i has been assigned to, $\mu(i)$. To give an example of justified envy, let us take either of the earlier scenarios A or B and suppose that student j 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 j has.

²In the context of student-college matching, a matching that is free from justified envy is called 'fair' by Balinski and Sönmez (1999) and is called 'weakly fair' by Svensson (1994) in a paper on the allocation of indivisible goods.

Though justified envy is a property applicable to students, we may use it to help us to define a property applicable to matchings. Relative to a given school choice problem, a matching may be *free from justified envy*. This is defined as follows.

Free from justified envy. Relative to a given school choice problem (I, S, q, P, \succeq) , a matching μ is free from justified envy if no student has justified envy at μ .

As yet, these definitions make no reference to mechanisms. But now we can use this property, that applies to matchings, as the basis for the following properties that apply to mechanisms. It is here that we come to a fork in the road, so to speak, a division, that is, into the ‘deterministic’ and ‘random’ frameworks. While the concept of a matching is common to both frameworks, they differ when we come to discuss mechanisms. First, let us define a property applicable to deterministic mechanisms. Recall that a deterministic mechanism is a function that associates a matching to each school choice problem.

Eliminates justified envy. A deterministic mechanism M eliminates justified envy if, for every problem (I, S, q, P, \succeq) , the matching $M(I, S, q, P, \succeq)$ is free from justified envy.

A random mechanism, on the other hand, is a function that associates a probability distribution over the set of matchings to each school choice problem. We may define a property applicable to random mechanisms as follows.

Eliminates *ex post* justified envy. A random mechanism M eliminates *ex post* justified envy if, for every problem (I, S, q, P, \succeq) , the probability distribution $M(I, S, q, P, \succeq)$ assigns a probability of zero to every matching that is not free from justified envy.

All of these definitions refer to the standard formulation of the school choice problem. But these properties can also be defined for the extended formulation of the problem. For example, we may define justified envy as follows. Relative to a given extended school choice problem (I, S, q, P, C, f, \leq) and a given matching μ , a student i has justified envy if there exists a student j such that (i) $f(j, \mu(j)) < f(i, \mu(i))$ and (ii) $\mu(j)P_i\mu(i)$. Here, part (i) requires that the criteria satisfied by the pair $(i, \mu(i))$ are more important than the criteria satisfied by the pair $(j, \mu(j))$. In other words, i has higher priority for j 's school than j has. So this definition is perfectly equivalent to the previous definition of justified envy.

Next, I will make an argument for a revision of the concept of justified envy, and this will *require* us to use the extended formulation of the school choice problem. This argument could not be considered within the standard framework.

6.1. Strongly justified envy

For a deterministic mechanism to eliminate justified envy, it must assign student i to s_1 in both scenarios A and B. And for a random mechanism to eliminate *ex post* justified envy, it must assign student i to s_1 with a probability of one in both scenarios A and B. However, I argue that the standard definition of justified envy makes sense when applied to scenario A but not when applied to scenario B. 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 c_1 , 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 c_1 to justify her envy. Indeed, if we swap the assignments of the two students who are assigned to s_1 and s_2 then 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*. A student may have strongly justified envy relative to a given matching and a given extended school choice problem. Of course, the fact that this definition refers to an *extended* school choice problem is crucial. It is impossible to define this concept relative to the standard formulation of the school choice problem.

Strongly justified envy is a property that applies to students. Naturally, though, we can use this property as the basis of definitions of properties that apply to matchings, to deterministic mechanisms and to random mechanisms. We can do this in just the same way that we developed the basic property of justified envy into properties applicable to matchings and mechanisms.³ Strongly justified envy is defined as follows.

Strongly justified envy. Relative to a given matching μ and a given extended school choice problem (I, S, q, P, C, f, \leq) , a student i has strongly justified envy if there exists a student j such that

- (i) $f(j, \mu(j)) < f(i, \mu(j))$,
- (ii) $\mu(j) P_i \mu(i)$,
- (iii) $f(i, \mu(j)) - f(i, \mu(i)) \neq \emptyset$ and
- (iv) $\mu(i) P_s$ implies $f(i, \mu(j)) - f(i, s) \neq \emptyset$ for all $s \in S$.

Part (i) of this definition requires that i has higher priority for $\mu(j)$ than j has, where $\mu(j)$ is the school that j has been assigned to. Part (ii) requires that i prefers j 's school to the one that i has been assigned to. So parts (i) and (ii) are simply the same conditions for justified envy that we saw earlier. A student's justified envy becomes strong if parts (iii) and (iv) are also satisfied. Part (iii) requires that the school $\mu(i)$ that i has been assigned to does not satisfy all of the criteria for i that $\mu(j)$ does. Part (iv) then extends this further; it requires that *none* of the schools that are *worse* than $\mu(i)$ (according to i 's own preference) satisfy all of the criteria for i that $\mu(j)$ does. To help explain strongly justified envy, let us consider some examples.

Let us return to our earlier scenario A first. As before, c_1 denotes one priority criterion, 'lives within walking distance', and c_2 denotes the other criterion, 'has a sibling at this school'. In this scenario, if i is assigned to s_2 then i has strongly justified envy toward whomever is assigned to s_1 . In this case, (i) and (ii) are satisfied since i is at the top of the priority ordering for s_1 and i prefers s_1 to s_2 . Condition (iii) is also satisfied since i has a sibling at s_1 and not at s_2 , that is, $f(i, s_1) - f(i, s_2) = \{c_2\}$.

³These corresponding properties might be called 'freedom from strongly justified envy', 'elimination of strongly justified envy' and 'elimination of *ex post* strongly justified envy', for example.

Condition (iv) is also satisfied because every school that i likes less than s_2 , which in this case is just s_3 , also fails to satisfy all of the criteria for i that s_1 satisfies. That is, we have $f(i, s_1) - f(i, s_3) = \{c_1, c_2\}$. So we see that in scenario A, if i is assigned to s_2 then i not only has justified envy (conditions (i) and (ii)) but strongly justified envy (all four conditions (i)–(iv)).

By contrast, in scenario B, i does not have strongly justified envy when assigned to s_2 . While conditions (i) and (ii) are satisfied in this case, condition (iii) is not. In this scenario, i does not have a sibling at any school. We have $f(i, s_1) = \{c_1\}$ and $f(i, s_2) = \{c_1\}$. Thus, $f(i, s_1) - f(i, s_2) = \emptyset$, contrary to condition (iii).

Scenarios A and B do not reveal any motivation for condition (iv) in the definition of strongly justified envy. In the following subsection we will consider another example that is intended to show why that requirement is included.

6.2. Motivating condition (iv)

To help motivate condition (iv), it is useful to consider an example in which conditions (i)–(iii) are satisfied and (iv) is not. Suppose that there are four schools, s_1, s_2, s_3 and s_4 . In the preference ordering of student i , s_1 comes first, s_2 second, s_3 third and s_4 last. There is just one priority criterion in this example: ‘this student requires bilingual teaching support and this school provides excellent support’. Schools s_1 and s_3 both offer excellent bilingual teaching support, while s_2 and s_4 provide a basic level of support. Student i has need of this support and another student j does not, and for this reason i is higher than j is in the priority ordering for s_1 . Nevertheless, suppose that j has been assigned to s_1 and i is assigned to one of the other schools, all of which i likes less than s_1 . So i will certainly have justified envy toward j , but we must consider each one of those three other schools to determine when i has *strongly* justified envy toward j .

In Table 1 we see which of the four conditions of strongly justified envy are satisfied when i is assigned to each of the three other schools. If i is assigned to s_2 then only three of the four conditions of strongly justified envy are satisfied. If i is assigned to s_4 then i envies j and this envy is strongly justified. This is indicated by the ‘x’ in all four columns in the bottom row of the table.

Condition (i) simply requires that i has higher priority for s_1 than j has. So this does not depend on which school i is assigned to. An ‘x’ in column (ii) means, ‘ i likes this school less than s_1 ’. An ‘x’ in column (iii) means, ‘this school does not have excellent bilingual support’. And finally, an ‘x’ in column (iv) means, ‘no school worse than this one offers excellent bilingual support’.

When i is assigned to s_3 then conditions (i) and (ii) are satisfied and so i has justified envy towards j , but i 's envy is not strongly justified in this case. This is because s_3 offers excellent bilingual teaching support just as s_1 does. Hence, the justification for i 's envy is undermined and condition (iii) is not satisfied.

Condition (iv) becomes critical when i is assigned to s_2 . If conditions (i)–(iii) are satisfied then why should we not regard this as a case of strongly justified envy? After all, by being assigned to s_2 , student i is being deprived of excellent bilingual teaching support, which i could benefit from and which j does not need. Yet, i 's envy is not strongly justified. To check whether condition (iv) is satisfied we look at the schools that i likes less than the one that i has been assigned to. So in this case we

Table 1. The four conditions of strongly justified envy

School	Bilingual support	(i)	(ii)	(iii)	(iv)
s_2	Basic	x	x	x	
s_3	Excellent	x	x		x
s_4	Basic	x	x	x	x

look at the rows of the table that are below the s_2 row. We find that one of those less-preferred schools, s_3 , offers excellent bilingual teaching support. So we have $f(i, s_1) - f(i, s_3) = \emptyset$. This is why condition (iv) is not satisfied. But why should it matter that s_3 offers this support? The reason is as follows. Student i 's preference for s_2 over s_3 implies that i regards s_2 as having attributes that more than compensate for the lack of excellent bilingual teaching facilities. These implicit, compensating attributes of s_2 weaken the justification for i 's envy of j , so that it is not strongly justified envy.

To help motivate condition (iv) further, consider the following. Suppose that a court has decided that a plaintiff is entitled to take ownership of a plot of land from his neighbour on the basis that this plot will give him access to a certain river. And suppose that this neighbour offers an alternative plot of land, and that the plaintiff actually *prefers* this alternative plot to the original one even though it does not give him access to the river. Then, surely, the neighbour may satisfy the plaintiff's claim by transferring ownership of the alternative plot. Analogously, we may respect i 's language-based priority by assigning i to s_3 , which has excellent bilingual teaching, or, crucially, by assigning i to a school that i *prefers* to s_3 , such as s_2 , even if that preferred school does not have excellent bilingual teaching.

6.3. Relevance to mechanism design

We have noted that strongly justified envy, though a property that applies to students, can be used to define properties applicable to deterministic and random mechanisms. However, it may not be immediately obvious why strongly justified envy is relevant to mechanism design. After all, mechanisms that eliminate justified envy also eliminate strongly justified envy. Since we can eliminate both kinds of justified envy, why should we ever be interested in mechanisms that merely eliminate strongly justified envy and that permit justified envy to arise?

Let us focus on random mechanisms. For scenario B, the s_1 and s_2 columns of priority matrix (3) are identical. That is, these two schools are identical over the priority criteria that are satisfied by each student. Furthermore, all three students share the same preference ordering over the schools, with s_1 the most preferred and s_2 the second-most preferred of every student. Given that s_2 satisfies the same priority criteria for i as s_1 does, I think that i should be assigned to s_2 with some non-zero probability, so that students j and k have some chance of being assigned to s_1 . For this reason, I object to expected matching (2) for scenario B. Moreover, I object to any expected matching for scenario B that assigns i to s_1 with a probability of one. Let us say that there is *unjust exclusivity* if i is assigned to s_1 with a probability of one in scenario B.

However, in scenario B, a matching is free from justified envy if and only if it assigns i to s_1 . Thus, a random mechanism eliminates *ex post* justified envy if and only if it assigns i to s_1 with a probability of one. Hence, the elimination of *ex post* justified envy implies unjust exclusivity. By contrast, a matching is free from strongly justified envy if and only if it assigns i to s_1 or to s_2 . Therefore, expected matchings (4)–(6), where we do not have unjust exclusivity, are all consistent with the elimination of strongly justified envy. So we have reason to be interested in mechanisms that permit justified envy while eliminating strongly justified envy; they can eliminate unjust exclusivity in scenario B.

I think that the discussion in this section demonstrates that adaptations of normative concepts such as justified envy can be thought-provoking and contestable, and not merely technical. This concludes my argument for the extended formulation of the school choice problem. In the next section I discuss the relationship of this paper to some recent literature.

7. Literature on diversity and slot-specific priorities

Socio-economic and racial diversity in schools is an important matter of public policy in the USA 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 *et al.* (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 (2016) 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 places (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.

8. 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 normative concepts can be sensitive to the additional information in the extended school choice problem I proposed a property called strongly justified envy. Crucially, it would be impossible to define this new concept 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.

Author ORCID.  Conal Duddy [0000-0002-7801-9305](https://orcid.org/0000-0002-7801-9305)

Acknowledgements. I thank Ashley Piggins, Bill Zwicker, Bettina Klaus and two anonymous reviewers for many helpful comments.

References

- Abdulkadiroğlu 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' Oneapp*. NBER Working Paper Series 23265.
- Abdulkadiroğlu A., Y.-K. Che and Y. Yasuda** 2015. Expanding 'choice' in school choice. *American Economic Journal: Microeconomics* 7, 1–42.
- Abdulkadiroğlu A. and T. Sönmez** 2003. School choice: a mechanism design approach. *American Economic Review* 93, 729–747.
- Balinski M. and T. Sönmez** 1999. A tale of two mechanisms: student placement. *Journal of Economic Theory* 84, 73–94.
- Bó I.** 2016. Fair implementation of diversity in school choice. *Games and Economic Behavior* 97, 54–63.
- Echenique F. and M. B. Yenmez** 2015. How to control controlled school choice. *American Economic Review* 105, 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, 669–689.
- Ergin H. and T. Sönmez** 2006. Games of school choice under the Boston mechanism. *Journal of Public Economics* 90, 215–237.

- Gale D. and L. S. Shapley** 1962. College admissions and the stability of marriage. *American Mathematical Monthly* **69**, 9–15.
- Hafalir I. E., M. B. Yenmez and M. A. Yildirim** 2013. Effective affirmative action in school choice. *Theoretical Economics* **8**, 325–363.
- Kesten O.** 2010. School choice with consent. *Quarterly Journal of Economics* **125**, 1297–1348.
- Kesten O. and M. U. Ünver** 2015. A theory of school-choice lotteries. *Theoretical Economics* **10**, 543–595.
- Kojima F.** 2012. School choice: impossibilities for affirmative action. *Games and Economic Behavior* **75**, 685–693.
- Kominers S. D. and T. Sönmez** 2016. Matching with slot-specific priorities: theory. *Theoretical Economics* **11**, 683–710.
- Morrill T.** 2013. An alternative characterization of top trading cycles. *Economic Theory* **54**, 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**, 234–243.
- Roth A. E.** 2008. What have we learned from market design? *Economic Journal* **118**, 285–310.
- Sönmez T. and T. B. Switzer** 2013. Matching with (branch-of-choice) contracts at the United States Military Academy. *Econometrica* **81**, 451–488.
- Svensson L.-G.** 1994. Queue allocation of indivisible goods. *Social Choice and Welfare* **11**, 323–330.

Conal Duddy lectures in Economics at the National University of Ireland, Galway. His current scholarship focuses on social choice and voting.