The Equitable Top Trading Cycles Mechanism for School Choice *

Rustamdjan Hakimov† Onur Kesten‡

This version: December 2017

Abstract

A particular adaptation of Gale's top trading cycles procedure applied to school choice, the so-called TTC mechanism, has attracted much attention both in theory and practice due to its superior efficiency and incentive features. We discuss and introduce alternative adaptations of Gale's original procedure that can offer improvements over TTC in terms of equity, along with various other distributional considerations. Instead of giving all the trading power to those students with the highest priority for a school, we argue for the distribution of the trading rights of all slots of each school: This allows them to trade in a thick market where additional constraints can be accommodated by choosing an appropriate pointing rule. We propose a particular mechanism of this kind, the *Equitable Top Trading Cycles (ETTC)* mechanism, which is also Pareto efficient and group strategy-proof just like TTC and eliminates avoidable justified envy situations. ETTC generates significantly fewer number of justified envy situations than TTC, both in simulations and the lab.

Keywords: school choice; stability; top trading cycles.

JEL classification: C78, C79, D61, D78, I20

*This work supersedes the second part of Kesten's 2004 paper entitled "Student placement to public schools in the US: Two new solutions". We thank the editor, four anonymous referees as well as Nick Arnosti, Yan Chen, Haluk Ergin, Dorothea Kübler, Oguzhan Gencturk, Fuhito Kojima, Herve Moulin, Muriel Niederle, Szilvia Pápai, Parag Pathak, Tayfun Sönmez, William Thomson, Utku Ünver, Al Roth, Serkan Zorba for helpful suggestions, and participants at Rochester, Carnegie Mellon, Harvard, Rice, Urbana-Champaign, Michigan, Sabanci, and various conferences for a number of useful discussions. Financial support from WZB is gratefully acknowledged. All remaining errors are our own.

†WZB, Reichpietschufer 50, 10785, Berlin, Germany; e-mail: rustamdjan.hakimov@wzb.eu

‡Corresponding author. Carnegie Mellon University, 5000 Forbes Ave. Pittsburgh, PA 15217 USA; e-mail:okesten@andrew.cmu.edu

1 Introduction

School choice programs that give families the flexibility to express preferences over schools have become increasingly popular both across the U.S. and worldwide. Despite such popularity, market designers have yet to reach a consensus on the "right" assignment mechanism. One major cause of this lack of agreement can be attributed to the existing trade-offs among the desirable features of mechanisms.

Three most important criteria to evaluate an assignment mechanism are: (1) stability/fairness in assignments (i.e., the extent to which student priorities can be accommodated), (2) student welfare, and (3) immunity to strategic action. In this context, however, there are apparent tensions between the three requirements. Even fairness and welfare are in conflict: there is no mechanism which is both stable[1] and Pareto efficient (Roth, 1982b). This tension automatically forces the designer to make a decision between the two properties. If one values stability over Pareto efficiency, then the choice is clear. The celebrated student-proposing *deferred acceptance (DA)* mechanism of Gale and Shapley (1962) is not only stable but it also produces the stable allocation that is most favorable to each student. Moreover, this mechanism is strategy-proof[2] (Dubins and Freedman, 1981; Roth, 1982a).[3]

On the other hand, an alternative mechanism is available should Pareto efficiency be the preferred feature between the two. Abdulkadiroğlu and Sönmez (2003) proposed an intuitive adaptation of Gale's 'top trading cycles' procedure to school choice. The *top trading cycles (TTC)* mechanism[4] is not only Pareto efficient but it is also strategy-proof.[5] TTC has also been regarded as a viable assignment mechanism to be used in practice. The task force for the Boston school district recommended the use of TTC as opposed to the eventually adopted DA.[6] Similarly, the San Francisco school district announced plans to implement TTC in 2011. Most recently, the Recovery School District in New Orleans adopted TTC in 2012, before abandoning it after one year.[7]

The top trading cycles idea was originally introduced in the context of a housing market (Shapley and Scarf, 1974). A housing market consists of a set of distinct objects, each with unit supply, and each owned by a distinct agent. The procedure works as follows: each agent points to the agent who owns his best choice object. Since the number of agents is finite, there is at least one cycle. In each cycle, the intended trades are performed and these agents and objects are removed. Then the same procedure is iteratively applied to the reduced market until no object remains in the market. In a housing market this procedure has been shown to possess quite appealing properties. It yields the unique *core* allocation of the housing market (Roth

[1]An allocation is stable if there is no unmatched student-school pair (i, s) such that student i prefers school s to his assignment, and school s either has not filled its quota or prefers student i to at least one student who is assigned to it.

[2]That is, no student ever gains by misrepresenting his preferences.

[3]Today two major school districts, Boston and NYC, have been implementing DA in student assignment for a decade now after abandoning their existing plans.

[4]The Abdulkadiroğlu and Sönmez adaptation of the original procedure is a substantial generalization of Papai's (2000) fixed-endowment inheritance rules.

[5]Pápai (2013) introduces a unified framework to characterize a class of rules that combine DA with TTC.

[6]See Kesten (2006) for a property-specific theoretical comparison of TTC and DA.

[7]In each case the decision was made in collaboration with economists. The details of the subsequent stages of the San Francisco plan, however, are not publicly available.

and Postlewaite, 1977) and it is strategy-proof (Roth, 1982b). Moreover, this procedure stands out as the unique strategy-proof, Pareto efficient, and individually rational mechanism (Ma, 1994). For the case of the unit supply of each object, large classes of assignment mechanisms based on this procedure have been characterized by (group-)strategy-proofness and efficiency, in conjunction with various other auxiliary requirements (Pápai, 2000; Pycia and Ünver, 2017; and Dur, 2013). The procedure has already proven useful for other applications, such as on-campus housing (Abdulkadiroğlu and Sönmez, 1999), assignment problems with transfers (Miyagawa, 2001), kidney exchange (Roth et al., 2005), random assignment problems (Kesten, 2009), the assignment of landing slots among flights (Schummer and Vohra, 2013), and tuition and worker exchanges (Dur and Ünver, 2015).

In all of these applications, objects are in unit supply and there is no ambiguity about the adaptation of Gale's top trading cycles procedure, as introduced by Pápai (2000) and further developed by Pycia and Ünver (2017). However, there is not a unique generalization of the procedure in the context of school choice, where an object may now be in multiple supply (i.e., each school has several slots) and where students have priorities over schools as opposed to the ownership rights in a housing market. Therefore, an adaptation of the procedure calls for a suitable interpretation of student priorities. In TTC, the trading market is generated by assigning each school to the student with the current highest priority for that school and keeping track of the remaining slots at each school. Trades among these top priority students are then carried out according to the top trading cycles procedure and the remaining slots are adjusted. Once a student has been placed in a school, the student is removed and the next highest priority student enters the market to be part of a trading cycle that is formed in a similar fashion. Despite its compelling welfare and incentive features, TTC may still introduce otherwise avoidable stability violations. We illustrate this point via a simple example.

Let the set of agents be $I = \{i_1, i_2, i_3\}$ and the set of schools be $S = \{s_1, s_2\}$ where school s_1 has one slot and school s_2 has two slots. The priorities for the schools and the students' preferences are given as follows:

$\succ_{s_1}$	$\succ_{s_2}$
i_1	i_3
i_2	i_2
i_3	i_1

P_{i_1}	P_{i_2}	P_{i_3}
$\boxed{\underline{\underline{s_2}}}$	$\underline{s_1}$	$\boxed{s_1}$
s_1	$\boxed{s_2}$	$\underline{s_2}$

When we apply TTC to this problem, student i_1 who has the highest s_1-priority, exchanges one slot at school s_1 in return for one slot at school s_2 from student i_3, who has the highest s_2-priority. The resulting allocation is indicated in boxes above. This allocation is Pareto efficient. However, the priority of student i_2 for school s_1 is violated by student i_3; i.e., i_2 has justified envy over i_3. On the other hand, notice that a Pareto efficient and stable allocation is also available for this problem.

In this paper, we argue that it may be possible to avoid the kind of justified envy situations observed in the above example under TTC, at no cost in terms of welfare or incentives. Observe that TTC gives student i_3 ownership over both slots of school s_2 before student i_2 enters the market. But then student i_1 has no choice but to trade with student i_3, which in turn leads to the violation of the priority of student i_2 for school s_1. However, had student i_1 traded his right to one slot at school s_1 with student i_2 for his right to one slot at school s_2, there would not be any priority violations. Indeed, such a trade would have led to the Pareto efficient and stable allocation underlined in the above profile.[8]

We argue that alternative adaptations of Gale's top trading cycles procedure can make a difference in terms of its practical desirability by improving upon its stability and making it possible to integrate other additional considerations into the trading process relative to the original TTC.[9] In fact, while expressing his view on TTC, the former superintendent of Boston, Thomas Payzant, underscored the need for such flexibility in trades in his memorandum to the school committee, dated May 25, 2005:

> *"There may be advantages to this approach...It may be argued, however, certain priorities—e.g., sibling priority—apply only to students for particular schools and should not be traded away."*

Since TTC allows only the current highest priority students to participate in the trading process, the number of participants in the market is at most the cardinality of the set of schools at any instant of this algorithm. Such a thin market may, however, entail justified envy for non-participants, as shown above, or may not be able to accommodate the trading constraints concerning specific priorities such as those due to sibling status. Our alternative approach consists of two parts. The first entails the construction of a much thicker trading market than TTC by increasing the number of "active" participants at any given instant of the trading market. We advocate for assigning all slots of each school x to all the q_x students with the highest $x-$priority, giving one slot to each student and endowing them with equal trading power subject to any additional constraints that may be desired. This would generate a sufficiently thick trading market, leading to a multitude of trading opportunities in which various considerations including fairness, affirmative action (e.g., diversity), class composition, etc., may also be incorporated into the trading process while making placements. A key observation is that in a thick market with competing participants, *being pointed to* is what leverages an agent to be part of a trading cycle, directly impacting the composition of the final allocation. Therefore, once market thickness is guaranteed, the second part of the approach is the choice of an appropriate "pointing rule" that specifies *who can point to whom*, which can be determined depending on the desideratum.

At any instant of the TTC algorithm, for example, each active participant of the market is endowed with a slot from a distinct school. Such thinness of the market automatically precludes the use of additional criteria for making assignments and already pins down all possible trades. To exemplify how the alternative approach works, consider the following pointing rule to address

[8]The introduction of avoidable justified envy situations in TTC may be further exacerbated when working with coarse priorities, which often require the use of random tie-breaking. Such randomization is known to lead to avoidable welfare losses under DA (Erdil and Ergin, 2008; Abdulkadiroğlu et al., 2009).

[9]Whether a Pareto efficient mechanism can actually be practically desirable inevitably depends on how school districts view stability, especially on whether some degree of instability can be tolerated. The much debated Boston mechanism remains an example of an unstable assignment rule that is widely used in practice, although the school choice reforms have considerably diminished its popularity. It may be argued that the assessment of the efficiency vs. stability trade-off by practitioners ultimately hinges on the grounds on which priorities are initially assigned, e.g., screening vs. lottery numbers.

superintendent Payzant's remark.[10] For simplicity, assume that the *a priori* priority assignments are independent of students' sibling status (as opposed to the current practice in Boston public schools). Coupled with the construction of the thick market where each slot in the district is up for trade, first let each student point only to a student (or students) holding a slot from his favorite school—note that unlike in TTC there are now multiple such students—and second, point only to that student (or those students) whose favorite school at this instant is one where he himself has a sibling. Such a pointing rule, which can lexicographically use other criteria (e.g., other priorities) during cases of ties, indeed ensures that a student is favored in trading due to his sibling status only when he is to be assigned to his sibling's school and is treated the same way as the rest of the students otherwise.

Whereas generalized trading mechanisms can readily attain efficiency by restricting to trades among the participants' most desired choices, achieving strategy-proofness becomes challenging due to the possibility of new participants joining the market: These participants may be endowed with a slot from a school that an existing participant may already be endowed with. Therefore, maintaining strategy-proofness requires particular attention to whether the pointing rule depends on the exogenous (e.g., priorities) or endogenous (e.g., preferences) specifications of a problem, as well as on the organization of the entry of new participants into the market. As a specific illustration of stability considerations that have been central in school choice, we propose and study a particular manifestation of the above ideas as a competitor to TTC. We call this mechanism the *equitable top trading cycles mechanism (ETTC)*. Unlike TTC, ETTC chooses the Pareto efficient and stable allocation in the above example. ETTC is also Pareto efficient and strategy-proof just like TTC (Propositions 1 & 2). Kesten (2010) showed that there is no strategy-proof mechanism that always selects the Pareto efficient and stable allocation when it exists. Therefore, ETTC cannot fully avoid justified envy situations. Nevertheless, we show that ETTC considerably improves the stability aspects of TTC without paying any cost when attention is restricted to pairwise cycles (Propositions 3 & 4). We also show that there is in general no specific type of top trading cycles mechanism, strategy-proof or not, that always eliminates justified envy among students whenever efficiency and stability are compatible.

Given the popularity of TTC as an efficient and strategy-proof mechanism, it is important to see how the two mechanisms perform in the controlled environment of the lab. The concern is that ETTC might induce different behavior in subjects, and thus theoretical comparison of the mechanisms might find no support in experiments, once the strategies of experimental subjects are considered. To this end, we designed experiments to compare the two. Ours is not the first attempt to evaluate the performance of TTC in the lab. TTC has been experimentally compared to other mechanism in the context of on-campus housing and school choice. In the former context, Chen and Sönmez (2002) experimentally compare TTC with a popular real-life mechanism, the so-called random serial dictatorship with squatting rights and show that TTC leads to superior welfare gains.[11] In the latter context, Chen and Sönmez (2006) find that TTC and

[10]Note that this issue cannot be resolved for TTC by simply giving higher priority to students who have siblings at some or at all schools. The reason is that, aside from the size of the market in which it operates, the pointing rule of TTC exclusively considers each student's priority at his *endowed* school and is agnostic about his priority at his *intended* assignment.

[11]For the same context, Guillen and Kesten (2012) compare TTC to another, much less popular real-life mechanism, which turns out to be equivalent to Gale-Shapley's (1962) deferred acceptance mechanism.

DA both provide higher performance in terms of efficiency relative to the much debated Boston mechanism primarily due to the significant differences in truth-telling rates. Our experiments also adopt much of their design idea. Calsamiglia et al. (2010) test TTC vs. DA and Boston when participants' preferences may be constrained in length. They show that constraining the preference list leads to efficiency losses. Guillen and Hing (2014) show that truth-telling under TTC can be distorted by third-party advice. Guillen and Hakimov (2015), however, show in a field experiment that advice has a strong positive effect on participants' truthful reporting, while disclosure of mechanics of the mechanism has detrimental effect of truthful reporting for a subsample of students. Pais and Pintér (2008) study the role of information in student placement mechanisms and find that the three mechanisms achieve their best performances when participants do not have information about either the preferences of the other participants or the priorities of the schools. Additionally, they report that TTC outperforms Boston and DA when participants hold partial or full information about each other's preferences and priorities. Guillen and Hakimov (2016) show that in the presence of full information, students tend to misrepresent their preferences even in a simple environment under TTC, especially in the case of uncertainty about other participants' behavior in the market. These findings, together with other experimental considerations, such as the minimization of favorable default choices, led to our choice of a complete information environment for the experiment.

The main purpose of our experiment lies in comparing TTC and ETTC from the point of view of incentives, welfare, and stability. We work with three environments: (1) *a designed environment* which mimics that of Chen and Sönmez (2006); (2) *a random-correlated environment* which generates a high degree of correlation among preferences; and (3) *a random-uncorrelated environment*. We observe similar rates of misrepresentation by subjects under both mechanisms. Nevertheless, ETTC produces significantly less justified envy in treatments with some correlation of the preferences. In addition to the usual comparison with respect to a number of justified envy situations, we employ the idea of Guillen and Kesten's (2012) ordinal efficiency test to compare the allocations by the two mechanisms with respect to fairness dominance. We refer to this as the ordinal fairness test. Given the same input of stated preferences, we determine the more plausible allocation from the criterion of fairness. Then we compare the number of dominations of one mechanism vs. the other. The test shows that ETTC is significantly more likely to generate less justified envy situations than TTC in all environments. We also couple our lab experiment with computational exercises, assuming full preference revelation. In this simulation, the advantage of ETTC becomes more prominent and the ETTC allocation is more likely to fairness dominate the TTC allocation, even in the case of the uncorrelated preferences of students.

In a follow-up work to the first version of the current paper in Kesten (2004), Morril (2015) proposed two simple variations of TTC, which he calls Clinch and Trade (CT) and First Clinch and Trade (FCT), to mitigate some fairness distortions under TTC. The idea behind both variations is to first allow a student to form a self-cycle (hence the term "clinch") for some school a if he is among the top q_a students and then to run TTC for the reduced market.[12] Morrill's rules do not allow the trade of a slot by any agent other than the one who currently has the

[12]Unlike TTC and ETTC, CT is bossy and its outcome is sensitive to the order in which cycles are processed. FCT, which is nonbossy, is equivalent to running TTC for a given problem after preferences and priorities are augmented to incorporate handling of self-cycles first.

highest priority, a feature that is different from ETTC but common with the original TTC,[13] since lower-priority agents who qualify (i.e., are within the quota) can only take the slot for themselves. This limits the scope of Morrill's rules compared to ETTC and can be seen using the above example: both CT and FCT would yield the same result as TTC as there are no self-cycles, whereas ETTC gives the stable and efficient outcome. While the self-cycle feature is readily embedded in the ETTC algorithm,[14] ETTC as well as the generalized trading cycles mechanisms we discuss also have the two distinctive features of working through a thick market (where all slots are up for trade) and utilizing specific pointing rules to eliminate justified envy among other potential goals.

The rest of the paper is organized as follows. The next section introduces the formal model. Section 3 describes TTC and its properties. Section 4 delineates the generalization of the top trading cycles idea and provides theoretical comparisons of TTC and ETTC. Section 5 offers the experimental analysis, and section 6 concludes. All proofs and experimental instructions are relegated to the appendix.

2 School Choice Problem

In a school choice problem, a certain number of students are to be placed in a certain number of schools. Each school has a certain number of available slots, and the total number of slots is no less than the number of students.[15] Let $I = \{i_1, i_2, \ldots, i_n\}$ denote the set of students. A generic element in I is denoted by i. Let $S = \{s_1, s_2, \ldots, s_m\}$ denote the set of schools. A generic element in S is denoted by s. Let q_s be the number of available slots at school s, or the *quota* of s.[16] Each student has strict preferences over all schools. Let P_i denote the preferences of student i. Let R_i denote the at-least-as-good-as relation associated with P_i. Formally, we assume that R_i is a linear order, i.e., a complete, transitive, and anti-symmetric binary relation on S. That is, for any $s, s' \in S$, $s \; R_i \; s'$ if and only if $s = s'$ or $s \; P_i \; s'$. A strict priority order of all students for each school is exogenously given. Let $\succ_s$ denote the priority order for school s.

Since priorities are pre-specified, a **school choice problem**,[17] or simply a problem, is a preference profile $P = (P_i)_{i \in I}$. Let $\mathcal{R}$ be the set of all problems. An **allocation** μ is a list of (final) placements such that each student is placed in a single school and the number of students

[13] Indeed, once self-cycles are carried out, the trading under CT and FCT is exactly the same as that of TTC. This implies that under both variations, the number of market participants in each round is essentially as many as those under TTC and the pointing rule is that of TTC.

[14] The idea of carrying out self-cycles first and then executing larger exchanges also appears in Kesten (2009), in an adaptation of the top trading cycles idea to the random assignment problem.

[15] This assumption is not needed for any of the results in this paper. In the absence of this assumption, the identical analysis can be carried out by introducing the "null school" that represents not being placed in any *real* school and that has unlimited seats.

[16] It is straightforward to capture outside options in this model such as private schooling by assuming the existence of a school with a quota of at least $|I|$. See Kesten and Kurino (2016) for the implications of such an assumption on the trade-offs in school choice.

[17] The student placement problem is also closely related to the "house allocation problem" in which there is a set of objects collectively owned by the society. In that problem, however, the quota for each house is one. See for example, Pápai (2000), Abdulkadiroğlu and Sönmez (1998, 1999), Ehlers et al. (2002), Ehlers (2002), Ehlers and Klaus (2003), and Kesten (2009).

placed in a particular school does not exceed the quota of that school. Formally, it is a function $\mu : I \to S$ such that for each $s \in S$, $|\mu^{-1}(s)| \leq q_s$. Given $i \in I$, $\mu(i)$ denotes the placement of student i at μ and given $s \in S$, $\mu^{-1}(s)$ denotes the set of students placed in school s at μ. Let $\mathcal{M}$ be the set of all allocations. An allocation μ is **non-wasteful** if no student prefers a school with an unfilled quota to the school he is placed in, i.e., for all $i \in I$, $s \; P_i \; \mu(i)$ implies $|\mu^{-1}(s)| = q_s$. An allocation μ is **Pareto efficient** if there is no other allocation that makes all students at least as well off and at least one student better off, i.e., there is no $\alpha \in \mathcal{M}$ such that $\alpha(i) \; R_i \; \mu(i)$ for all $i \in I$ and $\alpha(j) \; P_j \; \mu(j)$ for some $j \in I$.

The central equity notion is "stability." We say that **student** i **justifiably envies student** j **for school** s at a given allocation μ (or, alternatively, the *priority of student* i *for school* s *is violated*) if i would rather be placed in s, to which some student j who has a lower $s-$priority than i has been placed in, i.e., $s \; P_i \; \mu(i)$ and $i \succ_s j$ for some $j \in \mu^{-1}(s)$. An allocation is **stable** (or fair) if it is non-wasteful and no student's priority for any school is violated.

A **school choice mechanism**, or simply a mechanism φ, is a systematic procedure that chooses an allocation for each profile of preference reports. Formally, it is a function $\varphi : \mathcal{R} \to \mathcal{M}$. Let $\varphi(P)$ denote the allocation chosen by φ for P and let $\varphi_i(P)$ denote the placement of student i at this allocation. A mechanism is Pareto efficient (stable) if it always selects Pareto efficient (stable) allocations. A mechanism φ is **strategy-proof** if it is a dominant strategy for each student to truthfully report her preferences. Formally, for every $P \in \mathcal{R}$, every student $i \in I$, and every misreport P'_i by i, $\varphi_i(P) \; R_i \; \varphi_i(P'_i, P_{-i})$.

3 Top Trading Cycles Mechanism

Abdulkadiroğlu and Sönmez (2003) propose what they call the **top trading cycles mechanism (TTC)**. TTC is based on *Gale's top trading cycles procedure* proposed in the context of "housing markets" (Shapley and Scarf, 1974). In a housing market, there is a set of "indivisible objects" (e.g., houses) each of which is initially *assigned*[18] to a different agent among a set of "agents." Gale's top trading cycles procedure works as follows:[19] Each agent points to the agent who is assigned to his best choice object. Since the number of agents is finite, there is at least one cycle. Then in each cycle, the corresponding trades are performed (i.e., each agent in a cycle receives the object assigned to the agent he points to), and these agents and objects are removed. Some agents may not be able to participate in a cycle and therefore remain in the market. The same procedure is then applied to the new market and so on. The algorithm terminates when there are no agents left. This procedure yields the unique core allocation of this housing market (Roth and Postlewaite, 1977). The core mechanism for the housing market context has been shown to be the unique strategy-proof and Pareto efficient mechanism that ensures that no agent receives a worse object than his initial assignment (Ma, 1994).

Gale's top trading cycles procedure cannot be directly applied to the school choice context. Since there may now be multiple copies (slots) of a particular object (school), this difference in

[18]Throughout, we use the term "assignment" to refer to interim (tentative) pairings of agents (students) with objects (slots of schools) under a top trading cycles procedure, while reserving the term "placement" to mean a final outcome.

[19]The procedure we describe here is not the same as the one proposed by Gale. But, the two are equivalent. For reasons to be made clear shortly, we adopt this alternative procedure.

the models necessitates further modification of the procedure. TTC is one such adaptation of the procedure. For a given problem, the outcome of TTC can be found via the following algorithm:[20]

Step 1: Each student who has the highest priority for a school is assigned to all slots of that school. (A student may be assigned to the slots of different schools.) Each student points to the student (possibly himself) who is assigned to (all slots of) his best choice. There is at least one cycle. Each student in a cycle is placed in the school that was assigned to the student he is pointing to. Since each student who is part of a cycle is already placed in a school, he is removed and the number of available slots at that school is decreased by one.

In general,

Step k, $k \geq 2$: All of the remaining slots of each school that were assigned to a student who was part of a cycle at step k-1 are assigned to the student with the highest priority for that school among the remaining students. (A student may be assigned to the slots of different schools.) Each student points to the student (possibly himself) who is assigned to (all remaining slots of) his best choice among the remaining schools. There is at least one cycle. Each student in a cycle is placed in the school that was assigned to the student he is pointing to. Since each student who is part of a cycle is already placed in a school, he is removed and the number of available slots at that school is decreased by one.

The algorithm terminates when no student is left.

Since it is based on Gale's top trading cycles procedure, TTC inherits that procedure's desirable properties. The first such property is Pareto efficiency.

Proposition (Abdulkadiroğlu and Sönmez, 2003) *The top trading cycles mechanism is Pareto efficient.*

However, TTC is not stable.[21] (We return to this aspect of TTC shortly.) The second important property TTC has is strategy-proofness.

Proposition (Abdulkadiroğlu and Sönmez, 2003) *The top trading cycles mechanism is strategy-proof.*

Proposition (Roth, 1982) *A Pareto efficient and stable allocation may not always exist and if it exists, it is unique.*

[20]The algorithm we give here is equivalent to the one proposed by Abdulkadiroğlu and Sönmez (2003). Since this version of the algorithm will make it easier for us to compare it with the alternative adaptation of Gale's top trading cycles procedure that we propose later in the paper, we work with this equivalent algorithm. This alternative algorithm appears in Kesten (2006).

[21]Kesten (2006) gives a sufficient and necessary condition for the equivalence of DA and TTC, or to restore the stability, or the resource monotonicity, or the population monotonicity of TTC.

4 Generalized Top Trading Cycles Mechanisms

Abdulkadiroğlu and Sönmez (2003) have adopted Gale's top trading cycles procedure to school choice and propose a Pareto efficient and strategy-proof mechanism based on this procedure. Despite its appealing properties, this mechanism leaves room for improvement as far as fairness is concerned. To elaborate, in the TTC algorithm all slots of a given school are assigned to the student with the highest priority for that school. Since all students who have this school as a best choice have to point to this student, this particular student is given all the trading power of this school's slots. As we have illustrated in the introduction, such "excessive" power may, however, result in the justified envy of students who may have lower priority for that school but higher priority for the school this student is placed in. Of course, due to the incompatibility between stability and Pareto efficiency, it is not always possible to totally avoid cases of justified envy. Nonetheless, we shall show that it may still be possible to considerably reduce the number of these cases by considering alternative adaptations of Gale's top trading cycles procedure.

The key idea of our alternative approach is the construction of a much thicker trading market by increasing the number of "active" participants at any given instant of the trading market. Since TTC allows only the current highest priority students to be part of the trading process, there are at most $|S|$ participants at any instant of this algorithm. Such a thin market may, however, entail justified envy for non-participants. Instead of assigning all q_x slots of a given school x to the highest $x-$priority student, we propose assigning all slots to all the q_x students, with the highest $x-$priority giving one slot to each and endowing them with equal trading power. This would, for example, lead to an initial market with $\sum_x q_x$ active participants. And, in order to maintain the thickness of the market throughout the process, whenever a student is removed from the market, any unallocated endowments will be "inherited" by the next highest priority student remaining in the market.

An important subtlety that arises with a market containing multiple students who are assigned a slot from the same school is determining the terms of trade. A *pointing rule* specifies for each school-student pair which school-student pair(s) should be pointed to among those who contain the favorite school currently remaining in the market at any instant of the algorithm. The choice of the pointing rule will shape the ensuing incentive and fairness properties of the resulting trading mechanism.

While there is a rich set of pointing rules one can conceive, by and large we classify them into two groups. Consider any step t of the trading process applied to a given problem. Let (i, x) be a pair whose[22] currently remaining favorite school in the market is some $y \in S\backslash\{x\}$. Let $X_t^{\to y}$ be the set of such pairs (i.e., those who are endowed with a slot from x and whose remaining favorite school at step t is some $y \in S\backslash\{x\}$) and let Y_t be the set of pairs who are endowed with a slot from y (i.e., $Y_t = \cup_z Y_t^{\to z}$).[23]

- *Pointing rules that depend only on the exogenous parameters.* Formally, a rule of this

[22]With a slight abuse of language, the favorite school of a student in a given school-student pair will also be referred as the favorite school of the pair.

[23]Note that, for notational brevity, we denote schools by lowercase letters and the set of pairs containing a slot from a particular school with the corresponding uppercase letter.

kind is a (possibly multivalued)[24] mapping $r : (X_t^{\to y}, Y_t, E) \rightrightarrows Y_t$ where E captures an exogenously given constraint such as the priority structure. These pointing rules in effect describe which pair(s) in Y_t will be pointed to by a member in $X_t^{\to y}$, depending on the characteristics of the students contained in the two sets. Specifically, rules of this kind are agnostic about the preferences of the students in Y_t. For example, such a rule may be the one based on the lottery draw used to break ties in priorities in school districts such as Boston, where many students fall in the same priority class. For example, each pair in $X_t^{\to y}$ points to that pair in Y_t which contains the student with the best lottery draw or the student with the highest average priority for the schools which that student is not assigned at a student-school pair at the current step of the procedure i.e., students are dynamically re-ranked at each step based on their average priority at other schools to increase the likelihood of a stable allocation. Alternatively, the pointing rule may be chosen to fulfill certain affirmative action considerations and to promote diversity in student composition, as in San Francisco; e.g., each pair in $X_t^{\to y}$ points to those pairs in Y_t which contain a minority student.

- *Pointing rules that depend on the priority structure and the preferences.* Formally, a rule of this kind is a (possibly multivalued) mapping $r : (X_t^{\to y}, Y_t, E, P) \rightrightarrows Y_t$. These pointing rules in effect describe which pair(s) in Y_t will be pointed to by a member in $X_t^{\to y}$, depending not only on the two sets, but possibly also on the preferences of the students in Y_t. For example, this kind of rule may allow for considering the priority of students in Y_t for their favorite school(s) remaining in the market.

TTC is an example of the former type of pointing rules. Indeed, at any step t of TTC each pair $(i, x) \in X_t$ points to the pair in Y_t containing the student with the highest $y-$priority.[25] In the next section we present yet another example of these rules that specifically aims to establish fairness among pairwise exchanges. An important advantage of these pointing rules is that because such rules depend only on exogenous specifications, the top trading cycles mechanisms they induce are readily strategy-proof, in addition to being Pareto efficient.[26]

On the other hand, by allowing for dependence on the internal specifications of a problem, the latter type of pointing rules may render more ground in terms of stability, possibly at the expense of strategy-proofness, while also maintaining Pareto efficiency. In the sequel we also discuss the construction of a top trading cycles mechanism based on an intuitive example of this kind of pointing rule.

[24] A multivalued pointing rule (or correspondence) allows for multiple overlapping nested cycles to form, which in turn can offer more flexibility in accommodating additional distributional considerations of the planner. Section 4.1 discusses an example of such a correspondence that generates cycles with minimal justified envy.

[25] Formally, the pointing rule underlying TTC is as follows. For any given sets $X_t^{\to y}$ and Y_t of any step t, $r^{TTC}(X_t^{\to y}, Y_t, \succ_y) = \{(j, y) \in Y_t : j \succsim_y j' \text{ for any } (j', y) \in Y_t\}$.

[26] More precisely, this statement is true for pointing rules that are functions. In cases where the pointing rule is a correspondence, two or more cycles may be nested and an additional exogenous criterion also needs to be used for cycle selection. We shall discuss cycle selection rules subsequently.

4.1 Equitable Top Trading Cycles

Recall that in TTC each student-school pair points to the pair that contains the highest priority student for the school contained in the latter pair. We propose an alternative adaptation of Gale's top trading cycles procedure in which *each student-school pair points to the pair that contains the highest priority student for the school contained in the former pair.* This proposal is based on a "dual" pointing rule of that of TTC which ensures that whenever a cycle forms between two student-school pairs, the students included in that cycle have the highest priority for their favorite schools among their competitors at that step of the trading market.

Here is a description of our proposal. At the first step, for each school, slots are assigned to students one by one following their priority order to form student-school pairs in a thick market. A student can be contained in more than one student-school pair. We denote a student-school pair by (i, s), where i is a student and, with a slight abuse of notation, s denotes one slot from school s. Each student-school pair (i, s) points to the student-school pair (i', s') such that (i) school s' is the best choice of student i and, (ii) student i' is the student with the highest priority for school s among the students who are assigned to a slot from school s'. If there is already a student-school pair at which student i is assigned to one slot from his best choice school, then all student-school pairs containing him point to that student-school pair. Since the number of student-school pairs is finite, there is at least one cycle. In each cycle, corresponding trades are performed; i.e., if a student-school pair (i, s) is pointing to the pair (i', s') in a cycle, then student i is placed in school s' and he is removed, as well as the slot to which student i' is assigned.

It is possible that some student-school pairs that contain the same student appear in the same cycle or in different cycles. In such a case, the extra slots of that school (to which different student-school pairs containing him are pointing in other cycles) remain to be "inherited" by the remaining students.[27] An important twist of our algorithm is that the slots that remain to be inherited at the end of a step t, $t \geq 1$, are not necessarily inherited at the very next step by the remaining students. The inheritance of slots of a school s does not take place until all students who are contained in a student-school pair including school s at step t are removed.[28] Immediately after the last student who was contained in a student-school pair with a slot from school s at step t is removed, all slots of school s which thus far remain to be inherited are inherited by the remaining students, one by one, following the priority order for school s, i.e., these students are assigned those slots to again form student-school pairs. At each step, student-school pairs again point to each other in the way described above. Corresponding trades are carried out in each cycle and some slots remain to be inherited at the appropriate subsequent step. The procedure continues in a similar fashion. The following algorithm summarizes this procedure for a given problem:

In general,

[27] The idea of "inheritance of slots" we use here is in part inspired by Pápai (2000), who introduces and characterizes quite a large family of rules.These rules, which she calls "endowment inheritance rules," are also based on Gale's top trading cycles algorithm.

[28] As will be clear shortly, this restriction ensures that our mechanism is strategy-proof.

Step k, $k \geq 1$: **Inheritance:** *If $k = 1$, then all slots are available for inheritance and assigned to students one by one following schools' priority orders to form student-school pairs. Otherwise, for each school s such that (i) there are slots of school s which remained to be inherited from previous steps, and (ii) there are no existing pairs who were assigned to a slot of school s at a previous steps of the algorithm, i.e., the market has completely run out of the slots of school s, the slots which remained to be inherited are assigned to the remaining students one by one following the priority order for school s to form new student-school pairs.*

Pointing: *Each student-school pair (i, s) points to the student-school pair (i', s') such that (i) school s' is the best choice of student i and, (ii) student i' is the student with the highest priority for school s among the students who are assigned to a slot from school s'. If student i is already assigned to one slot from his best choice school, then all student-school pairs containing him point to that student-school pair.*

Trading: *There is at least one cycle. In each cycle, corresponding trades are performed, and all student-school pairs that participate in a cycle are removed. It is possible that student-school pairs containing the same student–say, student i–appear in the same cycle or in different cycles. In such a case, student i is placed in his best choice and the other slots of that school (to which the student-school pairs containing him are pointing in those other cycles) remain to be inherited. For each student-school pair (i, s) that participates in a cycle, the slots assigned to student i in other student-school pairs that do not participate in a cycle also remain to be inherited.*

Remark 1: *Note that slots remain to be inherited in one of two ways: (1) More than one student-school pair containing the same student participates in a cycle or cycles. That student is then given one slot from his best choice and the other slots of his best choice (to which he is pointing in any other cycle) remain to be inherited; (2) A student-school pair participates in a cycle and there are other student-school pairs containing the same student which do not participate in a cycle: Then the slots assigned to him in those other student-school pairs remain to be inherited.*

We call the mechanism that associates the outcome of the above algorithm to each problem the **equitable top trading cycles mechanism (ETTC)**. Next, we give two examples to illustrate the dynamics of this algorithm:

Example 1. Let $I = \{i_1, i_2, i_3, i_4, i_5, i_6, i_7, i_8\}$ and $S = \{s_1, s_2, s_3, s_4\}$, where the schools s_1 has one slot, s_2 has three slots, and s_3 and s_4 have two slots each. The priorities for the schools and the preferences of the students are given as follows:

$\succ_{s_1}$	$\succ_{s_2}$	$\succ_{s_3}$	$\succ_{s_4}$
i_1	i_4	i_1	i_6
i_2	i_2	i_5	i_2
i_5	i_3	i_3	i_5
i_7	i_8	i_7	i_3
i_6	i_7	i_4	i_7
i_3	i_1	i_2	i_8
i_8	i_5	i_6	i_4
i_4	i_6	i_8	i_1

P_{i_1}	P_{i_2}	P_{i_3}	P_{i_4}	P_{i_5}	P_{i_6}	P_{i_7}	P_{i_8}
$\boxed{s_2}$	$\boxed{s_3}$	$\boxed{s_1}$	$\underline{s_1}$	$\boxed{s_4}$	$\boxed{s_2}$	$\boxed{s_4}$	s_1
s_1	s_1	$\underline{s_3}$	$\boxed{s_3}$	s_4	s_4	s_2	s_3
s_3	s_2	s_4	s_2	s_3	s_2	s_1	$\boxed{s_2}$
s_4	s_4	s_2	s_4	s_1	s_1	s_3	s_4

Step 1: The student-school pairs are: (i_1, s_1), $(i_4, s_2), (i_2, s_2), (i_3, s_2), (i_1, s_3)$, $(i_5, s_3), (i_6, s_4)$, and (i_2, s_4). Note that i_1 and i_2 have slots of two schools each. Student i_1 has slots of s_1 and s_3, and student i_2 has slots of s_2 and s_4.

We determine which student-school pair points to which pair. Consider, for example, the student-school pair (i_1, s_1). The best choice for student i_1 is school s_2; hence, student i_1 will point to one of the student-school pairs that contain school s_2, which are $(i_4, s_2), (i_2, s_2)$, and (i_3, s_2). Since student i_2 has a higher s_1-priority than i_4 and i_3, student-school pair (i_1, s_1) points to (i_2, s_2).

Note that (i_1, s_3), containing the same student, i_1, points to a different student-school pair, namely (i_3, s_2), as i_3 has a higher s_3-priority than i_2 and i_4.

This leads to the fact that student i_1 can potentially participate in two different cycles, or can participate twice in one cycle.

A similar situation occurs with student-school pairs containing student i_2. The best choice for student i_2 is school s_3; hence, student-school pair (i_2, s_2) will point to one of the student-school pairs that contain school s_3, which are (i_1, s_3) and (i_5, s_3). Since student i_1 has a higher s_2-priority than student i_5, thus student-school pair (i_2, s_2) points to (i_1, s_3).

Note that (i_2, s_4), containing the same student, i_2, points to a different student-school pair, namely (i_5, s_3), as i_5 has a higher s_4-priority than student i_1.

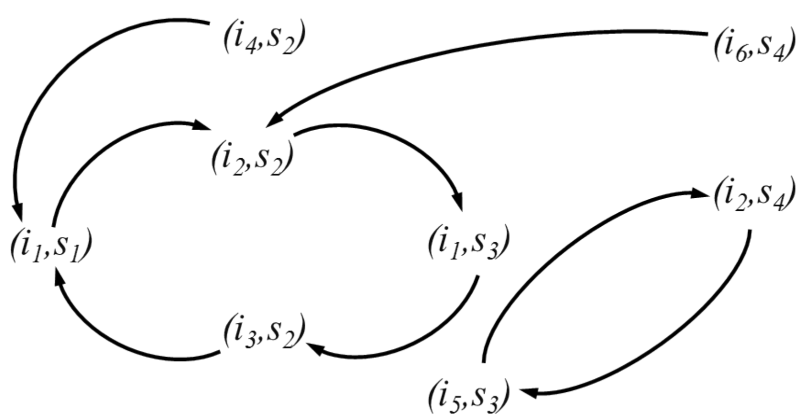

Figure 1: Step 1 of the ETTC for the Example 2.

Next we identify the cycles of step 1 (see Figure 1). Two cycles form in this step: one four-way

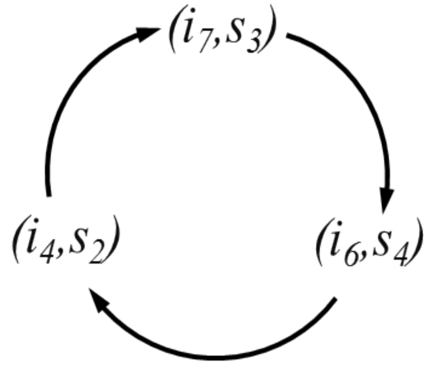

Figure 2: Step 2 of the ETTC for the Example 2.

cycle $\{(i_1, s_1), (i_2, s_2), (i_1, s_3), (i_3, s_2)\}$, and one pairwise cycle $\{(i_5, s_3), (i_2, s_4)\}$. Corresponding trades are implemented: that is, student i_1 is placed in school s_2, student i_2 is placed in school s_3, student i_3 is placed in s_1, and student i_5 is placed in s_4.

Note that the cycles of the first step involve two cases that never appear in TTC: once a student participates in the same cycle twice (student i_1), and once a student participates in two different cycles (student i_2).

Step 2: There is one slot at school s_2 and one slot at school s_3 to be inherited from the previous step. Since there is still a student who was assigned an s_2 slot (namely student i_4) previously, there is no inheritance of the s_2 slot at this step. Since no student who is assigned an s_3 slot is left, i.e., the market has completely run out of slots of s_3, the slot of s_3 is inherited by student i_7. Thus, the remaining student-school pairs are (i_4, s_2), (i_7, s_3), and (i_6, s_4) (see Figure 2). The only cycle at this step is a three-way exchange cycle, as student-school pair (i_4, s_2) points to (i_7, s_3), (i_7, s_3) points to (i_6, s_4), and (i_6, s_4) points to (i_4, s_2). The corresponding trade places student i_4 in school s_3, student i_7 in school s_4, and student i_6 in school s_2.[29]

Step 3: There is one slot of school s_2 to be inherited from the previous step. Since there are no students who is assigned an s_2 slot, this slot is inherited by student i_8. This is the only remaining student-school pair, thus it points to itself, and student i_8 is placed in s_2.

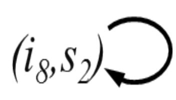

Figure 3: Step 3 of the ETTC for the Example 2.

The ETTC allocation is indicated in boxes. No student has justified envy in this allocation.

[29]Note that if we allowed for the inheritance of the s_2 slot at this step, student i_8 would inherit a slot of s_2. Thus, (i_6, s_4) would point to (i_8, s_2) instead of (i_4, s_2). This would lead to the following three-way cycle: (i_8, s_2) points to (i_7, s_3), (i_7, s_3) points to (i_6, s_4), and (i_6, s_4) points to (i_8, s_2). This in turn would lead to a different allocation.

Now consider the TTC allocation, which is underlined. The only difference in the two allocations concerns students i_3 and i_4. In TTC allocation students i_3 and i_8 have justified envy of student i_4 in school s_1. ⧫

As with TTC, ETTC is also Pareto efficient.

Proposition 1 *The equitable top trading cycles mechanism is Pareto efficient.*

As with TTC, ETTC is also group strategy-proof.

Proposition 2 *The equitable top trading cycles mechanism is group strategy-proof.*

ETTC not only shares the same compelling properties with TTC, it also has more to offer in terms of fairness. To make the argument more transparent, we give examples that compare ETTC with TTC in terms of stability aspects. The advantage of ETTC over TTC is most apparent when we consider cycles consisting of two student-school pairs.

Example 2. ETTC vs. TTC when there are cycles consisting of two student-school pairs: Let $I = \{i_1, i_2, \ldots, i_7\}$ and $S = \{s_1, s_2\}$, where school s_1 has three slots, and school s_2 has four slots. The priorities for the schools and the preferences of the students are given as follows:

$\succ_{s_1}$	$\succ_{s_2}$
i_1	i_4
i_2	i_5
i_3	i_6
i_7	i_7
i_5	i_3
i_6	i_1
i_4	i_2

P_{i_1}	P_{i_2}	P_{i_3}	P_{i_4}	P_{i_5}	P_{i_6}	P_{i_7}
$\boxed{s_2}$	$\boxed{s_2}$	$\boxed{s_2}$	$\underline{s_1}$	$\boxed{s_1}$	$\boxed{s_1}$	$\boxed{s_1}$
s_1	s_1	s_1	$\boxed{s_2}$	s_2	s_2	$\underline{s_2}$

Here, students i_4, i_5, i_6, and i_7 have identical preferences, and are competing for a slot at school s_1. According to the priority order $\succ_{s_1}$, it is student i_7 who "deserves" a slot at school s_1 before any other student among the four students.

When we apply the TTC algorithm to this problem, three cycles form in the first three steps. In these cycles, students i_1, i_2, and i_3 each trade one slot of school s_1 for one slot of school s_2 with students i_4, i_5, and i_6, respectively. Student i_7 inherits the last slot of school s_2 and forms a self-cycle. The TTC allocation is the underlined allocation above. Note that at this allocation, student i_7 has justified envy of all three students who have been placed in school s_1. Let us now apply the ETTC algorithm to the same problem.

The first step of the ETTC algorithm is depicted in Figure 4. The only cycle is $\{(i_3, s_1), (i_7, s_2)\}$. Student i_7 is placed in school s_1 and student i_3 in school s_2. Then these student-school pairs are

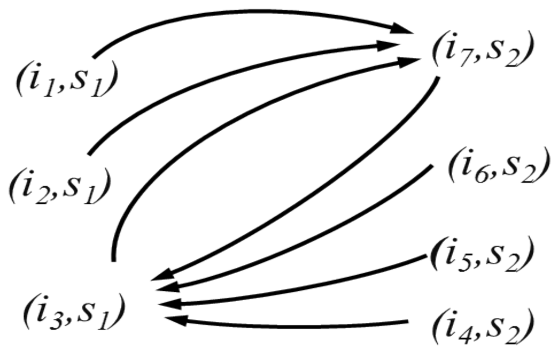

Figure 4: The first step of ETTC algorithm for Example 3.

removed. In the second step, all student-school pairs containing an s_1 slot point to the student-school pair (i_5, s_2) because student i_5 has the highest priority for school s_1 among the remaining students who are assigned to a slot of school s_2. All student-school pairs containing a slot of school s_2 point to the student-school pair (i_1, s_1) because student i_1 has the highest priority for school s_2 among the remaining students who are assigned to a slot of school s_1. Then the only cycle is $\{(i_5, s_2), (i_1, s_1)\}$. Student i_5 is placed in school s_1 and student i_1 to school s_2. Continuing in a similar way, we obtain the allocation marked with rectangles. Note that this allocation is stable.[30] Now student i_7 is better off, whereas student i_4 is worse off compared to the outcome of TTC. Unlike ETTC, TTC severely violates the priorities of student i_7 by giving all the trading power for slots of school s_2 to student i_4. ♦

One can also observe that the magnitude of the priority violations caused by TTC in the way described in Example 2 grows with the number of available slots. Another issue that may also impact the relative stability difference between the two mechanisms is that student preferences are often correlated in practice,[31] and this salient feature further restricts the scope of large exchanges. For example, in the extreme case of perfect correlation–i.e., when preferences are identical across all students–only self-cycles form under either mechanism.

The next two results establish the superiority of ETTC over TTC for each problem when only pairwise cycles are allowed to form.

Proposition 3 *Suppose there are two schools. Let i be any student who is entitled to one slot at a school $s \in S$, i.e., he is one of the students in the top q_s priority group. Then, under ETTC student i never has justified envy. This is not the case under TTC.*

Proposition 4 *Suppose there are two schools. If ETTC selects an unstable allocation for a problem, then TTC also selects an unstable allocation for the same problem. However, the converse*

[30] Of course, this is not to say that it will be the case in general. In fact, there is no strategy-proof and Pareto efficient mechanism that selects the Pareto efficient and stable matching whenever such a matching exists (Kesten, 2010).

[31] Motivated by this observation, Abdulkadiroğlu et al. (2011) study a stylized model of school choice where students have identical preference rankings.

is not necessarily true.[32]

Remark 2. *The two-school restriction in Propositions 3 and 4 renders a clear analytical comparison of the two mechanisms possible. Nevertheless, these results are not meant to say that ETTC improves upon TTC only when there are two schools. Indeed, the simulation results for more general cases, reported in the subsequent experimental analysis, also indicate a smaller number of priority violations on average under ETTC relative to TTC.*

In problems where larger cycles form, ETTC is unfortunately not necessarily more stable than TTC when compared problem-specifically rather than on average. The following example illustrates this point.

Example 3. Let $I = \{i_1, i_2, i_3, i_4\}$ and $S = \{s_1, s_2, s_3\}$, where the school s_1 has two slots, and school s_2 and s_3 have one slot each. The priorities for the schools and the preferences of the students are given as follows:

$\succ_{s_1}$	$\succ_{s_2}$	$\succ_{s_3}$
i_1	i_4	i_3
i_2	i_2	i_1
i_3	i_1	i_2
i_4	i_3	i_4

P_{i_1}	P_{i_2}	P_{i_3}	P_{i_4}
$\underline{s_3}$	$\boxed{s_3}$	$\underline{\boxed{s_2}}$	$\underline{\boxed{s_1}}$
$\boxed{s_1}$	$\underline{s_1}$		

The ETTC allocation, induced by a three-way exchange among pairs (i_2, s_1), (i_3, s_3), and (i_4, s_2), is indicated in boxes. Note that student i_1 has justified envy over student i_2 in school s_3. On the other hand, the underlined allocation, induced by a three-way exchange among pairs (i_1, s_1), (i_3, s_3), and (i_4, s_2), is the outcome of TTC, where no student has justified envy. Intuitively, to resolve the competition between students i_1 and i_2 for school s_3, ETTC uses the priority order for s_2 and TTC that for s_1. Since priority order for s_1 is consistent with the two students' priority ranking at s_3, TTC yields the stable allocation. Observe that, if instead student i_2 had higher s_3-priority than i_1, then the two mechanisms would yield the same allocations, with that of ETTC being stable and that of TTC being unstable. ⧫

Remark 3: *A common tool in practice is to use a single lottery draw to break the ties, when priorities are coarse, i.e. when many students fall in the same priority class based on criteria such as walk-zone. The situation in Example 3 leads to an important practical point when priorities are generated in this fashion. For example, suppose that s_2 and s_3 are two non-walk zone schools for i_1 and i_2. Then, if student i_1 has higher s_2-priority than i_2 based on the lottery draw, he must also have higher s_3-priority than i_2. But this correlation would then eliminate the kind of justified envy observed under ETTC as in the above example. This suggests that when priority orders are correlated, problems where TTC leads to fewer justified envy situations than ETTC become less likely.*

[32] The conclusion of Proposition 4 does not extend to arbitrary problems. However, we conjecture that a similar result would still be obtained *on average* if problems are randomly generated. While showing this result formally is beyond the scope of the current paper, our simulation results confirm this conjecture.

4.2 Cyclewise Equitable Top Trading Cycles Mechanisms

One plausible idea is to establish fairness among arbitrary sized trades that possibly involve more than two student-school pairs. To fix ideas, consider a potential three-way cycle that will involve a first student-school pair chosen from those pairs containing school a, a second pair chosen from those containing school b, and a third chosen from those containing school c. In order to avoid possible priority violations in the resulting allocation, fairness would require that the first pair we choose contain the highest $b-$priority student, the second pair contain the highest $c-$priority student, and the third contain the highest $a-$priority student. This reasoning motivates an interesting pointing rule that belongs to the second type of these rules.

For simplicity, first suppose that (any student in) any student-school pair containing the same school has the same favorite school, e.g., every student-school pair of the kind $(\cdot, x)$ will point to a pair containing school y. We require that *each student-school pair* (e.g., all pairs of the kind $(\cdot, a)$ in the example) *point to the student-school pair holding his favorite school* (e.g., one of the kind $(\cdot, b)$ in the example) *which contains the current highest priority student for the school which is the favorite school of the latter kind of pairs* (e.g., highest $c-$priority student contained in pairs of the kind $(\cdot, b)$ in the example).

The rationale behind such a pointing rule is clear. In a top trading cycles procedure for a student-school pair, being pointed to is tantamount to gaining trading power. Thus, when determining which pair needs to be pointed to among a set of pairs that are competing for a particular school, stability necessitates that the decision should be based on students' priority for that particular school.

One subtlety that is missing in the above discussion is that two pairs that contain the same school may have different favorite schools at a given instant of the algorithm. For example, suppose that while the favorite school of pair (j, y) is school z', the favorite school of pair (k, y) is a different school z''. Meanwhile, suppose that y is the favorite school of pair (i, x) at this instant. Then the question is: which of the two pairs will be pointed to by pair (i, x)?

One solution to this question is to consider a pointing *correspondence,* whereby a pair possibly points to multiple pairs. For instance, in the above example we can require (i, x) to point to the highest $z'-$priority pair among those pairs containing school y whose favorite school is z', *and* point to the highest $z''-$priority pair among those pairs containing school y whose favorite school is z''. On the other hand, pointing correspondences give rise to cycle selection issues, as we may now have overlapping cycles once a pair is allowed to point to multiple pairs. Therefore, one can imagine a rich inventory of cycle selection methods, including those based on cycle size, composition (e.g., affirmative action considerations), fairness across cycles, etc.

Although pointing correspondences may allow for superior fairness gains compared to the first type of pointing rules that depend only on the exogenous specifications, the preceding discussion leads to two observations about such correspondences. First, they involve critical choice decisions depending on the desideratum of the market designer.[33] And second, although the Pareto efficiency of the resulting mechanisms is ensured, the use of cycle selection rules makes these mechanisms vulnerable to strategic behavior.[34]

[33] Similar choice decisions are the subject of Roth et al. (2005) in the context of kidney exchange, where the authors propose and discuss a wide range of cycle and chain selection rules.

[34] In fact, it is possible to construct examples showing that strategy-proofness is lost regardless of the cycle

Given that pointing correspondences may be able to improve upon ETTC on the stability front, a curious question at this point is whether there can be any pointing rule/correspondence such that the associated mechanism selects the stable and efficient allocation when it exists. The following example leads to a negative answer.

Example 4. (No top trading cycles mechanism, strategy-proof or not, selects the stable and efficient allocation) Let $I = \{i_1, i_2, i_3\}$ and $S = \{s_1, s_2, s_3\}$, where each school has one slot. The priorities for the schools and the preferences of the students are given as follows:

$\succ_{s_1}$	$\succ_{s_2}$	$\succ_{s_3}$
i_1	i_2	i_1
i_3	i_3	i_2
i_2	i_1	i_3

P_{i_1}	P_{i_2}	P_{i_3}
$\boxed{s_2}$	$\underline{s_1}$	$\boxed{s_1}$
s_3	$\boxed{s_3}$	$\underline{s_3}$
s_1	s_1	s_2

The school-student pairs that form are (i_1, s_1), (i_1, s_3), and (i_2, s_2). Regardless of the pointing rule, a single cycle forms: the pairwise cycle $\{(i_1, s_1), (i_2, s_2)\}$. Therefore, any top trading cycles mechanism selects the underlined allocation above. However, the unique Pareto efficient and stable allocation is the one marked with boxes. ♦

5 Experimental design

We designed an experiment to compare the performances of the TTC and ETTC mechanisms, with particular attention paid to the question of the elimination of justified envy situations. The importance of experimental comparison comes from the fact that different mechanisms could lead to different participants behavior, and thus distort the theoretical comparison. Though the main insights of experiments are participants strategies, they allow us to make comparisons of the allocations, even given partially suboptimal reports. We implemented two treatments which differ only in the mechanism which is used to allocate students to schools. We ran three environments in each treatment. In each environment, there are 10 students and 10 school slots.

5.1 Designed environment

There are three schools in this environment. Schools A and B have three slots, while school C has four slots. The preferences are generated by the algorithm which is similar to the one used in Chen and Sönmez (2006) and approximates the preference formation in the real market. The payoff of each student depends on the quality of the school,[35] the proximity of the school and the presence of siblings in the school at the time of application.

Thus each student's utility ranking of the schools is generated by the following utility function:

selection rule. Of course, this is not to say that these mechanisms are easily manipulable, especially in a large market where participants lack complete information. A thorough investigation of pointing correspondences and the ensuing top trading cycles mechanisms is beyond the scope of the present paper.

[35] There is only one good school in this environment, namely school A. Schools B and C are of equally low quality.

$$u^i(s) = u^i_q(s) + u^i_p(s) + u^i_s(s) + u^i_r(s)$$

$u^i_q(s)$ represents the quality utility for student i at school s. It equals 40 for school A, and 20 for schools B and C for all students.

$u^i_p(s)$ represents the proximity utility for student i at school s. It equals 10 for school s, if the student lives in the district of school s. Students 1, 4, and 7 live in the district of school A. Students 2, 5, and 8 live in the district of school B. Students 3, 6, 9, and 10 live in the district of school C.

$u^i_s(s)$ represents the utility from having a sibling at school s. It equals 20 if a student has a sibling in the school. Student 1 has a sibling at school B, students 4 and 7 have siblings at school C. [36]

$u^i_r(s)$ represents a random utility which captures diversity of tastes. It is drawn from the uniform distribution from the interval [0,20]. It reflects a variety of preferences depending on abilities and interests.

The monetary payoffs for subjects were determined based on the resulting ranking. They received 15 euros for getting a slot in their most preferred school, 10 euros for the second most preferred school and 5 euros for the least preferred school.

Table 1 presents the preferences of students in this designed environment.

ID	1	2	3	4	5	6	7	8	9	10
The most preferred school	B	C	B	C	A	A	C	A	A	B
2nd most preferred school	A	A	A	B	C	C	A	B	B	A
The least preferred school	C	B	C	A	B	B	B	C	C	C

Table 1: Preferences of students in the designed environment

In all environments each school has a strict priority order of students of other districts. In the designed environment these priority orders are generated so as to insure that an allocation with no justified envy is possible under mechanisms of interest. As for students who live in the school district, they have the highest and equal priority for that school. The weak priority of students in the district is transformed into a strict priority by a random draw in the experiment. Students' draws are transformed into a queue used for tie-breaking when the allocation is calculated in the experiment. Note that draws are only used to tie-break the otherwise coarse priorities of district students, for instance, to tie-break priority of students 1, 4, and 7 in school A in the designed environment. The strict priority of non-district students is not affected by the random draws.

Table 2 presents the priority orders of schools in the designed environment.

5.2 Random-correlated environment

It is often the case in real life that different students want different schools depending on their abilities.[37] Given their abilities, students' preferences are strongly correlated, at least for the

[36] We introduced this utility in order to create preferences where the student's most preferred school does not correspond to her district school, as this situation is not interesting from an experimental perspective.

[37] Chen and Sönmez (2006) differentiate between two schools: one is stronger in the arts, another one in sciences. Here we employ a similar idea.

School	A	B	C
The highest priority	**1,4,7**	**2,5,8**	**3,6,9,10**
The second highest	10	10	7
.	8	9	2
.	9	1	1
.	3	7	8
.	2	4	5
.	5	6	4
No priority	6	3	

Table 2: Priority orders of schools in the designed environment

most preferred school. The random-correlated treatment is constructed to approximate this structure, assuming there is no utility of proximity to school.[38]

There are five schools in this environment – A, B, C, D, and E. Each school has only two slots. Students 1 and 2 live in the district of school A, students 3 and 4 live in the district of school B, students 5 and 6 live in the district of school C, students 7 and 8 live in the district of school D, and students 9 and 10 live in the district of school E.

Six students (1–6) prefer either school D or E, with the other one coming second in preferences, and the other four students (7–10) preferring either school A or school C. Note that it is insured that students 1–6 do not live in the district of schools D and E, as well as that students 7–10 do not live in the districts of schools A and C. The rankings from the third to the last (5th) are generated randomly. In the experiment, subjects received 15 euros for getting a slot in their most preferred school, 12.50 euros for the second most preferred school, 10 euros for the third most preferred school, 7.50 euros for the fourth most preferred school, and 5 euros for the least preferred school. [39]

Table 3 presents the preferences of students in the random-correlated environment.

ID	1	2	3	4	5	6	7	8	9	10
The most preferred school	D	D	E	E	D	E	C	C	C	A
2nd most preferred school	E	E	D	D	E	D	A	A	A	C
3rd most preferred school	B	A	B	C	A	A	E	E	E	E
4th most preferred school	C	B	A	B	C	C	B	B	B	B
The least preferred school	A	C	C	A	B	B	D	D	D	D

Table 3: Preferences of students in the random-correlated environment

Unlike in the designed environment, the schools' priority orders of non-district students are generated in such a way that a stable allocation is not feasible under both mechanisms. Table 4

[38]One could think of this environment as a game, where only the most ambitious students from districts with relatively bad schools are competing to get slots in the best schools of the city.

[39]The payoffs for the schools are different in all environments, but we keep the payoff for the most preferred school the same – 15 euros. The payoff for the least preferred school is also almost the same (5 or 4.5 euros). Payoffs for the other schools are generated in order to keep the same difference in payoffs between first and second schools, second and third schools, and so on.

presents the priority orders of schools in the random-correlated environment.

School	A	B	C	D	E
The highest priority	**1, 2**	**3, 4**	**5, 6**	**7, 8**	**9, 10**
The second highest	10	6	10	9	3
.	8	5	4	10	2
.	5	8	1	3	5
.	6	7	7	2	1
.	9	10	8	1	6
.	3	2	9	4	8
.	4	1	2	6	4
No priority	7	9	3	5	7

Table 4: Priority orders of schools in the random-correlated environment

5.3 Random-uncorrelated environment

In order to check the robustness of the results we create a random-uncorrelated environment. There are four schools in this environment – A, B, C and D. Schools A and B have two slots each, and schools C and D have three slots each. Students 1 and 2 live in the district of school A, students 3 and 4 live in the district of school B, students 5, 6, and 7 live in the district of school C, students 8, 9, and 10 live in the district of school D.

The preferences of each student are generated randomly. In situations when the district school has the highest payoff, the preferences of the students were regenerated. The resulting preferences are presented in Table 5. In the experiment, subjects received 15 euros for getting a slot in their most preferred school, 11.50 euros for the second most preferred school, 8 euros for the third most preferred school, and 4.50 euros for the fourth most preferred school.

ID	1	2	3	4	5	6	7	8	9	10
The most preferred school	D	C	C	A	D	A	B	C	C	B
2nd most preferred school	A	A	B	B	B	D	C	A	D	D
3rd most preferred school	B	D	D	C	A	C	A	D	A	C
The least preferred school	C	B	A	D	C	B	D	B	B	A

Table 5: Preferences of students in the random-uncorrelated environment

The priority orders for students from other districts were generated randomly and presented in Table 6

Both mechanisms are implemented as one-shot games of complete information. Each subject knows her own payoff table, the preference tables of other participants, and the priority orders of schools. [40]

The experimental design allows us to test three hypotheses based on the theoretical properties of TTC and ETTC:

[40]Except for the tie-breaking of the district priority.

School	A	B	C	D
The highest priority	**1,2**	**3,4**	**5,6,7**	**8,9,10**
The second highest	7	8	9	7
.	9	6	2	4
.	10	5	8	1
.	4	10	4	3
.	8	9	3	2
.	6	7	1	6
.	5	2	10	5
No priority	3	1	.	.

Table 6: Priority orders of schools in the random-uncorrelated environment

Hypothesis 1: Participants of the experiment choose to state their true preferences for allocations under both TTC and ETTC as both mechanisms are strategy proof.

Hypothesis 2: TTC and ETTC should not differ with respect to the efficiency criteria, as both mechanisms are Pareto-efficient.

Hypothesis 3: On average, the number of justified envy situations generated by ETTC should be lower than those of TTC.[41]

5.4 Experimental procedures

The experiment was run at the experimental economics lab at the Technical University of Berlin. We recruited subjects from our pool with the help of ORSEE by Greiner (2015). In total, 14 sessions were conducted – that is, seven sessions per treatment. In total, 140 subjects participated in the experiment. Each session included three environments which were played in a random order. No feedback was provided between environments, and thus we assume no learning effects. Only one of the three environments was payoff-relevant for subjects and it was determined at the end of the experiment by a random draw. All subjects played in a complete information environment, and were thus aware of the preferences of other players, as well as of the priorities of the schools. The experiment was paper-based, and sessions lasted approximately 70 minutes, with 40–45 minutes used for instructions and the public solution of an example. The average payoff, including a participation fee, was 15.32 euros.

In the experiment, subjects were randomly assigned a seat in the lab which corresponded to a subject ID. The experimenter read the instructions aloud. Subjects were allowed to ask questions, which were answered publicly. The subjects then had an additional 20 minutes to go through the explanation of the mechanism and the solution to the example. After all subjects had completed the reading, an experimenter went through an example solution in public. Subjects then had 10 minutes to solve an allocation problem. Before the decision sheets were distributed, the subjects drew a number from the bag to determine the initial queue, which was used to tie-break the

[41]Note that given truthful submissions the number of justified envy outcomes is 0 in ETTC and 3.3 in TTC in Designed environment, 8 in ETTC and 9.2 in TTC in Random-correlated environment, and 1 in ETTC and 2.7 in TTC in Random-uncorrelated environment. To calculate the number of justified envy outcomes for TTC, we use 10 random tie-breaking orders, and report the average of 10 allocations.

priority orders of students in the school of their district (one for each of the environments).[42] Subjects were not made aware of the number they had drawn, nor were they aware of the procedure which mapped the drawn numbers into the queue. The decision sheets were then distributed for one of the three environments. After subjects made their decisions, the sheets were collected, and the decision sheets for the next environment were distributed. After all the decisions for all the environments were completed, subjects filled out the questionnaire, while the experimenter entered the data to calculate the outcomes. The feedback on the allocations was given, together with the payment, to every subject in private. The instructions which were used in the experiment can be found in the Appendix.

6 Experimental Results

We compare the outcomes generated by TTC and ETTC to evaluate the mechanisms. There are three important criteria for comparison: truthful preference revelation, efficiency of allocation, and equity of allocation.

We first look at the proportions of truthful preference revelations. As under both TTC and ETTC, slots in the district schools are guaranteed for student: the truthful preference revelation requires that reported rankings coincide with truthful preferences from the top choice up to the district school of the participant.

	Designed environment	Random-correlated environment	Random-uncorrelated environment
TTC (1)	57%	30%	37%
ETTC (2)	56%	27%	40%
P-value (3)	0.865	0.708	0.728

Notes: The report is counted as truthful if reported preferences and truthful preferences coincide for choices from the top choice up to the district school of the participant. Rows (1) and (2) present proportion of truthful reporting. Row (3) shows two-sided p-values for the test of equality of proportions between treatments.

Table 7: Truth-telling rates

Result 1 (Truthful Preference Revelation): *In all environments, the differences in proportions of truthful preference revelation under TTC and under ETTC are not statistically significant.*

Table 7 presents the proportions of truthful preference revelation for each mechanism in each of the environments. The test of equality of proportions shows that proportions of truthful preference revelation under ETTC are not significantly different from that under TTC in

[42]Note that the lottery draws did not influence the strict priority of students below the top q priority.

all environments: z=0.170 (p=0.805) for the designed environment, z=0.374 (p=0.708) for the random-correlated environment, and z=0.3473 (p=0.728) for the random-uncorrelated environment. The proportion of truthful preference revelation is the highest in the designed environment, 57% under TTC and 56% under ETTC, and the lowest in the random-correlated environment: 30% and 27% correspondingly. Thus, neither TTC nor ETTC induced full truthful preference revelation for all participants, and therefore we reject Hypothesis 1. The proportions of truthful preference revelation in designed and random-uncorrelated environments for TTC are similar to those in Chen and Sönmez (2006) (50% and 43%, respectively). Low truth-telling rates in the random-correlated environment can be explained by high correlation of preferences within two most preferred schools, high number of schools and relatively low loss in payoffs by missing one rank reached. The variation shows that participants in the experiment react to the environment in the expected way.

The second criterion we consider is efficiency. We follow the cardinal concept of utility. Therefore we use the following formula to calculate the efficiency of allocation:

$$Efficiency = \frac{Actual\ sum\ of\ payments\ to\ participants}{Sum\ of\ payments\ to\ participants\, if\ they\ all\ state\ the\ truth}$$

Thus, we refer to the allocation with full preference revelation as being efficient. As we implemented sessions as one-shot games, we follow Chen and Sönmez (2002) in using the recombinant estimation technique proposed by Mullin and Reiley (2006) and modified by Abrevaya (2008). By using the recombinant technique we assume that every session and every environment is an independent drawn from an identical distribution. This assumption is restrictive, but we are making it due to the design feature of the one-shot game, without feedback to participants between environments.[43] Due to a relatively high number of sessions for each treatment, the full recombination is too demanding as it includes $(7)^{10}$ recombinations. That is why we take 20,000 random recombinations of subjects from different sessions, keeping the ID fixed and calculating the outcomes of interest. Thus, for every ID of a combination, we randomly determine a natural number from 1 to 7 from the uniform distribution which corresponds to the number of the session. We repeat this procedure 20,000 times, and control for the absence of repetitions afterwards. Based on these 20,000 combinations, we calculate a recombinant estimator for the expected value of the allocation outcome of interest. For the allocation of TTC, the initial queue of subjects used for tie-breaking the priority order of district students is crucial. To address this issue we calculate 10 different allocations for each combination for TTC based on 10 random initial queues. Note that ETTC does not require a tie-breaking the priority order for district students. Thus we produce 200,000 allocations for each environment for TTC, and 20,000 allocations for each environment for ETTC. Following the results of Abrevaya (2008), we estimate asymptotic standard errors and asymptotic variance for these outcomes. Based on asymptotic variance, we construct a Z-test for the significance of the difference between treatments.

Result 2 (Efficiency): *In all environments, the difference in efficiency under TTC and ETTC is not statistically significant.*

Table 8 presents average efficiency and the results of the recombinant estimation of mean efficiency in different environments. A two-sided Z test based on recombinant estimation of

[43] Thus, we abstract from any possible session specific shocks.

asymptotic variance shows that the efficiency of the allocation under ETTC is not significantly different from that under TTC in all environments: z=0.380 (p=0.70) for the designed environment, z=0.248 (p=0.39) for the random-correlated environment, and z=0.064 (p=0.95) for the random-uncorrelated environment. Thus we can support Hypothesis 2.

Mechanism		TTC	ETTC	p-value
Designed environment	Average efficiency (7 sessions)	81.77%	82.27%	0.85
	Mean efficiency	86.38%	82.37%	0.70
	Asymptotic standard error	0.026	0.028	
Random-correlated environment	Average efficiency (7 sessions)	91.64%	87.33%	0.33
	Mean efficiency	91.37%	89.30%	0.39
	Asymptotic standard error	0.015	0.019	
Random-uncorrelated environment	Average efficiency (7 sessions)	84.97%	85.31%	0.84
	Mean efficiency	84.51%	84.65%	0.95
	Asymptotic standard error	0.001	0.001	

Note: Average efficiency corresponds to the average of the data from seven experimental sessions for each treatment. Mean efficiency corresponds to recombination data, which includes 20,000 allocations for ETTC values and 200,000 allocations for respective TTC values (the same recombinations given 10 different random draws). The corresponding p-value is the two-sided p-value of the Wilcoxon ranksum test for the equality of average efficiency between treatments. Mean efficiency corresponds to the average efficiency of allocations reached after recombinations. The corresponding p-values are based on the two-sided Z test with asymptotic standard errors.

Table 8: Mean efficiency

The main interest of comparison of mechanisms lies in the fairness comparison. Using allocations from recombination we calculate two fairness indicators:

1. *The number of justified envy situations* depicts the total number of student pairs where one student has justified envy of another student in any school with respect to their true preferences. If one student has justified envy of several students in one or more schools, each of the cases is counted toward the sum of justified envy situations.

2. *The percent of students with justified envy* depicts the percent of students who have justified envy of at least one other student with respect to their true preferences.

Result 3 (Number of justified envy situations): *In the designed and random-correlated environments, ETTC produces significantly less justified envy situations than TTC does. In the random-uncorrelated environment, the difference in the number of justified envy situations produced by ETTC and TTC is not statistically significant.*

Mechanism		TTC	ETTC	p-value
Designed environment	Average number of justified envy situations (7 sessions)	6.14	2.00	0.04
	Mean number of justified envy situations	4.76	2.89	0.00
	Asymptotic standard error	0.66	0.58	
Random-correlated environment	Average number of justified envy situations (7 sessions)	9.71	9.00	0.43
	Mean number of justified envy situations	9.62	8.43	0.00
	Asymptotic standard error	0.61	0.16	
Random-uncorrelated environment	Average number of justified envy situations (7 sessions)	4.14	3.14	0.10
	Mean number of justified envy situations	3.80	3.42	0.22
	Asymptotic standard error	0.42	0.49	

Note: The justified envy situations are calculated with respect to the true preferences of the subjects. Average number of justified envy situations corresponds to the averages of the data from seven experimental sessions for each treatment without recombination. The corresponding p-values are the two-sided p-values of the Wilcoxon ranksum test for the equality of the average number of justified envy situations between treatments. Mean number of justified envy corresponds to the average number of justified envy situations in allocations reached after recombinations, which include 20,000 allocations for ETTC values and 200,000 allocations for respective TTC values (the same recombinations given 10 different random draws). The corresponding p-values are based on the two-sided Z test with asymptotic standard errors.

Table 9: Mean number of justified envy situations

Table 9 presents the the average number of justified envy situations and results of the recombinant estimation of the mean number of justified envy situations under both mechanisms in all environments. In the designed environment, the mean number of justified envy situations under ETTC is 2.89, which is significantly less than 4.76 under TTC. The two-sided Z test based on recombinant estimation of asymptotic variance yields z=3.723 (p=0.00). In the random-correlated environment the relation is 8.43 to 9.62 respectively, and the difference is statistically significant: z=3.01 (p=0.00). As for the random-uncorrelated environment, the expected number of justified envy situations is 3.42 under ETTC, and is 3.80 under TTC but the difference is not statistically significant: z=1.23 (p=0.22).

Result 4 (Percent of students with justified envy): *In the designed and random-correlated environments under ETTC, a significantly lower percent of students have justified envy than under TTC. In the random-uncorrelated environment, the difference in the percent of students with justified envy under ETTC and TTC is not statistically significant.*

Mechanism		TTC	ETTC	p-value
Designed environment	Average percent of students with justified envy (7 sessions)	35.7	15.7	0.04
	Mean percent of students with justified envy	38.5	22.5	0.00
	Asymptotic standard error	4.8	4.2	
Random-correlated environment	Average percent of students with justified envy (7 sessions)	52.9	38.6	0.02
	Mean percent of students with justified envy	70.4	56.1	0.00
	Asymptotic standard error	2.9	3.6	
Random-uncorrelated environment	Average percent of students with justified envy (7 sessions)	30.0	25.7	0.60
	Mean percent of students with justified envy	29.7	29.8	0.98
	Asymptotic standard error	2.9	4.5	

Note: The justified envy situations are calculated with respect to the true preferences of the subjects. Average percent of students with justified envy corresponds to the averages of the data from seven experimental sessions for each treatment without recombination. The corresponding p-values are the two-sided p-values of the Wilcoxon ranksum test for the equality of the average percent of students with justified envy between treatments. Mean percent of students with justified envy corresponds to the average percent of students with justified envy in allocations reached after recombinations, which include 20,000 allocations for ETTC values and 200,000 allocations for respective TTC values (the same recombinations given 10 different random draws). The corresponding p-values are based on the two-sided Z test with asymptotic standard errors.

Table 10: Mean percent of students with justified envy situations

Table 10 presents the the average percent of students with justified envy and results of the recombinant estimation of the mean percent of students with justified envy under both mechanisms in all environments. After recombination, in the designed environment, on average 22.5% students have justified envy of at least one student under ETTC, and 38.5% have it under TTC. The Z-test for equal means yields z=3.733 (p=0.00). For the random-correlated environment, the mean percents of students with justified envy are 56.1 and 70.4 respectively. And the difference is statistically significant: z=3.924 (p=0.00). In the random-uncorrelated environment the mean percents of students with justified envy are 29.8 under ETTC and 29.7 under TTC, and the difference is not statistically significant: z=0.021 (p=0.98).

Combining result 3 and result 4, we can conclude that our Hypothesis 3 is only partially supported by the experimental results. It holds for the case of the designed and the random-correlated environments, but is rejected in the case of the random-uncorrelated environment.

In spite of the fact that in the random-uncorrelated environment ETTC produces a result similar to TTC from the perspective of the fairness criteria, one should not underestimate the benefits of ETTC, as the result is driven by the high rate of preference misrepresentations in the form of district school bias in the random-uncorrelated environment in both mechanisms. It can be explained by the fact that, by design, stating the district school as a top choice in the random-uncorrelated environment leads to a much lower average loss in payoffs than in other environments. Note that seven out of 10 subjects have the district school as the second-best school in their preferences in the random-uncorrelated environment, while only three out of 10 in the designed environment do, and none in the random-correlated environment do. Thus, in the random-uncorrelated environment it is likely that preferences are reported in a way that at most one local student in each school will apply to another school as the top choice. This fact diminishes the advantage of ETTC relative to TTC because it is identical to the case when each school has only one slot to allocate,[44] where TTC and ETTC are equivalent.[45]

Next, we address the following question: if one cares about fairness, given the stated preferences, which mechanism (TTC or ETTC) should be chosen? Thus we switch attention from the true preferences of subjects, to their ***stated*** preferences. Indeed, a designer cares about fairness with respect to stated preferences, because students can have a truly justified envy only by stated preferences (for the purpose of filing a complaint or appeal). Using the recombinations of stated preferences from both treatments, we create 20,000 different preference profiles in each of the three environments.[46] We use both mechanisms to calculate the allocation for each preference profile and 10 random initial queues and evaluate the allocation with respect to the stated preferences. In total, we calculate 200,000 allocations for each mechanism in each environ-

[44]Since all other district students who applied to their home school can be excluded from the procedure, their tentative assignments are finalized.

[45]Recent paper by Abdulkadiroglu et al. (2017) found no significant difference in the number of justified envy situations between TTC and ETTC given submitted reports in New Orleans Recovery School District. This raises an important question of what features of the preferences one has to assume to draw policy conclusions and how similar are the preferences of students in different school choice programs. It might be that in New Orleans most of the students prefer the school where they have top-q priority, which makes TTC and ETTC closer to each other. Our experiments show that the degree of the advantage of ETTC to TTC varies depending on the environment, and thus it is crucial to study the features of each school choice program before the decision on implementation of one mechanism or the other.

[46]Alternatively, one can interpret the subsequent analysis as simulations.

ment. We run the procedure, similar to the ordinal efficiency test by Guillen and Kesten (2012). We calculate the number of times TTC dominates ETTC with respect to the fairness criterion (TTC dominations)[47] and the number of times ETTC dominates TTC (ETTC dominations). We calculate it for both equity criteria: the number of justified envy situations and the percent of students with justified envy. Finally, we use the Wilcoxon signed-rank test for the equality of dominations.

Result 5 (Fairness dominance): *ETTC is more likely to generate less justified envy situations and lower percent of students with justified envy than TTC in all environments, given the students' stated preferences.*

For the number of justified envy situations, the Wilcoxon signed-rank test rejects the hypothesis of equality in the number of TTC and ETTC dominations in all environments. In the designed environment, the number of TTC dominations is 16,990 out of 200,000 and the number of ETTC dominations is 111,924, leading to z=272.405 (p=0.00). In the random-correlated environment, the number of TTC dominations is 33,815 and the number of ETTC dominations is 84,562, leading to z=152.338 (p=0.00). In the random-uncorrelated environment, the number of TTC dominations is 20,231 out of 200,000 and the number of ETTC dominations is 66,294, leading to z=159.079 (p=0.00).

For the percent of students with justified envy, the Wilcoxon signed-rank tests of equality in the number of TTC and ETTC dominations yield p-values$<$0.01 in all environments as well.

The histograms for the differences of fairness criteria in a random environment are presented in Figure 5.

Result 5 is one of our main results from the perspective of policy. We showed that, given Pareto efficiency and strategy proofness of the allocation procedure, and considering the frequency of justified envy situations while choosing between competing mechanisms, one should prefer ETTC to TTC.

Next we address the question of equity from the individual's perspective. Note that here we switch our attention back to the justified envy with respect to the ***true preferences***. We assume that the revelation of true preferences is an exogenous decision, independent of the mechanism. Thus we interpret it as individual-based. This assumption is in line with the results of the experiment as there is no significant difference in the proportion of truth revelation across mechanisms. We, therefore, focus on the question: if one could choose between TTC and ETTC, which mechanism would she prefer on average when controlling for the rank reached and truthful revelation of preferences?[48]

Result 6 (Individual fairness): *The probability of a student having justified envy of any other student, controlling for preferences revelation and rank reached, is lower under ETTC than under TTC.*

[47] TTC creates less justified envy situations than ETTC.

[48] We assume that from an individual perspective the equity is less important than the rank reached and we control for it explicitly in the regression. Thus, we assume that any subject would prefer an allocation to the second most preferred school, with justified envy to an allocation to the third most preferred school without justified envy.

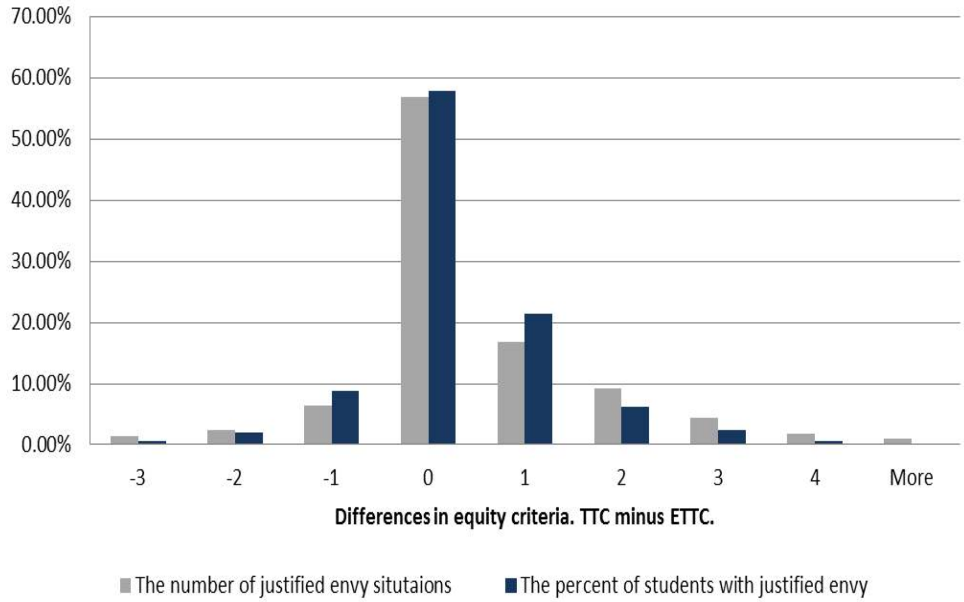

Figure 5: Difference in equity criteria in random environment

Models (1) and (2) in Table 11 present marginal effects at the mean for the probit regressions of the probability of having justified envy based on the data from sessions, without recombination. Rank reached variable represents the ordinal rank of the school a subject was placed to in the real preferences of the subject. The higher is the rank (or, the less preferred school the subject is placed to), the higher is the probability of having justified envy. Submitting the truthful preference list decreases the probability of having justified envy, but the coefficient is only significant at the 10% level. Finally, and most importantly, the coefficient of ETTC dummy is negative and significant. Thus, on average, ETTC decreases the probability of a subject having justified envy by 12 percentage points.[49]

However, to have robust estimation of the effect of the ETTC mechanism on the probability of justified envy, we construct a dataset on an individual level in the following way: we take 100 allocations for each of the participants in the experiment, based on a different recombination of the rest of the group, and calculate the individual result for the participant in both mechanisms. For TTC we use three different random draws of for tie-breaking the priority for each of the recombinations. It results in a database with 21,000 allocations for ETTC (70 subjects, times three environments, times 100 allocations) and 63,000 allocations for TTC. It contains the following variables: a dummy for stating the truth, a dummy for having justified envy situation, and the rank reached by each participant.

Using the recombinant data, we calculate the probability of having justified envy for a subject, given a distribution of possible strategies of the opponents. Thus, for a given subject and her action, we calculate the expected probability of having justified envy. Using our recombination

[49]Note, that controlling for the specific environments in model (2) of Table 11 leads to insignificant coefficients. This might appear as contradiction for the data presented in Table 10, but could be explained by the fact that we control for rank reached, which partially captures the differences in the environments.

Table 11: Marginal effects of probit regression of the dummy for justified envy

	(1)	(2)
Rank reached	0.27***	0.28***
	(0.04)	(0.04)
Truth dummy	-0.10*	-0.10*
	(0.05)	(0.05)
ETTC dummy	-0.12***	-0.12**
	(0.05)	(0.05)
Designed env.		0.05
		(0.06)
Random-correl. env.		-0.03
		(0.05)
Observations	420	420
No. of clusters	140	140
log(likelihood)	-184.70	-183.88

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Standard errors are in parentheses and are clustered at the individual level.

Note: Justified envy is defined with respect to the true preferences of subjects. The table presents marginal effects of probit regressions. Marginal effects of every variable are calculated at the mean of all other covariates. In models (1) and (2) the data from each session are used, without recombination. Thus for each subject we observe three outcomes, one in each of the environments. Variable "Rank reached" represents ordinal rank of the school a subject was placed to in the true preferences of the subject. Variable "Designed env." is a dummy for the designed environment, while the variable "Random-correl. env." is a dummy for random-correlated environment.

Table 12: OLS regression of the proportion of the allocations with justified envy for each student

	Proportion of allocations with justified envy
Av. rank reached	0.26***
	(0.02)
Truth dummy	-0.07**
	(0.03)
ETTC dummy	-0.05**
	(0.02)
Designed env.	0.02
	(0.03)
Random-correl. env	-0.002
	(0.03)
Observations	420
Clusters	140
R-squared	0.43

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table presents the coefficients OLS regressions of the the proportion of the allocations with justified envy for each student in the recombinant data, calculated for each environment separately. Standard errors are clustered on the level of each subject. Variable “Av. rank reached” represents the average ordinal rank of the school a subject was placed to in the true preferences of the subject, averaging among all 7,000 allocations for ETTC for each environment, and 21,000 allocations for TTC for each environment. Variable ”Designed env.” is a dummy for the designed environment, while the variable ”Random-correl. env.” is a dummy for random-correlated environment.

data on individual level, we calculate the fraction of allocations in which the subject has justified envy. Note that the subject and her action are held constant. This amounts to assuming that the subject is playing against a distribution of opponents (their types in terms of priority and preferences are known, but their strategies in the game are unknown). We then regress the calculated probability on ETTC dummy, the average rank reached in these games and the dummy for revealing truthful rank-order list.

Table 12 presents the results of the estimation of the linear regression. In line with results from Table 12, ETTC significantly decreases the probability of having justified envy, while the higher the average rank the more likely is the subject to have justified envy. The interpretation of the ETTC coefficient is different, however: Under ETTC, given a strategy of a subjects, controlling for the average rank reached, the proportion of allocations with justified envy is on average five percentage points lower than under TTC. Thus we conclude that result 6 is robust for different approaches of comparing fairness of allocations under TTC and ETTC.

7 Conclusion

Notwithstanding the theoretical attractiveness of the original TTC as an efficient and strategy-proof assignment method, its popularity in practice remained limited, primarily due to its lack of stability. The current paper offers a direction to improve this mechanism's desirability in practice in various dimensions including, but not limited to stability.

Our contribution has been twofold. First, we have argued that whenever the indivisible goods are in multiple supply – such as in the school choice context – the Gale's original idea can also be used in a much more flexible way to meet other distributional goals. This is due to the additional freedom the planner has emanating from the multitude of trading opportunities in general TTC mechanisms. To explore this potential one needs to first generate a thick market where all supplies of the goods are concurrently up for trade and then specify a pointing rule suitable for the planner's objectives. Second, we have proposed and studied a particular application of this idea to school choice as a potentially more stable mechanism than TTC when dealing with pairwise exchanges which often arise under correlated preference profiles. We have shown that establishing full fairness when stability and efficiency are compatible, may not be possible via any mechanism employing the top trading cycles idea.[50]

The findings from our experiment are consistent with the theory, as ETTC on average provides better stability results. In realistic settings (designed and random-correlated environments) ETTC generates a significantly lower number of students with justified envy, as well as significantly less cases of justified envy. In the simulated settings, when allocations are evaluated with respect to stated choices, ETTC dominates TTC with respect to the fairness criteria in all environments, including a random-uncorrelated environment. We have shown that from an individual point of view, ETTC is less likely to lead to justified envy than TTC.

[50]In a general canonical assignment model in the absence of priorities, Liu and Pycia (2016) establish quite a general limit theorem that shows that all symmetric, asymptotically efficient, and asymptotically strategy-proof mechanisms are asymptotically equivalent. Although our model is finite and deterministic, and the presence of priorities violates symmetry, their result implies that if an ex ante randomization over all priority structures were to be chosen, then all versions of top trading cycles will asymptotically coincide in a large market.

Notwithstanding our modeling assumption that schools' priorities are strict and exogenously given, ETTC can easily be adopted to the coarse priority structures which are commonly observed in reality.[51] ETTC has an additional advantage over TTC in such circumstances: the ETTC allocation does not depend on the random tie-breaking of the top (up-to-quota) priority students if the top q-groups across two different schools are disjoint.[52] In cases when the number of students in the top priority class is more than the number of slots in the school, the slots are guaranteed for these students under both TTC and ETTC. However, only ETTC treats all students with the guaranteed slots *equally*, whereas TTC would require the use of tie-breaking among the students within the same class, thus unnecessarily favoring some students over others. In other words, although random tie-breaking does not harm student welfare under TTC, as opposed to the case of the well-known DA (cf. Erdil and Ergin, 2008; Abdulkadiroğlu et al. 2009), it may nonetheless introduce an artificial loss of fairness. On the other hand, ETTC guarantees equal treatment of all students within the highest priority group, without

requiring random-tie breaking and can eliminate avoidable justified envy situations.

Under DA, allocations are tentative at each step and the students enrolled at each school can only be determined at the last step. By contrast, the trades made in each step under mechanisms based on top trading cycles are final, which in turn gives the planner a better sense of the resulting allocation as the market unfolds. Pointing rules that exploit this aspect of top trading cycles may be particularly useful when addressing concerns based on racial or gender diversity. For example, a pointing rule that takes into account the composition of the partial allocation that has formed until the current step can allow the planner to shift priority in that step to the currently underrepresented groups of students enabling him to achieve a more balanced distribution as the algorithm proceeds. We believe this type of history-dependent dynamic pointing rules, coupled with a thick market throughout, would be a worthwhile avenue to explore in future research.

8 The Appendix

Proof of Proposition 1

Lemma 1 *Given a problem and a step of the ETTC algorithm, if the best choice of a student i among the remaining slots is a slot at school s, then student i does not ever point to a pair which contains a slot from a different school until all slots of school s are allocated.*

Proof:

This simply follows from the fact that for each school s whose slots remain to be inherited from earlier steps, the to-be-inherited slots are assigned to the remaining students, immediately after the step at the end of which no student who was assigned to a slot of school s at an earlier step is left.

[51] For example, until recently in Boston, the top priority group of students have consisted of those students who live in the walk zone of the school and have a sibling at that school. This priority group is unlikely to fill all the quotas in any given school.

[52] That is, for any such problem, any shuffling of student positions for the top q_a-priorities of any school a yields exactly the same allocation under ETTC, a feature not present in TTC. This simply follows from the fact that the pointing rule of ETTC does not depend on students' own priorities.

Q.E.D.

A critical observation about the ETTC algorithm is that by Lemma 1, if a school still has vacant slots at a step of the ETTC algorithm, then there is a student-school pair containing a slot from that school at that step. Thus, the idea behind the Pareto efficiency of ETTC is the same as that of TTC. Given a problem, each student who leaves at the first step is placed in his best choice, hence he can not be made better off. Each student who leaves at the second step is placed in his best choice among the remaining schools, hence he cannot be made better off without making someone who left at the first step worse off. Continuing in this way, no student can be made better off without making someone who left at an earlier step worse off.

Q.E.D.

Proof of Proposition 2

By Pápai (2000), a mechanism is group strategy-proof if and only if it is strategy-proof and nonbossy. We first show that ETTC is strategy-proof.

Lemma 2 *Given a problem $P = (P_i)_{i \in I}$, suppose a student i is removed at a step t of the ETTC algorithm and, if he submits preferences P_i' instead of P_i, he is then removed at a step t'. Then the remaining students, available slots of schools, and to-be-inherited slots of schools at the beginning of step $\min\{t, t'\}$ are the same.*

Proof:

Given a problem $P = (P_i)_{i \in I}$, because no student-school pair containing student i participates in a cycle before step $\min\{t, t'\}$, the same cycles form until step $\min\{t, t'\}$ and the same students are placed in the same schools until step $\min\{t, t'\}$.

Q.E.D.

Lemma 3 *Given a problem, if a student-school pair (i, s) is pointing to another student-school pair (i', s') at some step of the ETTC algorithm, then student-school pair (i, s) continues pointing to the student-school pair (i', s') as long as student i' is not removed, i.e., he is not contained in any cycle.*

Proof:

Given a problem, if a student-school pair (i, s) is pointing to another student-school pair (i', s') at a step t, then school s' is the best choice of student i among the remaining schools. Furthermore, student i' has the highest priority for school s among those who are assigned a slot from school s'. By Lemma 1, student i will keep pointing to a pair containing a slot from school s' until no vacant slots at school s remain. Then, the only case where student-school pair (i, s) can point to another student-school pair (i'', s'') before student i' is removed is when (i) $s'' = s'$ and (ii) student i'' has higher priority than student i' for school s. But this is only possible if student i'' is assigned to a slot of school s through inheritance at some step t', $t' > t$. But no inheritance of slots of school s' takes place before student i is removed.

Q.E.D.

Given a problem $P = (P_i)_{i \in I}$ and a student i, let t be the step at which student i is removed and s be the school he is placed in. We will show that if student i submits preferences P'_i, he cannot be placed in a school which is better for him than school s. Let t' be the step at which student i is removed when he submits P'_i and $\widehat{s}$ be the school he is placed in at this step. We consider two cases:

Case 1. $t \geq t'$: Consider step t'. By Lemma 2, at the beginning of this step, the remaining students and slots are the same. Note that by Lemma 3, each student-school pair (i', s') that is pointing to a pair containing student i continues pointing to that pair as long as student i stays. Similarly, each student-school pair (i'', s'') that is pointing to the pair (i', s') continues pointing to that pair as long as student i' stays which is the case as long as student i stays, and so on. Consequently, at step t, student i has the opportunity to participate in any of the cycles he participates in under preferences P'_i. Since he is pointing to his best choice under P_i at step t, school s cannot be worse than school $\widehat{s}$ for student i.

Case 2. $t < t'$: By Lemma 2, at the beginning of step t, the remaining students and slots are the same. Since student i is placed in his best choice at this step, he cannot be placed in a better school at a later step.

Next, we show that ETTC is nonbossy. In addition to the above, suppose that, if student i submits preferences R'_i, his assignment remains school s at the new problem. Without loss of generality, assume that $t \leq t'$. We first contrast the market at the beginning of step t with that of step t'. By Lemma 2, at the beginning of step t, the remaining students and slots are the same for both problems. By Lemma 1, any student-school pair involved in a pointing chain that ends up with a pair containing student i at step t of the original problem remains in the same chain until step t' of the new new problem.[53] Let M be the set of students contained in all such chains, i.e., all the students included in any chain that ends up with a pair containing student i at step t of the original problem.

Let C be the set of cycles that form from step t to step t' at the new problem and I_C and S_C be the set of students and schools contained in these cycles, respectively. Since student i leaves the market at step t', he is clearly not contained in any such cycle. By the same logic, no student in M is contained in any such cycle either. This means the students in I_C who were assigned the slots from the schools in S_C have higher priority than those remaining in the market after step t. When student i leaves at step t, any slot assigned to him of a school in S_C will not be inherited by any student until all existing slots of that school in the market are allocated. This implies that at the orginal problem after i leaves at step t, the slots of schools in S_C will be inherited by exactly the same students in I_C and cycles in C will form in exactly the same way. But then the inheritance of slots and formation of the cycles by the remaining students should be identical for both problems, which means that all remaining students are placed in the same schools at both problems.

Q.E.D.

Proof of Proposition 3

[53]Formally, a (pointing) chain that ends up with a pair containing student i is an ordered collection of student-school pairs $((j_1, s_1), (j_2, s_2), \ldots, (j_k, s_k), (i, x))$ where each pair, except the last, points to the following pair in the order.

Example 2 establishes the latter statement in the proposition. Let $P = (P_i)_{i \in I}$ be a problem, i a student and s a school such that i is entitled to one slot at school s. Suppose by contradiction that student i has justified envy for a student j who is placed in a school s' under ETTC. Since i is entitled to one slot at school s, we have $s \neq s'$. Note that student i has higher priority than student j for school s'. This means student j does not inherit any slot of school s' but in fact there is a step t such that a student-school pair (j, s) forms a cycle with another student-school pair (k, s') where $k \neq i, j$. Note that since student i has higher priority than student j for school s', student-school pair (j, s) must be removed before step t. But this is possible only if student j is placed in school s'. A contradiction.

Q.E.D.

Proof of Proposition 4

Example 2 establishes the latter statement in the proposition. To prove the former statement, we shall show that if TTC selects a fair allocation for a problem, then ETTC also selects a fair allocation for the same problem. Let $P = (P_i)_{i \in I}$ be a problem, and α and β be the TTC and the ETTC allocations, respectively. Suppose α is stable. We show that $\alpha = \beta$. Let $S = \{s, s'\}$. Let $I^* = I^*_s \cup I^*_{s'}$ be the set of students who form a student-school pair at the first step of the ETTC algorithm, where I^*_s (resp. $I^*_{s'}$) is the set of students who are contained in a student-school pair with school s (resp. school s'). In addition, let I_1 be the set of students who are contained in a student-school pair of the first step of the ETTC algorithm that form a self-cycle. Note that $I_1 \subseteq I^*$. Clearly, $\alpha|_{I_1} = \beta|_{I_1}$. If there is no cycle that forms after the student-school pairs containing the students in I_1 are removed, then we proceed to the next paragraph. If there is such a cycle, call it cycle C_1. Let (i_1, s) and (j_1, s') be the two student-school pairs contained in cycle C_1. Note that student i_1 (resp. student j_1) has the highest priority for school s' (resp. for school s) among those students in $I^* \backslash I_1$ whose first choice is school s' (resp. school s). The existence of cycle C_1 implies that there also exists a cycle C'_1 that forms at some step of the TTC algorithm, such that two agents in $I^* \backslash I_1$ point to one another. Then, since α is fair and student i_1 (resp. student j_1) has the highest priority for school s' (resp. for school s) among those students in $I^* \backslash I_1$, this means $\alpha|_{\{i_1, j_1\}} = \beta|_{\{i_1, j_1\}}$. If there is no cycle that forms after cycle C_1, then we proceed to the next paragraph. If there is such a cycle, call it cycle C_2. Let (i_2, s) and (j_2, s') with $s \neq s'$ be the two student-school pairs contained in cycle C_2. Note that student i_2 (resp. student j_2) has the highest priority for school s' (resp. for school s) among those students in $I^* \backslash I_1$ whose first choice is school s' (resp. school s). Again, the existence of cycle C_1 implies that there also exists a cycle C'_2 that forms at some step of the TTC algorithm, such that two agents in $I^* \backslash (I_1 \cup \{i_1, j_1\})$ point to one another. As before, since α is fair, $\alpha|_{\{i_2, j_2\}} = \beta|_{\{i_2, j_2\}}$. Continuing in this way, we conclude that both ETTC and TTC agree on the placement of those students who are contained in a cycle until the inheritance of slots of some of the schools takes place for the first time.

Without loss of generality, let s be a school with n_s slots that remains to be inherited by the remaining students. (If there is no such school, then the result is trivial.) Note that at this step no student in I^*_s is left. From the preceding paragraph, the students who get the first $q_s - n_s$ slots are the same under ETTC and TTC. This implies there is a step of the TTC algorithm in which the student with the highest priority for school s among the remaining students inherits all the n_s slots of school s. Note that under the ETTC algorithm, these n_s slots are assigned to the remaining students (starting with the student with the highest priority for school s) one by one, following their priority order to form new student-school pairs. This means we are back where we started and can again use the same argument we used in the previous paragraph. Applying this argument iteratively, we conclude that $\alpha = \beta$.

Q.E.D.

INSTRUCTIONS – Mechanism TTC

This is an experiment in the economics of decision making. The instructions are simple, and if you follow them carefully and make good decisions, you might earn a considerable amount of money. In this experiment, we simulate a procedure to allocate students to schools. The procedure, payment rules, and student allocation method are described below. Do not communicate with each other during the experiment. If you have questions at any point during the experiment, raise your hand and the experimenter will help you.

Procedure

There are 10 participants in this experiment. You are participant 1. (This number is drawn by each participant before entering the room.) Each participant represents a student who wants to get a slot at a school. You will have to make three decisions in three different school markets. In all markets allocation will be determined by the same algorithm. This algorithm works as follows:

Allocation Method

Each participant is first tentatively assigned to the school within her respective district. Students of the school district have the highest priority in the school. Next, Decision Sheet rankings are used to determine mutually beneficial exchanges between two or more participants. The order in which these exchanges are considered is determined by a fair lottery. This means each participant has an equal chance of being the first in the line, the second in the line,..., as well as the last in the line. To determine this fair lottery, a participant will be asked to draw 10 numbers from an basket, one at a time. Each number corresponds to a participant ID number. The sequence of the draw determines the order in the lottery.

The specific allocation process is explained below:

1. Initially, all 10 slots are available for allocation.

2. All participants are ordered in a queue based on the order in the lottery.

3. Next, an application to the highest-ranked school on the Decision Sheet is submitted for the participant at the top of the queue.

- If the application is submitted to her district school, then her tentative assignment is finalized (thus, she is assigned to a slot at her district school). The participant and her assignment are removed from subsequent allocations. The process continues with the next participant in line.

- If the application is submitted to another school, the procedure progresses as follows: Say applicant Claudia's home district school is school A and she is applying to school B.

Then Claudia's application is submitted to school B. After that, one of the students who tentatively holds the slot at school B has to be chosen. In particular, among all students who tentatively hold a slot at school B, procedure chooses the one who is closest to the top of the queue. (So procedure follows the queue ordering while choosing among students of school B). This student is then moved to the top of the queue, directly in front of the requester (Claudia).

In general, whenever the queue is modified, the process continues similarly with the next student in the queue: An application of this student is submitted to her highest-ranked school with available slots.

- If the application is submitted to her district school, then her tentative assignment is finalized. The process continues with the next participant in line.
- If the application is submitted to another school – say, school S – then we adhere to the following procedure (explained above for Claudia): One of the students who tentatively hold a slot at School S needs to be chosen. In particular, among all the students who tentatively hold a slot at school S, the procedure chooses the student who is closer to the top of the queue. Then this student is moved to the top of the queue directly in front of the requester.

4. A mutually-beneficial exchange is obtained when a cycle of applications are made in sequence, which benefits all affected participants: e.g., I apply to Stefan's district school, Stefan applies to your district school, and you apply to my district school. In this case, the exchange is completed and the participants, as well as their assignments, are removed from subsequent allocations.

5. The process continues until all participants are assigned a school slot.

Example:

In order to understand the mechanism better, let us go through an example together.

If you have any questions about any step of the allocation procedure please feel free to ask at any point. There are six students (ID numbers from 1 to 6) on the market, and three schools (school A, school B, and school C) with two free slots each. Students 1 and 2 live in the district of school A, students 3 and 4 live in school district B, and, finally, students 5 and 6 live in school district C. It means that tentative assignments look as follows:

Tentative assignments of students (IDs)	School A	School B	School C
slot 1	1	3	5
slot 2	2	4	6

The lottery determined the following order (student IDs): 1-2-3-4-5-6

Students submitted the following school rankings in their Decision Sheets:

Student ID	1	2	3	4	5	6
Top choice	B	C	A	C	C	A
Middle choice	A	A	C	B	A	B
Last choice	C	B	B	A	B	C

This allocation method consists of the following steps:

Step 1. The queue is as follows: 1-2-3-4-5-6 (the initial queue order is always determined by the lottery.) Thus student 1 (the first in the queue) applies to school B. It is not her district school. Students 3 and 4 are tentatively assigned to school B. One of the two students needs to be chosen. Between the two students, student 3 is the closest to the top of the queue, that is why she moves to the top of the queue. And thus the queue is modified.

Step 2. The queue is as follows: 3-1-2-4-5-6. Thus student 3 (the first in the queue) applies to school A. This school is not her district school, but the cycle of beneficial exchange appears. Student 3 wants to attend student 1's school, and student 1 wants to attend student 3's school. The beneficial exchange is obtained. Allocations of students 1 and 3 are finalized and they are excluded from the queue, and also one slot in each of the schools A and B is excluded from the allocation process.

Finalized assignments	School A	School B	School C
slot 1	3	1	-
slot 2	-	-	-

Step 3. The queue is as follows: 2-4-5-6. Student 2 (the first in the queue) applies to school C. It is not her district school. Students 5 and 6 are tentatively assigned to the school C. One of the two students needs to be chosen. Between two students, student 5 is closer to the top of the queue – that is why she moves to the top of the queue. And thus the queue is modified.

Step 4. The queue is as follows: 5-2-4-6. Thus, student 5 (the first in the queue) applies to school C. This school is her district school: Thus student 5 is assigned to school C. Her allocation is finalized and she is excluded from the queue as well as the slot in school C.

Finalized assignments	School A	School B	School C
slot 1	3	1	5
slot 2	-	-	-

Step 5. The queue looks as follows: 2-4-6. Student 2 (the first in the queue) applies to school C again. It is not her district school. Student 6 is the only student left, who is tentatively assigned to school C. Thus student 6 moves to the top of the queue. And thus the queue is modified.

Step 6. The queue is as follows 6-2-4. Therefore student 6 (the first in the queue) applies to school A. This school is not her district school, but the cycle of beneficial exchange appears. Student 6 wants to attend student 2's school, and student 2 wants to attend student 6's school. The beneficial exchange is obtained. Allocations of students 2 and 6 are finalized and they are excluded from the queue, and also one slot in each of the schools A and C is also excluded from the allocation process.

Finalized assignments	School A	School B	School C
slot 1	3	1	5
slot 2	6	-	2

Step 7. There is only one student in the queue – student 4. She wants to apply to school C, but there are no more free slots there, and so she applies to the second choice, school B. It is her district school and she is assigned to the slot in school B.

Thus, final allocation of students is as follows:

Finalized assignments	School A	School B	School C
slot 1	3	1	5
slot 2	6	4	2

INSTRUCTIONS – Mechanism ETTC

This is an experiment in the economics of decision making. The instructions are simple, and if you follow them carefully and make good decisions, you might earn a considerable amount of money. In this experiment, we simulate a procedure to allocate students to schools. The procedure, payment rules, and student allocation method are described below. Do not communicate with each other during the experiment. If you have questions at any point during the experiment, raise your hand and the experimenter will help you.

Procedure

There are 10 participants in this experiment. You are participant 1. (This number is drawn by each participant before entering the room.) Each participant represents a student who wants to get a slot at a school. You will have to make three decisions in three different school markets. In all markets allocation will be determined by the same algorithm. This algorithm works as follows:

Allocation Method

Each participant is first tentatively assigned to the school within her respective district. Each school has a priority ordering of all other students. Students of the school district have the highest priority in the school. Next, Decision Sheet rankings are used to determine mutually beneficial exchanges between two or more participants. The order in which these exchanges are considered is determined by a fair lottery. This means each participant has an equal chance of being the first in the line, the second in the line,..., as well as the last in the line. To determine this fair lottery, a participant will be asked to draw 10 numbers from a basket, one at a time. Each number corresponds to a participant ID number. The sequence of the draw determines the order in the lottery.

The specific allocation process is explained below:

1. Initially, all 10 slots are available for allocation.

2. All participants are ordered in a queue based on the order in the lottery.

3. Next, an application to the highest-ranked school on the Decision Sheet is submitted for the participant at the top of the queue.

- If the application is submitted to her district school, then her tentative assignment is finalized (thus, she is assigned to a slot at her district school). The participant and her assignment are removed from subsequent allocations. The process continues with the next participant in line.

- If the application is submitted to another school, the procedure progresses as follows: Say applicant Claudia's home district school is school A and she is applying to school B. Then Claudia's application is submitted to school B. After that, one of the students who tentatively holds the slot at school B has to be chosen. In particular, among all students who tentatively hold a slot at school B, we choose the student with the highest priority at Claudia's district school, i.e., school A. (So we follow the priority ordering of Claudia's school, school A). This student is then moved to the top of the queue directly in front of the requester (Claudia).

Whenever the queue is modified, the process continues similarly: An application is submitted to the highest-ranked school with available slots for the participant at the top of the queue.

- If the application is submitted to her district school, then her tentative assignment is finalized. The process continues with the next participant in line.

- If the application is submitted to another school, say school S, then we follow the procedure, explained for Claudia: a participant with the highest priority in the district school of the requester among those who tentatively hold a slot at school S is moved to the top of the queue directly in front of the requester. This way, each participant is guaranteed an assignment which is at least as good as her district school based on the preferences indicated in her Decision Sheet.

4. A mutually-beneficial exchange is obtained when a cycle of applications are made in sequence, which benefits all affected participant; e.g., I apply to Stefan's district school, Stefan applies to your district school, and you apply to my district school. In this case, the exchange is completed and the participants as well as their assignments are removed from subsequent allocations.

5. The process continues until all participants are assigned a school slot.

Example:

In order to understand the mechanism better, let us go through an example together.

If you have any questions about any step of the allocation procedure, please feel free to ask at any point. There are six students (ID numbers from 1 to 6) on the market, and three schools (school A, school B, and school C) with two free slots each. Students 1 and 2 live in the school district A, students 3 and 4 live in school district B, and, finally, student 5 and 6 live in school district C. It means that tentative assignments are as follows:

Tentative assignments of students (IDs)	School A	School B	School C
slot 1	1	3	5
slot 2	2	4	6

The lottery determined the following order (student IDs): 1-2-3-4-5-6
Students submitted the following school rankings in their Decision Sheets:

Student ID	1	2	3	4	5	6
Top choice	B	C	A	C	C	A
Middle choice	A	A	C	B	A	B
Last choice	C	B	B	A	B	C

This allocation method consists of the following steps:

Step 1. The queue is as follows: 1-2-3-4-5-6 (the initial queue order is always determined by the lottery.) Thus student 1 (the first in the order) applies to school B. It is not her district school. Students 3 and 4 are tentatively assigned to school B. Student 1 has to choose between them. Student 4 has higher priority in school A (school of student 1) – that is why student 1 chooses student 4, and she moves to the top of the queue. And thus the queue is modified.

Step 2. The queue is as follows: 4-1-2-3-5-6. Thus student 4 applies to school C. This school is not her district school. Students 5 and 6 are tentatively assigned to school C. Student 4 has to choose between them. Student 5 has higher priority in school B (school of student 4), that is why student 4 chooses student 5, and she moves to the top of the queue. And thus the queue is modified.

Step 3. The queue is as follows: 5-4-1-2-3-6. Student 5 applies to the school C. It is her district school. Thus student 5 is assigned to school C. Her allocation is finalized, and she is excluded from the queue as well as the slot in school C.

Finalized assignments	School A	School B	School C
slot 1	-	-	5
slot 2	-	-	-

Step 4. The queue is as follows: 4-1-2-3-6. Thus student 4 applies to school C again. This school is not her district school. Student 6 is the only student left, who is tentatively assigned to school C, and she moves to the top of the queue. And thus the queue is modified.

Step 5. The queue is as follows: 6-4-1-2-3. Thus student 6 applies to school A. This school is not her district school. Students 1 and 2 are tentatively assigned to school A. Student 6 has to choose between them. Student 1 has higher priority in school B (school of student 6), that is why student 6 chooses student 1, and she moves to the top of the queue. And thus the queue is modified.

Step 6. The queue is as follows: 1-6-4-2-3. Thus student 1 applies to the school B again. It is not her district school, but the cycle of beneficial exchange appears. Student 1 wants to attend student 4's district school again, and at the same time student 4 wants to attend student 6's district school, and student 6 wants to attend student 1's district school. The beneficial exchange is obtained. Allocations of students 1, 4, and 6 are finalized and they are excluded from the queue, and also 1 slot in each of the schools A, B, and C is excluded from the allocation process.

Finalized assignments	School A	School B	School C
slot 1	6	1	5
slot 2	-	-	4

Step 7. The queue is as follows 2-3. Student 2 wants to apply to school C, but there are no free slots any more. Thus student 2 applies to her second choice – school A. It is her district school. Thus student 2 is assigned to the school A. Her allocation is finalized, and she is excluded from the queue, and also the slot in school A is excluded from the allocation process.

Finalized assignments	School A	School B	School C
slot 1	6	1	5
slot 2	2	-	4

Step 8. There is only one student in the queue – student 3. She wants to apply to school A, but there are no more free slots there, so she applies to the second choice – school C. But there are no more free slots there, either, and so she applies to school B. It is her district school and she is assigned to the slot in school B.

Thus, the final allocation of students is as follows:

Finalized assignments	School A	School B	School C
slot 1	6	1	5
slot 2	2	3	4

Questionnaire in Instructions(common for both mechanisms)

In order to check the level of understanding of the allocation procedure, we ask you to find out the students allocation for the following market: There are six students (ID numbers from 1 to 6) on the market, and three schools (school A, school B, and school C) with two free slots each. Students 1 and 2 live in the district of school A, students 3 and 4 live in the district of school B and, finally, students 5 and 6 live in the district of school C. It means that the tentative assignments are as follows:

The lottery determined the following order (student IDs): 1-2-3-4-5-6

Students submitted the following school rankings in their decision sheets:

You have 10 minutes to perform the task. If you have any questions, raise your hand and we will come to you. After 10 minutes, you must submit your answer sheet and then the experimenter will go through the solution on the board.

Tentative assignments of students (IDs)	School A	School B	School C
slot 1	1	3	5
slot 2	2	4	6

Student ID	1	2	3	4	5	6
Top choice	B	C	A	C	B	B
Middle choice	A	A	C	B	C	A
Last choice	C	B	B	A	A	C

Decision Sheets (common for both mechanisms)[54]

Decision sheet. Market TP1. Student ID1.

There are three schools in the market: school A, school B, and school C. School A and B have three slots each. School C has four slots.

Students 1, 4, and 7 live in the district of school A.
Students 2, 5, and 8 live in the district of school B.
Students 3, 6, 9, and 10 live in the district of school C.

Recall: You are student 1.

Your payoff amount depends on the school slot you hold at the end of the experiment. Payoff amounts for market 1 are outlined in the following table.

School	A	B	C
Payoff, EUR	10	15	5

You will be paid 10 euros if you hold a slot at school A in the market 1 at the end of the experiment.
You will be paid 15 euros if you hold a slot at school B in the market 1 at the end of the experiment.
You will be paid 5 euros if you hold a slot at school C in the market 1 at the end of the experiment.
The preferences of other students are as follows:

ID	1	2	3	4	5	6	7	8	9	10
top	B	C	B	C	A	A	C	A	A	B
middle	A	A	A	B	C	C	A	B	B	A
bottom	C	B	C	A	B	B	B	C	C	C

Note that the preferences above correspond to payoffs. They do not necessarily coincide with stated preferences, which are used for the allocation.

The priorities of students in schools are as follows:

[54] We use the preferences of student 1 and the market 1, for example.

School	A	B	C	D
High priority	1	3	5	8
	2	4	6	9
	7	8	7	10
	9	6	9	7
	10	5	2	4
	4	10	8	1
	8	9	4	3
	6	7	3	2
	5	2	1	6
No priority	3	1	10	5

Note that the framed blue highlights correspond to the students of the district. They have equally high priority.
Please write down your ranking of the schools (A through C) from your first choice to your last choice. Please rank ALL schools. You have five minutes to think about the rankings, and then we will proceed to the next market

First choice	Second choice	Third choice

Your experimenter Name(ID) ....
This is the end of market 1. Please submit your rankings to the experimenter.

References

Abdulkadiroglu, A., Y.-K. Che, P. A. Pathak, A. E. Roth, and O. Tercieux (2017): "Minimizing Justified Envy in School Choice: The Design of New Orleans' OneApp," Tech. rep., National Bureau of Economic Research.

Abdulkadiroğlu, A., Y.-K. Che, and Y. Yasuda (2011): "Resolving Conflicting Preferences in School Choice: The Boston Mechanism Reconsidered," *American Economic Review*, 101, 399–410.

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, 1954–1978.

Abdulkadiroğlu, A. and T. Sönmez (1998): "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems," *Econometrica*, 66, 689–701.

——— (1999): "House Allocation with Existing Tenants," *Journal of Economic Theory*, 88, 233–260.

——— (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729–747.

Abrevaya, J. (2008): "On Recombinant Estimation for Experimental Data," *Experimental Economics*, 11, 25–52.

Calsamiglia, C., G. Haeringer, and F. Klijn (2010): "Constrained School Choice: An Experimental Study," *American Economic Review*, 100, 1860–1874.

Chen, Y. and T. Sönmez (2002): "Improving Efficiency of On-campus Housing: An Experimental Study," *American Economic Review*, 92, 1669–1686.

——— (2006): "School Choice: An Experimental Study," *Journal of Economic Theory*, 127, 202–231.

Dubins, L. E. and D. A. Freedman (1981): "Machiavelli and the Gale-Shapley Algorithm," *American Mathematical Monthly*, 88, 485–494.

Dur, U. M. and M. U. Ünver (2015): "Two-sided Matching via Balanced Exchange: Tuition and Worker Exchanges," *working paper*.

Ehlers, L. (2002): "Coalitional Strategy-Proof House Allocation," *Journal of Economic Theory*, 105, 298–317.

Ehlers, L. and B. Klaus (2003): "Resource-Monotonicity for House Allocation Problems," *International Journal of Game Theory*, 32, 545–560.

Ehlers, L., B. Klaus, and S. Pápai (2002): "Strategy-Proofness and Population-Monotonicity for House Allocation Problems," *Journal of Mathematical Economics*, 38, 329–339.

Erdil, A. and H. Ergin (2008): "What's the Matter with Tie-breaking? Improving Efficiency in School Choice," *American Economic Review*, 98, 669–689.

Gale, D. and L. S. Shapley (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–15.

Greiner, B. (2015): "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," *Journal of the Economic Science Association*, 114–125.

Guillen, P. and R. Hakimov (2017a): "The Effectiveness of Top-Down Advice in Strategy-Proof Mechanisms: a Field Experiment," *European Economic Review*.

——— (2017b): "Not Quite the Best Response: Truth-Telling, Strategy-Proof Matching, and the Manipulation of Others," *Experimental economics*, 20(3), 670–686.

Guillen, P. and A. Hing (2014): "Lying Through Their Teeth: Third Party Advice and Truth Telling in a Strategy Proof Mechanism," *European Economic Review*, 70, 178–185.

Guillen, P. and O. Kesten (2012): "Matching Markets with Mixed Ownership: The Case for a Real-life Mechanism," *International Economic Review*, 53(3), 1027–1046.

Kesten, O. (2006): "On Two Competing Mechanisms for Priority-based Allocation Problems," *Journal of Economic Theory*, 127, 155–171.

——— (2009): "Coalitional Strategy-Proofness and Resource Monotonicity for House Allocation Problems," *International Journal of Game Theory*, 38, 17–21.

——— (2010): "School Choice with Consent," *Quarterly Journal of Economics*, 125(3), 1297–1348.

Kesten, O. and M. Kurino (2016): "Do Outside Options Matter in Matching? A New Perspective on the Trade-offs in Student Sssignment," Tech. rep., working paper.

Liu, Q. and M. Pycia (2016): "Ordinal Efficiency, Fairness, and Incentives in Large Markets," Tech. rep., working paper.

Ma, J. (1994): "Strategy-proofness and the Strict Core in a Market with Indivisibilities," *International Journal of Game Theory*, 23, 75–83.

Miyagawa, E. (2001): "House Allocation with Transfers," *Journal of Economic Theory*, 100, 329–355.

Mullin, C. H. and D. H. Reiley (2006): "Recombinant Estimation for Normal-form Games, with Applications to Auctions and Bargaining," *Games and Economic Behavior*, 54, 159–182.

Pais, J. and Á. Pintér (2008): "School Choice and Information: An Experimental Study on Matching Mechanisms," *Games and Economic Behavior*, 64, 303–328.

Pápai, S. (2000): "Strategyproof Assignment by Hierarchical Exchange," *Econometrica*, 68, 1403–1433.

Pápai, S. (2013): "Individual Preference Restrictions in Hedonic Coalition Formation Games," Tech. rep., working paper.

Pycia, M. and M. U. Ünver (2017): "Incentive Compatible Allocation and Exchange of Discrete Resources," *Theoretical Economics*, 12, 287–329.

Roth, A. E. (1982a): "Incentive Compatibility in a Market with Indivisible Goods," *Economics letters*, 9, 127–132.

——— (1982b): "The Economics of Matching: Stability and Incentives," *Mathematics of operations research*, 7, 617–628.

Roth, A. E. and A. Postlewaite (1977): "Weak Versus Strong Domination in a Market with Indivisible Goods," *Journal of Mathematical Economics*, 4, 131–137.

Roth, A. E., T. Sönmez, and M. U. Ünver (2005): "Pairwise Kidney Exchange," *Journal of Economic Theory*, 125, 151–188.

Schummer, J. and R. V. Vohra (2013): "Assignment of Arrival Slots," *American Economic Journal: Microeconomics*, 5, 164–185.

Shapley, L. and H. Scarf (1974): "On Cores and Indivisibility," *Journal of Mathematical Economics*, 1, 23–37.