



**Respecting priorities versus respecting preferences in school choice:
When is there a trade-off ?**

Estelle Cantillon
ECARES, Université libre de Bruxelles

Li Chen
University of Gothenburg

Juan Pereyra
Universidad de la Republica Uruguay

December 2022

ECARES working paper 2022-39

Respecting priorities versus respecting preferences in school choice: When is there a trade-off?*

Estelle Cantillon[†], Li Chen[‡] and Juan Pereyra[§]

December 2022

Abstract

A classic trade-off that school districts face when deciding which matching algorithm to use is that it is not possible to always respect both priorities and preferences. The student-proposing deferred acceptance algorithm (DA) respects priorities but can lead to inefficient allocations. The top trading cycle algorithm (TTC) respects preferences but may violate priorities. We identify a new condition on school choice markets under which DA is efficient and there is a unique allocation that respects priorities. Our condition generalizes earlier conditions by placing restrictions on how preferences and priorities relate to one another only on the parts that are relevant for the assignment. We discuss how our condition sheds light on existing empirical findings. We show through simulations that our condition significantly expands the range of known environments for which DA is efficient.

Keywords: Matching, envyfreeness, efficiency, priorities, preferences, matching algorithms

JEL Codes: C78, D47, D78, D82

*Funding from FNRS PDR grant 23671793, ANII FMV-1-2021-1-166576, and Tore Browaldhs Stiftelse BFh18-0007 is gratefully acknowledged. We thank Péter Biró, Lars Ehlers, Bettina Klaus, David Manlove, Parag Pathak, Camille Terrier and seminar and conference audiences for comments and suggestions. The simulations were enabled by resources in project SNIC 2022/22-675 provided by the Swedish National Infrastructure for Computing (SNIC) at UPPMAX, partially funded by the Swedish Research Council through grant agreement no. 2018-05973.

[†]FNRS, Université Libre de Bruxelles (ECARES) and CEPR, email: Estelle.Cantillon@ulb.be.

[‡]University of Gothenburg, email: li.chen@economics.gu.se.

[§]Department of Economics, FCS, Universidad de la República Uruguay, email: juan.pereyra@cienciassociales.edu.uy.

1 Introduction

Many countries promote parental choice for schools. How choice is actually implemented greatly varies across countries and school districts and is often hotly debated by stakeholders. A cornerstone of this debate is the choice of the algorithm used to allocate students to schools, when capacity is limited. Indeed, different algorithms will typically result in different final allocations, even when submitted preferences and priorities are the same.

In their classic paper, [Abdulkadiroğlu and Sönmez \(2003\)](#) identify desirable properties of matching algorithms in the school choice context. These include efficiency, which can be interpreted as a measure of the extent to which the algorithm respects students’ preferences, and envyfreeness, a measure of the extent to which priorities are respected. Unfortunately, no algorithm exists that always produces an envyfree and efficient final allocation ([Balinski and Sönmez, 1999](#)), and there is a trade-off between respecting students’ preferences and respecting schools’ priorities. Two strategyproof mechanisms satisfy one of the desiderata: the student-proposing deferred acceptance (DA) algorithm,¹ originally proposed by [Gale and Shapley \(1962\)](#), always yields an envyfree allocation; the top trading cycles (TTC) algorithm, first described by [Shapley and Scarf \(1974\)](#), is efficient. Other non-strategyproof mechanisms are used in practice. They are often variants of the school-proposing DA, which always produces an envyfree matching but is not efficient ([Roth, 1984](#)), or of the immediate acceptance (IA) algorithm, also known as the Boston mechanism, which produces an envyfree allocation in equilibrium ([Ergin and Sönmez, 2006](#)), but may not be efficient.

How big is the trade-off in practice is an empirical question, and existing evidence is mixed. [Table 1](#) summarizes some of the available evidence. [Abdulkadiroğlu, Pathak, and Roth \(2009\)](#) and [Che and Tercieux \(2019\)](#) have documented a significant trade-off between respecting students’ preferences and schools’ priorities in the NYC High School markets. In Budapest, [Ortega and Klein \(2022\)](#) have documented large violations of priorities when using TTC. Using estimated preferences, [Calsamiglia, Fu, and Güell \(2020\)](#) and [De Haan, Gautier, Oosterbeek, and Klaauw \(forthcoming\)](#) also find significant differences across algorithms. This contrasts with the evidence from the Boston Public Schools system where [Abdulkadiroğlu, Pathak, Roth, and Sönmez \(2006\)](#) and [Pathak \(2017\)](#) have found very little difference between the outcomes of DA and TTC. Likewise, we were able to look at data from the allocation of seats in elementary schools in the city of Ghent (Belgium) and found the school-proposing DA close to be efficient. Evidence from New Orleans is somewhere in between.

These examples raise the question of when the choice of algorithms actually matters: When should a school district spend time weighting the choice among algorithms? When is

¹When not specified, we refer to the student-proposing DA algorithm as DA in short.

Table 1: The tension between efficiency and envyfreeness across school districts

School district	Algorithms compared	% students with Pareto improving trade	% students with justified envy	Special features
Boston, all levels (Abdulkadiroğlu, Pathak, Roth, and Sönmez, 2006; Pathak, 2017)	Student-proposing DA, TTC		6.8	Guaranteed placement and sibling priority, catchment area
Budapest, secondary (Biró, 2012; Ortega and Klein, 2022)	Student-proposing DA, TTC		64	Combination of school grades, centralized exam and own school test/interview
Ghent elementary (own source)	school-proposing DA, TTC	< 1	9.2	Sibling and staff priority, distance as tie-breaker
New Orleans - elementary to middle school (Abdulkadiroğlu, Che, Pathak, Roth, and Tercieux, 2020)	school-proposing DA, TTC		13	Sibling priority, catchment area
New York, high school (Abdulkadiroğlu, Pathak, and Roth, 2009)	Student-proposing DA, TTC	5.45*	44*	Mix of schools and of priority and ranking criteria

* indicates inferred values when these statistics were not directly available. There are no commonly accepted measures of the efficiency and envyfreeness trade-off. We consider two ordinal measures. A first measure is the fraction of students in the student-proposing DA with an available Pareto improving trade. A student is said to have a Pareto improving trade if there exists one or more students, and exchanges of seats among them, such that all are better off based on submitted preferences. A second measure, taking TTC allocation as starting point, is the fraction of students who have justified envy, i.e. they are not accepted at a school that they prefer to their assigned school, even though a student with lower priority got in (Doğan and Ehlers (2022) provides foundations for this measure).

the choice of the algorithm a second-order issue?

We identify a new condition on preferences and priorities under which there is no trade-off between the two goals: there is an envyfree and efficient allocation, and in fact it is also unique. Our condition, Generalized Mutually Best Pairs (GMBP), generalizes existing conditions identified in the literature. Like some of the existing conditions, it seeks to capture the degree to which priorities and preferences are congruent, but it restricts attention to the parts of the preferences and priorities that actually matter for the allocation (we call it the *simplified market*). Some students who have priority in a school may never need to consider it because they can access a preferred school for sure. Reversely, some schools that students have listed in their preferences may not be attainable anyways. Roughly speaking, our GMBP condition is satisfied if, in the simplified market, we can *sequentially* match students to the best school in their preference lists for which their priority qualifies them for one of the available seats. If a market satisfies the GMBP condition, then there is no trade-off between efficiency and envyfreeness. There is a unique envyfree allocation and it is efficient. It is reached at equilibrium by the student-proposing DA, the school-proposing DA, the IA algorithm or any other mechanism that produces an envyfree allocation. Moreover, if we can sequentially match students to their best available schools in the original market (without simplifying), then all the algorithms – including TTC – produce the same allocation that is both efficient and envyfree.

We discuss how this condition naturally arises in existing school choice environments. An expanding empirical literature has shown that parents typically value school quality and proximity (Abdulkadiroğlu, Agarwal, and Pathak, 2017; Fack, Grenet, and He, 2019). We show that such preferences, when combined with priorities based on distance, meet our GMBP condition. Likewise, to the extent that preferences for academic quality correlate with academic performance, priorities based on academic performance will also meet the condition. We show through simulations that our condition significantly expands the set of known environments for which there is no trade-off between efficiency and envyfreeness. Our results shed light on the differences in performance across school environments described in Table 1.

Positioning within the literature. A number of papers have explored the trade-off between envyfreeness and efficiency in school choice markets, starting from the seminal papers by Balinski and Sönmez (1999) and Abdulkadiroğlu and Sönmez (2003). The DA and TTC algorithms are natural starting points here: DA maximizes efficiency among algorithms that produce envyfree outcomes (Gale and Shapley, 1962; Balinski and Sönmez, 1999), and TTC performs well (and under some circumstances best) on envyfreeness, within the class of efficient and strategyproof mechanisms (Abdulkadiroğlu, Che, Pathak, Roth, and Tercieux, 2020; Doğan and Ehlers, 2022). When TTC does not involve any violation of priorities, it

yields the same outcome as DA.

One approach has been to identify domain restrictions on the set of priorities such that DA is efficient (Ergin, 2002; Ehlers and Erdil, 2010; Erdil and Kumano, 2019; Ishida, 2019) or TTC yields the same outcome as DA (Kesten, 2006; Ishida, 2019), for any preference profile.² Heo (2019) explores the flip-side question of domain restrictions on preferences such that DA is efficient or TTC yields the same outcome as DA, for any priority profile.

Of course, whether efficiency and envyfreeness conflict depends on the *combination* of preferences and priorities. Moreover, in practice, priorities are not entirely independent of preferences and are instead a partial reflection of what school districts (or schools) view as legitimate preferences. For example, many school districts give priorities to siblings or to students living close to schools, and parents often prefer to send their children to nearby schools or to the school where a sibling is already educated, everything else equal.

Our GMBP condition captures this insight and generalizes existing conditions, placed on how preferences and priorities relate, known to guarantee that DA is efficient, such as Salonen and Salonen (2018)’s single peakedness, Clark (2006)’s no crossing condition or Reny (2021)’s student-oriented preferences.³

The trade-off between efficiency and envyfreeness is conceptually connected to the issue of uniqueness of stable matchings which has been explored in one-to-one two-sided markets (Alcalde, 1994; Eeckhout, 2000; Clark, 2006; Niederle and Yariv, 2009; Legros and Newman, 2010; Romero-Medina and Triossi, 2013; Lee and Yariv, 2014; Gutin, Neary, and Yeo, 2021). Intuitively, given the lattice structure of stable matchings (Knuth, 1997), uniqueness of stable matchings suggests that preferences on both sides (priorities and preferences in the school choice context) are “sufficiently compatible” that starting from one side or the other when using deferred acceptance does not impact the final outcome. Methodologically, our procedure of removing irrelevant choices of students to define the simplified market, is related to the simplification of Gutin, Neary, and Yeo (2021). The main difference is that our simplification process only removes irrelevant choices from students’ preferences and update schools’ priorities accordingly. We do not remove irrelevant students for schools, because doing so impedes efficiency.

Though uniqueness of stable matchings does not in itself guarantee efficiency (in a school choice context) nor the equivalence between DA and TTC, we argue in Section 4 that most existing conditions that have been identified to imply uniqueness rely, in their definition or their proof, on some form of mutually best pairs condition. They therefore, once appro-

²The second question is more restrictive as DA can be efficient without producing the same allocation as TTC, whereas TTC yields the same allocation as DA when it is envyfree.

³Clark (2006) originally solely focused on the uniqueness of stable matching in one-to-one two-sided matching but Salonen and Salonen (2018) show that his condition also ensures that DA is efficient (and that TTC yields the same outcome as DA).

privately extended to one-to-many matching contexts, also imply that DA is efficient. This connects the two strands of the literature.

2 Model

There is a wide variety of algorithms and procedures used in practice to match students to schools. We focus on *direct* and *priority-based* mechanisms which elicit preference rankings from students (this is the “direct” part) and use priorities to assign students to schools when demand exceeds capacity (the “priority” part).

Let I denote the set of students ($|I| = n$), and S the set of schools ($|S| = m$). Let school s_{m+1} represent the outside option for the students. A market is defined by students’ preferences over schools, schools’ priorities over students and schools’ capacities, and is denoted by $\mathcal{E} = (\succ, P, q)$, where \succ are students’ preferences over $S \cup \{s_{m+1}\}$, P are schools’ priorities, and $q = (q_1, \dots, q_m) \in \mathbb{N}^m$ are school capacities (without loss of generality, we let the capacity of s_{m+1} to be equal to n). We assume that preferences and priorities are strict linear orders (no indifference). In addition, the priority order of each school only contains those students who consider the school as acceptable, i.e. preferred to their outside options.

An allocation is a mapping $\mu : I \rightarrow S \cup \{s_{m+1}\}$ that describes to which schools students are assigned, with the understanding that $\mu(i) = s_{m+1}$ means that student i is not assigned. An allocation is feasible if it does not allocate more students to a school than its capacity, $|\mu^{-1}(s)| \leq q_s$, for every $s \in S$. A feasible allocation μ is *Pareto efficient* if it is not Pareto dominated by any other feasible allocation μ' , that is, if there is no μ' such that $\mu'(i) \succeq_i \mu(i)$ for all $i \in I$ and $\mu'(i) \succ_i \mu(i)$ for some $i \in I$. A feasible allocation is blocked by a pair (i, s) if i prefers s to her assignment $\mu(i)$, and either s has some empty seats under μ , or there is a lower priority student j who is assigned to s under μ , that is, formally, $s \succ_i \mu(i)$, and $|\mu^{-1}(s)| < q_s$ or $i P_s j$ for some $j \in \mu^{-1}(s)$. An allocation is said to be *envyfree* or *stable* if it is not blocked by a pair.⁴ A direct priority-based mechanism is a function that maps student preferences \succ , schools’ capacities q and priorities P into a feasible allocation.

We first consider two mechanisms at the center of the debate between envyfreeness and efficiency. The first one is the student-proposing DA, first proposed by [Gale and Shapley \(1962\)](#). The mechanism is strategyproof, i.e., it is a weakly dominant strategy for students to submit their true preferences, and it always produces an envyfree allocation based on the submitted preferences and priorities. The allocation it produces Pareto-dominates all other envyfree allocations ([Gale and Shapley, 1962](#); [Balinski and Sönmez, 1999](#)). The second algorithm, TTC, was first described in [Shapley and Scarf \(1974\)](#) and adapted to the school

⁴Given our modeling choice for the outside option, this definition implies that an envyfree allocation is also individually rational since students can always block with their outside option.

choice problem by [Abdulkadiroğlu and Sönmez \(2003\)](#). The mechanism is strategyproof and always produces an efficient allocation based on submitted preferences. However, it may violate priorities.

The student-proposing DA and TTC serve as natural benchmarks to measure the tension between envyfreeness and efficiency: student-proposing DA maximizes efficiency among algorithms that produce envyfree allocations and TTC performs well (and under some circumstances best) on envyfreeness, within the class of efficient and strategyproof mechanisms ([Abdulkadiroğlu, Che, Pathak, Roth, and Tercieux, 2020](#); [Doğan and Ehlers, 2022](#)).⁵

Additionally, two common mechanisms used in practice are the school-proposing DA and the Immediate Acceptance algorithm. Unlike its student-proposing counterpart, the school-proposing DA offers seats to students in order of school priorities. It always produces an envyfree allocation but is neither efficient nor strategyproof: in the Nash equilibrium of the school-proposing DA, students may be tempted to misreport their preferences to get a better (still envyfree) allocation ([Roth, 1982](#)).

Like student-proposing DA, the IA mechanism starts with students' preferences but it allocates seats first to first choices (possibly using school priorities if demand is higher than the number of seats), before considering second choices. The IA mechanism is efficient based on submitted preferences but it is not strategyproof: some students may have incentives to submit preferences different from their true preferences in order to get a preferred assignment. [Ergin and Sönmez \(2006\)](#) show that the set of Nash equilibrium outcomes is equal to the set of feasible envyfree allocations.

In the rest of the paper, we assume that students play the weakly dominant strategy of truth telling under both student-proposing DA and TTC and that they play Nash equilibrium strategies under the school-proposing DA and the IA mechanisms. All mechanisms are described formally in the Appendix.

3 The Generalized Mutually Best Pairs condition

We are interested in analyzing the set of markets – conditions on priorities and preferences – for which there is no conflict between envyfreeness and efficiency. Since the DA allocation Pareto dominates all other envyfree allocations, the most direct formal translation of this question is to ask when DA is efficient. A slightly more demanding requirement is to ask when DA yields the same allocation as TTC. This requirement is more demanding because

⁵Specifically, TTC minimizes envy within the class of efficient and strategyproof mechanisms when all schools have unit capacity. This is no longer true with capacities larger than one. [Morrill \(2015\)](#) and [Hakimov and Kesten \(2018\)](#) have proposed variants of [Abdulkadiroğlu and Sönmez \(2003\)](#)'s TTC version for multiple units that reduce envy. Because none of these variants would change our results, we work with the better known [Abdulkadiroğlu and Sönmez \(2003\)](#)'s variant.

DA can be efficient, and yet yield a different allocation than TTC. In that case, TTC is efficient (as always) but not envy-free.

Our condition builds on [Eckhout \(2000\)](#)'s sufficient condition for uniqueness in one-to-one two-sided matching environments. Specifically, in a context where all schools have unit capacity, [Eckhout \(2000\)](#)'s condition comes down to requiring that there exists a reordering of schools and students such that (unit capacity) schools and students are **mutually best pairs**, i.e. for each $i \in S$, $s_i \succ_i s_k$ for all $k > i$ and $i P_{s_i} k$ for all $k > i$.

In an environment with unit capacity, this condition trivially ensures that DA is efficient and leads to the same allocation as TTC. To see this, consider first student 1 and school 1. School 1 is student 1's top choice and since they have the highest priority in that school, they will be assigned to that school for sure under DA.⁶ Next, consider student 2. They may prefer school 1 to school 2 but under the mutually best pairs condition, school 2 is their top or second top (after school 1) preferred school. Hence, since student 2 is also school 2's preferred student (except possibly for student 1 who is already assigned), student 2 will be assigned to school 2 under DA. As the process continues, it is clear that, under DA, student i will be allocated to school s_i . This outcome is also the outcome that arises from TTC by iterated elimination of unit length cycles (school 1 pointing to student 1 and the reverse, etc).

[Eckhout \(2000\)](#) finds that this condition is sufficient for the set of envyfree allocations to be singleton and we have just seen that it is also sufficient for DA to be efficient and for DA to yield the same allocation as TTC. The next example suggests that this condition might be over-restrictive, even in the unit capacity school context.

Example 1. Consider a market with unit-capacity schools. We describe students' preferences by listing, for each student, their acceptable schools in their preferences orders. Similarly for schools' priorities. The preferences and priorities are as follows:

$$\begin{array}{ll}
 i_1 : s_1 \ s_3 \ s_2 & s_1 : i_2 \ i_1 \\
 i_2 : s_2 \ s_1 & s_2 : i_1 \ i_2 \\
 i_3 : s_3 & s_3 : i_1 \ i_3
 \end{array}$$

It is easily checked that DA and TTC yield the same allocation, $(i_1, s_1), (i_2, s_2), (i_3, s_3)$, where (i, s) means that student i is assigned to school s , and that there is a unique envyfree allocation. Yet, this market does not satisfy [Eckhout \(2000\)](#)'s condition because there is no mutually best pair to start with.

Now consider school s_2 in student i_1 's preferences. School s_2 is an irrelevant choice for student i_1 since that student is ranked first by school s_3 and they prefer school s_3 to school

⁶We use "they" as the gender-neutral third person singular pronoun.

s_2 . Hence, student i_1 will never be assigned to school s_2 in any envyfree allocation and we might as well remove them from the priority of school s_2 .

Building on this intuition, we can define for every school market an associated “simplified school market” where students’ preferences are truncated up to the best school they can get for sure (their “safe school”) and priorities are updated accordingly. In the example above this gives, after two additional rounds of truncations:

$$\begin{array}{ll} i_1 : s_1 & s_1 : i_1 \\ i_2 : s_2 & s_2 : i_2 \\ i_3 : s_3 & s_3 : i_3 \end{array}$$

This simplified market satisfies [Beckhout \(2000\)](#)’s mutually best pairs condition. We will formally show in [Lemma 1](#) below that, by construction, it admits the same set of envyfree allocations as the original market. Therefore, we can leverage the mutually best pairs condition to conclude that the original market admits a unique envyfree allocation and that DA is efficient.⁷ \square

Following [Example 1](#), we will consider, from now on, simplified markets where irrelevant choices have been removed from students’ rank order lists through **iterated** truncations at the best “safe school”, and priority lists have been updated accordingly.

Definition 1 (Simplified market associated with \mathcal{E}). Consider any school choice market $\mathcal{E} = (\succ, P, q)$. Its associated simplified market is given by $\mathcal{E}^* = (\succ^*, P^*, q)$ where

1. \succ_i^* is a truncation of \succ_i at the first school s on \succ_i that ranks student i among its top q_s priority students according to P_s^* .
2. P_s^* is a selection of P to students who rank school s according to \succ^* .

Because the simplification process only removes irrelevant choices from students’ rank-order lists, the set of envyfree allocations in the simplified market is the same as in the original market as the next lemma shows.

Lemma 1. *The sets of envyfree allocations in \mathcal{E}^* and its associated simplified version, \mathcal{E} , are the same.*

Proof. We first show that all envyfree allocations in \mathcal{E} are also envyfree in \mathcal{E}^* . Consider the envyfree allocation μ in \mathcal{E} and suppose it is not envyfree in \mathcal{E}^* . This means there exists (i, s) such that $s \succ_i^* \mu(i)$ and $|\mu^{-1}(s)| < q_s$ or $i P_s^* j$ for some $j \in \mu^{-1}(s)$. But this means that

⁷In general, truncation will change the outcome of TTC so the fact that TTC yields the same outcome as DA in the simplified market does not necessarily mean that this property also applies in the original market.

$s \succ_i \mu(i)$ since \succ^* is a truncation of \succ and either $|\mu^{-1}(s)| < q_s$ or $iP_s j$ for some j (since P^* is a selection of P). A contradiction.

Consider now μ , an envyfree allocation in \mathcal{E}^* and suppose it is not envyfree in \mathcal{E} . This means there exists (i, s) such that $s \succ_i \mu(i)$ and $|\mu^{-1}(s)| < q_s$ or $iP_s j$ for some $j \in \mu^{-1}(s)$. But since \succ^* is a truncation of \succ , $s \succ_i^* \mu(i)$ and either $|\mu^{-1}(s)| < q_s$ or $iP_s^* j$ for some $j \in \mu^{-1}(s)$ must hold. A contradiction. \square

Our second example illustrates the added difficulty arising from allowing schools to have multiple seats.

Example 2. Consider the following market with preferences, capacities and priorities as follows:

$$\begin{array}{ll}
 i_1 : s_2 & s_1 \\
 i_2 : s_1 & \\
 i_3 : s_1 & s_2 \\
 i_4 : s_1 & s_2 & s_3
 \end{array}
 \qquad
 \begin{array}{ll}
 s_1 : i_1 & i_2 & i_3 & i_4 & (\text{capacity} = 2) \\
 s_2 : i_2 & i_1 & i_4 & i_3 & (\text{capacity} = 1) \\
 s_3 : i_4 & & & & (\text{capacity} = 1)
 \end{array}$$

Note first that the school choice environment cannot be further simplified: there is no irrelevant choices in students' rank order lists. Furthermore, this market does not satisfy the mutually best pairs condition, even if we divide the schools into mini-schools of unit capacity that inherit the priorities of the original schools as is typically done (see e.g. [Abdulkadiroğlu and Sönmez \(2003\)](#)). Yet, DA is efficient and produces the same allocation as TTC, namely, $(i_1, s_2), (i_2, s_1), (i_3, s_1), (i_4, s_3)$.

The issue here is that only school s_1 's top priority student is considered when checking the mutually best pairs condition, even though school s_1 can admit 2 students.

One way to address this issue is to consider schools' top q students when identifying mutually best pairs. So, in this case, student i_2 and school s_1 are mutually best pairs since student i_2 is one of school s_1 's top 2 students. Once student i_2 is removed from schools' priority lists, student i_1 and school s_2 become mutually best pairs, leaving student i_3 to be mutually best pair with school s_1 , and finally student i_4 and school s_3 are a mutually best pair. \square

We are now ready to introduce our generalized mutually best pairs condition. We do this in two steps. We first define the Sequential Mutual Best Pairs (MBP) condition, which is the generalisation of Eeckhout's condition to environments with multi-unit capacity:

Definition 2 (Sequential Mutually Best Pairs condition). A market $\mathcal{E} = (\succ, P, q)$ satisfies the Sequential Mutually Best Pairs condition if there is a reordering of students (i_1, i_2, \dots) and an associated list of schools $S, (s_{(1)}, s_{(2)}, \dots)$, where $s_{(i)} \in S$ stands for the

school associated with student i and the same school does not appear more times than its capacity, such that:

1. $s_{(1)} \succeq_{i_1} s$ for all $s \in S \setminus \{s_{(1)}\}$ and i_1 is among the top $q_{s_{(1)}}$ students in school $s_{(1)}$'s priority list,
2. (for $k > 1$), $s_{(k)} \succeq_{i_k} s$ for all $s \in S^k = \{s \in S : q_s^k > 0\}$, where $q_s^k = q_s - \sum_{l=1}^{k-1} 1_{\{s_{(l)}=s\}}$ is the remaining capacity of school s by the time we reach student i_k , and student i_k is among the top $q_{s_{(k)}}^k$ students in school $s_{(k)}$'s priorities, among students i_k, i_{k+1}, \dots

In words, a market satisfies the Sequential MBP condition if we can *sequentially* match students to the best school in their preference lists for which their priority qualifies them for one of the remaining seats. The sequential Mutually Best Pairs condition describes environments where [Salonen and Salonen \(2018\)](#)'s Iterated Best Match algorithm converges (produces a non-wasteful matching, in their terminology).

To illustrate Definition 2 with Example 2, we can reorder students and schools as follows:

$$\begin{array}{cccc} i_2 & i_1 & i_3 & i_4 \\ s_1 & s_2 & s_1 & s_3 \end{array}$$

Our Generalized Mutually Best Pairs (GMBP) condition imposes the Sequential Mutually Best Pairs condition only on the simplified version of the original market and is thus a relaxation of the sequential MBP condition.

Definition 3 (Generalized Mutually Best Pairs condition). A market $\mathcal{E} = (\succ, P, q)$ satisfies the Generalized Mutually Best Pairs condition if its simplified market \mathcal{E}^* satisfies the Sequential Mutually Best Pairs condition.

While the sequential MBP condition serves to prevent unrealistic choices in the top of students' preferences to derail the MBP condition, the generalized MBP removes uninterested students from the top of schools' priorities. The two actions together help us to focus on the relevant set of possible allocations. We defer the detailed discussion of which school market environments are likely to satisfy the GMBP condition until Section 4. Let us simply note for now that the condition captures a sense of compatibility between students' preferences and school priorities (mutually best pairs), once we account for the feasible choice set of students (the mutually best condition is checked sequentially, on the relevant set of alternatives). We have the following result:

Proposition 1 (No trade-off between efficiency and envyfreeness). *Suppose $\mathcal{E} = (\succ, P, q)$ satisfies the generalized mutually best pairs condition. Then, there is a unique envyfree allocation and it is efficient. It is produced by the student-proposing DA, school-proposing DA and IA.*

Proof. We first establish that there is a unique envyfree allocation in \mathcal{E}^* (by Lemma 1, this implies that there is a unique allocation in \mathcal{E} as well). The proof proceeds by iteration. Consider the reordering of students and schools associated with the application of the GMBP condition. Student 1 is matched with school $s_{(1)}$ at every envyfree allocation, otherwise the student blocks the allocation with $s_{(1)}$. Remove this student from the market. Likewise, student 2 is matched with school $s_{(2)}$ at every envyfree allocation. Otherwise, this student will block the allocation. Indeed, the only case where student 2 may not be assigned to $s_{(2)}$ without blocking the assignment with $s_{(2)}$, is when student 1 is assigned to it. But student 1 is allocated to $s_{(1)}$ at every envyfree allocation.

A similar argument applies to student 3. Student 3 will be assigned to $s_{(3)}$ at every envyfree allocation. Otherwise, they will block the assignment, unless student 1 or 2 is assigned to that school. But as we have just shown, these students are assigned to $s_{(1)}$ and $s_{(2)}$, respectively, at every envyfree allocation.

This argument can be repeated for the rest of the students, establishing that there exists a unique envyfree allocation in the market. Clearly, this allocation is also efficient.

The claim follows from the fact that DA, the school-proposing DA and IA all produce an envyfree allocation in equilibrium (Roth, 1984; Ergin and Sönmez, 2006). \square

Proposition 1 does not claim that DA produces the same allocation as TTC. The reason is that the simplification process removes trading opportunities for TTC. However, if the original market satisfies the GMBP condition, then we can prove the stronger claim that TTC, student-proposing DA, school-proposing DA and IA, all produce the same allocation. In other words, the choice of the algorithm is second order.

Proposition 2 (Irrelevance of the algorithm). *If $\mathcal{E} = (\succ, P, q)$ satisfies the sequential mutually best pairs condition, then TTC, DA, the school-proposing DA and IA yield the same allocation and this allocation is both efficient and envyfree.*

Proof. We simply need to show that TTC yields the same outcome as DA. Note that in the execution of the TTC, whenever a student points to a school, the school has k remaining seats, and the student is among the top k students in the school's priority, they will be assigned to that school. This is precisely what happens in the execution of the TTC when the market satisfies the sequential mutually best pairs condition. Indeed, in the first step student 1 points to $s_{(1)}$ and they are among the top $q_{s_{(1)}}$ students in the school's priority. When we remove student 1 from the market, then student 2 points to $s_{(2)}$ and they are among the top $q_{s_{(2)}}^2$ students in the school's priority. And so on, and so forth, for the rest of the students. \square

A natural question that arises is to what extent GMBP is also necessary. The answer is

negative. There are school choice markets where efficiency and envyfreeness are compatible, yet they do not satisfy GMBP as Example 3 illustrates.

Example 3. Consider the following market with preferences and priorities as follows (all schools have unit capacity):

$$\begin{array}{ll}
 i_1 : s_3 & s_1 \\
 i_2 : s_2 & \\
 i_3 : s_1 & s_2 \quad s_3
 \end{array}
 \qquad
 \begin{array}{ll}
 s_1 : i_1 & i_3 \\
 s_2 : i_2 & i_3 \\
 s_3 : i_3 & i_1
 \end{array}$$

Note first that the school choice environment cannot be further simplified: there is no irrelevant choice in students' preferences. Furthermore, this market does not meet the generalized mutually best pair condition: while i_2 and s_2 are mutually best, the process stops there. Yet, DA is efficient and produces the same outcome as TTC, namely, $(i_1, s_3), (i_2, s_2), (i_3, s_1)$. Interestingly, this example is also one where the set of envyfree allocations is not unique: the school-proposing DA produces $(i_1, s_1), (i_2, s_2), (i_3, s_3)$. \square

4 The GBMP condition in stylized school choice environments

This section discusses, in the context of stylized school choice environments, how our condition compares with existing conditions that guarantee the efficiency of DA. The exercise helps crystallize the relationship among existing conditions and build intuition for the way in which our condition enlarges the set of known environments for which DA is efficient.

To do this, we turn to cardinal representations of preferences and priorities. Let u_{is} denote a cardinal utility representation for student i 's preference for school s , with the convention that $s \succ_i s'$ if and only if $u_{is} > u_{is'}$. Likewise, let π_{is} describe student i 's priority at school s with the convention that $i P_s j$ if and only if $\pi_{is} > \pi_{js}$. Table 2 shows to what extent different school choice environments are covered by existing conditions and our GMBP condition.

Ergin (2002) was the first to investigate when DA is efficient in a school choice context. He identified acyclicity of priorities as a necessary and sufficient condition for DA to be efficient for any profile of preferences.⁸

⁸Acyclicity requires (in addition to a scarcity condition) that there are no three students i, j, k and two schools, s and s' , such that $i P_s j P_s k P_s' i$. Ergin (2002) shows that this is equivalent to requiring, that for every pair of schools and for every student ranked below the sum of the two schools' capacities in one school, that student's position differs at most by one across the two schools' lists of priorities. Ehlers and Erdil (2010) and Erdil and Kumano (2019) extend this result to the case where priorities are coarse and to quotas, respectively.

While his condition does not place any restriction on preferences, it is restrictive. Of all the stylized environments we consider in Table 2, it is only met in school choice environments where all schools use the same priorities, e.g., because seats are allocated on the basis of a test score or on the basis of a single tie-breaking rule.⁹ All other conditions identified in the literature happen to rely on the existence of a mutually best pair condition, either in their definitions or in their proofs.¹⁰ Where the conditions differ is in the domain over which this mutually best pairs condition is required to hold.

The strongest condition requires the mutually best pair condition to hold everywhere. For one-to-one matching environments it comes down to Alcalde (1994)’s α -reducibility condition as adapted by Clark (2006) and, equivalently, to Niederle and Yariv (2009)’s aligned preferences. In many-to-one environments, it is equivalent to Salonen and Salonen (2018)’s single-peaked preferences.

Environments where school priorities and student preferences depend on the same student-school match quality (e.g. distance, religion, academic inclination) satisfy this condition. Such preferences and priorities take the form $u_{is} = \pi_{is} = d_{is}$, where d_{is} represents the student-school match quality. The condition also holds if, in addition, schools value student quality or students value school quality as long as they value these characteristics identically (no heterogeneity). In this case, preferences take the form $u_{is} = d_{is} + v_s$, and priorities take the form $\pi_{is} = d_{is} + g_i$. To verify that this environment satisfies the MBP condition everywhere, consider a $I \times N$ matrix, with elements $\phi_{is} = d_{is} + v_s + g_i$. The matrix rows represent students’ preferences. Indeed, $\phi_{is} > \phi_{is'}$ if and only if $u_{is} = d_{is} + v_s > u_{is'} = d_{is'} + v_{s'}$ (the g_i term drops out). Likewise, the matrix columns represent school priorities: $\phi_{is} > \phi_{i's}$ if and only if $\pi_{is} = d_{is} + g_i > \pi_{i's} = d_{i's} + g_{i'}$ (the v_s term drops out). Assuming no identical student-quality match quality, a property of this matrix is that there is always an element that is maximal, for every submatrix. In other words, the MBP condition holds everywhere.

Note that the richer and more realistic environments where students are heterogeneous in their valuation of quality (e.g. Abdulkadiroğlu, Agarwal, and Pathak (2017)) do not meet any of the conditions generically. So, for example, if $u_{is} = d_{is} + \alpha_i v_s$ and priorities are distance-based $\pi_{is} = d_{is}$, we can easily construct a situation where a student prefers the school further away because they value quality more, blocking the construction of a mutually best pair.

Another environment where the MBP condition holds everywhere is when students prefer the school where their sibling, if any, goes, and where schools prioritize students with siblings

⁹Such priorities are common in university admissions, e.g. in China (Chen and Kesten, 2017), Germany (Grenet, He, and Kübler, 2021), Spain (Arenas and Calsamiglia, 2020) and Turkey (Arslan, 2019), but also in secondary education, e.g. in Mexico (Chen and Pereyra, 2019), Romania (Pop-Eleches and Urquiola, 2013) and Singapore (Teo, Sethuraman, and Tan, 2001) or for exam schools in many countries.

¹⁰This includes Alcalde (1994), Clark (2006), Eeckhout (2000), Legros and Newman (2010), Niederle and Yariv (2009), Reny (2021), Rong, Tang, and Zhang (2020), and Salonen and Salonen (2018).

and use a single tie-breaking rule for all others. To see this, consider any subset of schools and students. If there is a student with a sibling in one the schools, then they will be mutually best pairs. If no student has a sibling, then all schools will be ranking students in exactly the same way, and one of these schools (at least) will be the preferred choice of one of the students (we are back in the environment of the first row of Table 2).

Table 2: Comparing conditions across school choice environments

Preferences	Priorities	Application	Ergin acyclicity	MBP everywhere	Sequential MBP	GMBP
Any	Identical priorities (e.g. based on test scores)	University and high school admissions in several countries	✓	✓	✓	✓
$u_{is} = d_{is}$	$\pi_{is} = d_{is}$	Match-quality priorities and preferences		✓	✓	✓
$u_{is} = d_{is} + v_s$	$\pi_{is} = d_{is} + g_i$	Match-quality + common priorities and preferences		✓	✓	✓
prefers school with sibling, no restriction otherwise	$\pi_{is} = \mathbf{1}_{(i=\text{sibling})} + \varepsilon_i$ ($\varepsilon_i \in [0, 1]$)	Sibling priorities, single tie-breaking rule for rest		✓	✓	✓
Prefers one of the catchment school, with some exception	$\pi_{is} = \mathbf{1}_{i=\text{in catchment}} + \varepsilon_i$ ($\varepsilon_i \in [0, 1]$)	Guaranteed admission in catchment area school, single tie-breaking otherwise		only in case of run-away / attraction catchment areas	possible	more likely

The second strongest condition requires the mutually best pair condition to hold sequentially. In one-to-one settings, this is [Eeckhout \(2000\)](#)'s condition. In one-to-many settings, it corresponds to environments for which [Salonen and Salonen \(2018\)](#)'s Iterated Best Match process converges to a non-wasteful matching. The advantage of such condition relative to the MBP condition everywhere is that it ignores unrealistic preferences that some students may have and that prevent the MBP condition to hold.

The fifth row of Table 2 provides an illustration of the type of environments this condition allows. In the example, priority is given to students living in the catchment area of the schools and these students have essentially guaranteed access to at least one of them (in case of excess demand, a single tie-breaking rule is used). Most students prefer to go to a school in their catchment area but some prefer a school outside of their catchment area.

We can first check that if out-of-catchment-area choices are asymmetric, in the sense that the catchment areas that students select for their out-of-catchment-area first choices are not the same as the catchment areas that students seek to leave (e.g. there are popular

catchment areas and run-away catchment areas), then the market satisfies the MBP condition everywhere. To see this, take any two students and the schools which they list first. If at least one of them makes a within-catchment-area choice, then at least one of them has top priority at their preferred school. If both make an out-of-catchment-area choice, the same priority ordering applies to them because of the single tie-breaking rule and therefore at least one of them is part of a mutually best pair.

As soon as there is one student in one catchment area (say A) who prefers a school in another catchment area (say B), and the reverse, the “MBP condition everywhere” fails (just take the submarket made of these two students and the two schools they prefer).

However, it may still satisfy the sequential MBP. Indeed, through the sequential match of mutually best pairs, some schools, which the students making out-of-catchment-area choices listed, will reach capacity and will therefore exit the consideration set for these students. The sequential process will converge to a reduced set of students making out-of-catchment-area choices. This school choice market will satisfy the sequential MBP if there are no cross-catchment-area choices.

We can use this example to illustrate what the generalized MBP condition allows for, on top of the sequential MBP condition. Consider a student (say student a) in catchment area A who lists a school in catchment area B as first choice. That student will be ranked below every student in catchment area B who listed that school as acceptable, even if they are not interested in that school because they are sure to get another school in catchment area B if they ask for it (safe school). These students are hampering the formation of mutually best pairs and the generalized MBP, by truncating students’ preferences at their safe school and updating priority lists accordingly, will correct this and make the formation of mutually best pairs easier.

We are unaware of generic environments where the generalized MBP condition is always met but the sequential MBP condition is not. However, this example illustrates how GMBP enables the identification of an expanded set of school choice environments where DA is efficient. The next section quantifies this statement in select stylized environments.

5 Quantification

Having shown qualitatively how the GMBP condition expands on existing conditions, we turn in this section to a quantitative assessment of the GMBP condition. Specifically, we generate a large number of school choice markets and check (1) to what extent DA is efficient in those markets, and (2) to what extent these markets satisfy the sequential and generalized MBP conditions.

Building on the recent empirical literature in school choice that estimates preferences

for schools (e.g. Hastings, Kane, and Staiger (2009); Abdulkadiroğlu, Agarwal, and Pathak (2017); Calsamiglia, Fu, and Güell (2020); Fack, Grenet, and He (2019); Pathak and Shi (2021)), we assume that cardinal utilities underlying student preferences take the following form:

$$u_{is} = \lambda(\delta d_{is} + (1 - \delta)v_s) + (1 - \lambda)\varepsilon_{is},$$

where, as before, d_{is} stands for the student-school match quality (distance, religion, academic inclination), v_s captures characteristics of the school that are valued equally by all students, and ε_{is} is an idiosyncratic component capturing individual taste.¹¹

The first term, $\delta d_{is} + (1 - \delta)v_s$, captures the structural part of preferences, driven by match quality and school characteristics. When $\delta = 1$, preferences are driven by match quality, which can be understood as the result of horizontal differentiation between schools, whereas when $\delta = 0$, schools are vertically differentiated in the eyes of students. The parameter λ determines to what extent idiosyncratic preferences matter. When $\lambda = 1$, there are no idiosyncratic factors beyond match quality. When $\lambda = 0$, preferences are entirely idiosyncratic and independent of priorities.

We consider the following priority structure for schools:

$$\pi_{is} = \alpha(\beta d_{is} + (1 - \beta)g_i) + (1 - \alpha)\eta_{is},$$

where g_i captures priorities based on student characteristics and single tie-breaking, and η_{is} captures residual priorities based on idiosyncratic factors or multiple tie-breaking. The parameter α measures the degree of structure on priorities, whereas the parameter β measures to what extent priorities are driven by match quality rather than student characteristics valued equally by all schools (e.g. grades).

We carry out simulations as follows. For each school market environment—characterized by the vector of parameters $(\lambda, \delta, \alpha, \beta)$, the number of students n , the number of schools m , and the school capacities q (we assume all schools have identical capacities), we draw 1,000 realizations of the vector of variables $(d_{is}, v_s, \varepsilon_{is}, g_i, \eta_{is})$, where all variables are independently drawn from the uniform distribution on $[0, 1]$. This generates 1,000 market realizations for every specific school market environment. We set $n = 1,000$, $m = 50$ and $q = 20$.¹²

Table 3 reports the percentage of market realizations where DA was found to be efficient and which satisfied the sequential MBP and generalized MBP, respectively. We fix the values for λ and α , and average over all market realizations and values for δ and β , where δ and β vary from 0 to 1 in 0.05 increments. The table presents the results for λ taking the values

¹¹The main difference with empirically estimated preferences is that our coefficients on the match quality and school characteristics are assumed to be common across students.

¹²We played with different market sizes and school sizes, with no qualitative changes in the results.

Table 3: Percentage of markets where DA is efficient, sequential MBP and GMBP are satisfied

Preferences	Priorities	(1)	(2)	(3)
		DA is efficient	Sequential MBP	GMBP
$\lambda = 1$	$\alpha = 1$	41.15	36.79	40.86
	$\alpha = 0.95$	21.76	12.71	20.54
$\lambda = 0.75$	$\alpha = 1$	14.95	8.54	14.07
	$\alpha = 0.95$	3.83	0	2.93
$\lambda = 0.5$	$\alpha = 1$	7.71	4.80	6.54
	$\alpha = 0.95$	1.15	0	0.37
$\lambda = 0.25$	$\alpha = 1$	6.08	4.76	5.1
	$\alpha = 0.95$	0.88	0	0.14
$\lambda = 0$	$\alpha = 1$	5.17	4.76	4.80
	$\alpha = 0.95$	0.76	0	0

Notes: The numbers indicate the percentage of school choice markets for which DA is efficient (column (1)) and that satisfy, respectively, the sequential MBP condition (column (2)) and the GMBP condition (column (3)). Percentages are computed first on the basis of 1,000 independent draws of variables $(d_{is}, v_s, \varepsilon_{is}, g_i, \eta_{is})$ from the uniform distribution $[0, 1]$, for a given value of δ and β , then averaged over the δ and β parameters taking values from 0 to 1 in 0.05 increments.

of 1, 0.75, 0.5, 0.25, 0, and α taking the values of 1 and 0.95.¹³


Three observations stand out from the Table. First, the ability of DA to generate efficient allocations vary strongly across environments and, in particular, drops sharply with students' idiosyncratic preferences (low λ) and schools' idiosyncratic priorities (low α), beyond match quality. This is when the trade-off between respecting preferences and respecting priorities is largest.

Second, the GMBP condition outperforms the sequential MBP condition in identifying the environments for which DA is efficient. The gap is especially big as soon as we allow for idiosyncratic preferences ($\lambda < 1$) or priorities ($\alpha < 1$). Recall from Example 1 that the GMBP condition outperforms the sequential MBP condition in situation where the market does not have mutually best pairs to start with, but by removing the irrelevant schools for students, new mutually best pairs emerge. This situation is more likely to occur when schools priorities and students preferences go in opposite direction: schools prefer more students who prefer them the least. Smaller values for λ or α allow for this possibility.


The third insight is that, for a given value of λ , the GMBP condition is able to identify a large proportion of the school choice environments for which DA is efficient especially when

¹³For α taking smaller values such as 0.75, 0.5, 0.25 and 0, DA is inefficient in all simulated markets. Neither sequential nor generalized MBP is satisfied.


Figure 1: Share of markets where DA is efficient (green), the sequential MBP is satisfied (blue), and GMBP is satisfied (red), as a function of δ and β




(a) $\lambda = 1, \alpha = 1$



(b) $\lambda = 1, \alpha = 0.75$




(c) $\lambda = 1, \alpha = 0.5$



(d) $\lambda = 1, \alpha = 0.25$

Figure 1: Share of markets where DA is efficient (green), the sequential MBP is satisfied (blue), and GMBP is satisfied (red), as a function of δ and β (continued)



(e) $\lambda = 1, \alpha = 0$

Notes: Simulations for a school market with 1,000 students, 50 schools and school capacities of 20. Shares are averages over 1,000 independent draws of variables $(d_{is}, v_s, \epsilon_{is}, g_i, \eta_{is})$, for a given value of the $\lambda, \alpha, \beta,$ and δ parameters.

α takes high values. Figure 1 provides a closer look at this property. Specifically, Figure 1 maps, for specific values of $(\lambda, \alpha, \delta, \beta)$, the proportion of markets, taken over the 1,000 draws of $(d_{is}, v_s, \epsilon_{is}, g_i, \eta_{is})$, where DA is efficient and the sequential and generalized MBP conditions are satisfied. Confirming the results in Table 3, we see that the proportion of markets such that DA is efficient is highest when λ and α are equal to 1 (no idiosyncratic components of preferences and priorities beyond match quality). Low values for δ (i.e. strong vertical differentiation of schools) and high values for β (i.e. greater emphasis on match quality in school priorities) fosters the efficiency of DA. A higher value of β allows school priorities to “align” with preferences through the match quality, even when students value more factors such as school quality than proximity in distance, and the GMBP condition is capable of identifying most of these markets when DA is efficient.

6 Discussion

When is there a trade-off between efficiency and envyfreeness or, equivalently, when is there a trade-off between preferences and priorities? Our results confirm the conjecture put forward by Pathak (2016) according to which “correlation between preferences and priorities induced by proximity may, in turn, result in less scope for Pareto-improving trades across priority

groups that involve situations of justified envy. This pattern may then result in a small degree of inefficiency in DA”, and clarify to which extent the conjecture holds and to what extent it generalizes. Specifically, our Generalized Mutually Best Pairs condition maximally captures the set of environments where priorities and preferences are sufficiently congruent that DA is efficient, and there is no trade-off between efficiency and envyfreeness.

Our results shed light on the empirical evidence presented in Table 1. A small trade-off was found in (primarily) elementary school markets with priorities to siblings, staff, and some measure of distance (Boston, Ghent, New Orleans). The elementary school level is exactly the education level where one would also expect parents to place greater emphasis on proximity or selecting the same school as the older sibling. In other words, markets where preferences and priorities are congruent. Such markets are well captured by the functional forms in rows 2 and 3 of Table 2 and high values of λ and α in Table 3.

On the other hand, the secondary and high school markets of Budapest and New York City, respectively, are characterized by a higher level of idiosyncratic school-specific priorities and, presumably, a higher level of horizontal differentiation across schools (different specialisation tracks). Our numerical results suggest that DA is less likely to be efficient in such markets. The high level of justified envy found in the data suggests that the trade-off between envyfreeness and efficiency might indeed be big in those school markets.

In our simulations, we found that preference and priority congruence, as captured by the Generalized Mutually Best Pairs condition, covers a large fraction of the environments for which DA is efficient. We nevertheless know from Example 3 that there can be situations where preferences and priorities are not congruent and yet, DA is efficient. In Example 3, this happened because, while preferences and priorities conflicted, students had sufficiently different preferences that they could nevertheless get their first choice. Priorities were toothless. This is reminiscent of Che and Tercieux (2019)’s finding that the main source of the efficiency – envyfreeness trade-off (in their setting) is the excess competition for seats. It is also perhaps not surprising that Example 3 also happens to be an example where the set of envyfree allocations is not a singleton.

What lessons can policy-makers draw from our analysis? A first general lesson is to understand their markets – what drives student preferences – and assess to what extent school priorities are likely to be congruent with those, for the parts of the market where there is excess demand. If this is the case, the choice of the algorithm is likely to be second order. If not, extensive evaluation of different designs might be useful. A second lesson is that they can use their discretionary power, when available, to increase the probability that their markets meet the GMBP condition. An obvious example is the choice of the tie-breaking rule when priorities are weak. Our results here echo Ashlagi and Nikzad (2020)’s recommendation that popular schools use a single tie-breaking rule: fostering preferences

and priority congruence is especially valuable for schools with excess demand.

References

- ABDULKADIROĞLU, A., N. AGARWAL, AND P. A. PATHAK (2017): “The welfare effects of coordinated assignment: Evidence from the New York City high school match,” American Economic Review, 107(12), 3635–89.
- ABDULKADIROĞLU, A., Y.-K. CHE, P. A. PATHAK, A. E. ROTH, AND O. TERCIEUX (2020): “Efficiency, justified envy, and incentives in priority-based matching,” American Economic Review: Insights, 2(4), 425–42.
- ABDULKADIROĞLU, A., P. PATHAK, A. ROTH, AND T. SÖNMEZ (2006): “Changing the Boston School Mechanism: Strategyproofness as Equal Access,” Discussion paper, mimeo.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2009): “Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match,” American Economic Review, 99(5), 1954–78.
- ABDULKADIROĞLU, A., AND T. SÖNMEZ (2003): “School choice: A mechanism design approach,” American economic review, 93(3), 729–747.
- ALCALDE, J. (1994): “Exchange-proofness or divorce-proofness? Stability in one-sided matching markets,” Economic design, 1(1), 275–287.
- ARENAS, A., AND C. CALSAMIGLIA (2020): “The Design of University Entrance Exams and its Implications for Gender Gaps,” Unpublished Manuscript.
- ARSLAN, H. A. (2019): “An empirical analysis of college admissions with endogenous entrance exam scores,” Available at SSRN 3346459.
- ASHLAGI, I., AND A. NIKZAD (2020): “What matters in school choice tie-breaking? How competition guides design,” Journal of Economic Theory, 190, 105120.
- BALINSKI, M., AND T. SÖNMEZ (1999): “A tale of two mechanisms: student placement,” Journal of Economic theory, 84(1), 73–94.
- BIRÓ, P. (2012): “Matching practices for secondary schools–Hungary,” MiP country profile, 6.
- CALSAMIGLIA, C., C. FU, AND M. GÜELL (2020): “Structural estimation of a model of school choices: The boston mechanism versus its alternatives,” Journal of Political Economy, 128(2), 642–680.

- CHE, Y.-K., AND O. TERCIEUX (2019): “Efficiency and stability in large matching markets,” Journal of Political Economy, 127(5), 2301–2342.
- CHEN, L., AND J. S. PEREYRA (2019): “Self-selection in school choice,” Games and Economic Behavior, 117, 59–81.
- CHEN, Y., AND O. KESTEN (2017): “Chinese college admissions and school choice reforms: A theoretical analysis,” Journal of Political Economy, 125(1), 99–139.
- CLARK, S. (2006): “The uniqueness of stable matchings,” Contributions in Theoretical Economics, 6(1).
- DE HAAN, M., P. GAUTIER, H. OOSTERBEEK, AND B. V. D. KLAAUW (forthcoming): “The performance of school assignment mechanisms in practice,” Journal of Political Economy.
- DOĞAN, B., AND L. EHLERS (2022): “Robust minimal instability of the top trading cycles mechanism,” American Economic Journal: Microeconomics, 14(4), 556–82.
- ECKHOUT, J. (2000): “On the uniqueness of stable marriage matchings,” Economics Letters, 69(1), 1–8.
- EHLERS, L., AND A. ERDIL (2010): “Efficient assignment respecting priorities,” Journal of Economic Theory, 145(3), 1269–1282.
- ERDIL, A., AND T. KUMANO (2019): “Efficiency and stability under substitutable priorities with ties,” Journal of Economic Theory, 184, 104950.
- ERGIN, H., AND T. SÖNMEZ (2006): “Games of school choice under the Boston mechanism,” Journal of public Economics, 90(1-2), 215–237.
- ERGIN, H. I. (2002): “Efficient resource allocation on the basis of priorities,” Econometrica, 70(6), 2489–2497.
- FACK, G., J. GRENET, AND Y. HE (2019): “Beyond truth-telling: Preference estimation with centralized school choice and college admissions,” American Economic Review, 109(4), 1486–1529.
- GALE, D., AND L. S. SHAPLEY (1962): “College admissions and the stability of marriage,” The American Mathematical Monthly, 69(1), 9–15.
- GRENET, J., Y. HE, AND D. KÜBLER (2021): “Decentralizing centralized matching markets: Implications from early offers in university admissions,” arXiv preprint arXiv:2107.01532.

- GUTIN, G. Z., P. R. NEARY, AND A. YEO (2021): “Unique Stable Matchings,” .
- HAKIMOV, R., AND O. KESTEN (2018): “The equitable top trading cycles mechanism for school choice,” International Economic Review, 59(4), 2219–2258.
- HASTINGS, J., T. J. KANE, AND D. O. STAIGER (2009): “Heterogeneous preferences and the efficacy of public school choice,” NBER Working Paper, 2145, 1–46.
- HEO, E. J. (2019): “Preference profiles for efficiency, fairness, and consistency in school choice problems,” International Journal of Game Theory, 48(1), 243–266.
- ISHIDA, W. (2019): “On two competing mechanisms for priority-based allocation problems with general weak priorities,” Available at SSRN 3098685.
- KESTEN, O. (2006): “On two competing mechanisms for priority-based allocation problems,” Journal of Economic Theory, 127(1), 155–171.
- KNUTH, D. E. (1997): Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms, vol. 10. American Mathematical Soc.
- LEE, S., AND L. YARIV (2014): “On the efficiency of stable matchings in large markets,” .
- LEGROS, P., AND A. NEWMAN (2010): “Co-ranking mates: Assortative matching in marriage markets,” Economics Letters, 106(3), 177–179.
- MORRILL, T. (2015): “Two simple variations of top trading cycles,” Economic Theory, 60(1), 123–140.
- NIEDERLE, M., AND L. YARIV (2009): “Decentralized matching with aligned preferences,” Discussion paper, National Bureau of Economic Research.
- ORTEGA, J., AND T. KLEIN (2022): “Improving Efficiency and Equality in School Choice,” QMS Research Paper, 2.
- PATHAK, P. A. (2016): “What really matters in designing school choice mechanisms,” .
- (2017): “What really matters in designing school choice mechanisms,” Advances in Economics and Econometrics, 1, 176–214.
- PATHAK, P. A., AND P. SHI (2021): “How well do structural demand models work? Counterfactual predictions in school choice,” Journal of Econometrics, 222(1), 161–195.

- POP-ELECHES, C., AND M. URQUIOLA (2013): “Going to a better school: Effects and behavioral responses,” American Economic Review, 103(4), 1289–1324.
- RENY, P. J. (2021): “A simple sufficient condition for a unique and student-efficient stable matching in the college admissions problem,” Economic Theory Bulletin, pp. 1–3.
- ROMERO-MEDINA, A., AND M. TRIOSSI (2013): “Acyclicity and singleton cores in matching markets,” Economics Letters, 118(1), 237–239.
- RONG, K., Q. TANG, AND Y. ZHANG (2020): “On stable and efficient mechanisms for priority-based allocation problems,” Journal of Economic Theory, 187, 105014.
- ROTH, A. E. (1982): “The economics of matching: Stability and incentives,” Mathematics of operations research, 7(4), 617–628.
- (1984): “Misrepresentation and stability in the marriage problem,” Journal of Economic theory, 34(2), 383–387.
- SALONEN, H., AND M. A. SALONEN (2018): “Mutually best matches,” Mathematical Social Sciences, 91, 42–50.
- SHAPLEY, L., AND H. SCARF (1974): “On cores and indivisibility,” Journal of mathematical economics, 1(1), 23–37.
- TEO, C.-P., J. SETHURAMAN, AND W.-P. TAN (2001): “Gale-shapley stable marriage problem revisited: Strategic issues and applications,” Management Science, 47(9), 1252–1267.

A Appendix

A.1 Description of the mechanisms

Student-proposing Deferred Acceptance (DA)

Step 1: Each student i proposes to the best school according to \succ_i . Each school s provisionally accepts the q_s -highest ranked students, according to P_s , among those students that have proposed to s , and rejects the others.

Step k : Each student i , who has not been previously accepted, proposes to the best school according to \succ_i , among those schools that have not yet rejected i . Each school s provisionally accepts the q_s -highest ranked students, according to P_s , among those students that have proposed to s along steps 1 to $k + 1$, and rejects the others.

The algorithm terminates at the step where no rejections are made and provisional acceptances become definitive by matching each school s to the set of students provisionally accepted at this step.

School-proposing Deferred Acceptance (DA)

Step 1: Each school s proposes to the q_s -highest ranked students according to P_s . Each student i provisionally accepts the best school according to \succ_i , among those schools that have proposed to i , and rejects the others.

Step k : Each school s , which has been previously rejected, proposes to the next highest priority students according to P_s up to capacity q_s , among those students that have not yet rejected s . Each student s provisionally accepts the best school, according to \succ_i , among those schools that have proposed to i along steps 1 to $k + 1$, and rejects the others.

The algorithm terminates at the step where no rejections are made and provisional acceptances become definitive by matching each school s to the set of students provisionally accepted at this step

Top Trading Cycles (TTC)

Step 1: Each student i points to the best school according to \succ_i . If no school is acceptable, i points to s_{m+1} and is removed from the problem. Each school s points to the best student according to P_s , and s_{m+1} points to all students.

Since the sets of students and schools are finite, there exists at least one cycle which is of the form $i_1 \rightarrow s_1 \rightarrow \dots \rightarrow i_K \rightarrow s_K \rightarrow i_1$ or $i \leftrightarrow s_{m+1}$, where $x \rightarrow y$ means “ x points to y ”.¹⁴ Each student i in a cycle is matched to the school s that they point to (if $i \rightarrow s$), in which case i and a seat in s are removed from the problem. If i points to s_{m+1} . They remain unmatched and i is removed from the problem.

Step k . Each remaining student i points to the best school according to \succ_s , among the schools that still have empty seats. Each school s with an empty seat, points to the best student, according to P_s , among the remaining students, and s_{m+1} points to all of these students. There is at least one cycle. Each student i in a cycle of the form $i_1 \rightarrow s_1 \rightarrow \dots \rightarrow i_K \rightarrow s_K \rightarrow i_1$ is matched to the school s that they point to, and i and a seat in s are removed from the problem. If i points to s_{m+1} . One remains unmatched and i is removed from the problem.

The algorithm terminates when each student i is either matched to a school or to s_{m+1} .

¹⁴There may be many cycles, although each student and school $s \neq s_{m+1}$ can be part of at most one cycle.

Immediate Acceptance (IA)

Step 1: Each student i proposes to the best school according to \succ_i . Each school s accepts the q_s -highest ranked students, according to P_s , among those students that have proposed to s , and rejects the others. Accepted students are definitive matched to the school. Schools' capacities are reduced by the number of students accepted.

Step k : Each student i , who has not been previously accepted, proposes to the best school according to \succ_i , among those available schools that have not yet rejected i . Each school s accepts the q_s -highest ranked students, according to P_s , among those students that have proposed to s , and rejects the others. Accepted students are definitive assigned to the school. Schools' capacities are reduced by the number of students accepted.

The algorithm terminates at the step where no rejections are made.