

# A Unified Theory of School Choice

Peter Doe

April 12, 2025

## Abstract

In school choice, policymakers consolidate a district’s objectives for a school into a priority ordering over students. They then face a trade-off between respecting these priorities and assigning students to more-preferred schools. However, because priorities are the amalgamation of multiple policy goals, some may be more flexible than others. This paper introduces a model that distinguishes between two types of priority: a between-group priority that ranks groups of students and must be respected, and a within-group priority for efficiently allocating seats within each group. The solution I introduce, the unified core, integrates both types. I provide a two-stage algorithm, the DA-TTC, that implements the unified core and generalizes both the Deferred Acceptance and Top Trading Cycles algorithms. This approach provides a method for improving efficiency in school choice while honoring policymakers’ objectives.

## 1 Introduction

School choice programs<sup>1</sup> allocate school seats using *priorities* (such as sibling status or proximity): when a school is oversubscribed, these priorities determine which students are admitted. The interpretation of these priorities is fundamental to the program’s design, and there are two leading interpretations. The stronger interpretation posits that priorities are the right to attend the school ahead of those with lower priorities. In the weaker interpretation, priorities reflect “better opportunities” to attend a school, all else equal (Abdulkadiroğlu and Sönmez, 2003).

These interpretations underpin the central trade-off in school choice between fairness and efficiency. *Fairness* means no student is denied admission to a school in favor of a lower-priority student; *efficiency* means students are assigned to schools in such a way that no student can be

---

<sup>1</sup>What economists refer to as *school choice*—the ability to choose between public schools—is frequently called “open enrollment” in the media. Although, in common parlance, “school choice” refers to voucher or tax credit programs for defraying the costs of private schooling or homeschooling, I will use the standard language in economics.

improved without harming another student. By requiring the stronger interpretation, the designer arrives at a fair match because no student prefers a school that a lower-priority student is assigned to. However, allowing for the weaker interpretation shifts the set of permissible assignments to include an efficient match. The main dilemma is that no algorithm delivers a match that is both fair and efficient.

Despite the appeal of efficiency, policymakers consistently choose fairness.<sup>2</sup> Empirically, the cost is substantial: for instance, [Abdulkadiroğlu, Pathak and Roth \(2009\)](#) finds that the assignments of 4300 eighth-grade students in New York City could have been improved without harming any other students. In response to this, recent research focuses on weaker versions of fairness to allow for efficiency gains.

However, an examination of why policymakers side with fairness reveals that they are frequently concerned with the violations of some priorities but not with others. For instance, a priority derived from a high test score may have a stronger interpretation than a priority assigned based on proximity to the school. Because economists have not considered how policymakers may treat some priorities differently than others, this omission has precluded a promising approach to reconciling fairness and efficiency. The key to understanding how policymakers view priorities lies in the two-step process that is typically used to construct priorities.

In the first step, the policymaker identifies a set of student characteristics to prioritize and uses these characteristics to partition students into priority groups.<sup>3</sup> Then, the policymaker assigns a priority over the priority groups that reflects the policymaker's objectives for the school; this is the *between-group* priority. In the second step, within each group, another *within-group* priority is provided to break ties, which could depend upon other student characteristics or a random lottery. A student's priority at a school is lexicographically determined first by her between-group priority and then by her within-group priority.

Policymakers understand that because priorities arise from different sources, they require different interpretations. Allocating a seat to a student in a lower priority group ahead of a student in a higher priority group may be unacceptable, but assignments within groups may be more subtle. The prominent case of Boston Public Schools (BPS)—the first district to carefully consider its assignment mechanism—illustrates this difference.

BPS created two priority groups: students with a sibling attending the school and all other

---

<sup>2</sup>New Orleans Recovery School District and the Common App is the only example I know of that sought an efficient outcome in some rounds of the assignment process, but this was abandoned after the first year.

<sup>3</sup>For example, these characteristics frequently include whether the student has a sibling attending the school, whether the student is within a particular geographic region, the student's test scores, etc., and a priority group consists of the students with the same set of characteristics.

students, with the former given higher between-group priority than the latter. Within each priority group, BPS gave higher within-group priority to students residing in the school's walk zone, and broke any remaining ties with a random lottery. When comparing the efficient algorithm with the fair algorithm, Superintendent Payzant implies that the between-priority derived from having a sibling present should be treated differently from the within-priorities:

[The efficient system] presents the opportunity for the priority of one student at a given school to be “traded” for the priority of a student at another school, [...] There may be advantages to this approach, [...] It may be argued, however, that certain priorities – e.g., sibling priority – apply only to students for particular schools and should not be traded away. (Abdulkadiroğlu et al., 2006)

The implication from his silence concerning walk-zone priority is that, in some cases, walk zone priorities *could* be traded away. Sibling priority has the stronger interpretation while walk zone priority has the weaker interpretation of a better opportunity to attend the school.<sup>4</sup>

Despite the different interpretation of priorities within these two steps, economists usually remain agnostic about the source of the priority. The nuances between priority groups are dismissed, and all priorities are treated equally. The existing options afforded to policymakers do not allow them to express these nuances. This leaves policymakers in a quandary because a more efficient match is better than not, but violating the stronger priorities may be impermissible. When BPS faced this decision, they erred on the side of respecting priorities and implemented the fair mechanism even though only a subset of their priorities required the stronger interpretation.

In this paper, I introduce a model that explicitly distinguishes between two types of priority—between-group and within-group—and I propose the unified core to connect the two. In my model, each school's priority consists of two layers. The first layer is the between-group priority, a weak preference that specifies which students belong to which priority group and how those groups are ordered. The second layer is the within-group priority, a strict preference that refines the between-group priority.<sup>5</sup> Accordingly, only the between-group priority is interpreted as a

<sup>4</sup>Readers familiar with this story will note that my exposition of the priority groups is slightly different than that of Abdulkadiroğlu et al. (2006). Formally, BPS designated five priority groups: continuing students, sibling-walk, sibling, walk, and all others. However, as Payzant's quote illustrates, sibling priority and walk zone priority are normatively different. Conveniently, the order on the priority groups allows for walk zone priority to be viewed as a tiebreaker within the sibling priority group. If BPS ranked these groups differently, my model would not apply directly. For ease of exposition, I exclude the continuing students from the model; continuing students can be easily added to the model by including an additional priority group.

<sup>5</sup>In most applications, the within-group priority consists of a random lottery. However, I refrain from referring to it as a *tiebreaker* because the within-group priority may include substantive components, such as geographic location (as in

right to attend the school ahead of lower-priority students, and it cannot be violated. On the other hand, the within-group priority has the milder interpretation of an opportunity to attend the school; it is a means of allocating the school seat, but it is not inviolable. The interpretation of the within-group priority is weak enough to guarantee efficiency within a group while still using the priority to allocate the seat.

The unified core connects the two cores that underlie the fair and efficient algorithms, and is implemented by carefully combining the algorithms that implement the fair and efficient cores. The core underlying the fair algorithm is the *fair* core of [Gale and Shapley \(1962\)](#) (henceforth GS), which allows a student to claim a seat at a school if she has a higher priority than a student currently there. This stems from the strong interpretation of priorities and treats the school as a strategic player and the priority as the school's preference. The result is a fair match because any unfair match can be blocked by the student whose priority is violated.

The core underlying the efficient algorithm is from [Rong, Tang and Zhang \(2022\)](#) (henceforth RTZ),<sup>6</sup> who develop their *efficient* core to provide a foundation for the use of several efficient algorithms in school choice. RTZ's efficient core allows a student to take possession of a school provided no higher-priority students can veto that action. In effect, the efficient core allows for a form of tradable priorities, where a low-priority student can receive a priority from another student so long as no higher-priority students are harmed. As is paralleled in many economic models, the power to trade freely (in this case, trade school seats) guarantees that an efficient allocation is reached.

Whether policymakers prefer the outcomes of the fair core or the efficient core ultimately depends on the source of the students' priorities. If the priority reflects different priority groups, then the fair core is appealing; if it reflects within-group tiebreakers such as a random lottery, the efficient core allows more students to attend the schools that they prefer. However, neither the fair core nor the efficient core allows for combinations of the two priorities.

The unified core resolves this tension and connects the two cores by allowing the policymaker to explicitly designate two kinds of priorities. By treating the between-group priorities as in the fair core while the within-group priorities as in the efficient core, the unified core generalizes both cores. Students can always claim a seat at a school if they have a higher between-group priority than a currently assigned student, as in GS. Within-group priorities, however, may be traded amongst students within the same priority class, as in RTZ. This guarantees a form of "within-group" efficiency. The way I implement this mirrors RTZ: a student can take a seat provided no

---

the case of BPS). The emphasis, however, is that these priorities have a weaker interpretation.

<sup>6</sup>It is worth mentioning here that RTZ's core identifies precisely the *just* assignments of [Morrill \(2015\)](#).

within-group interrupters in her priority group can block her.

The main challenge in crafting a solution for this model is integrating the differing normative implications of the two priorities. Unlike the model of [Erdil and Ergin \(2008\)](#) (henceforth EE),<sup>7</sup> the within-group priorities are not merely random tiebreakers. Within-group priority may encode substantive differences (as in the case of BPS) that must be respected, just not in the same fashion as the between-group priority. The difficulty lies in creating a unified model that allows for both interpretations, one between groups and the other within groups.

The power of the unified core lies in its ability to interpolate between the frameworks of GS and RTZ. When the between-group priority is a strict ranking where each student forms a singleton priority group, the unified core corresponds to the fair core of GS. On the other hand, when the between-group priority is indifferent over all students (there is a single priority group), then the unified core corresponds to the efficient core of RTZ. At intermediate stages, the unified core blends both RTZ and GS in a principled manner.

To make the unified core practical, I introduce a novel two stage algorithm—the DA-TTC algorithm—that always produces a match in the unified core. In the first stage, I allow students to freely apply to their most-preferred schools as in the Deferred Acceptance algorithm of GS. Schools tentatively accept the highest-priority applicants and immediately reject the rest. The process continues until no more students are rejected, at which point the first stage terminates and returns the match  $\mu_1$ .

In the second stage, I adapt Gale’s Top Trading Cycles algorithm but with a restriction on the allowable trades ([Shapley and Scarf, 1974](#)). Essentially, students may only trade their seats (that they have from  $\mu_1$ ) to other students in the same priority group. I iteratively find groups of students who can, by trading their assigned schools among themselves, be matched to their most-preferred school out of the ones currently available. This algorithm is the crux of this paper, and is presented in [Section 3.2](#). I label the resulting match  $\mu_2$ . Once all students have been removed, the second stage terminates and returns the match  $\mu_2$ .

My main result establishes that  $\mu_2$  belongs to the unified core. The advantage of my approach is a practical algorithm to finding an element of the unified core. The DA-TTC is also computationally feasible.<sup>8</sup>

A key advantage of the unified core in applications is the simplicity of its definition.<sup>9</sup> First,

<sup>7</sup>I discuss the differences in models in [Section 1.1](#), and the differences in our techniques in [Section 4](#). In short, EE uses only between-group priorities; however, I must use more care because of the within-group priorities. This prevents a direct application of the Stable Improvement Cycles algorithm.

<sup>8</sup>It runs in polynomial time.

<sup>9</sup>Indeed, Payzant expressed reservation concerning the TTC because of the complexity of the algorithm. I return to

there are no between-group priority violations, so the fairness property of the fair core is preserved. However, when there is a within-group priority violation, the explanation is straightforward: rectifying the violation would harm a student with higher within-group priority. By focusing on a reasonable outcome, the definition of the unified core lends itself to applications in school choice.

The rest of this paper is structured as follows. First, I close out this section with a discussion of the relevant papers, emphasizing the differences between the unified core, related concepts, and the methods used. [Section 2](#) introduces the model, the bulk of which is devoted to the distinction between within-group and between-group priorities. In [Section 3](#), I develop my main result ([Theorem 1](#)) which shows that a match in the unified core can always be found using an efficient algorithm. I conclude in [Section 4](#) with a discussion of the result and further avenues of research.

## 1.1 Related Literature

This paper connects several strands of research.

First, as mentioned previously, the model in this paper has substantial overlap with [Erdil and Ergin \(2008\)](#) (henceforth, EE). The key difference between my model and EE's model is that I include the within-group priorities in the model primitives, while they take these priorities as randomly generated from the between-group priorities. Although mathematically similar, because the within-group priorities may carry normative implications, the outcomes in the unified core are restricted. In EE, any student within the same priority group can be assigned to the school so long as the match is Pareto optimal among the matches that are stable with respect to the between-group priorities. In contrast, in the unified core, a student of higher within-group priority can claim a seat so long as no higher priority students in the same group are harmed. The within-group priority disciplines how seats are allocated within priority groups.

Second, ever since [Abdulkadiroğlu and Sönmez \(2003\)](#), researchers have been trying to modify the Deferred Acceptance algorithm to accommodate efficiency gains ([Ehlers and Morrill, 2020](#); [Kesten, 2010](#); [Trojan, Delacrétaz and Kloosterman, 2020](#); [Reny, 2022](#)). These approaches, however, take the priorities as given and weaken the definition of the core. The unified core complements this approach by weakening the priorities and then adapting the core to this framework. The upshot of the unified core is that the definition of the core is simpler, but the downside is that the preferences must have a two-stage structure to make any efficiency gains.

Third, a separate approach has been to provide foundations for the use of the Top Trading Cy-

---

this topic in the discussion.

cles algorithm in school choice (Abdulkadiroğlu and Che, 2010; Morrill, 2013; Abdulkadiroğlu et al., 2017; Chen, Chen and Hsu, 2021; Rong, Tang and Zhang, 2022; Dur and Paiement, 2024). These axiomatic definitions ground this work, and I extend our understanding of these mechanisms by integrating them into the fairness framework of GS.

The paper closest in spirit to mine is Abdulkadiroğlu (2011). The main difference between his paper and mine is the restriction they place on between-priorities: in their model, between-group priorities either express indifference or a strict preference over all students. Additionally, his focus is on student optimality rather than on the core. He uses the within-priorities in a similar way, but only in his algorithm; the unified core provides the underpinnings for his solution.

## 2 Model

In this section I present the formal model of a *school choice problem* and my key definitions. In the first subsection, I introduce the standard primitives, except that I replace a school's one priority with two priorities: the between-group priority and within-group priority. In the second subsection, I develop the theory of the unified core. In the third subsection, I turn to the solutions of GS, RTZ, and EE, and highlight the differences and advantages of the unified core.

### 2.1 The Primitives

There is a finite set of *students*  $I$  and a finite set of *schools*  $S$ , collectively called *agents*. Each school has a *capacity*  $q_s \geq 1$  of seats. A *match* is a function  $\mu : I \rightarrow I \cup S$  with the following two properties:

1. if  $i \in I$ , then  $\mu(i) \in S \cup \{i\}$ ; and
2. if  $s \in S$ , then  $|\mu^{-1}(s)| \leq q_s$ .

The first requirement guarantees that the market is two-sided: students are matched to schools or unmatched (this is denoted as being matched to oneself). The second requirement guarantees that capacities are not exceeded.

Every student has a strict preference  $\succ_i$  over  $S \cup \{i\}$ ; the associated weak preference is  $\succeq_i$ . A match  $\mu$  is *individually rational* if for every  $i \in I$ ,  $\mu(i) \succeq_i i$ . When considering a set of students  $A \subseteq I$ , I say that  $A$  *prefers*  $\nu$  to  $\mu$  if for every  $i \in A$ ,  $\nu(i) \succeq_i \mu(i)$  and the comparison is strict for some student in  $A$ .

Schools, however, come equipped with two orders over  $I$ , the between-group priority and the

within-group priority.<sup>10</sup> The *between-group priority* order is a weak preference  $\succsim_s$ .<sup>11</sup> It represents the priorities that require the stronger interpretation as a right to attend the school ahead of others. When  $\succsim_s$  does not rank two students strictly, it means that they are in the same priority group. Formally, denote the *priority group* of student  $i$  at  $s$  as

$$[a]_s \equiv \{j \in I : j \sim_s i\}$$

The *within-group priority* order is a strict preference  $\succ_s^*$ . It represents the priorities that are allowed the weaker interpretation of a better opportunity to attend the school. In practice, the priority that a policymaker uses is the lexicographic combination of  $\succsim_s$  and  $\succ_s^*$ : first,  $\succsim_s$  is used to partition students into groups, and then  $\succ_s^*$  is used to break ties within groups. Rather than define a combined priority, I assume that  $\succ_s^*$  is a refinement of  $\succsim_s$ ; that is, if  $a \succ_s b$ , then  $a \succ_s^* b$ . That is,  $\succ_s^*$  is what most researchers refer to as school priority, whereas for me it is the result of the policymaker breaking ties within priority groups.

Because this model is many-to-one, it is necessary to construct the school's between-group priority over groups of students.<sup>12</sup> Toward this end, I make the standard assumption that the school's between-group priority over groups of students is *responsive* (Roth, 1985). In words, this means that for equally-sized groups, the first has greater between-group priority than the second if every student in the first can be paired with a student in the second such that the first has greater between-group priority. For unevenly sized groups, the smaller group never has greater between-group priority than the larger group, but the larger group has greater between-group priority if there is a subset of the larger group that is the same size as the smaller group and has greater between-group priority.

Formally, I construct  $\succsim_s$  for groups of students in the following way. For  $A, A' \subseteq I$  such that  $|A| = |A'|$ , I write  $A \succsim_s A'$  if the students in  $A$  and  $A'$  can be indexed such that  $a_1 \succsim_s a'_1, a_2 \succsim_s a'_2, \dots, a_{|A|} \succsim_s a'_{|A|}$ . When  $A \succsim_s A'$  and  $A' \succsim_s A$ , I write  $A \sim_s A'$ .<sup>13</sup> If  $A \succsim_s A'$  but not  $A' \succsim_s A$ , then I write  $A \succ_s A'$ . When  $|A| \neq |A'|$ , I write  $A \succsim_s A'$  if  $|A| > |A'|$  and there is some  $\tilde{A} \subseteq A'$  such that  $\tilde{A} \succsim_s A$ .

## 2.2 Unified Core

With the primitives in hand, I turn to a discussion of the core of the model. As introduced earlier, policymakers view between-group and within-group priorities differently. To accommodate both

<sup>10</sup>Implicit here is the assumption that every student is acceptable to every school. Extending this model to incorporate unacceptable students is straightforward, but adds unnecessary notation.

<sup>11</sup>The irreflexive portion is  $\succ_s$  (the strict preference relation) and the reflexive portion is  $\sim_s$  (the indifference relation).

<sup>12</sup>Extending the within-group priority in a similar manner is unnecessary in my analysis.

<sup>13</sup>Note that  $\succsim_s$  may not be a complete relation over groups of students: there are some coalitions  $A$  and  $A'$  such that neither  $A \succsim_s A'$  nor  $A' \succsim_s A$ .



types of priorities, I introduce two separate types of blocks. In both cases the conditions on student preferences are the same: the coalition of students must prefer the new match to the old one. What changes is the definition of what the coalition can enforce.

The between-group priorities resemble the preferences in the GS two-sided model. However, within my framework I focus on students as the active players. To accommodate this, [Definition 1](#) allows a student to block a match if her between-group priority has been violated. This definition mirrors the standard blocking definition of GS. It rules out any violations of the between-group priority order.

**Definition 1.** The coalition  $A$  can *between-group enforce* match  $\nu$  over match  $\mu$  if for every  $i \in A$ , either there is a student  $j \in \mu^{-1}(\nu(i))$  such that  $i \succ_s j$  or  $|\mu^{-1}(\nu(i))| < q_{\nu(i)}$ .<sup>14</sup>

The definition of blocking is standard: a coalition can *between-group block* match  $\mu$  if it can between-group enforce some match  $\nu$  that the coalition prefers to  $\mu$ . Note that if a match is not individually rational, then it is between-group blocked by some student who wishes to be unmatched.

Between-group blocks are useful in school choice models, but do not fully model the priorities given by schools. As seen in EE, using between-group blocks alone can provide substantial welfare improvements for students. However, the within-priorities may also carry normative value as Payzant implied. To address this, I turn to a treatment of within-group priorities.

The within-group priorities are modeled as in the RTZ framework. In their model, a coalition of students can enforce a match to some schools only if no group of students *outside* the coalition could veto such a match by asserting their stronger claims. I import this definition into my model by restricting attention only to students within the same priority group. To formalize this, I first define the *within-group upper contour set*:

$$U_s(i) = \{j \in [i]_s : j \succeq_s^* i\}$$

In words,  $j$  is in  $U_s(i)$  if  $j$  has a higher within-group priority at  $s$ . When considering whether  $i$  can claim a seat at  $s$ ,  $U_s(i)$  are the students who could prevent this match because they have a stronger within-group claim to  $s$ . For example, suppose  $q_s = 1$ ; if  $i$  wishes to match to  $s$ , then  $i$  must guarantee that students in  $U_s(i)$  are not harmed by this action. I call students in  $U_s(i)$  but not in  $A$  the *within-group interrupters* of  $i$ . A sufficient condition is that there are no within-group interrupters. When  $q_s > 1$ , it is not necessary to eliminate every within-group interrupter: if  $n$

<sup>14</sup>Traditionally in GS-style models, the schools are included in the coalition; however, in RTZ-style models (see next paragraph), this is not the case. I present this student-only version for use here. The downside is that the definition of enforcement relies on  $\mu$ ; the upside is fewer cases in the definition of the unified core.

students from  $A$  match to  $s$ , then  $q_s - n$  within-group interrupters do not need to be included because the within-group interrupters (even when acting together) are unable to prevent the coalition members from claiming their seats at the school.<sup>15</sup> [Definition 2](#) formalizes this intuition.

**Definition 2.** A coalition  $A \subseteq I$  can *within-group enforce* match  $\nu$  over match  $\mu$  if for every school  $s \in \nu(A)$ , the following two conditions hold:

1.  $\nu^{-1}(s) \sim_s \mu^{-1}(s)$
2. the sum of the number of within-group interrupters across all students in  $A$  that are matched to  $s$  is less than or equal to  $q_s - |\nu^{-1}(s)|$ .

Again, the definition of blocking is symmetric: a coalition can *within-group block* match  $\mu$  if it can within-group enforce some match  $\nu$  that the coalition prefers to  $\mu$ .

The main difference between within-group blocks and between-group blocks is their treatment of schools. Schools must strictly benefit (according to the between-priority) from between-group blocks, but may be indifferent in a within-group block. The within-group block compensates by placing a stronger condition upon the student coalition.

With both types of blocks in hand, I turn to defining the unified core in [definition 3](#).

**Definition 3.**  $\mu$  is in the *unified core* if it is neither between-group blocked nor within-group blocked.

In [Section 3](#), I provide a two-stage algorithm that always finds a match in the unified core. Before introducing the algorithm, I first discuss the two most-closely related solutions.

### 2.3 Comparison with Related Models

In this subsection, I briefly introduce the models of GS and RTZ.

In the fairness framework of GS, a match  $\mu$  is *fair* if there is no student  $i$  and  $s$  such that  $s \succ_i \mu(i)$  and either there is a  $j \in \mu^{-1}(s)$  such that  $i \succ_s^* j$  or  $|\mu^{-1}(s)| < q_s$ . Note that when the between-group priority orders all students strictly, then the set of fair matches and the unified core are the same.

In contrast, the efficiency framework of RTZ uses priorities to allocate school seats as if they were objects. The within-group priority represents a form of ownership of the school and allows for a form of trading. The key difference between the efficient core and the unified core is that the

---

<sup>15</sup>The reader may wonder why I (and RTZ) include this slackness in the definition of enforceability. Without it, a student with top priority might trade *multiple* seats away. The intuition is that, if every within-group interrupter is included, then including the highest-priority student might necessitate (through a chain of enforceability claims) including two or more lower-priority students who are matched to the school. An example demonstrating this is available upon request.

efficient core assumes that there is a single priority group: any group of students can use their priorities to trade their schools. This is made by a slight change in definitions: the *efficient upper contour set* is

$$U_s^*(i) = \{j \in I : j \succ_s^* i\}$$

and for a coalition  $A$  and match  $\nu$ , the *efficient interrupters* of  $i \in A$  are the students in  $U_{\nu(i)}^*(i)$  but not  $A$ . The definition of *efficient enforcement* is then the same as within-group enforcement except that the condition  $\nu^{-1}(s) \sim_s \mu^{-1}(s)$  is dropped and the efficient interrupters are counted instead of within-group interrupters. The definition of efficient blocking and the efficient core is similar. When the between-group priority places all students within the same priority group, the efficient core and the unified core are the same.

The power of the unified core is its ability to interpolate between the efficient and fair cores. By allowing for the between-group priority to range between a strict ranking of all students and complete indifference, the unified core captures the nuances of policymakers' objectives and the two-stage nature of priorities. In [example 1](#), I show how the unified core models what the fair and efficient cores cannot.

**Example 1.** Consider a school choice problem with six students and six schools, each school having one seat. The students' preferences and the schools' between- and within-priority orders are as follows (unlisted preferences/priorities are irrelevant):

$\succ_1$	$\succ_2$	$\succ_3$	$\succ_{1'}$	$\succ_{2'}$	$\succ_{3'}$
$a$	$b$	$a$	$a'$	$b'$	$a'$
$b$	$a$	$c$	$b'$	$a'$	$c'$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

  

$\succ_a$	$\succ_b$	$\succ_c$	$\succ_{a'}$	$\succ_{b'}$	$\succ_{c'}$
1, 2, 3	1, 2	3	2'	1'	3'
$\vdots$	$\vdots$	$\vdots$	3'	2'	$\vdots$
			1'	$\vdots$	
			$\vdots$		

  

$\succ_a^*$	$\succ_b^*$	$\succ_c^*$	$\succ_{a'}^*$	$\succ_{b'}^*$	$\succ_{c'}^*$
2	1	3	2'	1'	3'
3	2	$\vdots$	3'	2'	$\vdots$
1	$\vdots$		1'	$\vdots$	
$\vdots$			$\vdots$		

Notice that this example consists of two parallel problems side-by-side. In the first, the students 1 and 2 are the highest-priority students at the school the other most-prefers (schools  $a$  and  $b$ , respectively), and in the second, students  $1'$  and  $2'$  are the highest-priority at the school the other

most-prefers (schools  $a'$  and  $b'$ , respectively). The fair core is the match:

$$\mu^{\text{GS}} = \begin{pmatrix} 1 & 2 & 3 & 1' & 2' & 3' \\ b & a & c & b' & a' & c' \end{pmatrix}$$

The fair core does not incorporate the difference between the between-group and within-group priorities; hence it is less efficient than allowed. However, the efficient core is also unsatisfactory:

$$\mu^{\text{RTZ}} = \begin{pmatrix} 1 & 2 & 3 & 1' & 2' & 3' \\ a & b & c & a' & b' & c' \end{pmatrix}$$

The difficulty with RTZ is that it does not recognize that student  $3'$  can claim a seat at school  $a'$  ahead of student  $1'$ .

The unified core, however, correctly identifies the difference between these priority groups:

$$\mu^{\text{UC}} = \begin{pmatrix} 1 & 2 & 3 & 1' & 2' & 3' \\ a & b & c & b' & a' & c' \end{pmatrix}$$

### 3 Analysis

In this section, I present my main result ([Theorem 1](#)). It states that the unified core is non-empty for every school choice problem and provides an algorithm for finding one such match.

I use a two-stage algorithm to prove the result. In the first stage, I leverage the Deferred Acceptance algorithm (DA) to produce a match that is not between-group blocked ([Gale and Shapley, 1962](#)). In the second stage, I build a variation of the Top Trading Cycles algorithm (TTC) to remove any within-group blocks from the previous match while not adding any between-group blocks ([Shapley and Scarf, 1974](#)). The match produced is in the unified core.

#### 3.1 Removing between-group blocks: The Deferred Acceptance Algorithm

In this subsection, I present the first stage of the DA-TTC: the Deferred Acceptance algorithm from GS. In the DA, students initially start unmatched. In the first round, students “propose” to their favorite school. Every school tentatively accepts students up to its capacity (picking the highest according to its within-group priority) and immediately rejects the rest. In subsequent rounds, students who are not tentatively accepted propose to their favorite school among those that have not rejected them. Every school considers the new proposals simultaneously with the students it has tentatively accepted, and again rejects all but the top students up to its capacity. This process continues until every student is matched or has been rejected by every acceptable school. I formally write this in [Algorithm 1](#).

**Algorithm 1** Deferred Acceptance Algorithm

---

every student  $i$  points to the  $\succsim_i$ -highest agent

**while** more than  $q_s$  students point to some school  $s$  **do**

$s$  rejects all but the  $\succ_s^*$ -top  $q_s$  students who proposed to  $s$

every student  $i$  points to the  $\succsim_i$ -highest agent *who has not rejected him yet*

**end while**

$\mu_1$  assigns each student to the agent he last proposed to

---

The DA was originally designed by GS to construct a fair match. Fairness is a stronger requirement than not being between-group blocked, this result applies to my model. [Lemma 1](#) translates the standard result into my model. The proof is standard and is relegated to the Appendix.

**Lemma 1.**  $\mu_1$  is not between-group blocked by any coalition.

### 3.2 Removing within-group blocks: The Top Trading Cycles Algorithm

In this subsection, I present the second stage of the DA-TTC: the Top Trading Cycles algorithm of [Shapley and Scarf \(1974\)](#). In the TTC, each student “owns” a seat at the school she is matched to in  $\mu_1$ . Students are allowed to trade the seats that they own with others, but only within priority groups.

The way that I implement this is by first restricting which students and schools are active, and then restrict which schools a student can “point” to. At the start of the TTC, each student  $i$  is *active* if she is in the lowest priority group of  $\mu_1(i)$  among the students matched to  $\mu_1(i)$  (i.e. for every  $j \in \mu_1^{-1}(\mu_1(i))$ ,  $j \succsim_{\mu_1(i)} i$ ). Throughout the TTC, each school  $s$  is active if at least one student in  $\mu_1^{-1}(s)$  is active (at the start of the TTC, every school such that  $|\mu_1^{-1}(s)| \geq 1$  is active by construction).

In every round of the TTC, for each student, construct the list of *admissible* schools. A school  $s$  is admissible to student  $i$  if there is some active  $j \in \mu_1^{-1}(s)$  such that  $i \sim_s j$ . Each student points to her most-preferred admissible school, and each active school points to the highest-priority active student in  $\mu_1^{-1}(s)$  (by definition, a school is only active if there is an active student in  $\mu_1^{-1}(s)$ , so this is well-defined). Because every active student points to an active school, and every active school points to an active student, a cycle must form. Every student on the cycle is assigned a seat at the school she points to and becomes inactive. The process is repeated until no students (and hence no schools) are active. [Algorithm 2](#) provides the formal definition.

**Theorem 1.**  $\mu_2$  is in the unified core.

---

**Algorithm 2** Top Trading Cycles Algorithm
 

---

```

initialize  $\mu_2 \leftarrow \mu_1$ 
for all  $i \in I$  do            $\triangleright$  activates students in the lowest-priority group at their assigned school
    initialize  $i$  as active if  $j \succsim_{\mu_1(i)} i$  for every  $j \in \mu_1^{-1}(\mu_1(i))$ 
end for
for  $s \in S$  do
    initialize  $s$  as active if there is an active student in  $\mu_1^{-1}(s)$ 
end for
while there is an active student do
    for every active student  $i$  do            $\triangleright$  students point to most-preferred admissible school
         $i$  points to her  $\succ_i$ -most preferred active school in  $\{s : \text{exists active } j \in \mu_1^{-1}(s) \text{ s.t. } i \sim_s j\}$ 
    end for
    each active school points to the highest-priority active student in  $\mu_1^{-1}(s)$ 
    a cycle  $i_1, s_1, \dots, i_n, s_n$  exists
    for all  $1 \leq k \leq n$  do            $\triangleright$  removes students in cycle from problem, then restarts
        set  $\mu_2(i_k)$  to  $s_k$ 
        deactivate  $i_k$ 
        if there are no active students in  $\mu_1^{-1}(s_k)$ , then deactivate  $s$ 
    end for
end while
  
```

---

The full proof of [Theorem 1](#) is available in the Appendix. Here, I outline the key points in the argument. A preliminary step is to show that  $\mu_2$  is not between-group blocked; given that  $\mu_2$  is a Pareto improvement of  $\mu_1$  for the students but is equivalent under the between-priorities, this is not difficult to show.<sup>16</sup> The rest of the proof supposes that  $A$  within-group blocks  $\mu_2$  with  $\nu$ , and then finds a contradiction.

The major difficulty with applying a standard proof that the TTC is in the core is that, for a school  $s$  in  $\nu(A)$ , not every student in  $\nu^{-1}(s)$  is necessarily in  $A$ . Put differently, using the TTC usually requires a well-defined set of owners, which the unified core lacks. The crux of the proof is a construction of a bijection  $T_s$  between students in  $\nu^{-1}(s) \cap A$  and students in  $\mu_2^{-1}(s) \cap A$ , which identifies an owner for each relevant seat. I construct  $T_s$  by addressing students in both  $\nu^{-1}(s) \cap A$  and  $\mu_2^{-1}(s) \cap A$  separately from those in only  $\nu^{-1}(s) \cap A$ . For those in both, I make  $T_s$  the identity function. For those in only  $\nu^{-1}(s) \cap A$ , I leverage the  $\nu^{-1}(s) \sim_s \mu_2^{-1}(s)$  of within-group enforcement requirement to find an equally-sized set of students in  $(\mu_2^{-1}(s) \cap A) \setminus \nu^{-1}(s)$ . I further show that if  $i$  is active in the TTC, then  $T_{\nu(i)}(i)$  is also active in the TTC and  $T_{\nu(i)}(i) \in U_{\nu(i)}(i)$ . Hence, for active students, the  $T_s$  essentially identifies which student owns which seat at  $s$ .

With the bijection  $T_s$  in hand, I then can use standard methods. I construct a cycle of students  $i_1, \dots, i_n$  in  $A$  where  $\nu(i_k) = \mu_2(i_{k+1})$  as in most proofs involving the TTC. I construct it such that least one student in this cycle must strictly prefer  $\nu$  to  $\mu_2$ . But I also show that every  $i_k$  is active in the Top Trading Cycles algorithm, so each must be deactivated only after the next student. This then leads to a situation in which  $i_k$  must be deactivated strictly before  $i_k$  is, a contradiction.

## 4 Discussion

In this section, I place [Theorem 1](#) and the DA-TTC in conversation with existing results and point toward several avenues for future research. I first compare this result with that of EE, highlighting the complexities arising from the within-group blocks. I then turn to the question of finding a constrained efficient match. Finally, I turn to a more general discussion of the DA and TTC in school choice, and highlight the difficulties present in bridging these two algorithms.

With the proof of the non-emptiness of the unified core in hand, I turn to a discussion of the differences between the DA-TTC and the Stable Improvement Cycles algorithm (DA-SIC) of EE. For context, EE considers the matches that are not between-group blocked. The DA-SIC starts by running the DA using the within-group priorities to establish a baseline match that is not between-group blocked. The SIC then checks for a cycle of students who each prefer the next

<sup>16</sup>But the importance of this cannot be overstated; as I discuss in [Section 4](#), having to find a Pareto improvement of  $\mu_1$  has restricted the kinds of algorithms available in school choice.

student's match such that, if the students exchange seats, no between-group blocks are created; this is the eponymous *stable* improvement cycle. The output of the SIC stage is a student-optimal match among those that are not between-group blocked. The critical difference between the SIC and the TTC is that the SIC allows for trades *across* priority groups while the TTC does not. The following example with three students and two schools (each with one seat) illustrates this:

$\succ_1$	$\succ_2$	$\succ_3$
$a$	$b$	$a$
$b$	$a$	$b$

$\succ_a$	$\succ_b$
2	1
1, 3	2, 3

$\succ_a^*$	$\succ_b^*$
2	1
3	3
1	2

$$\mu^{\text{DA}} = \begin{pmatrix} 1 & 2 & 3 \\ b & a & 3 \end{pmatrix}$$

At the match  $\mu^{\text{DA}}$ , EE identifies a stable improvement cycle of  $1 \rightarrow a$  and  $2 \rightarrow b$ . However, allowing this pair to trade schools violates student 3's within-priority at both schools: after the trade, the coalition of  $\{3\}$  can within enforce a match to either  $a$  or  $b$ . In effect, the stable improvement cycle is stable with respect to  $\succ$ , but disregards  $\succ^*$ .<sup>17</sup> This allows for these trades between priority groups. This difference is not a mere technical nuance; in the case of BPS, it is precisely trades like this that concern the policymaker.

Thus far, I have dealt only with existence rather than constrained-efficiency. The major drive of EE is to find a student-optimal match subject to the stability constraints, which they show can be found by iteratively eliminating stable improvement cycles. Given the similarity between the DA-TTC and the DA-SIC, the natural question arises as to whether an analogous method can be used to eliminate “unified” stable improvement cycles. Unfortunately, this approach seems doomed to failure. Although for the between-priority the stable improvement cycle can be imported as-is, checking whether a cycle respects within-priority is significantly more complex. The crux is that a student's within-priority can be violated when a higher-priority student trades her seat to a lower-priority agent. However, in a cycle, whether such a “trade” can be found depends upon the matches of students outside the cycle. Put differently, a stable improvement cycle may involve trades across priority groups, but then a trade within the priority group must be found to protect the former trade. The problem is compounded by the possibility that the latter trade may itself be the result of another simultaneous trade. Whether a constrained efficient match can be found using this technique or others is an open question.

<sup>17</sup>There is a second difference between the DA-TTC and the DA-SIC. Consider a school choice problem with one school (with one seat) and two students, both of whom desire the seat at the school. If the students are in the same priority group, then the DA-SIC could assign either student to the school. The DA-TTC, however, is constrained to only assign the higher-priority student. Put differently, the DA-TTC can emulate the TTC because the DA-TTC makes use of the within-group priority; the DA-SIC cannot.



There are several open questions about algorithms similar to the DA-TTC. The DA-TTC is a part of a growing set of algorithms that can be roughly described as “DA baseline, then trade to improve.” Starting with [Abdulkadiroğlu and Sönmez \(2003\)](#), several authors have used this formula, such as EE, [Kesten \(2010\)](#), and [Doe \(2024\)](#). The difficulty with combining these two families of algorithms is maintaining monotonicity on a set of matches. [Echenique and Oviedo \(2004\)](#) shows how the DA can be viewed as a monotonic function on matches; however, the TTC does not share this property ([Echenique, Goel and Lee, 2024](#)). When combining these two algorithms, care must be taken to guarantee that the trades made in the TTC do not upset the match found in the DA, which is why the previous papers allow students only to trade the school they received from the DA rather than any school they have the highest-priority for. Going the other way around is tougher: if the designer first allows students to trade schools and then attempts to remove fairness violations, a student may trade for a school that is taken away; with one student-school pair removed from the group of trading students, the outcome of the TTC shifts unpredictably. The only paper to tackle both simultaneously is [Abdulkadiroğlu \(2011\)](#). Is this approach to connecting the DA and TTC the only viable option? Can conditions be placed upon priorities to ameliorate the non-monotonicity of the TTC in the presence of the DA?

Another obstacle to implementing the TTC is the complexity of the algorithm. Payzant alludes to this; after his previous quote, he continued: “Moreover, Top Trading Cycles is less transparent – and therefore more difficult to explain to parents – because of the trading feature executed by the algorithm, which may perpetuate the need or perceived need to ‘game the system.’” Concerns like Payzant’s have spurred a growing body of literature to understand the complexity of the TTC ([Gonczarowski and Thomas, 2024](#); [Leshno and Lo, 2017](#)). In this paper, I instead focused on properties of the match rather than the process used to reach the match. The properties of a match may prove more transparent to stakeholders than the properties of an algorithm.

I close with a brief note about the DA and TTC. Within the literature on school choice, the vast majority of algorithms stem from the DA and the TTC. The monotonicity of the DA is well-suited for the lattice structure of the stable matches, and the TTC flexibly handles ownership economies ([Gusfield and Irving, 1989](#); [Papai, 2000](#)). However, as I note in the previous paragraphs, these algorithms are difficult to combine. Additionally, given the complexity of the TTC and how few districts have attempted to implement it, relying on TTC-style algorithms to improve the efficiency of the DA may not be successful in applications. Developing general algorithms—as in [Abdulkadiroğlu \(2011\)](#)—to probe the efficiency-fairness frontier is a research direction of primary importance.

## References

- Abdulkadiroğlu, Atila.** 2011. “Generalized Matching for School Choice.”
- Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 2003. “School Choice: A Mechanism Design Approach.” *American Economic Review*, 93(3): 729–747.
- Abdulkadiroğlu, Atila, and Yeon-Koo Che.** 2010. “The Role of Priorities in Assigning Indivisible Objects: A Characterization of Top Trading Cycles.” *Working Paper*.
- Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth.** 2009. “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match.” *American Economic Review*, 99(5): 1954–1978.
- Abdulkadiroğlu, Atila, Parag Pathak, Alvin Roth, and Tayfun Sonmez.** 2006. “Changing the Boston School Choice Mechanism.” National Bureau of Economic Research w11965, Cambridge, MA.
- Abdulkadiroğlu, Atila, Yeon-Koo Che, Parag Pathak, Alvin Roth, and Olivier Tercieux.** 2017. “Minimizing Justified Envy in School Choice: The Design of New Orleans’ OneApp.” National Bureau of Economic Research w23265, Cambridge, MA.
- Chen, Siwei, Yajing Chen, and Chia-Ling Hsu.** 2021. “New axioms for top trading cycles.” arXiv:2104.09157 [econ].
- Doe, Peter.** 2024. “Matching With Pre-Existing Binding Agreements: The Agreeable Core.” *Working Paper*.
- Dur, Umut, and Scott Paiement.** 2024. “A characterization of the top trading cycles mechanism for the school choice problem.” *Mathematical Social Sciences*, 129: 93–100.
- Echenique, Federico, and Jorge Oviedo.** 2004. “A Theory of Stability in Many-to-many Matching Markets.” *SSRN Electronic Journal*.
- Echenique, Federico, Sumit Goel, and SangMok Lee.** 2024. “Stable allocations in discrete exchange economies.” *Journal of Economic Theory*, 222: 105921.
- Ehlers, Lars, and Thayer Morrill.** 2020. “(Il)legal Assignments in School Choice.” *The Review of Economic Studies*, 87(4): 1837–1875.
- Erdil, Aytok, and Haluk Ergin.** 2008. “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice.” *American Economic Review*, 98(3): 669–689.

- Gale, D., and L. S. Shapley.** 1962. "College Admissions and the Stability of Marriage." The American Mathematical Monthly, 69(1): 9–15. Publisher: Mathematical Association of America.
- Gonczarowski, Yannai A., and Clayton Thomas.** 2024. "Structural Complexities of Matching Mechanisms." 455–466. Vancouver BC Canada:ACM.
- Gusfield, Dan, and Robert W Irving.** 1989. The stable marriage problem: structure and algorithms. MIT press.
- Kesten, Onur.** 2010. "School Choice with Consent \*." Quarterly Journal of Economics, 125(3): 1297–1348.
- Leshno, Jacob D, and Irene Lo.** 2017. "The Simple Structure of Top Trading Cycles in School Choice: A Continuum Model."
- Morrill, Thayer.** 2013. "An alternative characterization of top trading cycles." Economic Theory, 54(1): 181–197.
- Morrill, Thayer.** 2015. "Making just school assignments." Games and Economic Behavior, 92: 18–27.
- Papai, Szilvia.** 2000. "Strategyproof Assignment by Hierarchical Exchange." Econometrica, 68(6): 1403–1433.
- Reny, Philip J.** 2022. "Efficient Matching in the School Choice Problem." American Economic Review, 112(6): 2025–2043.
- Rong, Kang, Qianfeng Tang, and Yongchao Zhang.** 2022. "The Core of School Choice Problems." Working Paper.
- Roth, Alvin E.** 1985. "The college admissions problem is not equivalent to the marriage problem." Journal of Economic Theory, 36(2): 277–288.
- Shapley, Lloyd, and Herbert Scarf.** 1974. "On cores and indivisibility." Journal of Mathematical Economics, 1(1): 23–37.
- Troyan, Peter, David Delacrétaz, and Andrew Kloosterman.** 2020. "Essentially stable matchings." Games and Economic Behavior, 120: 370–390.

## A Omitted Proofs

**Lemma A.1.** *Every student weakly prefers  $\mu_1$  to being unmatched.*

*Proof.* Observe that no student rejects a proposal from herself (because students do not make rejections). This implies that no student prefers being unmatched more than  $\mu_1$ .  $\square$

**Proof of Lemma 1:**

Suppose (toward a contradiction) that  $\mu_1$  is between-group blocked by a coalition  $A$  between-enforcing a match  $\nu$  which it prefers.

First, by lemma A.1, there is some school in  $\mu_1(A)$ .

Second, observe that a school only rejects students if it receives more proposals (cumulatively) than its capacity  $q_s$ . Therefore, if  $i$  strictly prefers a school to  $\mu_1(i)$ , then that school is filled to capacity at  $\mu_1$ .

The previous two points imply the existence of a school  $s$  and two students  $i$  and  $j$  with the following properties:

1.  $s \in \nu(A)$ ; and
2.  $i \in A \cap \nu^{-1}(s)$  but  $i \notin \mu_1^{-1}(s)$ ; and
3.  $j \in \mu_1^{-1}(s)$  but  $j \notin \nu^{-1}(s)$ ; and
4.  $i \succ_s j$ .

But then notice that  $i$  must have proposed to  $s$  but have been rejected. However,  $j$  must have proposed to  $s$  but *not* have been rejected. If  $j$  proposed to  $s$  before  $s$  rejects  $i$ , then  $s$  rejects  $j$  before rejecting  $i$ , a contradiction. If  $j$  proposed to  $s$  after  $s$  rejects  $i$ , then  $s$  must have rejected  $i$  in favor of a higher within-group priority student  $i^*$ ; hence,  $j$  is also rejected, a contradiction.

Therefore, there is no coalition which between-group blocks  $\mu_1$ .  $\square$

**Lemma A.2.** *If  $A \sim_s A'$ , then for every  $i \in I$ :  $||[i]_s \cap A| = |[i]_s \cap A'|$*

*Proof.* Suppose (toward a contradiction) that for some  $i \in I$ :  $||[i]_s \cap A| \neq |[i]_s \cap A'|$ . Without loss of generality, let  $i$  be in the  $\succsim$ -highest priority group such that  $||[i]_s \cap A| \neq |[i]_s \cap A'|$ . Again, without loss of generality, let  $||[i]_s \cap A| > |[i]_s \cap A'|$ . Consider that  $A \sim_s A'$  implies  $A' \succsim_s A$ . But then  $A'$  and  $A$  cannot be indexed such that  $a'_k \succsim_s a_k$  because every student in a  $\succsim$ -more preferred priority group must be paired with a student in that same priority group and the students in the  $[i]_s$  priority group are imbalanced. This is a contradiction.  $\square$

**Lemma A.3.** *The following statements are true:*

1. *If  $A$  within-group blocks  $\mu_2$  by within-enforcing  $v$ , then  $A$  within-group blocks  $\mu_1$  by within-enforcing  $v$  as well.*
2. *If  $A$  between-group blocks  $\mu_2$  by between-enforcing  $v$ , then  $A$  within-group blocks  $\mu_1$  by within-enforcing  $v$  as well.*

*Proof.* I show that  $A$  can within-group enforce or between-group enforce  $v$  over  $\mu_1$  (respectively). Because every student weakly prefers  $\mu_2$  to  $\mu_1$ , it follows that  $A$  prefers  $v$  to  $\mu_1$  and hence within-group blocks or between-group blocks  $\mu_1$  with  $v$ .

For the first statement, suppose that  $A$  within-group blocks  $\mu_2$  by within-enforcing  $v$ . I show that  $A$  can within-group enforce  $v$  over  $\mu_1$ . To see this, notice that  $U_s(i)$  is independent of  $\mu_2$ . Hence, the sum of within-group interrupters across all students in  $v^{-1}(s) \cap A$  is the same for both  $\mu_2$  and  $\mu_1$ . Additionally, by definition, for every  $s \in v(A)$ :  $\mu_2^{-1}(s) \sim_s v^{-1}(s)$ . By the construction of  $\mu_2$ :  $\mu_2^{-1}(s) \sim_s \mu_1^{-1}(s)$ . Hence,  $v(s) \sim_s \mu_1^{-1}(s)$ . Therefore,  $A$  can within-group enforce  $v$  over  $\mu_1$ .

For the second statement, suppose that  $A$  between-group blocks  $\mu_2$  by within-enforcing  $v$ . I show that  $A$  can between-group enforce  $v$  over  $\mu_1$ . To see this, notice that for every school  $s$ ,  $\mu_2 \sim_s \mu_1$ . Hence, if  $v^{-1}(s) \succ_s \mu_2$ , then  $v^{-1}(s) \succ_s \mu_1$ . Therefore,  $A$  can between-group enforce  $v$  over  $\mu_1$ .

This completes the proof.  $\square$

**Lemma A.4.** *If  $A$  within-group blocks  $\mu_2$  with  $v$ , then  $i \in A$  strictly prefers  $v(i)$  to  $\mu_2(s)$  only if  $s$  is filled to capacity at  $\mu_2$ .*

*Proof.* Suppose (toward a contradiction) that  $i \in A$  strictly prefers  $v(i) \equiv i$  to  $\mu_2(s)$  but  $|\mu_2^{-1}(s)| < q_s$ .

By lemma A.3,  $A$  within-group blocks  $\mu_1$  with  $v$ . By the construction of  $\mu_2$ ,  $|\mu_2^{-1}(s)| = |\mu_1^{-1}(s)|$ . By the construction of  $\mu_1$ ,  $s$  only rejects  $i$  if more than  $|q_s|$  students point to  $s$ . But less than  $q_s$  other students point at  $s$ , so  $s$  does not reject  $i$ . Therefore,  $i$  weakly prefers  $\mu_1$  to  $i$ . This contradicts that  $i$  strictly prefers  $s$  to  $\mu_2$ . Therefore,  $|\mu_2^{-1}(s)| = q_s$ .  $\square$

**Proof of Theorem 1:** By lemma A.3,  $\mu_2$  is not between-group blocked; otherwise,  $\mu_1$  would be between blocked, a contradiction to lemma 1.

Now consider within-group blocks. Suppose (toward a contradiction) that  $\mu_2$  is within-group blocked by a coalition  $A$  within-enforcing a match  $v$ .

**First**, for every school  $s \in v(A)$ , I define an injection  $T_s : v^{-1}(s) \cap A \rightarrow \mu_1^{-1}(s) \cap A$  in the following manner. Let  $v^{-1}(s) \cap A = \{i_1, i_2, \dots, i_n\}$  be indexed such that  $i_{k+1} \succ_s^* i_k$  for every  $k$ . Notice that by the construction of  $\mu_1$ , if  $i_k \in \mu_1^{-1}(s)$ , then for  $k' > k$ ,  $i_{k'} \in \mu_1^{-1}(s)$ . Therefore, there is a unique  $m$

(possibly taking the value of 0 or  $n$ ) such that  $i_m \notin \mu_1^{-1}(s)$  but  $i_{m+1} \in \mu_1^{-1}(s)$ . I define  $T_s$  piecewise based on the index  $k$ :

- For  $m+1 \leq k \leq n$ , let  $T_s(i_k) \equiv i_k$ .

Observe that  $T_s(i_k) \in \mu_1^{-1}(s)$  by the definition of  $m$ . Notice that  $T_s(i_k) \in A$  because  $i_k \in v^{-1}(i) \cap A$ . This piece of  $T_s$  is clearly injective.

- For  $1 \leq k \leq m$ , consider the following argument:

Because  $A$  within-group enforces  $v$  at  $\mu_2$ , it follows that  $v^{-1}(s) \sim_s \mu_2^{-1}(s)$ . By construction,  $\mu_1^{-1}(s) \sim_s \mu_2^{-1}(s)$ . By construction of  $\mu_1$ , every student in  $\mu_1^{-1}(s)$  is strictly  $>_s^*$ -preferred to every student in  $i_1, \dots, i_m$ . Hence by [lemma A.2](#),  $i_1, \dots, i_m \in [i_m]_s$ .

Again, because  $A$  within-group enforces  $v$  over  $\mu_2$ , there are at most  $q_s - |v^{-1}(s)|$  within-group interrupters for student  $i_m$ . Because  $|v^{-1}(s)| = q_s$  by [lemma A.4](#), it follows that there are no within-group interrupters of  $i_m$ . Hence,  $U_s(i_m) \subseteq A$ . But every student in  $\mu_1^{-1}(s)$  is  $>_s^*$ -preferred to  $i_m$ . Thus,  $\mu_1^{-1}(s) \cap [i_m]_s \subseteq U_s(i_m)$ . Therefore,  $\mu_1^{-1}(s) \cap [i_m]_s \subseteq A$ .

Finally, note that  $|\mu_1^{-1}(s) \cap [i_m]_s| = |v^{-1}(s) \cap [i_m]_s|$  by [lemma A.2](#). Let  $M$  be the greatest index such that  $i_M \in [i_m]_s$ , and note that  $\{i_1, \dots, i_M\} \subseteq v^{-1}(s) \cap [i_m]_s$ . But then I have

$$|\mu_1^{-1}(s) \cap [i_m]_s| \geq |\{i_1, \dots, i_m, \dots, i_M\}|$$

However, because  $\{i_{m+1}, \dots, i_M\} \subseteq \mu_1^{-1}(s)$ , this can be rewritten:

$$\left| \left( \mu_1^{-1}(s) \cap [i_m]_s \right) \setminus \{i_{m+1}, \dots, i_M\} \right| + |\{i_{m+1}, \dots, i_M\}| \geq |\{i_1, \dots, i_m, \dots, i_M\}|$$

This further implies

$$\left| \left( \mu_1^{-1}(s) \cap [i_m]_s \right) \setminus \{i_{m+1}, \dots, i_M\} \right| \geq |\{i_1, \dots, i_m\}| \quad (1)$$

Define  $T_s(i_k)$  as the  $k^{\text{th}}$   $>_s^*$  least-preferred student in  $(\mu_1^{-1}(s) \cap [i_m]_s) \setminus \{i_{m+1}, \dots, i_M\}$ . By [equation \(1\)](#),  $T_s(i_k)$  is well-defined and is an injection on this piece. Because  $T_s(i_k) \notin \{i_{m+1}, \dots, i_M\}$ , this piece of  $T_s$  has no overlap with the first piece. Finally, by the above argument,  $T_s(i_k) \in A$ .

Hence, there is a well-defined injection  $T_s : v^{-1}(s) \cap A \rightarrow \mu_1^{-1}(s) \cap A$ .

**Second**, iteratively construct the following sequence of students. Let  $i_1$  be some student in  $A$  such that  $v(i_1) \neq \mu_2(i_1)$ . Because  $A$  prefers  $v$  to  $\mu_2$ , such a student exists. Let  $i_{k+1} \equiv T_{v(i_k)}(i_k)$ .

**Third**, I claim that  $i_k$  is well-defined,  $i_k \in A$ , and  $\mu_1(i_k) \neq v(i_k)$ . I show this by induction. For the base case when  $k = 1$ , I make three observations:

1. That  $i_1$  is well-defined is noted previously.

2.  $i_1 \in A$  by definition.
3. Because  $i_1$  strictly prefers  $v(i_1)$  to  $\mu_2(i_1)$  and  $i_1$  weakly prefers  $\mu_2$  to  $\mu_1$  by construction, it follows that  $v(i_1) \neq \mu_1(i_1)$ .

For the inductive step, suppose that every student with index less than  $k$  is well-defined, is in  $A$ , and is not matched to the same school in  $v$  and  $\mu_1$ . I make three observations:

1. Note that  $v(i_{k-1}) \in v(A)$ . Similarly, note that  $v(i_{k-1}) \in S$  because  $\mu_2$  is individually rational. Hence,  $T_{v(i_{k-1})}$  is well-defined. Therefore  $i_k$  is well-defined.
2.  $i_k \in A$  because the domain of  $T_{v(i_{k-1})}$  is a subset of  $A$ .
3. Because  $v(i_{k-1}) \neq \mu_2(i_{k-1})$  and  $i_{k-1}$  weakly prefers  $\mu_2$  to  $\mu_1$ , it follows that  $v(i_{k-1}) \neq \mu_1(i_{k-1})$ . Thus,  $i_{k-1} \notin \mu_1^{-1}(v(i_{k-1}))$ . By the construction of  $T_{v(i_{k-1})}$ , it follows that  $\mu_1(i_k) = v(i_{k-1})$  and  $v(i_k) \neq v(i_{k-1})$ . Therefore,  $\mu_1(i_k) \neq v(i_k)$ .

**Fourth**, because each  $T_{v(i_k)}$  is an injection and each student  $i_k$  is matched to only one school in  $\mu_1$ , it follows that if  $i_k \neq i_l$ , then  $T_{v(i_k)} \neq T_{v(i_l)}$ . Because there are a finite number of students, it follows that there is a (minimal) index  $n$  such that  $i_1 = i_{n+1}$ .

**Fifth**, consider the cycle  $i_1, \dots, i_n$ . Because  $T_{v(i_{k-1})}(i_{k-1}) \neq i_k$ , it follows that  $j \succ_{\mu_1(i_k)} i_k$  for every  $j \in \mu_1^{-1}(\mu_1(i_k))$ . Hence, every student in the cycle is active the Top Trading Cycles stage. Notice that each  $i_k$  must be deactivated weakly later than  $i_{k+1}$  because  $i_k$  points to  $\mu_1(i_{k+1})$ . However,  $i_1$  must be deactivated *strictly* later than  $i_2$  because  $v(i_1) \neq \mu_1(i_2)$ . Therefore,  $i_1$  must be deactivated strictly before  $i_1$ , a contradiction.

**Therefore**, there is no coalition that within-group blocks  $v$ . □