Implementing Diversity in School Choice

Peter Troyan and Daniel Fragiadakis

Department of Economics, Stanford University

1 Introduction

In recent years, school choice policies have become popular in cities across the United States to give parents the opportunity to choose which school their child will attend, and economists have been important in the design of many of the mechanisms used for the actual implementation of these policies (e.g., in New Orleans, Denver, Boston, and New York City, among others).[1] In conjunction with giving parents more choice, many school districts also set socioeconomic diversity as an important goal, which is often achieved by imposing lower and upper quotas on the numbers of each type of student that can be enrolled at a school. For example, the Cambridge, Massachusetts school district divides students into two socioeconomic classes (low SES and high SES) and requires that each SES class comprise 45-55% of the students at every school.[2] Similarly, Montclair, New Jersey divides the city into three zones based on socioeconomic data and attempts to equalize the number of students from each zone at each school.[3] In New York City, "Educational Option" schools must have 16 percent of students score above grade level on a standardized reading test, 68 percent must score at grade level, and 16 percent must score below grade level.

While some papers such as [4] and [1] consider the problem of diversiy in school choice, most work considers only type-specific *upper quotas*. Such upper quotas can be deficient, because they can still result in completely segregated schools.[4] One simple (and often used) approach is to impose upper quotas that are artificially lower than the true upper quotas (what we will call "artificial caps"). By imposing sufficiently stringent artificial caps, the school district can ensure that all of the true *lower and upper* quotas are satisfied. However, this solution may waste seats in the sense that it will be possible to reassign some students and make them better off, while still satisfying all diversity constraints (and without harming other students).

In this paper, we propose mechanisms that guarantee all diversity constraints will be satisfied and outperform imposing artificial caps. In addition, we show that our mechanisms perform well on other dimensions that are important in matching markets in general, and school choice markets in particular: namely, strategyproofness, fairness, and nonwastefulness. To the best of our knowledge, we are the first to provide (nontrivial) strategyproof mechanisms that satisfy all type-specific lower and upper quotas.

[1] See, for example, [2] and [3].

[2] The exact percentages may change from year to year.

[3] Due to recent court rulings, many school districts can no longer use race as a factor in making assignments, and so use other criteria such as geography or socioeconomic status.

[4] In addition, [8] shows that these mechanisms can actually be Pareto inferior for minority students, the supposed beneficiaries.

2 Model

There is a set $I = \{1, \ldots, I\}$ of students and a set $S = \{s_1, \ldots, s_M\}$ of M schools. Each student is of exactly one type in the set $\Theta = \{\theta_1, \ldots, \theta_T\}$. Each school has a capacity of q_s seats, and type specific lower and upper quotas $L_{s\theta}$ and $U_{s\theta}$, respectively, for each type.[5] Each school has a **priority relation** $\succ_s$ on the set of students, and each student has a **preference relation** P_i on the set of schools. A matching is a mapping $\mu : I \cup S \to 2^{I \cup S}$ such that $\mu(i) \in S$ for all $i \in I$; $\mu(s) \subseteq I$ for all $s \in S$; and $i \in \mu(s) \iff \mu(i) = s$. We say a matching is **feasible** if all capacities and quotas are satisfied.

If, at a matching μ, there is some student-school pair (i, s) such that $s \succ_i \mu(i)$ and the matching μ' in which i is moved to s and all other assignments remain the same is feasible, then μ is **wasteful**. If μ is not wasteful, we say it is **nonwasteful**.

If at matching μ, there is a student i such that $s \succ_i \mu(i)$ and $i \succ_s j$ for some $j \in \mu(s)$, we say that i **envies** j. If no student has any envy, we call the matching **fair**.

A mechanism χ is a function that, for each vector of reports of the students $P = (P_i)_{i \in I}$, outputs a matching. A mechanism is **strategyproof** if no student can ever improve her assignment by lying about her preferences, no matter what the other students report.

Strategyproofness, nonwastefulness, and fairness are three important desiderata in school choice. However, it is simple to show that once lower quotas are introduced, the set of fair and nonwasteful matchings may be empty.[6] Due to this impossibility result, we must weaken either fairness or nonwastefulness.

[5] We of course must impose some consistency conditions on the capacities and quotas and number of students of each type, to ensure that we can feasibly assign all students to a school. For brevity, we do not discuss these conditions here.

[6] See, for example, [5] and [6].

	q_s	$L_{s\ell}$	L_{sh}	$U_{s\ell}$	U_{sh}	Priorities
A	1	0	0	1	1	$h_1 \succ_A h_2 \succ_A \ell_1$
B	1	0	0	1	1	$h_1 \succ_B h_2 \succ_B \ell_1$
C	2	1	1	1	2	$h_1 \succ_C h_2 \succ_C \ell_1$

Table 1: Capacities, type-specific quotas, and priorities for the three schools. The market consists of two students of type h, labeled h_1 and h_2, and one student of type ℓ, labeled ℓ_1.

3 An illustrative market

For the remainder of the paper, we use an example of a simple market to illustrate our ideas. The market consists of three students, $I = \{h_1, h_2, \ell_1\}$. Two students, h_1 and h_2, are of high socioeconomic status, while one student, ℓ_1, is of low socioeconomic status. The set of schools is $S = \{A, B, C\}$.

The standard approach found in [4] and [1] would be to run the deferred acceptance algorithm on the above market using only the upper quotas. However, this may unfortunately result in an infeasible assignment. To see this, consider the following preferences: $P_{h_1} : A, B, C$, $P_{h_2} : B, A, C$ and $P_{\ell_1} : C, A, B$.[7]

The standard solution when there are only upper quotas is to use the *deferred acceptance (DA) algorithm*. In round 1, each student applies to her first choice school, and the school *tentatively* accepts applicants according to its priority relation up to its quotas and capacities, rejecting the rest. All students rejected then apply to their next choice, and the schools consider students tentatively held from the previous round and all new applicants, again tentatively accepting students according to its priority relation. The algorithm continues until either no student is rejected or all students have applied to all schools.

When we run DA on the reports $P = (P_{h_1}, P_{h_2}, P_{\ell_1})$ using only the upper quotas from Table 1, the resulting matching is

$$\mu_1 = \begin{pmatrix} A & B & C \\ h_1 & h_2 & \ell_1 \end{pmatrix}$$

Note that this assigns no type h students to C, which violates the lower quotas at C, and so is not feasible.

Artificial caps DA

One easy solution to the lack of feasibility is to artificially cap the number of h students at A by lowering U_{Ah} from 1 to 0. Now, if we run DA with these quotas, the output is

$$\mu_2 = \begin{pmatrix} A & B & C \\ \emptyset & h_1 & h_2, \ell_1 \end{pmatrix},$$

which does satisfy all upper and lower quotas. By using this artificial cap DA algorithm (ACDA), we can **guarantee** that the true lower and upper quotas will be satisfied for *any* report of the students. Such an approach is taken by many real-world markets (see, for example, [4] or [7]).

[7] We use the notation $P_i : x, y, z$ to mean that student i prefers x to y to z.

While simple, this mechanism is problematic. This can be seen by considering a profile of reports $\tilde{P} = (P_{h_1}, \tilde{P}_{h_2}, P_{\ell_1})$, where students h_1 and ℓ_1 reports are unchanged, but student h_2 reports $\tilde{P}_{h_2}; C, B, A$. The resulting matching after running DA with artificial caps is

$$\mu_3 = \begin{pmatrix} A & B & C \\ \emptyset & h_1 & h_2, \ell_1 \end{pmatrix}.$$

Note here that, while μ_3 satisfies all of the quotas, we can in fact make h_1 better off by assigning him to A, and the resulting matching would still satisfy all of the quotas.

This wastefulness is a general feature of artificial caps type mechanisms. These mechanisms eliminate some seats ex-ante, without regard for student demand. Doing so ensures that all of the lower quotas will be filled, but can potentially result in large efficiency losses. Our proposed mechanisms will rectify this by only reducing quotas when necessary.

Deferred acceptance with dynamic quotas (DADQ)

The idea behind DADQ can be illustrated by our example: rather than reducing the type h quota at A from the outset as in ACDA, we only do so when necessary. More specifically, we start by running the algorithm with the original upper quotas from Table 1. If the resulting matching satisfies all of the lower quotas, the algorithm finishes. If not, we reduce a quota at some school, and then run the algorithm again. We continue until the resulting matching is feasible.

To illustrate, let us describe the algorithm for our example when the submitted preference profile is P. In round 1, the resulting matching is μ_1, just as above. Since the lower quota at C is not filled, we then lower the type h quota at A from 1 to 0, and rerun DA. The new matching is μ_2, which does satisfy all of the quotas, and so the algorithm finishes.

While the output in the above is the same as ACDA, the benefit of DADQ is revealed when we consider the preference profile $\tilde{P}$. In round 1 of the algorithm, the matching is

$$\mu_4 = \begin{pmatrix} A & B & C \\ h_1 & \emptyset & h_2, \ell_1 \end{pmatrix}.$$

Since this satisfies all quotas, the algorithm ends and the final matching is μ_4, which Pareto dominates μ_3, the matching obtained from ACDA in the previous section. This Pareto dominance is general, as we discuss in Section 4.

Multistage DA

We also consider an alternative mechanism which we call multistage DA (MSDA) (the tradeoffs between DADQ and MSDA will be discussed in Section 4). MSDA is also run in several stages. The idea behind multistage DA is to first “reserve” a number of students equal to the sum of the minimum quota seats remaining. Then, run standard DA using only the upper quotas on the remaining students. Since we have the reserved students, we know that no matter how the students in the first stage are allocated, we will certainly have enough agents to fill any remaining quotas. This process is then repeated until all students have a seat.

To illustrate, consider our example with preferences P. Based on the starting lower quotas, we must reserve 1 type h student and 1 type ℓ student. Say that we reserve h_2 and ℓ_1. Then, we run DA on student h_1, who applies (and is accepted at) A. Now, we reduce the quota and capacity at A by 1. After this assignment, there the number of lower quota seats is exactly equal to the number of agents of each type, and so we run DA on the remaining students, but only allow them to apply to schools with lower quota seats remaining. The final matching is matching μ_4 from above.

4 Discussion

The market given above is a particularly simple example, but all of the mechanisms can be extended to markets of arbitrary size. Some complexities arise, because for DADQ, we must decide in what order to reduce the quotas, while for MSDA, we must decide which students to reserve at each stage. Strategyproofness is preserved by requiring that these decisions not depend on the submitted preferences of the students.

Recall the properties of interest from Section 2: strategyproofness, fairness, and nonwastefulness. The DADQ and MSDA mechanisms illustrate the necessary tradeoff (because of the impossibility result) between fairness and nonwastefulness: DADQ is fair (in the sense that it eliminates all envy among students of the same type), but may waste seats, while MSDA is nonwasteful, but will not be fair.

While DADQ cannot be nonwasteful (since we already know it is fair), it does not give up on nonwastefulness entirely. In particular, note that it only eliminates seats (i.e., lowers quotas) when necessary, and so in fact Pareto dominates ACDA. Similalrly, while MSDA cannot be fair (since we know it is nonwasteful), it will satisfy a weaker fairness property, which we define in the full paper.

The full proofs of the above fairness and nonwastefulness properties are quite intuitive. However, the last key property we discuss, strategyproofness, proves much more difficult. Strategyproofness is regarded to be an important property in many matching markets, and in school choice markets in particular.[8] Because of the importance of strategyproofness, and the lack of strategyproof mechanisms for problems with lower quotas (besides ACDA-type mechanisms, which, as we noted, have other problems), we took strategyproofness as integral to the design of our mechanisms, and in fact, under certain (broad) conditions, both of our mechanisms will be strategyproof.

For MSDA to be strategyproof, we simply require that the order in which students are "reserved" not depend on the submitted preferences. Then, since within each stage DA is strategyproof, and an agent cannot affect the stage in which he participates, the MSDA mechanism overall is strategyproof.

Proving that DADQ is strategyproof is significantly more difficult. It may seem at first that this would not even be true, because the report of an agent can affect which quotas are filled at the end of a given stage, and thus also will affect the stage at which the algorithm ends. While this is true, we show that, as long as the choice functions of the schools get weakly smaller at each stage,[9] the mechanism will be strategyproof. Intuitively, any individual agent's report can only affect the number of seats filled at any school by at most 1. Consider an agent i who can potentially manipulate, and say that if he reports truthfully, the algorithm ends in stage k, when the quota and capacity vectors (for all schools and types) have been lowered to q^k and U^k, respectively. It is clear that the student then never wants to submit a (false) report that causes the algorithm to end at a stage $k' > k$, since at later stages, there are simply less seats for everyone. What is more difficult to show is that agent i does not want to submit a report that ends the algorithm early, at some $k' < k$, where he "locks-in" a seat that would have been cut if the algorithm had continued to stage k. The general intuition is that, if i is able to obtain a seat at some school s by reporting some $P_i'(\neq P_i)$ and causing the algorithm to end in stage $k' < k$, then under the true report P_i (or, indeed, any report), when i applies to s, all of the lower quotas will be filled, and so the algorithm will end before he is rejected from s. Since he cannot be rejected from s, there is no risk to i from reporting the schools preferred to s truthfully.

5 Conclusion

This paper provides the first (nontrivial) strategyproof mechanisms for matching markets with type-specific lower and upper quotas. We identify a fairness/nonwastefulness tradeoff and introduce two mechanisms, one for each side of this tradeoff: the deferred acceptance with dynamic quotas mechanism (DADQ) is strategyproof and fair, while the mulitstage deferred acceptance mechanism (MSDA) is strategyproof and nonwasteful. DADQ still performs well with respect to nonwastefulness, Pareto dominating the simple and commonly used solution of imposing stringent artificial caps. MSDA, on the other hand, will still satisfy a weaker definition of fairness.

The main application of our model is to school choice markets in which diversity is a large concern, of which there are numerous examples. However, we would like to end by noting that our mechanisms can be applied to any markets where lower quotas are relevant. Particular examples include the Japanese hospital residency market studied in [7], the military cadet matching market studied in [10], or, in general, any organization which assigns members to projects, with each project having a minimum staffing requirement (such as a firm and its employees).

[8] Strategyproofness prevents parents from needing to play a complicated preference revelation game, and was a crucial property in the redesign of Boston's school choice mechanism. See also [9].

[9] For example, when the quota U_{Ah} is lowered to force a type h student to apply to C and fill the lower quota there, school A is not allowed to reassign this seat to a type ℓ student. If it was allowed to do so, student ℓ_1 may have an incentive to prolong the algorithm in order to lower the U_{Ah} quota, allowing him entry to school A.

References

[1] Atila Abdulkadiroğlu. College admissions with affirmative action. *International Journal of Game Theory*, 33:535–549, 2005.

[2] Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. The new york city high school match. *American Economic Review, Papers and Proceedings*, 95:364–367, 2005.

[3] Atila Abdulkadiroğlu, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez. The boston public school match. *American Economic Review, Papers and Proceedings*, pages 364–367, 2005.

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

[5] Lars Ehlers, Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. School choice with controlled choice constraints: Hard bounds versus soft bounds. 2011. working paper, Carnegie Mellon University.

[6] Daniel Fragiadakis, Atsushi Iwasaki, Peter Troyan, Suguru Ueda, and Makoto Yokoo. Strategyproof matching with minimum quotas. 2013. working paper, Stanford University.

[7] Yuichiro Kamada and Fuhito Kojima. Efficient matching under distributional concerns: Theory and application. 2013. unpublished manuscript.

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

[9] Parag A. Pathak and Tayfun Sönmez. School admissions reform in chicago and england: Comparing mechanisms by their vulnerability to manipulation. *American Economic Review*, 103:80–106, 2012.

[10] Tayfun Sönmez. Bidding for army career specialties: Improving the rotc branching mechanism. *Journal of Political Economy*, forthcoming, 2013.