Revising the school choice problem

Duddy, Conal

National University of Ireland, Galway

29 March 2015

Revising the school choice problem

Conal Duddy*

March 29, 2015

Abstract

In a school choice problem each school has a priority ordering over students. These priority orderings depend on criteria such as whether a student lives within walking distance or has a sibling already at the school. We argue that by including just the priority orderings in the problem, and not the criteria themselves, we lose crucial information. This loss of information results in mechanisms that discriminate between students in ways that are difficult to justify. We propose an alternative school choice problem and adaptations of the Gale-Shapley student optimal stable mechanism and the top trading cycles mechanism.

1 Introduction

In some 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 what is called the *school choice problem*. There are three mechanisms that are central in the literature on the school choice problem, a literature that begins with Abdulkadiroğlu and Sönmez (2003). These

*Discipline of Economics, National University of Ireland Galway, University Road, Galway, Ireland (e-mail: conal.duddy@gmail.com). I thank Ashley Piggins for helpful comments. This research was co-funded by the Irish Research Council and the European Commission via a postdoctoral fellowship.

are the Boston mechanism, the Gale-Shapley student optimal stable mechanism and the top trading cycles mechanism.

Each of these mechanisms assigns students to schools based on preference orderings submitted by the students and the priority orderings of the schools. 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.

There is no consensus as to which mechanism is best because there are inescapable tensions in the school choice problem between objectives such as efficiency, strategy-proofness and the avoidance of what is called *justified envy*. The debate about this matching problem has centered around trade-offs between these important normative principles.

However, this paper is not about those trade-offs. Here we argue that the school choice problem itself must be modified if we are to design mechanisms that treat students fairly. We define the school choice problem formally in the next section and argue against its suitability in Section 3. We propose an alternative problem in Section 4. Finally, in Section 5 we show how standard mechanisms can be adapted to this new version of the school choice problem.

2 The school choice problem

A 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. This list of five items is based on the one provided by Ergin and Sönmez (2006).

A *matching* assigns each student to a school. We require that the number of students assigned to a school does not exceed the capacity of that school, and that each student is assigned to exactly one school. A *mechanism* is a systemic way of generating a matching for a school choice problem. A mechanism may involve some randomization. Given such a mechanism we can associate an *expected matching* to a school choice problem prior to the resolving of any lotteries. An expected matching describes each student's probability of being matched to each school.

There is another form of the school choice problem called the *controlled school choice problem*. In the controlled problem, school enrollments 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 socioeconomic groups. For simplicity, we leave aside these kinds of exogenous constraints here. For analysis of controlled school choice problems see Kojima (2012), Hafalir, Yenmez and Yildirim (2013) and Ehlers et al. (2014).

3 Motivating a revision

To help motivate a revision of the school choice problem let us consider a simple scenario involving two neighborhoods and three schools. Two of the schools, call them s_1 and s_2, are located in the Oak Hill neighborhood while s_3 is in the Elm Hill neighborhood.

There are twenty available places at each school, and there are sixty students to be assigned to these places. Twenty of these students live in Oak Hill and forty live in Elm Hill. All sixty students share the same preference ordering over the schools; they all like s_1 the most and s_3 the least. The twenty Oak Hill students have priority for both s_1 and s_2 because they live within walking distance of those schools and the Elm Hill students do not.

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 all twenty available places at s_1 to Oak Hill students. The other two standard mechanisms, the Gale-Shapley mechanism and the top trading cycles mechanism, are more complex and we will postpone a discussion of them until Section 5. For now it suffices to say that they satisfy a principle called "Mutual Top", 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 top" students). Thusly those two mechanisms agree with the Boston mechanism in this case; the twenty Oak Hill students should be assigned to s_1 and the forty Elm Hill students should be split between s_2 and s_3.

Crucially, however, the only difference between schools s_1 and s_2 is that one is more desirable than the other. Twenty Elm Hill students must travel to Oak Hill for their schooling and, seemingly without justification, they are all assigned to the inferior Oak Hill school.

We do not object to there being some inequality of access to more desirable schools in general. Such inequality is often justifiable. We may quite reasonably prefer students to attend schools within walking distance of their homes, and this will inevitably entail unequal access to better schools. But in this scenario the degree of inequality is needlessly high. After all, we could have given students from Elm Hill a non-zero chance of being matched to s_1, by way of some kind of lottery, while still ensuring that all Oak Hill students are assigned to s_1 or s_2.

This scenario motivates a change to the important concept of justified envy that we briefly mentioned earlier. This change, in turn, requires a revision of the school choice problem itself.

3.1 Justified envy

Justified envy occurs when a student i prefers some other school s to the school he or she has been assigned to and i has higher priority for s than a student who has been assigned to s. In the above scenario, if an Elm Hill student is assigned to s_1 while an Oak Hill student is assigned to s_2 then we have a case of justified envy. The Oak Hill student, and her parents, may well feel aggrieved that she has been deprived of a place at a school located in her neighborhood in favor of a student from another area. Indeed, this is the basis on which the standard mechanisms allocate all twenty places at s_1 to Oak Hill students. Yet, we argue, this family's grievance surely loses any justification it seemed to have once we see that the school that the student has been assigned to, s_2, is *also* located in her neighborhood. As we have noted, the schools differ only in their desirability.

To help us to define an alternative concept that we call *strongly justified envy* we introduce some notation and terms here. A key feature of this paper is that we associate a set of "priority factors" to each student-school pair. A priority factor could be "lives within walking distance" or "has a sibling at the school". Given a student-school pair (i,s) let the corresponding set of priority factors be denoted by $\phi(i,s)$. This set may be empty.

Given a student i and schools s and s', if $\phi(i,s)$ is a subset of $\phi(i,s')$ then we say that s' is substitutable for s with respect to i. Note that a school is always substitutable for itself. Let $\omega(i,s)$ denote i's least preferred school that is substitutable for s with respect to i. We now define strongly justified envy as follows.

Strongly justified envy. A student i has been assigned to school s', another student has been assigned to school s despite having lower priority for s than i has, and i prefers $\omega(i,s)$ to s'.

Suppose that student i prefers s to s' but has been matched to s'. It transpires that a student with lower priority for s has been matched to s. This case of justified envy may fail to be "strong" for one of two possible reasons. One is that s' is substitutable for s with respect to i. This would rule out i preferring $\omega(i,s)$ to s'. The second possible reason is that s' is *better* than a school that is substitutable for s. In other words, i strictly prefers s' to $\omega(i,s)$. In the above scenario, an Oak Hill student attending s_3 would regard any Elm Hill student at s_1 with strongly justified envy, but an Oak Hill student at s_2 would not.

4 An alternative problem

In this section we propose an alternative definition of the problem. The seven items that constitute an *alternative 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 F of priority factors,
6. a weak ordering over 2^F, the power set of F, and
7. a mapping ϕ from $I \times S$ to 2^F.

Items 1–4 are unchanged from the original problem. The weak ordering over 2^F is a ranking of priority factors, and combinations of priority factors, by their importance. For example, "has a sibling at the school" may be considered more important than "lives within walking distance".

Items 5–7 induce a priority ordering for each school. Let $\succeq$ be the weak ordering over 2^F. Then i is ranked equal to or above j in the priority ordering for school s if and only if $\phi(i,s) \succeq \phi(j,s)$.

Crucially, the alternative school choice problem allows a mechanism to take account of the "substitutability" of schools with respect to each student. In the next section we show how this can lead to fairer mechanisms.

5 Solutions

Given that items 5–7 induce a priority ordering for each school, we can derive a standard school choice problem from an alternative school choice problem. Therefore we can apply any standard mechanism to an alternative problem. In Section 3, however, we argued that the standard mechanisms may discriminate between students in ways that are difficult to justify. In this section we adapt the Gale-Shapley student optimal stable mechanism and the top trading cycles mechanism to the alternative school choice problem. We leave the adaptation of the Boston mechanism as an open problem.

5.1 The Gale-Shapley mechanism

When school priority orderings are strict the Gale-Shapley student optimal stable mechanism, or more simply the Gale-Shapley mechanism, generates a matching as follows. A matching is said to be *stable* if it is free from justified envy. A school is deemed *possible* for a student if assigning that student to that school is consistent with a stable matching. Each student is assigned to his or her most-preferred possible school. Gale and Shapley (1962, Theorem 2) prove that this matching is always feasible. The matching generated by the Gale-Shapley mechanism is thusly Pareto-superior to all other stable matchings.

When there are ties in school priority orderings, as there typically are in practice, the matter is less straightforward. For analysis of the complex issues that arise in this case see Erdil and Ergin (2008) and Abdulkadiroğlu, Che and Yasuda (2015).

Here we propose one simple adaptation of the Gale-Shapley mechanism to the alternative school choice problem. In what follows we do not assume that the school priority orderings (as induced by items 5–7) are strict. We define this modified Gale-Shapley mechanism as an iterative procedure as follows. At each step k we draw an unassigned student i at random and we assign i to the highest school s in i's preference ordering such that: (i) there is a free place at s and (ii) assigning i to s, combined with all previous assignments made in steps 1 to $k-1$, is consistent with a complete matching that is free from strongly justified envy.

By way of an example, let us apply this mechanism to the scenario we discussed in Section 3. We noted that the three standard mechanisms generate the following expected matching. The first row shows the probability that a representative Oak Hill student will be assigned to each of the three schools, and the second row shows the respective probabilities for an Elm Hill student.

	s_1	s_2	s_3
Oak	1	0	0
Elm	0	1/2	1/2

When we instead apply the modified version of the Gale-Shapley mechanism we generate the following expected matching.

	s_1	s_2	s_3
Oak	$1/3$	$2/3$	0
Elm	$1/3$	$1/6$	$1/2$

The first twenty students drawn at random will be assigned to s_1. None of those assignments are inconsistent with freedom from strongly justified envy. So every one of the sixty students has a probability of one third of being assigned to s_1. Any Oak Hill student not assigned to s_1 will be assigned to s_2. Freedom from strongly justified envy requires this. Thusly each Oak Hill student is assigned to s_2 with a probability of two thirds. The last twenty of the forty Elm Hill students drawn at random will be assigned to s_3. So each Elm Hill student is assigned to s_3 with a probability of one half. It follows that each Elm Hill student is assigned to s_2 with a probability of one sixth.

This second expected matching is much fairer than the first and it fully respects the entitlement of Oak Hill students to attend a local school. Let us also note that if, say, one of the Oak Hill students had a sibling already attending s_1 then the probability of that student being assigned to s_1 would be one. This is because s_2 would not be substitutable for s_1 with respect to that student. Assigning such a student to s_2 would result in strongly justified envy.

Of course, by construction, this method always generates a matching that is free from strongly justified envy. It is also clear that a matching generated by this method will not be Pareto-inferior to any other matching that is free from strongly justified envy. However, like the original mechanism, it may generate a matching that is Pareto-inferior to a matching that *does* exhibit strongly justified envy.

This method also fails to be strategy-proof. Erdil and Ergin (2008) provide an example (their Example 2) of a standard school choice problem for which no strategy-proof mechanism can generate a matching that is stable and constrained efficient. A stable matching is constrained efficient if it is not Pareto-inferior to any other stable matching. We can construct an alternative school choice problem in which no school is substitutable for another with respect to any student and where student preference orderings and the induced school priority orderings match those of Erdil and Ergin's example. In the absence of any substitutability of schools, the concepts of justified envy and strongly justified envy coincide. By recasting their example in this way we can see that no mechanism that eliminates strongly justified envy and achieves constrained efficiency (vis-à-vis strongly justified envy) can be strategy-proof for every alternative school choice problem.

5.2 Top trading cycles

The top trading cycles algorithm originates in a paper by Shapley and Scarf (1974). They investigate a model of trading in goods that are indivisible. They attribute the idea for the algorithm to David Gale. A standard school choice problem corresponds to the kind of scenario Shapley and Scarf analyze when each school has just one place available. Abdulkadiroğlu and Sönmez (2003) adapt the original top trading cycles algorithm to the case where each school may have multiple seats available. Given a standard school choice problem, their method works in the following way to determine a matching.

By random selection let us resolve any ties in the priority orderings of the schools so that all priority orderings become strict. Each student points to his or her most preferred school of those that are available. Each school points to its highest priority student among those not yet matched to a school. There will be at least one cycle. Each student involved in a cycle is assigned to the school he or she is pointing to. The matched students are removed and any school that has now reached its capacity is removed, and the process is then repeated. The procedure ends when all students have been assigned to schools.

There are multiple other ways of adapting the original top trading cycles mechanism to problems in which schools have capacities greater than one. Compelling adaptations have been proposed by Morrill (2014) and Hakimov and Kesten (2014). While there are convincing normative arguments in favor of the newer adaptations, we will prioritize simplicity here and base our proposal on the mechanism proposed by Abdulkadiroğlu and Sönmez (2003), which remains the simplest and most well-known of the adaptations.

Our modified version of the top trading cycles mechanism for the alternative school choice problem works as follows. Let us take the school priority orderings (induced by items 5–7) and use randomization to break ties and make those priority orderings strict. Following this, we distribute school seats to students in the following way to generate a kind of initial endowment of seats before trading takes place. Some students may be endowed with multiple seats and others with none. The total number of seats at each school is equal to the capacity of that school. In each step we begin by selecting at random a seat that is not yet in any student's initial endowment. Suppose we select a seat at school s. We allocate it to the highest priority student for s who has not already been allocated a seat at a school that is substitutable for s with respect to that student. If every student already has a seat at a school that is substitutable for s then we choose at random a student who has not already been allocated a seat at s and we allocate the seat to him or her. We repeat this step until all seats are allocated to students. At this point we have determined initial endowments for the students but we have not yet actually assigned any student to a school. We say that a student is "holding" the seats in his or her endowment.

As in the method proposed by Abdulkadiroğlu and Sönmez (2003), we have each school point at just one student at a time. Initially, each school points to a student chosen at random from among those who hold a seat at that school. Once a school points to a student it continues to point to that same student for as long as he or she remains in the procedure. Only when that student is removed does the school point to another student (again chosen at random from among those holding a seat at the school). Each student always points to his or her most preferred school of those that have not yet reached capacity. Then, just as in the standard algorithm, any student involved in a cycle is assigned to the school he or she is pointing to. When a student is assigned to a school s the particular seat that the student is matched to is the one held by the student that s is pointing to. Those matched students and seats are removed from the procedure. If a removed student had been holding more than one seat then the unmatched seats are immediately redistributed to the endowments of the remaining students by a continuation of the procedure we used to determine the initial endowments. A school is removed once all of its seats have been matched to students. The procedure ends when all students have been assigned to schools.

By way of an example, let us apply this method to the scenario discussed in Section 3. In that scenario the Oak Hill students have higher priority than the Elm Hill students for both s_1 and s_2. Each time a seat at one of those schools is selected at random, during the initial endowment phase, it will be given to an Oak Hill student until all Oak Hill students have been allocated a seat at either s_1 or s_2. The substitutability of those two schools means that none of those students will be allocated seats at both schools. So each Oak Hill student will receive a seat at s_1 with a probability of one half, and a seat at s_2 with a probability of one half. Twenty seats at Oak Hill schools will be allocated to Elm Hill students. The twenty seats at s_3 will also be allocated to Elm Hill students. Since all sixty students share the same preference ordering it is easy to see that no trading will take place. That is, every cycle that occurs will be a trivial cycle in which a school and a student point directly to one another.

Prior to the resolving of lotteries to determine the order in which schools seats are initially allocated and to break ties in the school priority orderings, the students face the following expected matching.

	s_1	s_2	s_3
Oak	1/2	1/2	0
Elm	1/4	1/4	1/2

This is less egalitarian than the expected matching that resulted from the modified version of the Gale-Shapley mechanism, though more egalitarian than the one generated by the standard mechanisms. As an aside, let us observe that in this expected matching we see that, for all sixty students, the *conditional* probability of a student being assigned to s_1 *given* that he or she is assigned to an Oak Hill school is one half. The same is true of s_2. On the basis of this equality in conditional probabilities one may argue that this expected matching is fair even though an Elm Hill student has less chance of attending s_1 than an Oak Hill student has.

In modifying the top trading cycles mechanism we retained an important feature of the original. The rule that determines which student each school points to does not take any account of the preferences of students. This is a simple way of ensuring that the resulting mechanism is strategy-proof. This is why we ignored the popularity of s_1 when distributing seats to students and instead treated s_1 and s_2 symmetrically, as is reflected in the expected matching.

To see that this method is strategy-proof one may apply the very same intuition that Abdulkadiroğlu and Sönmez (2003) provide for their Proposition 4. Also, since in every iteration we only match students to schools that are their most preferred of the remaining schools, this method always generates a matching that is Pareto-efficient. However, while it does take account of the substitutability of schools, it does not eliminate strongly justified envy. This is unsurprising given that a central fact in the literature on matching problems is that the Gale-Shapley mechanism ensures stability but not efficiency, while the top trading cycles mechanism ensures efficiency but not stability.

References

Abdulkadiroğlu, Atila, Yeon-Koo Che and Yosuke Yasuda. 2015. "Expanding "choice" in school choice." *American Economic Journal: Microeconomics* 7 (1): 1–42.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. "School choice: A mechanism design approach." *The American Economic Review* 93 (3): 729–747.

Ehlers, Lars, Isa E. Hafalir, M. Bumin Yenmez and Muhammed A. Yildirim. 2014. "School choice with controlled choice constraints: Hard bounds versus soft bounds." *Journal of Economic Theory* 153: 648–683.

Erdil, Aytek, and Haluk Ergin. 2008. "What's the matter with tie-breaking? Improving efficiency in school choice." *The American Economic Review* 98 (3): 669–689.

Ergin, Haluk, and Tayfun Sönmez. 2006. "Games of school choice under the Boston mechanism." *Journal of Public Economics* 90: 215–237.

Gale, David, and Lloyd S. Shapley. 1962. "College admissions and the stability of marriage." *The American Mathematical Monthly* 69 (1): 9–15.

Hafalir, Isa E., M. Bumin Yenmez and Muhammed A. Yildirim. 2013. "Effective affirmative action in school choice." *Theoretical Economics* 8 (2): 325–363.

Hakimov, Rustamdjan and Onur Kesten. 2014. “The equitable top trading cycles mechanism for school choice.” WZB Discussion Paper SP II 2014-210.

Kojima, Fuhito. 2012. “School choice: Impossibilities for affirmative action.” *Games and Economic Behavior* 75 (2): 685–693.

Morrill, Thayer. 2014. “Two simple variations of top trading cycles.” *Economic Theory* DOI: 10.1007/s00199-014-0820-4.

Shapley, Lloyd S. and Herbert Scarf. 1974. “On cores and indivisibility.” *Journal of Mathematical Economics* 1 (1): 23–37.