



SEII

SCHOOL EFFECTIVENESS
& INEQUALITY INITIATIVE

Discussion Paper #2017.01

Minimizing Justified Envy in School Choice: The Design of New Orleans' OneApp

Atila Abdulkadiroglu
Yeon-Koo Che
Parag A. Pathak
Alvin E. Roth
Olivier Tercieux

March 2017

MIT Department of Economics
77 Massachusetts Avenue, Bldg E52-300
Cambridge, MA 02139

National Bureau of Economic Research
1050 Massachusetts Avenue, 3rd Floor
Cambridge, MA 02138

Minimizing Justified Envy in School Choice: The Design of New Orleans' OneApp
Atila Abdulkadiroglu, Yeon-Koo Che, Parag A. Pathak, Alvin E. Roth, and Olivier Tercieux
SEII Discussion Paper #2017.01
March 2017

ABSTRACT

In 2012, New Orleans Recovery School District (RSD) became the first U.S. district to unify charter and traditional public school admissions in a single-offer assignment mechanism known as OneApp. The RSD also became the first district to use a mechanism based on Top Trading Cycles (TTC) in a real-life allocation problem. Since TTC was originally devised for settings in which agents have endowments, there is no formal rationale for TTC in school choice. In particular, TTC is a Pareto efficient and strategy-proof mechanism, but so are other mechanisms. We show that TTC is constrained-optimal in the following sense: TTC minimizes justified envy among all Pareto efficient and strategy-proof mechanisms when each school has one seat. When schools have more than one seat, there are multiple possible implementations of TTC. Data from New Orleans and Boston indicate that there is little difference across these versions of TTC, but significantly less justified envy compared to a serial dictatorship.

* We are grateful to Neil Dorosin, Gabriela Fighetti, and John White for assistance with this study. Thanks also to seminar participants at the 2016 ASSA meetings for helpful comments. Vira Semenova provided excellent research assistance. Pathak is grateful to the W. T. Grant Foundation and the National Science Foundation for support. This paper supersedes "The Role of Priorities in Assigning Indivisible Objects: A Characterization of Top Trading Cycles" cited by others as Abdulkadiroglu, Atila and Yeon-Koo Che (2010) or Abdulkadiroglu, Atila, Yeon-Koo Che, and Olivier Tercieux (2010). Abdulkadiroglu, Pathak, and Roth are members of the scientific advisory board of the Institute for Innovation in Public School Choice. IIPSC was involved in designing OneApp in New Orleans. Abdulkadiroglu, Pathak, and Roth also advised Boston Public Schools and New York City's Department of Education on designing their student assignment systems, discussed herein. This article does not represent the views of the New Orleans Recovery School District or any other school district. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

1 Introduction

Research on priority-based resource allocation focuses on two central mechanisms: one based on Gale and Shapley (1962)'s deferred-acceptance (DA) algorithm and the other based on Gale's top trading cycles (TTC) algorithm (Shapley and Scarf 1974). This focus is due in part to the fact that both mechanisms are strategy-proof as direct mechanisms: truthful reporting is a weakly dominant strategy for the proposing side in DA and for individuals in TTC. The main difference between the mechanisms is that DA produces an outcome that is not Pareto efficient but eliminates justified envy. An allocation eliminates justified envy if there is no blocking pair; that is, if no individual prefers another assignment over her assignment and has a higher priority at the preferred assignment. TTC is Pareto efficient but does not eliminate justified envy. There is no mechanism that is both Pareto efficient and eliminates justified envy.¹

Abdulkadiroğlu and Sönmez (2003) proposed resolving this trade-off in school assignment mechanism design as follows: if the designer prefers the elimination of justified envy over efficiency, then use DA. If the designer prefers efficiency over elimination of justified envy, then use TTC.² Since Gale and Shapley (1962) showed that DA produces a justified envy-free matching that is not Pareto dominated by any other justified envy-free matching, the argument for DA rests on a firm foundation. DA is constrained-optimal in the sense that it weakly Pareto dominates any other matching with no justified envy.

Without a similar constrained-optimality result for TTC, the formal argument for TTC is not on as solid footing. In particular, there are other strategy-proof and Pareto efficient mechanisms for priority-based resource allocation. For example, a serial dictatorship mechanism, which assigns each student in a given order to her most preferred choice among the remaining schools, is strategy-proof and Pareto efficient. Importantly, priorities have no role in this mechanism, in the sense that changing priorities does not alter the outcome for a given ordering of students. Under DA, priorities have a clear role since there is no justified envy. However, the role of priorities in TTC is much less clear, since TTC allows for justified envy but uses priorities to compute allocations.

In this paper, we provide a constrained optimality result for TTC that is analogous to the result for DA. To compare efficient allocations, we measure the extent to which two allocations eliminate justified envy. Our approach is inspired by recent literature examining problem-wise

¹This fact is a consequence of Roth (1982a), who shows that an allocation that is free of justified envy need not be Pareto efficient. Kesten (2010) shows a related result: there is no Pareto efficient and strategy-proof mechanism that selects the Pareto efficient and justified-envy-free allocation whenever it exists. Ergin (2002) characterizes situations where there is no conflict between elimination of justified envy and Pareto efficiency.

²Abdulkadiroğlu and Sönmez (2003) state (p. 732): "... the choice between these two competing mechanisms depends on the structure and interpretation of the priorities. In some applications, policy makers may rank complete elimination of justified envy before full efficiency, then Gale-Shapley student optimal stable mechanism can be used in those cases. Efficiency may be ranked higher by others, and the top trading cycles mechanism can be used in such applications."

comparisons to contrast mechanisms.³ We say that mechanism φ_1 **has less justified envy** than mechanism φ_2 if, for any problem, the set of blocking pairs of mechanism φ_1 is a subset of the set of blocking pairs of mechanism φ_2 . Our main result states that there is no Pareto efficient and strategy-proof mechanism that has less justified envy than TTC, when each school has one seat. When minimizing justified envy is second only to efficiency, this result provides the missing foundation for recommending TTC.

Lack of a rigorous understanding of the roles priorities and preferences play in TTC may also explain its limited use in practice. For instance, from 2003 to 2005, a Student Assignment Task force in Boston, led by civic leaders, evaluated possible improvements to the school choice system. As part of their review, the task force recommended TTC over DA, and provided the following statement, which indicates how they perceived the mechanisms after reviewing the academic literature (Landsmark, Dajer and Gonsalves 2014):

[T]he Gale-Shapley algorithm [...] cuts down on the amount of choice afforded to families. The Top Trading Cycles algorithm also takes into account priorities while leaving some room for choice. [...] [C]hoice was very important to many families who attended community forums...

Based on this report and further consultation with the community and academic experts, the final recommendation was for DA (Abdulkadiroğlu, Pathak, Roth and Sönmez 2005a). TTC was faulted for allowing students to trade as stated in the final school committee report (BPS 2005):

[TTC's] trading shifts the emphasis onto the priority and away from the goals BPS is trying to achieve by granting these priorities in the first place.

The OneApp assignment system used for all public school children in New Orleans' Recovery School District (RSD) was based on TTC in its initial year in 2012 (Vanacore 2012). To the best of our knowledge, the RSD is the only place that TTC has been used in practice, in school assignment or elsewhere.^{4,5} However, after one year, officials abandoned TTC in favor of DA,

³For examples in different contexts, see Chen and Kesten (2016), Kesten (2006), Maskin (2016), and Pathak and Sönmez (2013).

⁴A team of economists led by Muriel Niederle and Clayton Featherstone proposed a TTC system for school choice for San Francisco, and the school board approved it in March 2010 (see San Francisco Unified School District (2010), Board Approves New Student Assignment System for San Francisco Schools, Press Release, March 10, 2010, <https://web.archive.org/web/20100331064212/http://portal.sfusd.edu/data/news/pdf/3%2010%2010%20SF%20School%20Board%20Approves%20New%20Student%20Assignment%20System.pdf>). However, it appears that the district staff never implemented the Board's decision; in any event, the school district has declined to make public its current algorithm. For a blow-by-blow description of events, browse the posts at Roth, Alvin E., Market Design blog, <http://marketdesigner.blogspot.com/search?q=Francisco+AND+school>.

⁵A version of TTC was initially proposed as a way to organize kidney exchange (see Roth, Sönmez and Ünver (2004)) but was never used in practice due in part to the difficulty of implementing long cycles. Instead, an incremental process has led to the use of long chains and cycles identified through integer programming (see Anderson, Ashlagi, Gamarnik and Roth (2015)).

in part due to its treatment of priorities. Motivated by these developments, we use data from New Orleans and Boston to report on how TTC compares to other efficient mechanisms and to DA. The empirical comparison also shows the relevance of our results when schools have more than one seat.

This paper is related to several other papers on axiomatic mechanism design characterizing strategy-proof mechanisms. The most closely related result is Ma (1994), who characterizes TTC as the unique strategy-proof, Pareto efficient, and individually rational mechanism in the housing market model of Shapley and Scarf. (Strategy-proofness was established in Roth (1982b)). In Section 5, we show how our main result generalizes Ma (1994). Pápai (2000) introduces and characterizes a wide class of mechanisms, known as hierarchical exchange rules, by Pareto efficiency, group strategy proofness, and re-allocation proofness. TTC is a hierarchical exchange rule defined by the school priority lists but Pápai (2000) does not make explicitly reference any specific priority and therefore is silent on how TTC respects priorities. In a similar vein, Pycia and Ünver (2016) introduce and characterize trading cycles with brokers and owners by Pareto efficiency and group strategy-proofness.⁶ Dur (2013) and Morrill (2015a) each characterize TTC by relying on additional axioms, which are distinct from our result. Following the circulation of Abdulkadiroglu and Che (2010), Morrill (2015a) developed a characterization of TTC based on a new fairness concept that he terms “just.” Essentially, this approach amounts to excluding some incidences of justified envy from the fairness consideration.⁷ By contrast, we consider all incidents of justified envy in our fairness consideration.

There are other notable characterizations of mechanisms used for school assignment. Alcalde and Barberá (1994) show that student-proposing DA is the only mechanism that eliminates justified-envy, is individually rational, is non-wasteful, and is strategy-proof. Balinski and Sönmez (1999) show that DA is the only matching mechanism that eliminates justified-envy, is individually rational, is non-wasteful, and respects improvements. Kojima and Manea (2010) characterize DA using Maskin monotonicity.

This paper also aligns with more recent studies motivated by TTC mechanisms’ fairness implications. Hakimov and Kesten (2014) and Morrill (2015b) propose TTC variants that improve fairness under some circumstances; we shall discuss these extensively. Che and Tercieux (2016a), Che and Tercieux (2016b) and Leshno and Lo (2017) study Top Trading Cycles in large economies under different asymptotics.⁸ They show that the amount of justified envy entailed

⁶See also Bade (2016). Pycia and Ünver (2011) extend their characterization when each discrete resource can be represented by a copy.

⁷Specifically, Morrill (2015a) declares a mechanism to be unjust (only) if i has justified envy toward j at j ’s assignment s and reassigning j could potentially lead to a “change” of assignment for some student k who has a higher priority at s than i . Importantly, this notion makes no distinction whether the “change” makes student k better or worse off.

⁸Che and Tercieux (2016a) and Che and Tercieux (2016b) consider a setting in which the number of students and the number of schools grow large, whereas Leshno and Lo (2017) study a setting in which the number of students and the number of seats at each school grow large while the number of schools remains fixed and finite.

by TTC does not vanish as the economy grows large.

The rest of this paper is organized as follows. Section 2 introduces the model and our method of comparing mechanisms in terms of their sets of blocking pairs. Section 3 describes a simple example we use to demonstrate our main theorem. Section 4 defines TTC and states our main result. Section 5 describes extensions when schools have multiple seats, when priorities are not strict, and for the housing market model of Shapley and Scarf. Section 6 uses data from New Orleans and Boston to compare assignment mechanisms and reviews design issues that led to the adoption of DA after one year in New Orleans. Section 7 concludes.

2 The Model

2.1 Notation

We define a priority-based resource allocation problem in terms of school assignment, following Abdulkadiroğlu and Sönmez (2003). There are a number of students, each of whom should be assigned a seat at one of a number of schools. Each student has a strict preference ordering over all schools as well as remaining unassigned, and each school has a strict priority ranking of all students. Each school also has a maximum capacity.

Formally, a school choice problem consists of:

1. a set of students $I = \{i_1, \dots, i_n\}$,
2. a set of schools $S = \{s_1, \dots, s_m\}$,
3. a capacity vector $q = (q_{s_1}, \dots, q_{s_m})$,
4. a list of strict student preferences $P = (P_{i_1}, \dots, P_{i_n})$, and
5. a list of strict school priorities $\succ = (\succ_{s_1}, \dots, \succ_{s_m})$.

For any student i , P_i is a strict preference relation over $S \cup \{i\}$ where $sP_i i$ means student i strictly prefers a seat at school s to being unassigned. For any student i , let R_i denote the “at least as good as” relation induced by P_i . For any school s , \succ_s is a complete, irreflexive and transitive binary priority relation over I .⁹ Thus, $i \succ_s j$ means that i has strictly higher priority at s than j . The model assumes that priorities are strict, and we examine the weak priority case in Section 5. Priority rankings are determined by the school district, and schools have no control over them. We also use the following shorthand notation: $P_{-I'} = (P_i)_{i \in I - I'}$ and $\succ_{-S'} = (\succ_s)_{s \in S - S'}$.

We fix the set of students, the set of schools and the capacity vector throughout the paper; hence the pair (P, \succ) denotes a school choice problem (or simply an **economy**).

⁹Our results follow under a more general assumption that \succ_s is a complete, irreflexive and transitive binary priority relation over $I \cup \{s\}$, where $i \succ_s s \succ_s j$ means that i is eligible for a seat at s and j is not eligible.

The outcome of a school choice problem is a **matching**. Formally, a matching $\mu : I \rightarrow S \cup I$ is a function such that

1. $\mu(i) \notin S \Rightarrow \mu(i) = i$ for any student i , and
2. $|\mu^{-1}(s)| \leq q_s$ for any school s .

We refer $\mu(i)$ as the assignment of student i under matching μ .

A matching μ **Pareto dominates** a matching ν , if $\mu(i)R_i\nu(i)$ for all $i \in I$ and $\mu(i)P_i\nu(i)$ for some $i \in I$. A matching is **Pareto efficient** if it is not Pareto dominated by any other matching.

A matching μ **eliminates justified envy** if

1. it is individually rational in the sense that there is no student i who prefers remaining unassigned to her assignment $\mu(i)$, and
2. there is no student-school pair (i, s) such that,
 - (a) student i prefers s to her assignment $\mu(i)$, and
 - (b) either school s has a vacant seat under μ or there is a lower priority student j who nonetheless received a seat at school s under μ .

If there is a student i who violates the first condition, we say the matching is **blocked by an individual**. If there is student-school pair (i, s) that violates the second condition, we say that student i **has justified envy** and (i, s) form a **blocking pair**. A matching **eliminates justified envy** if it is not blocked by either any individual or any pair. Eliminating justified envy is mathematically the same as stability for the college admissions model (Gale and Shapley 1962), when colleges have responsive and acceptant preferences (Roth 1985).¹⁰ The main difference involves interpretation: in college admissions, each school is a strategic agent, so stability has strategic implications, whereas in priority-based resource allocation eliminating justified envy is motivated by fairness considerations.

A **mechanism** selects a matching for each economy. If φ is a mechanism, let $\varphi(P, \succ)$ denote the matching selected by φ . Let $\varphi(P, \succ)(i)$ denote the assignment of i in matching $\varphi(P, \succ)$. A mechanism is Pareto efficient if it only selects Pareto efficient matchings, i.e., for any problem (P, \succ) , the matching $\varphi(P, \succ)$ is Pareto efficient. Similarly, a mechanism eliminates justified envy if it only selects justified envy-free matchings. Finally, a mechanism φ is **strategy-proof** if reporting true preferences is a dominant strategy for every individual in the preference revelation game induced by φ ; that is, for all P, \succ, i and P'_i

$$\varphi(P, \succ)(i)R_i\varphi(P'_i, P_{-i}, \succ)(i).$$

¹⁰A college's preferences are responsive if there is a linear order of the students according to which it wishes to fill its seats (up to its quota). Its preferences are acceptant if it would accept any student if a seat is available; i.e., it finds no student unacceptable.

2.2 Mechanisms

We describe three mechanisms for priority-based resource allocation. The simplest is a **serial dictatorship**, which is defined through a priority ordering of agents. For a given preference P , a serial dictatorship assigns the highest priority agent her first choice, the second highest priority agent her top choice among schools with remaining capacity, and so on. This mechanism is strategy-proof and Pareto efficient (see, e.g., Roth (1982a), Abdulkadiroğlu and Sönmez (1998), and Svensson (1999)).

Abdulkadiroğlu and Sönmez (2003) adapt Gale’s TTC for settings with priorities. Given a problem (P, \succ) , **TTC** finds a matching via the following algorithm:

- **Step 1:** *Assign each school a counter that keeps track of how many seats are still available at the school. Initially set the counters equal to the capacities of the schools. Each student points to her favorite school under her announced preferences. If a student has no acceptable school, they are removed from the market. Each school points to the student who has the highest priority for the school. Since the numbers of students and schools are finite, there is at least one cycle. A cycle $c = \{s_k, i_k\}_{k=1, \dots, K}$ is an ordered list of schools and individuals such that s_k points to i_k and i_k points to s_{k+1} for every k where $s_{K+1} = s_1$. Moreover, each school can be part of at most one cycle. Similarly, each student can be part of at most one cycle. Every student in a cycle is assigned a seat at the school she points to and is removed. Each school’s counter reduces by one in each cycle, and if the counter reduces to zero, the school is also removed. All other schools’ counters stay put.*

In general, at

- **Step k :** *Each remaining student points to her favorite school among the remaining schools. If a student has no acceptable school, they are removed from the market. Each remaining school points to the student with highest priority among the remaining students. There is at least one cycle. Every student in a cycle is assigned a seat at the school that she points to and is removed. Each school’s counter reduces by one in each cycle, and if the counter reduces to zero, the school is also removed. All other schools’ counters stay put.*

The algorithm terminates when no more students can be assigned seats. We denote the outcome of this mechanism as $TTC(P, \succ)$. This algorithm trades priorities of students among themselves starting with students with highest priorities. Abdulkadiroğlu and Sönmez (2003) show this mechanism is strategy-proof and Pareto efficient.

The last mechanism we define is based on Gale and Shapley’s (1962) celebrated DA algorithm, which assigns student as follows:

- **Step 1:** *Each student proposes to her first choice. Each school tentatively assigns its seats to its proposers one at a time, following their priority order. Any remaining proposers are rejected.*

In general, at

- **Step k :** *Each student who was rejected in the previous step proposes to her next choice. Each school considers the students it has been holding together with its new proposers and tentatively assigns its seats to these students one at a time, following their priority order. Any remaining proposers are rejected.*

The algorithm terminates when there are either no new proposals or each student has exhausted her rank order list. Dubins and Freedman (1981) and Roth (1982a) show that DA is strategy-proof. Gale and Shapley (1962) show that it Pareto dominates any other mechanism that eliminates justified envy. DA is widely used in school choice systems, including in New York City (Abdulkadiroğlu, Pathak and Roth 2005b, Abdulkadiroğlu, Pathak and Roth 2009), London and elsewhere (see, e.g., Pathak and Sönmez (2013) and Pathak (2011)).

2.3 Comparing Mechanisms

To isolate the way in which TTC uses priorities over other efficient and strategy-proof mechanisms, we need a way to compare efficient allocations based on other properties. A recent literature examines problem-wise comparisons to rank mechanisms by properties that are not satisfied for each problem (see, e.g., Chen and Kesten (2016), Kesten (2006), Maskin (2016), and Pathak and Sönmez (2013).) We follow this approach with our definition:

Definition 1. *A mechanism φ has **less justified envy** than ψ at \succ , if for any P and student-school pair (i, s) , if pair (i, s) blocks $\varphi(P, \succ)$, then pair (i, s) blocks $\psi(P, \succ)$.*

When we compare mechanisms for a given problem, this definition is not equal to a comparison of the count of blocking pairs across mechanisms. Rather, the concept involves a subset relationship: the set of individuals and schools that are in blocking pairs is at least as large under ψ as under for φ for each problem.

Two mechanisms can have less justified envy than each other if they have the same blocking pairs for each problem. Our next definition rules this out:

Definition 2. *A mechanism φ has **strictly less justified envy** than ψ if φ has less envy than ψ , but ψ does not have less envy than φ .*

Since the concept of less justified envy defines a preorder, our last definition describes the minimal element of that order.¹¹

Definition 3. *Given a class of mechanisms \mathcal{C} , φ is **justified envy minimal** in \mathcal{C} if there is no other mechanism ψ in \mathcal{C} that has strictly less envy than φ .*

¹¹The concept of less justified envy does not define a partial order since the preorder it defines violates anti-symmetry. To see this, let ψ and ψ' be two distinct, justified envy-free mechanisms. Clearly, ψ has less justified envy than ψ' , and ψ' has less justified envy than ψ . However, $\psi \neq \psi'$.

These definitions compare the sets of students who have justified envy under two different mechanisms. As will be clear, we shall perform this comparison within the class of mechanisms that are Pareto efficient and strategy-proof. In the rest of the paper, when not specified, the class of mechanisms, \mathcal{C} , are Pareto efficient and strategy-proof mechanisms.

Our problem-wise comparison is related to similar ideas for different contexts. The most closely related notion was examined by Chen and Kesten (2016), who say that a mechanism φ is more stable than ϕ if at any problem where ϕ is stable, φ is also stable, and there is a problem where ϕ is stable and φ is not.¹² This concept cannot be used to compare mechanisms that do not produce stable outcomes for any problem, while our notion does allow for comparisons in those cases. In that respect, our notion allows for more refined comparisons: if φ has less justified envy than ϕ , then φ is more stable than ϕ .¹³

3 Motivating Example

We begin with an example illustrating the tradeoff between efficiency and eliminating justified envy.¹⁴ Suppose there are three students $I = \{i_1, i_2, i_3\}$ and three schools $S = \{s_1, s_2, s_3\}$, each with one seat. Preferences and priorities are given by

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

DA produces matching

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix},$$

which eliminates justified envy but is Pareto inefficient. In particular, it is Pareto dominated by another matching:

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_3 \end{pmatrix},$$

which is Pareto efficient. TTC, by letting students i_1 and i_2 trade their top priorities respectively at s_1 and s_2 , implements this outcome. The trade leaves i_3 justifiably envying students i_1 and i_2 , as her priorities at s_1 and s_2 are violated.

¹²Hakimov and Kesten (2014) use this idea to compare the stability of two versions of TTC in a problem with two schools.

¹³Variations on problem-wise comparisons used to compare mechanisms by their vulnerability to manipulation proposed by Pathak and Sönmez (2013) are examined in Andersson, Ehlers and Svensson (2014), and Arribillaga and Massó (2016a) and Arribillaga and Massó (2016b). Aside from preference manipulation, Kesten (2006) shows that given a priority structure, if DA is not consistent for an economy, neither is TTC.

¹⁴As we already mentioned in the introduction, this tradeoff is well known. See, for instance, Roth (1982a) and Abdulkadiroğlu and Sönmez (2003).

In short, TTC *uses* priorities but does not protect *all* students' priorities. At the same time, it does protect the priorities of students i_1 and i_2 , as they are no worse off than their respective top priority schools. Our main question is whether it is possible to further reduce justified envy without sacrificing efficiency or strategy-proofness? Put simply, the answer is no. Consider any mechanism $\psi \neq TTC$, and suppose ψ does not match i_1 with s_2 . (The argument will be the same if ψ matches i_2 differently from TTC.) For student i_1 to remain free of justified envy under ψ , however, he cannot be assigned s_3 , so she must match with s_1 . Given this, student i_2 must match with s_2 , or else i_2 will have justified envy. Hence, the only way to reduce the amount of justified envy produced by TTC is to match students according to the DA matching above, which is Pareto inefficient.¹⁵ Summing up, if one requires Pareto efficiency, strategy-proofness, *and* respect for i_1 and i_2 's priorities, just as in the above assignment, then the matching *must* coincide with the TTC matching. This illustrates the “envy-minimizing role” of TTC in its use of priorities.

Take serial dictatorship, which is another prominent efficient and strategy-proof mechanism. If we require the priorities of i_1 and i_2 to be respected, the serial order must be either $i_1 - i_2 - i_3$ or $i_2 - i_1 - i_3$; both cases result in the same matching as the TTC matching. But if preferences were different, serial dictatorship would typically do worse than TTC under either serial order. Suppose the students' preferences are now:

P_{i_1}	P_{i_2}	P_{i_3}
s_3	s_2	s_3
s_2	s_3	s_2
s_1	s_1	s_1

(Their priorities are the same as before.) TTC will now assign

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix},$$

which is not only efficient but also free of justified envy; note that all students are assigned to their top priority schools. Under either of two serial orders, however, the serial dictatorship mechanism assigns:

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_3 & s_2 & s_1 \end{pmatrix},$$

which violates i_3 's priority.

Next, we will formalize the sense in which TTC “economizes” on the incidence of justified envy and showing that its main contender, serial dictatorship, fails to to economize, both theoretically and empirically.

¹⁵This argument shows that if ψ has less justified envy than TTC, then the two mechanisms must coincide at the preference profile we are considering. Our proof of Theorem 1 builds on this argument but must also use strategy-proofness to show that ψ and TTC actually coincide at other preference profiles.

4 TTC as Justified Envy Minimizer

4.1 Main result

Pareto efficiency and eliminating justified envy are two conflicting concepts since there is no mechanism that satisfies both properties. Our main result states the following constrained-optimality property of TTC:

Theorem 1. *Suppose each school has one seat. Let φ be a Pareto efficient and strategy-proof mechanism. If φ has less justified envy than TTC at \succ , then $\varphi(\cdot, \succ) = \text{TTC}(\cdot, \succ)$.*

This result implies that TTC is justified envy-minimal in the class of Pareto efficient and strategy-proof mechanisms. The proof, provided in the Appendix, proceeds by contradiction. We first posit that there is an efficient mechanism φ with less justified envy than TTC. That is, there is a problem where φ produces an efficient allocation that has less justified envy than the allocation produced by TTC. By comparing the outcome of φ to the steps of TTC, the proof shows that φ cannot have less justified envy than the TTC allocation without contradicting either the Pareto efficiency or strategy-proofness of φ .

We do not claim TTC to be the only justified-envy minimal mechanism in the class of Pareto efficient and strategy-proof mechanisms. However, we can identify a well-known mechanism in this class that is not justified-envy minimal. For each school s , let $f_s : \succ_s \rightarrow \succ_s$ be an arbitrary function that transforms its priority into another (possibly same or distinct) priority. Let $f = (f_s)_s$. Consider a class of mechanisms $\varphi(\cdot, \succ) = \text{TTC}(\cdot, f(\succ))$. In other words, the new mechanism is obtained by running a TTC under priorities that may differ from true priorities. Observe that φ is Pareto efficient and strategy-proof. A well-known mechanism in this class is serial dictatorship, which is obtained if for some ordering \succ over students, $f_s(\succ_s) = \succ$ for any \succ_s and s .

Proposition 1. *Suppose $f_s(\succ_s) \neq \succ_s$ for some school s . Then, the mechanism $\varphi(\cdot, \succ) = \text{TTC}(\cdot, f(\succ))$ is not justified-envy minimal.*

Proof. Assume $\succ'_s := f_s(\succ_s) \neq \succ_s$. This must mean that there are at least two students i_1 and i_2 such that $i_1 \succ_s i_2$, but $i_2 \succ'_s i_1$. Consider an environment where school s is the only school. Clearly, in this environment, there is trivially a strategy-proof, Pareto efficient, and justified envy-free mechanism. For instance, TTC is one such mechanism. Further assume that students i_1 and i_2 are the only students ranking school s as acceptable. It is obvious that in such a case, the “distorted” TTC fails to select an envy-free assignment. Hence, it cannot be justified envy-minimal.¹⁶ ■

¹⁶While it may look obvious that running TTC with the “wrong” priorities will fail to be justified envy-minimal, we provide an example in Appendix C showing that, if we are allowed more fine-tuning (for instance, if the priority modifications are allowed to depend on the preferences) a modified version of TTC can be justified envy-minimal in the class of strategy-proof and Pareto-efficient mechanisms.

This result shows that the envy-minimality of TTC is nontrivial. It also highlights the roles played by the priorities in economizing on justified envy: the envy-minimality fails even with minimally modified priorities.

4.2 “Tightness” of Theorem

Our characterization is tight in the sense that relaxing any of the assumptions yields a mechanism with less justified envy and the remaining properties. Indeed, when we drop the efficiency requirement, TTC does not minimize envy in the class of strategy-proof mechanisms. In that class, TTC produces matchings allowing for justified envy, while DA does not. Therefore, DA minimizes justified envy in that class.

In addition, TTC does not minimize envy in the class of Pareto efficient mechanisms, when we drop the strategy-proofness requirement. The following example provides a mechanism that is Pareto efficient and has strictly less justified envy than TTC.

Example 1 (A Pareto efficient but not strategy-proof mechanism with less justified envy than TTC). *Suppose there are three students $I = \{i_1, i_2, i_3\}$ and three schools $S = \{s_1, s_2, s_3\}$. Preferences and priorities are given by*

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

TTC yields the matching

$$\mu = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_3 \end{pmatrix}.$$

Matching μ is blocked by the pair (i_3, s_1) and the pair (i_3, s_2) . The following matching is Pareto efficient and justified envy-free:

$$\nu = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_3 & s_2 & s_1 \end{pmatrix}.$$

We build the mechanism φ as follows: for any problem other than the one considered above, let φ pick the same matching as TTC. For the above problem, φ picks ν , the efficient and justified envy-free matching. Clearly, φ is Pareto efficient and φ has less justified envy than TTC. Our main result implies that mechanism φ is not strategy-proof. Indeed, if the preferences are given as in the above instance, student i_1 is assigned to school s_3 . However, if i_1 misreports his preferences by switching school s_3 and s_1 , then we obtain the TTC outcome: i_1 is matched to s_2 and is better off.

Another assumption in our characterization result is that each school has a single seat. We discuss the case with multiple school seats in detail in the next section.

5 Extensions

5.1 Multiple school seats

When schools may have more than one seat, there are multiple possible implementations of TTC. We refer to the version defined in Section 2.2 as **TTC-Counters**, since each school starts with a counter.

Morrill (2015b) introduces a TTC variation known as **TTC-Clinch and Trade**. This mechanism is motivated by the idea that TTC-Counters allows an applicant i to use her priority at another school to “trade” into a school even when the applicant is among the highest q_s priority students at school s . These trades may lead to unnecessary situations of justified envy, which can be avoided by “clinching” or assigning applicants in this highest-priority group before forming and executing cycles. The mechanism works as follows:

- **Step 1:**

- 1a) *For each $i \in I$, if it is one of the q_s highest ranked students at i 's most preferred school s , then assign i to s , remove i , and set $q_s = q_s - 1$. Whenever we remove a student, we adjust the rankings of all schools accordingly. This is called clinching a school. Iterate the clinching procedure until no student has one of the q_s highest priorities at her most preferred school s .*
- 1b) *Have each student that remains point to her most preferred school that has capacity greater than zero. Have each school with available capacity point to the highest ranked student. Note that a cycle must exist. For every cycle that exists, assign the student to the school to which she is pointing, remove the student, and reduce the capacity of the school by one.*

- **Step k :**

- ka) *If the school to which i was pointing in Round $k - 1$ still has available capacity, then i continues to point to the same school. For the students whose favorite school was removed in the previous round, iterate the clinching process until no student has one of the q_s highest priority at her most preferred school s unless she was pointing to s at the end of round $k - 1$.*
- kb) *Have each remaining student point to her most preferred school that has capacity greater than zero. Have each school with available capacity point to the highest priority student. A cycle must exist. For every cycle that exists, assign the student to the school to which she is pointing, remove the student, and reduce the capacity of the school by one.*

The algorithm terminates when all students are assigned a seat. Morrill (2015b) shows that Clinch and Trade is strategy-proof and efficient.

Another TTC variation is motivated by the fact that in TTC-Counters, the highest priority student has “ownership” over all slots of a school. That is, in TTC-Counters, there are at most $|S|$ participants in each round, and the market may be made thickened by considering all school seats. Hakimov and Kesten (2014) define **Equitable TTC**, which assigns all slots of school s to all the q_s students with the highest priority, giving one seat to each student and endowing them with equal trading power. The key part of this extension requires specifying who can point to whom. The mechanism works as follows:

- **Step 1:** *For each school, all available slots are assigned to students one by one, following their priority order, to form student-school pairs. 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 her best choice school, then all student-school pairs containing her point to that student-school pair. 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 either the same or in different cycles. In such a case, student i is placed in her best choice, and the other slots of that school (to which the student-school pairs containing her 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.*
- **Step k :**
 - ka) *Inheritance: For each school s , such that (i) there are slots at school s that remain to be inherited from previous steps, and (ii) there are no existing pairs that were assigned to a slot at school s in previous steps of the algorithm, the slots that remained to be inherited from previous steps are assigned to the remaining students one by one, following the priority order for school s to form new student-school pairs.*
 - kb) *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 her best choice school, then all student-school pairs containing her point to that student-school pair. 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 her best choice, and the other slots of that school (to which the student-school pairs containing her 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.

Hakimov and Kesten (2014) shows Equitable TTC is Pareto efficient and strategy-proof. They also show that if there are two schools, if Equitable TTC selects an allocation with justified envy for a problem, then TTC also selects an allocation with justified envy, but the converse is not necessarily true.

To understand the differences between these variations on TTC, consider the following example.

Example 2. [Variations on TTC] Suppose there are three students $I = \{i_1, i_2, i_3\}$ and two schools $S = \{s_1, s_2\}$, with $q_{s_1} = 2$ and $q_{s_2} = 1$. The preferences and priorities are given by

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

TTC produces:

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_1 \end{pmatrix},$$

where during TTC, i_2 forms a cycle with i_1 that assigns her to s_1 , even though she has high enough priority to be assigned to s_1 without being in a cycle with another student. TTC-Clinch and Trade as well as Equitable TTC avoid this cycle and produce the efficient and justified envy-free matching:

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_1 & s_2 \end{pmatrix}.$$

This example might suggest other possible mechanisms that have strictly less justified envy than TTC. However, TTC-Clinch and Trade and TTC-Counters are not generally comparable in terms of justified envy in general, since the set of blocking pairs are non-empty and disjoint for both mechanisms. Roughly speaking, by clinching, an applicant obtains her top choice, resulting in other applicants improving their priority. These priority improvements, in turn, generate new trading opportunities that can result in justified envy.

None of TTC-Counters, TTC-Clinch and Trade, or Equitable-TTC is constrained-optimal with multi-unit capacity, as we next show.

Proposition 2. Suppose there is a school with more than one seat. Then,

- i) TTC-Counters, TTC-Clinch and Trade, and Equitable TTC are not justified-envy minimal within the set of Pareto efficient and strategy-proof mechanisms.*
- ii) For each of the three mechanisms, there exists a justified envy minimal, Pareto efficient, and strategy-proof mechanism that has strictly less justified envy than that mechanism.*

The proof of the first part of this result constructs a Pareto efficient and strategy-proof mechanism φ that produces the same outcome as TTC-Counters for nearly all problems. We define it to produce a different outcome than TTC-Counters in a problem that always has a Pareto efficient and justified envy-free matching. For instance, coming back to Example 2, one can check that the priorities satisfy the conditions identified in Ergin (2002); for any set of preferences, there will be a matching that is Pareto efficient and justified envy-free. Hence, in this economy, the strategy-proof student-proposing deferred acceptance mechanism selects a matching satisfying those two properties. Recall, however, that TTC-Counters produces a matching that is not justified envy-free in this example. Hence, we can build a mechanism that selects the outcome of the student-proposing deferred acceptance mechanism when priorities are given as in Example 2 and otherwise coincides with TTC-Counters. This mechanism is strategy-proof and Pareto efficient and has strictly less justified envy than TTC-Counters. The full argument for TTC-Counters—as well as TTC-Clinch and Trade and Equitable TTC—is provided in Appendix.

The second part of this result follows from the finiteness of I and S : for fixed number of schools, students, and school capacities, the number of Pareto efficient and strategy-proof mechanisms is finite, since the number of priorities and preferences is also finite. Hence, starting from any mechanism in this class, one can iteratively find a mechanism improving in the justified envy (which is a transitive order) until one reaches a justified-envy minimal mechanism in finite steps. While the existence of justified-envy minimal mechanisms can be shown, no simple constructive algorithm has been found so far. Whether such an algorithm exists remains an open question.

These observations motivate the empirical analysis in Section 6, which evaluates the performance of three variations of TTC compared to serial dictatorship and deferred acceptance.

5.2 Weak Priorities

The school choice problem was defined with strict priorities. In practice, priorities are often defined via coarse categories in practice, such as having a sibling enrolled at the school or residing in the school’s geographic catchment area. In this section, we relax the assumption that priorities are strict. Let \succsim be a profile of weak priority relations, where each element is complete and transitive but need not be irreflexive or anti-symmetric.

An assignment mechanism φ selects a matching for each problem (P, \succsim) . For example, one possible mechanism is to convert the weak priorities into strict priorities using a tie-breaker and

then apply one of the mechanisms described in Section 2.2. Formally, a tie-breaker r_s at school s is a function $r_s : I \rightarrow \mathbb{N}$ that assigns each student a number. Under \succsim_s , when two students have equal priority the tie-breaker constructs \succ_s by giving higher priority to the student with the lower value of the tie-breaker. Let τ be a collection of tie-breakers for each school; that is, $\tau = (r_s)_{s \in S}$. TTC_τ is a mechanism that applies the tie-breaker τ to construct the strict priority order, and then applies TTC to the resulting problem. That is, given (P, \succsim) , TTC_τ first applies tie-breaker τ to \succsim to construct \succ^τ and then computes $TTC(P, \succ^\tau)$.

For any given τ , TTC_τ is Pareto efficient and strategy-proof. However, even when each school has one seat, it is possible to show that TTC_τ does not minimize justified envy among Pareto efficient and strategy-proof mechanisms.

Example 3 (TTC_τ does not minimize justified envy for problem (P, \succsim)). *There are three students $I = \{i_1, i_2, i_3\}$ and two schools $S = \{s_1, s_2\}$, where $q_{s_1} = q_{s_2} = 1$. School s_1 has strict priorities given by $i_1 \succ_{s_1} i_2 \succ_{s_1} i_3$, while s_2 ranks i_2 last and is indifferent between i_1 and i_3 . Consider the tie-breaker τ , where $r(\cdot)$ is the same for each school and is defined by $r(i_1) = 2$, $r(i_2) = 3$, and $r(i_3) = 1$. With this tie-breaker, \succ^τ is given by:*

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

Suppose that preferences are as follows:

P_1	P_2	P_3
s_2	s_1	s_1
s_1	s_2	s_2

TTC yields the following matching

$$\left(\begin{array}{ccc} i_1 & i_2 & i_3 \\ s_2 & i_2 & s_1 \end{array} \right),$$

where i_2 is unassigned. In problem (P, \succsim) , the pair (i_2, s_1) blocks this matching.

Consider student-proposing DA with some other tie-breaking rule $\hat{\tau}$, where each school's tie-breaker is given by $\hat{r}(\cdot)$ with $\hat{r}(1) = 1$, $\hat{r}(2) = 2$, and $\hat{r}(3) = 3$. The corresponding priorities $\succ^{\hat{\tau}}$ are given by:

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

It is easy to check that this satisfies the acyclicity condition in Ergin (2002). Therefore, DA applied to $(P, \succ^{\hat{\tau}})$ is Pareto efficient.

To show that TTC_τ does not minimize envy in the class of Pareto efficient and strategy-proof mechanisms for problems (P, \succsim) , we build φ which applies student-proposing DA with tie-breakers $\hat{\tau}$ applied to \succsim whenever the priorities are given as above and applies TTC with tie-breaker τ applied to \succsim for any other priority relation. The resulting mechanism φ is Pareto efficient and strategy-proof and has less justified envy than TTC_τ .

A potentially unusual aspect of this example is that the tie-breaking rule depends on which mechanism we use. It may be more natural to compare mechanisms holding fixed the way tie-breakers are selected by the mechanism. For a mechanism φ for problem (P, \succ) , we can define φ_τ for problem (P, \succsim) as the mechanism that applies tie-breaking τ to \succsim to construct \succ^τ and then computes the outcome of φ for (P, \succ^τ) . With this definition in hand, the following proposition is an immediate consequence of our main theorem:

Proposition 3. *Suppose that each school has one seat. Let τ be a tie-breaking rule and φ_τ be a Pareto efficient and strategy-proof mechanism. If φ_τ has less justified envy than TTC_τ , then $\varphi_\tau = TTC_\tau$*

Proof. Suppose that φ_τ has less justified envy than TTC_τ . Then there exists (P, \succsim) such that $\varphi_\tau(P, \succsim)$ has less justified envy than $TTC_\tau(P, \succsim)$. Consider (P, \succ^τ) , and observe

$$\varphi(P, \succ^\tau) = \varphi_\tau(P, \succsim) \quad \text{and} \quad TTC(P, \succ^\tau) = TTC_\tau(P, \succsim).$$

Therefore, $\varphi(P, \succ^\tau)$ has less justified envy than $TTC(P, \succ^\tau)$, which contradicts Theorem 1 which states that if φ has less justified envy than TTC , φ must equal TTC . ■

5.3 The Shapley-Scarf Housing Market Model

As noted above, Gale initially proposed the TTC algorithm for the housing market model of Shapley and Scarf (1974). In the housing market model, the number of individuals is equal to the number of objects or “houses.” Each individual has a complete strict ranking over houses and is endowed with a house. Roth and Postlewaite (1977) showed that the market’s unique core is found via Gale’s TTC algorithm, since in each cycle an individual points to the owner of his most preferred house among the remaining houses. (Recall that for priority-based resource allocation, a cycle consists of individuals pointing to objects and objects pointing to individuals.) For this model, an allocation is individually rational if every individual is assigned a house that she weakly prefers to her initial endowment. Roth (1982b) shows that TTC is strategy-proof in this model, and Ma (1994) shows that Gale’s TTC is the only Pareto efficient, strategy-proof, and individually rational mechanism for this problem.

If we re-interpret the Shapley-Scarf economy such that the agent endowed with a house is first in that house’s priority relation, while all other agents have the equivalent priority just below the owner, then we can consider this setting as equivalent to our model with weak priorities. Let \succsim be as described above. A pair (i, s) blocks a matching μ in economy (P, \succsim) if and only if

$sP_i\mu(i)$, where s is the unique school at which i has top priority, which means that individual rationality is violated. Thus, if a Pareto efficient and strategy-proof mechanism is individually rational in a Shapley-Scarf economy, it must be that this mechanism has less justified envy than *TTC* at \succsim . Theorem 1 therefore implies the following¹⁷

Corollary 1 (Ma (1994)). *In a Shapley-Scarf economy, a mechanism is Pareto efficient, strategy-proof, and individually rational if and only if it is TTC.*

6 Comparing Mechanisms in New Orleans and Boston

In this section, we use data from New Orleans and Boston to compare Pareto efficient and justified envy-free allocations in practice.

6.1 New Orleans

The New Orleans Recovery School District was formed in 2003 to facilitate state control of schools in New Orleans. The RSD’s role expanded considerably following Hurricane Katrina in 2005, and by 2011-12, the district came to oversee the majority of schools in New Orleans. In 2011-12, the RSD included a total of 73 schools: 13 district-run and 60 charter schools. In 2011-12, the district launched a pioneering unified enrollment process known as OneApp, in an effort to make the process easier to navigate and improve the efficiency of the previous decentralized application process.¹⁸

Early on, officials identified three major priority groups: students with siblings currently enrolled in a school, students currently enrolled in schools that are closing, and students who live in the school’s geographic area. More background on OneApp is presented in Vanacore (2011), Vanacore (2012), and Harris, Valant and Gross (2015). There was little concern that RSD schools would form blocking pairs and try to enroll pupils from other RSD schools, as might motivate a strategic rationale for eliminating blocking pairs. When deciding on the assignment mechanism, the RSD relied on the Boston experience, whose school priority structure was most similar to New Orleans. At the time, Boston had a sibling and walk zone priority (see, e.g., Abdulkadiroğlu et al. (2005a) and Pathak and Sönmez (2008)).¹⁹ New York City’s experience with a new centralized assignment system was also relevant because the superintendent and

¹⁷Theorem 1 is proved for strict priorities. However, it is easily checked that it extends to the weak priorities considered in this section. Indeed, the proof of the theorem goes through as long as, when a school s points to a student i at some step of TTC, $i \succ_s j$ for any individual j who remains at this step. Under the priority relations considered in this section, this property is satisfied.

¹⁸16 schools managed by the Orleans Parish Board, a second school district in New Orleans were not part of the OneApp process in its first year, nor were five charter schools authorized by the state’s Board of Elementary and Secondary Education (BESE).

¹⁹In 2013, Boston Public Schools eliminated walk zones, and the current system only has a sibling priority. See Dur, Kominers, Pathak and Sönmez (2016) for more details.

many other RSD staffers had worked in New York City’s Department of Education when it adopted a new mechanism.

The discussion centered on the trade-off between efficiency and eliminating justified envy, and there were concerns about how to explain TTC to the public. Officials eventually sided with efficiency and selected TTC-Counters based on the desire for “as many students as possible get into their top choice school” (RSD 2012). Other versions of TTC had not yet been systematically investigated at the time of the decision and so were not considered. To address the concerns about explaining the algorithm, staff were trained in the basics of how the algorithm worked, and promotional materials were released to the public to provide an overview of the algorithm before it was used (see, e.g., the newspaper descriptions in the New Orleans Times-Picayune referenced in Vanacore (2011) and Vanacore (2012)).²⁰

We use data from elementary and high school applicants in grade PK and 9 in the 2011-12 school year, and report on the performance of five mechanisms: TTC-Counters, TTC-Clinch and Trade, Equitable TTC, Serial Dictatorship, and DA. Applicants are asked to rank up to eight choices; in practice, less than 5% of applicants rank all eight choices, so it seems reasonable to run the alternative algorithms without taking any behavioral response into account since truth-telling is a weakly dominant strategy in each mechanism without a constraint.²¹ We draw 100 sets of lottery numbers, one for each child, and ran the assignment algorithms for each lottery draw. For the serial dictatorship, we order students according to the realized lottery order. Table 1 reports the average across lottery draws and two grades.

Our analysis shows that all three versions of TTC produce nearly identical aggregate rank distributions and similar amounts of justified envy. Roughly 65% of applicants are assigned to their top choice, and 19% are unassigned. Of those who obtain their first choice across lottery draws, the standard deviation is 5.2, so there is virtually no difference once we account for simulation uncertainty. Across the three versions of TTC, about 13% of students and 10% of schools are involved in a blocking pair. Here again, the differences between columns are much smaller than the standard deviation across lottery draws. Likewise, the differences in the count of blocking pairs and instances of justified envy are similar for each version of TTC.²²

²⁰School choice design comprises not only algorithms but also communication packages.

²¹Strictly speaking, when there is a constraint on the number of choices an applicant can rank, none of the mechanisms is strategy-proof. But for the 95% of applicants who ranked fewer than eight choices, they could have ranked an additional choice, and so there is little reason to strategize among the choices they submitted. Haeringer and Klijn (2009), Calsamiglia, Haeringer and Klijn (2010), and Pathak (2016) provide additional analysis and discussion of constrained rank order lists.

²²Note that this conclusion differs from some of the simulations performed by Hakimov and Kesten (2014) and Morrill (2015b), which show substantial differences across TTC variants. One way to interpret the difference is that the simulation environments (e.g., the distribution of preferences and priorities) considered by these authors do not approximate real-world markets such as New Orleans and Boston. Another possibility is that the difference across different TTC variants may diminish in large markets; for instance, Leshno and Lo (2017) reach a conclusion consistent with our empirical findings, showing that in a large market the TTC-Clinch and Trade may not perform better (and sometimes considerably worse) than the TTC-Counters in terms of justified envy.

Further, each of the three TTC mechanisms has less justified envy than Serial Dictatorship.²³ About 18% of students have justified envy, and five more schools trigger justified envy under Serial Dictatorship compared to TTC. These differences are larger than the corresponding standard deviation due to lottery draws. Moreover, the number of blocking pairs is roughly 40% greater under Serial Dictatorship than under TTC. This evidence suggests that the forces identified in Theorem 1 and Proposition 1 carry through (even) when schools can have more than one seat; namely, TTCs (of different variants) perform well in terms of “economizing” on justified envy when schools have more than one seat. It is worth noting that the aggregate rank distribution under the Serial Dictatorship is similar to that under TTC, indicating that, based on this metric, there may be little cost to minimizing blocking pairs.

In 2012-13, the RSD switched to DA, citing three main reasons. First, there were on-going efforts to encourage participation from a broader set of schools in New Orleans, including screened schools in the Orleans Parish School Board (OPSB). OPSB schools include those that audition applicants for specialized music, arts, and specialized athletic programs, in which eliminating blocking may be valuable for strategic reasons.²⁴ Second, TTC was perceived as difficult for participants to understand and for staff to explain why a child was assigned a seat over a child who was not. Under DA, officials could explain that an applicant did not obtain an assignment at a higher ranked seat because another applicant with higher priority was assigned to that seat. Such an explanation was not available for TTC. Finally, eligible families were allowed to apply for non-public schools participating in Louisiana’s Scholarship Program (LSP), a voucher program that allows low-income students in low-performing schools to apply for private schools (for more details on this program, see Abdulkadiroglu, Pathak and Walters (2016)). The LSP expanded statewide during Year 2 of OneApp due to a state law, Act 2. The Louisiana Department of Education interpreted blocking pairs created by TTC as being potentially subject to legal challenge under Act 2.

The last column of Table 1 suggests that the switch to DA may have had little impact on the overall aggregate rank distribution of choices received by applicants. Slightly fewer students are assigned to their top choice under DA. The difference between DA and TTC-Counters is 10 students, which is greater than the standard deviation of first choice under TTC-Counters, which was 5.2. Slightly more students are assigned their second choice, and DA has somewhat fewer students unassigned. The benefit of this small reduction in top choice assignment is that no student is involved in a blocking pair.

²³Note also that Serial Dictatorship entails similar aggregate rank distribution as each TTC mechanism. The finding that all four Pareto efficient mechanisms perform similarly in aggregate rank distribution is consistent with Che and Tercieux (2017), who establish approximate payoff equivalence of Pareto efficient mechanisms in large allocation problems.

²⁴Direct-run OPSB schools participated in OneApp in 2012-13, but OPSB charter schools and charter schools run directly by the state under the purview of OPSB did not. In December 2012, the OPSB passed an amendment requiring charter schools to participate in OneApp when their contracts are reauthorized. The RSD encouraged all OPSB charters to participate.

6.2 Boston

Boston Public Schools operates a school choice process wherein applicants can apply to any school in their zone for elementary and middle school and city-wide for high school. Like in New Orleans, students are prioritized based on sibling and walk zone status, so that students with siblings at a school who reside in the walk zone have highest priority, students with siblings at a school have the next highest priority, and students in the walk zone have the next highest priority. Within each priority, students are ordered using their random number.²⁵ We work with data from four years, from 2009-2010 through 2012-2013, when the district employed the student-proposing deferred acceptance algorithm.²⁶ We focus on main round applicants for the elementary (grade K1 and K2), middle (Grade 6), and high school (Grade 9) transition points. For each applicant, we observe the priority of the choices they ranked.²⁷ Since the mechanism is based on DA, applicants can rank as many schools as they wish, and BPS advises families to list schools in order of true preference (BPS 2012), we simulate the outcome of other mechanisms holding the submitted preferences fixed.

Table 2 shows that the patterns in Boston parallel those in New Orleans. Namely, there is little difference across the three versions of TTC in terms of the aggregate rank of choices assigned and amount of justified envy. Furthermore, TTC has significantly less justified envy than the Serial Dictatorship and a similar aggregate rank distribution. Finally, DA assigns slightly fewer students to their first choice and has a small number who are unassigned. This difference is far smaller than the difference in the number of students with justified envy under TTC. These facts suggest our results may be relevant for other cities with similar priority structures as New Orleans and Boston.

7 Conclusion

This paper provided a formal basis for recommending TTC over other Pareto efficient and strategy-proof mechanisms for priority-based resource allocation. Our main result is a counterpart to the constrained-optimality result for DA and shows that there is a formal justification for resolving the trade-off between Pareto efficiency and the elimination of justified envy as initially envisioned by Abdulkadiroğlu and Sönmez (2003).

Using data from New Orleans and Boston, we investigated three TTC variations for settings where schools may have more than one seat. The aggregate rank distribution and set of blocking

²⁵Appendix E of Dur et al. (2016) provides additional details on the priority structure, and describes how the mechanism applies walk zone priority for only one-half of school seats.

²⁶Boston continues to employ DA, but changed the configuration of its schools zones after our time-period (see, e.g., Dur et al. (2016) for more details).

²⁷To implement TTC, it is necessary to compute the priority of choices that are not ranked. We impute sibling priority at unranked schools using a list of schools that siblings attend and walk zone priority using information linking residential locations to the set of walk zone schools.

pairs are similar across these versions of TTC. Further, all TTC variations have significantly less justified envy than a serial dictatorship, thereby providing an empirical analogue to our main result.

In the field, there is growing momentum in favor of DA over TTC (see, e.g., Pathak and Sönmez (2013) and Pathak (2016)). This trend may be driven by a first-mover advantage of DA, since New York City and Boston adopted DA in 2003 and 2005, and its familiarity as the core of the National Residency Matching Program for U.S. doctors (Roth and Peranson 1999). Our paper shows that the aggregate rank distribution of choices assigned under DA and TTC are similar in Boston and New Orleans, which may further shift the balance in favor of DA for markets similar to these two areas.²⁸

Moreover, officials in New Orleans echoed a theme in Boston: TTC seems harder to explain and participate in compared to DA. It remains to be seen whether there will be another setting where TTC is used for school assignment in the field. Part of the challenge in advocating for TTC is explaining how it reflects priorities, a difficulty that may be easier to surmount with our formal result. It's possible that TTC would have been selected in other settings with knowledge of this result, and at the very least, advocates now have a new argument in favor of TTC.

²⁸New York City high school assignment exhibits a more discernible difference between TTC and DA in aggregate rank distribution of assigned choices (see Abdulkadiroğlu et al. (2009) and Che and Tercieux (2016b)), and the pattern of difference is consistent with a theoretical result of Che and Tercieux (2016b) for a large-market model. New York City school assignment is unique in size (it involves more than 80,000 applicants annually) and variety of programs (it includes more than 750 programs of differing in types, including screened, limited unscreened, and Ed-opt, etc.). The extent to which differences in market characteristics affect the performance of alternative matching mechanisms is worth further investigation.

Table 1. Comparison of Mechanisms in New Orleans for Main Transition Grades (PK and Grade 9)

	TTC-Counters (1)	TTC-Clinch and Trade (2)	Equitable TTC (3)	Serial Dictatorship (4)	Student- Proposing Deferred Acceptance (5)
A. Choice Assigned					
1	772	770	771	777	762
2	126	129	127	121	137
3	46	47	47	44	51
4	18	18	18	17	19
5+	11	11	11	8	10
Unassigned	222	221	222	228	217
Total	1196	1196	1196	1196	1196
B. Statistics on Blocking Pairs					
Students with justified envy	158	157	159	213	0
Schools involved in blocking pairs	7	7	7	12	0
Blocking pairs (i,s)	228	224	215	308	0
Instances of justified envy (i, (j,s))	1111	1086	1100	6546	0

Notes: The main transition grades are PK and Grade 9 for 2012. TTC-counters defined in Abdulkadiroglu and Sonmez (2003). TTC-Clinch and Trade defined in Morrill (2015). Equitable-TTC defined in Hakimov and Kesten (2014). Students with justified envy means there exists a school s where (i,s) is a blocking pair. School involved in blocking pairs means there is a school s such that there exists student i such that (i,s) is a blocking pair. Blocking pairs (i,s) means there exists at least one applicant j such that $(i,(j,s))$ is a blocking instance. Instance of justified envy $(i,(j,s))$ means student i complains about student j 's assignment at s . The numbers represent averages of 100 different lottery draws for each grade, which are then averaged over grades PK and 9. The standard deviation across lottery draws in column 1 for first choice assigned is 5.2, for unassigned is 5.2, for students with justified envy is 8.5, for schools involved in blocking pairs is 0.4, for blocking pairs is 20.2, and for instances of justified envy is 110.7. Standard deviations are similar for the other columns.

Table 2. Comparison of Mechanisms in Boston for Main Transition Grades (K1, K2, 6, and 9)

	TTC-Counters (1)	TTC-Clinch and Trade (2)	Equitable TTC (3)	Serial Dictatorship (4)	Student- Proposing Deferred Acceptance (5)
A. Choice Assigned					
1	1240	1240	1240	1236	1227
2	322	323	323	315	336
3	134	134	134	132	138
4	56	55	55	51	57
5+	39	39	53	34	40
Unassigned	102	101	101	124	96
Total	1893	1893	1893	1893	1893
B. Statistics on Blocking Pairs					
Students with justified envy	129	126	125	280	0
Schools involved in blocking pairs	18	18	18	44	0
Blocking pairs (i,s)	160	156	157	369	0
Instances of justified envy (i, (j,s))	768	711	696	3650	0

Notes: Data cover four school years from 2009-2010 through 2012-13. TTC-counters defined in Abdulkadiroglu and Sonmez (2003). TTC-Clinch and Trade defined in Morrill (2015). Equitable TTC defined in Hakimov and Kesten (2014). Instance of justified envy (i,(j,s)) means student i complains about student j's assignment at s. Blocking pair (i,s) means there exists at least one applicant j such that (i,(j,s)) is a blocking instance. Students with justified envy means there exists a school s where (i,s) is a blocking pair. Schools involved in blocking pairs means there is a school s such that there exists student i such that (i,s) is a blocking pair. BPS precedence implemented as Walk-Open. The numbers represent averages of 100 different lottery draws for each grade, which are then averaged over grades K1, K2, 6, and 9. The standard deviation across lottery draws in column 1 for first choice assigned is 6.8, for unassigned is 5.2, for students with justified envy is 16.9, for schools involved in blocking pairs is 2.4, for blocking pairs is 21.6, and for instances of justified envy is 98.6. Standard deviations are similar for the other columns.

A Proof of Theorem 1

Proof. Let φ be a Pareto efficient and strategy-proof mechanism that has less justified envy than TTC at the profile of priority relations \succ . To the contrary, suppose that there exists P such that

$$\varphi(P, \succ) \neq TTC(P, \succ).$$

The proof will proceed by induction on the steps of TTC . Let $I_k(P, \succ)$ be the set of individuals who are matched in step k of $TTC(P, \succ)$.

Claim 1: For all $i \in I_1(P, \succ)$,

$$\varphi(P, \succ)(i) = TTC(P, \succ)(i).$$

Proof. Suppose that for some $i \in I_1(P, \succ)$,

$$\varphi(P, \succ)(i) \neq TTC(P, \succ)(i).$$

Let $c = \{s_k, i_k\}_{k=1, \dots, K}$ be the cycle in which i is matched with $TTC(P, \succ)(i)$ and $i = i_K$. We will abuse notation and say that $i \in c$ for some i if $i = i_1, \dots, i_K$. Note that every i_k trades s_k for s_{k+1} , which is i_k 's first choice.

Consider the alternative preference relation:

$$P'_{i_K} : s_1, s_K, \dots$$

Since we've only altered the preferences of individuals in c and each obtains her first choice, the TTC matching remains the same:

$$TTC(P'_{i_K}, P_{-\{i_K\}}, \succ) = TTC(P, \succ).$$

Since s_1 is i_K 's first choice and

$$\varphi(P, \succ)(i_K) \neq \underbrace{TTC(P, \succ)(i_K)}_{=s_1},$$

we obtain

$$s_1 P_{i_K} \varphi(P, \succ)(i_K). \tag{1}$$

Also, since s_K points to i_K , by construction of TTC , we know that

$$TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) R'_{i_K} s_K.$$

Therefore, (i_K, s_K) does not block $TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)$.

Since φ has less justified envy than TTC at \succ , (i_K, s_K) also does not block $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)$. In addition, since s_K points to i_K in cycle c , for all $j \neq i_K$, we have

$$i_K \succ_{s_K} j$$

so that

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) \in \{s_1, s_K\}$$

must hold. Then equation (1) and strategy-proofness of φ imply

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) = s_K,$$

for otherwise i_K would be able to manipulate φ in economy (P, \succ) by submitting P'_{i_K} to obtain s_1 . As a result, i_K obtains her second choice under P'_{i_K} .

Next, consider i_{K-1} and the alternative preference

$$P'_{i_{K-1}} : s_K s_{K-1} \dots$$

As before, the TTC matching remains the same:

$$TTC(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ) = TTC(P'_{i_K}, P_{-\{i_K\}}, \succ) = TTC(P, \succ).$$

Since s_K is i_{K-1} 's first choice and

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}) \neq \underbrace{TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1})}_{=s_K},$$

we obtain

$$s_K P_{i_{K-1}} \varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}). \quad (2)$$

Also, since s_{K-1} points to i_{K-1} , by construction of TTC , we know that

$$TTC(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)(i_{K-1}) R'_{i_{K-1}} s_{K-1},$$

it must be the case that (i_{K-1}, s_{K-1}) does not block $TTC(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)$. Since φ has less justified envy than TTC at \succ , (i_{K-1}, s_{K-1}) does not block $\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)$.

In addition, since s_{K-1} points to i_{K-1} in cycle c , for all $j \neq i_{K-1}$, we have

$$i_{K-1} \succ_{s_{K-1}} j.$$

Thus, using the fact that (i_{K-1}, s_{K-1}) does not block $\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)$, as before, we obtain

$$\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)(i_{K-1}) R'_{i_{K-1}} s_{K-1},$$

so that

$$\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)(i_{K-1}) \in \{s_K, s_{K-1}\}.$$

By strategy-proofness of φ and by (2), it must be that

$$\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)(i_{K-1}) = s_{K-1}.$$

Then

$$\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)(i_{K-2}) \neq \underbrace{TTC(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1}, i_K\}}, \succ)}_{=s_{K-1}}(i_{K-2}).$$

Repeating the argument recursively for i_{K-2} , and for every individual in the cycle c , we obtain that

$$\varphi(P'_c, P_{-c}, \succ)(i_k) = s_k,$$

where $P'_c = \{P'_{i_k}\}_{i_k \in c}$ and $P'_{i_k} : s_{k+1}, s_k, \dots$

To see this, note that at the last step of the recursive argument, we know that

$$\varphi(P'_c, P_{-c}, \succ)(i_1) = s_1.$$

This implies that

$$\varphi(P'_c, P_{-c}, \succ)(i_K) \in \{s_K, \emptyset\}.$$

Now since

$$TTC(P'_c, P_{-c}, \succ)(i_K) R'_{i_K} s_K,$$

we know that (i_K, s_K) does not block $TTC(P'_c, P_{-c}, \succ)$ and because φ has less justified envy than TTC, it does not block $\varphi(P'_c, P_{-c}, \succ)$. Since, by definition of TTC, for all $j \neq i$,

$$i_K \succ_{s_K} j,$$

we must have

$$\varphi(P'_c, P_{-c}, \succ)(i_K) R'_{i_K} s_K.$$

Therefore,

$$\varphi(P'_c, P_{-c}, \succ)(i_K) = s_K.$$

Applying the argument recursively, we get that $\varphi(P'_c, P_{-c}, \succ)(i_k) = s_k$ for all k . But this contradicts Pareto-efficiency of φ because every individual in the cycle will be better off if every i_k is matched with s_{k+1} without changing the matching of individuals in $I - c$.

Therefore, we have established Claim 1: for all $i \in I_1(P, \succ)$,

$$\varphi(P, \succ)(i) = TTC(P, \succ)(i).$$

□

Now, assume that

$$\varphi(P, \succ)(i) = TTC(P, \succ)(i)$$

for all i that is matched in the first $m - 1$ steps of TTC , i.e. $i \in \bigcup_{l=1, \dots, m-1} I_l(P, \succ)$.

Claim 2. For all $i \in I_m(P, \succ)$,

$$\varphi(P, \succ)(i) = TTC(P, \succ)(i).$$

Proof. Suppose to the contrary that for some $i \in I_m(P, \succ)$,

$$\varphi(P, \succ)(i) \neq TTC(P, \succ)(i).$$

Let $c = \{s_k, i_k\}_{k=1, \dots, K}$ be the cycle in which i is matched with $TTC(P, \succ)(i)$ in step m of $TTC(P, \succ)$ and $i = i_K$.

Note that every i_k trades s_k for s_{k+1} which is her first choice among remaining objects at step m of $TTC(P, \succ)$. Consider the alternative preference relation:

$$P'_{i_K} : s_1, s_K, \dots$$

By construction, the TTC matching remains the same, i.e.

$$TTC(P'_{i_K}, P_{-\{i_K\}}, \succ) = TTC(P, \succ).$$

Since s_1 is i_K 's first choice among remaining objects at step m of $TTC(P, \succ)$; that is,

$$\varphi(P, \succ)(i_K) \neq TTC(P, \succ)(i_K) = s_1.$$

By our inductive hypothesis, under $\varphi(P, \succ)$, all schools s assigned before step m of $TTC(P, \succ)$ have been assigned to

$$\varphi(P, \succ)(s) = TTC(P, \succ)(s).$$

That is, all schools assigned before step m are assigned to a student not in $I_m(P, \succ)$. Therefore, i_K does not receive any such schools s under $\varphi(P, \succ)$, i.e.

$$\varphi(P, \succ)(i_K) \neq s$$

for all schools s assigned before step m of $TTC(P, \succ)$. This yields

$$s_1 P_{i_K} \varphi(P, \succ)(i_K). \tag{3}$$

Also, by construction of TTC , we know that

$$TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) R'_{i_K} s_K,$$

so it must be the case that (i_K, s_K) does not block $TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)$. Since φ has less justified envy than TTC at \succ , (i_K, s_K) does not block $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)$.

In addition, by construction of TTC , we know that for all j that are not yet eliminated at step m of TTC ,

$$i_K \succ_{s_K} j.$$

Note that

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(s_K)$$

is not eliminated at step m of TTC . To see this, note that if $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(s_K)$ had been eliminated in TTC before step m , we would have by the inductive hypothesis

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(s_K) = TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)(s_K) = i_{K-1},$$

meaning that i_{K-1} would have been eliminated before step m in $TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)$, a contradiction.

Thus, since $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(s_K)$ is not eliminated at step m of TTC and $i_K \succ_{s_K} j$ for all j that are not eliminated yet at step m of TTC , we obtain that either

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(s_K) = i_K$$

or

$$i_K \succ_{s_K} \varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(s_K).$$

Thus, because (i_K, s_K) does not block $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)$, in each of these cases, we must have that

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) R'_{i_K} s_K,$$

so that

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) \in \{s_1, s_K\}.$$

By strategy-proofness of φ and (3), it must be that

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) = s_K.$$

Then,

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}) \neq \underbrace{TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1})}_{=s_K}.$$

As above, we repeat these arguments to arrive a contradiction with Pareto-efficiency of φ , which completes the proof of the claim. \square

Claim 2 shows the inductive step. As a result, if φ has less justified envy than TTC at \succ , then $\varphi = TTC$, completing the proof. \blacksquare

B Proof of Proposition 2

The example below shows that TTC-Counters, TTC-Clinch and Trade, and Equitable-TTC do not minimize justified envy in the class of Pareto efficient and strategy-proof mechanisms, when schools have multiple seats.

Consider TTC-Counters (as will become clear, the argument is similar for TTC-Clinch and Trade or Equitable-TTC). We build a Pareto efficient and strategy-proof mechanism φ that allows for strictly less justified envy than TTC-Counters. Suppose φ selects the same matching

as TTC-Counters except for the following instance: there are four individuals i_1, i_2, i_3 , and i_4 and three schools s_1, s_2, s_3 , where $q_{s_1} = 1$ and $q_{s_2} = q_{s_3} = 3$. In this instance, φ selects a matching that is free of justified envy and Pareto efficient, which can be computed by student-proposing DA.

Ergin (2002) characterizes priorities under which, for any individual preference, there exists a matching that is Pareto efficient and justified envy-free. Priorities must be such that there is no Ergin-cycle. A profile \succ has an Ergin-cycle if there are three individuals i_1, i_2 , and i_3 and two schools s_1 and s_2 such that the two conditions are satisfied:

1. *Loop condition.* $i_1 \succ_{s_1} i_2 \succ_{s_1} i_3$ and $i_3 \succ_{s_2} i_1$,

2. *Scarcity condition.* There are (possibly empty) disjoint sets N_{s_1} and $N_{s_2} \subseteq I \setminus \{i_1, i_2, i_3\}$ s.t. $N_{s_1} \subseteq U_{s_1}(i_2)$ and $N_{s_2} \subseteq U_{s_2}(i_1)$ and $|N_{s_1}| = q_{s_1} - 1$ and $|N_{s_2}| = q_{s_2} - 1$ where $U_{s_1}(i_2)$ and $U_{s_2}(i_1)$ are the strict upper contour set of i_2 and i_1 , respectively (i.e., $U_{s_1}(i_2) := \{\ell : \ell \succ_{s_1} i_2\}$ and $U_{s_2}(i_1) := \{\ell : \ell \succ_{s_2} i_1\}$).

In the instance described above, the scarcity condition in the definition of an Ergin-cycle can never be satisfied. To see this, observe that for a school $s \in \{s_2, s_3\}$,

$$|N_s| = 3 - 1 = 2,$$

while $N_s \subseteq I \setminus \{i_1, i_2, i_3\}$ implies that

$$|N_s| \leq 1,$$

since $|I| = 4$. Therefore, sets satisfying the scarcity condition do not exist. Hence, any profile of priority relations is Ergin-acyclic.

Finally, it is enough for our purpose to build some (P, \succ) where the set of blocking pairs of φ is a proper subset of the set of blocking pairs of TTC-Counters. Since φ eliminates justified envy, we only need to show that there is (P, \succ) under which TTC-Counters does not eliminate justified envy.

Consider the following profile of preferences and priority relations:

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

TTC-Counters produces:

$$\begin{pmatrix} i_1 & i_2 & i_3 & i_4 \\ s_3 & s_3 & s_1 & s_2 \end{pmatrix},$$

where (i_2, s_1) is a blocking pair.

This completes the argument for TTC-Counters. It is easy to check that TTC-Counters, TTC-Clinch and Trade, and Equitable TTC coincide to produce the same assignment for the

above profile of preferences and priority relations. Hence, the same argument can be used for TTC-Clinch and Trade and Equitable TTC.

C Another justified envy-minimal mechanism

When each school has a single seat, we build a mechanism different from TTC that is strategy-proof, Pareto-efficient, and justified envy-minimal. The mechanism is identical to TTC except at the following instance of priorities: we have three students i_1, i_2 and i_3 and two schools s_1 and s_2 each with a single seat. Priorities are given by

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

In essence, the mechanism will rank i_3 on top of each school's ranking and run standard TTC on these modified priorities (except for some preference profiles where there is a unique efficient and stable allocation where the original priorities will still be used to run TTC). This mechanism will be denoted TTC^* . Let us describe it precisely. For the instance of priorities described above and for each profile of preferences P , TTC^* selects a matching as follows.

Case A.

If under P there is an individual who ranks all schools as unacceptable then run TTC.

Case B.1.

If under P both i_1 and i_2 rank s_1 first and i_1 finds s_2 unacceptable then run TTC

Case B.2.

If under P both i_1 and i_2 rank s_2 first and i_2 finds s_1 unacceptable then run TTC

Case C.

If none of the above cases apply, move i_3 to the top of each school's ranking. Run TTC on the modified priorities.

Clearly, TTC^* is Pareto efficient. We prove below that it is strategy-proof.

Proposition 4. *TTC* is strategy-proof.*

Proof. Fix P falling into case A. If some student i deviates to P'_i , this cannot be profitable if we remain into case A or fall into Case B.1 or B.2 (since TTC is strategy-proof). Therefore, consider the case where we fall into Case C after i 's deviation. After the deviation, all individuals rank at least one school acceptable (since we are not in Case A anymore). Since at least one individual must rank all schools unacceptable before the deviation and since we are looking at

a single deviation by individual i , we conclude that P_i ranks all schools unacceptable. Hence, under TTC^* , i is unmatched under P and since P_i ranks all schools unacceptable, there cannot be any profitable deviation.

Fix P falling into case B.1 (and not in case A). If some student i deviates to P'_i , this cannot be profitable if we remain into case B.1 or fall into Case A or B.2 (since TTC is strategy-proof). Therefore, consider the case where we fall into Case C after i 's deviation. This must mean that before deviation both i_1 and i_2 rank s_1 first and i_1 finds s_2 unacceptable, though this is not the case anymore after deviation. Note that this must mean that i is either i_1 or i_2 . Further, since P falls into Case B.1, TTC^* runs standard TTC . Hence, i_2 gets matched to her top choice s_1 and so i_2 has no incentive to deviate (recall that each individual finds at least one school acceptable since we are not in Case A). Hence, let us consider $i = i_1$. The only way to reach (by a deviation of i_1) Case C is for i_1 to claim that s_2 is acceptable (while s_2 is not acceptable to i_1 under the original preferences P_i). Now, to complete the argument, we distinguish two cases. First, assume that i_3 ranks s_2 first. Then, since we fall into Case C after deviation, $\text{TTC}^*(P'_i, P_{-i})$ is given by

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ i_1 & s_1 & s_2 \end{pmatrix}.$$

In particular, i_1 cannot get s_1 (the only acceptable school under P_i) so the deviation to P'_i cannot be profitable. Similarly, in the other case where i_3 ranks s_1 first, $\text{TTC}^*(P'_i, P_{-i})$ is given by

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & i_2 & s_1 \end{pmatrix}.$$

Here again, i_1 fails to obtain s_1 and so the deviation is not profitable. The case under which P falls into Case B.2 can be treated in the same way.

Fix P falling into case C. If some student i deviates to P'_i , this cannot be profitable if we remain into case C (since TTC is strategy-proof). Therefore, consider the case where we fall into Case A. This must mean that P'_i ranks all schools unacceptable. Since TTC^* is individually rational, i 's deviation cannot be profitable. So consider the case where after deviation we fall into Case B.1 (and not into Case A). This must mean that before deviation, either i_1 or i_2 does not rank s_1 first, or i_1 finds s_2 acceptable while after deviation both i_1 and i_2 rank s_1 first and i_1 finds s_2 unacceptable. This must mean that i is either i_1 or i_2 . If $i = i_1$, this means that, i_2 ranks s_1 first (recall that before and after deviation, each individual has at least one acceptable school since we do not fall into Case A before and after deviation). So, in particular, after deviation, i_1 cannot get s_1 (s_1 will be allocated to i_2). Since after deviation, i_1 finds s_2 unacceptable, i_1 will not get s_2 either, and so she will remain unmatched. So the deviation cannot be profitable to i_1 .

Now, consider the other case where deviator $i = i_2$. This means that i_1 ranks s_1 first and ranks s_2 unacceptable. This also means that before deviation, i_2 ranks s_2 first while after deviation i_2 ranks s_1 first. To complete the proof, we distinguish two cases. First, assume

that i_3 ranks s_2 first. Then, at P , i_2 gets s_1 if she finds s_1 acceptable or remains unmatched. After deviation, i_2 ranks s_1 on top and i_2 must be getting s_1 after deviation, so this cannot be profitable. In the other case where i_3 ranks s_1 first, before deviation, i_2 is getting s_2 , which is her top choice. So the deviation cannot improve on this. The same reasoning holds if the deviation falls into Case B.2. ■

We fix any Pareto efficient and strategy-proof mechanism φ with less justified envy than TTC^* . We claim that $\varphi = TTC^*$.

Proposition 5. *Fix any P that falls into Case A, B.1 or B.2. $\varphi(P) = TTC^*(P)$.*

Proof. Fix any P falling into Case A. Some individual must rank all schools as unacceptable. It is easy to check that, in such a case, there is a unique efficient and stable allocation that is selected by TTC (with only two students, priorities are trivially Ergin-acyclic). Hence, because φ has less justified envy than TTC^* , φ must also select the unique efficient and stable allocation, and we obtain $\varphi(P) = TTC^*(P)$.

Now, fix any P falling into Case B.1. Both i_1 and i_2 rank s_1 first, and i_1 finds s_2 unacceptable. Here again, one can check that TTC selects the unique efficient and stable allocation, and we obtain $\varphi(P) = TTC^*(P)$. A similar reasoning holds for any P falling into Case B.2. ■

Proposition 6. *Fix any P that falls into Case C. $\varphi(P) = TTC^*(P)$.*

Proof. We assume that P falls into Case C and prove the above proposition in the four following claims.

Claim 1. Assume that $s_1 P_{i_1} s_2$ and $s_2 P_{i_2} s_1$.

$$\varphi(P) = TTC^*(P).$$

Proof. Clearly, under TTC^* , i_3 is never part of any blocking pair. Hence, because φ has less justified envy than TTC^* , we must have that i_3 is never part of any blocking pair under φ as well. Assume wlog that s_1 is i_3 's top choice (recall that because P falls into Case C, each individual finds at least one school acceptable).

In the sequel, we claim that i_3 is assigned its top choice s_1 under matching $\varphi(P)$. If i_3 is not assigned its top choice s_1 under φ , then in order to ensure that (i_3, s_1) does not block $\varphi(P)$, we must have that i_2 is matched to s_1 under $\varphi(P)$. Now, consider two cases. First, i_3 is matched to s_2 under $\varphi(P)$. In that case, i_2 and i_3 would be better off switching their assignments, a contradiction with Pareto efficiency of φ . In the other case, i_3 must be unmatched under $\varphi(P)$. If i_1 gets matched to s_2 under $\varphi(P)$, allowing i_2 and i_1 to switch their assignments would be beneficial to both of them, again a contradiction with Pareto efficiency of φ . If i_1 is not matched to s_2 under $\varphi(P)$ then s_2 is unmatched, and by assigning it to i_2 we Pareto-improve on $\varphi(P)$, a contradiction.

Thus, we proved that $\varphi(P)(i_3) = s_1 = \text{TTC}^*(P)(i_3)$. Now, let us complete the argument and show that $\varphi(P) = \text{TTC}^*(P)$. First, consider the case where i_1 finds s_2 acceptable. TTC^* yields the following matching

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & i_2 & s_1 \end{pmatrix}.$$

If under $\varphi(P)$, i_1 remains unmatched, then (i_1, s_2) would block $\varphi(P)$ while it does not block $\text{TTC}^*(P)$, a contradiction with our assumption that φ has less justified envy than TTC^* . Thus, $\varphi(P)(i_1) = s_2 = \text{TTC}^*(P)(i_1)$ and so we conclude that $\varphi(P) = \text{TTC}^*(P)$. Now, consider the second case where i_1 finds s_2 unacceptable. Recall that s_2 must be acceptable to i_2 , and so TTC^* yields the following matching

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ i_1 & s_2 & s_1 \end{pmatrix}.$$

Clearly, since we showed that i_3 gets matched to s_1 under $\varphi(P)$, i_1 remains unmatched under $\varphi(P)$. So by Pareto efficiency of φ , it must be that $\varphi(P)(i_2) = s_2 = \text{TTC}^*(P)(i_2)$. We conclude that $\varphi(P) = \text{TTC}^*(P)$. \square

Claim 2. Assume that $s_1 P_{i_1} s_2$ and $s_1 P_{i_2} s_2$.

$$\varphi(P) = \text{TTC}^*(P).$$

Proof. There are two cases.

Case 1. $s_1 P_{i_3} s_2$. Because P falls into Case C, each individual ranks at least one school acceptable (since P does not fall into Case A) and i_1 finds s_2 acceptable (since P does not fall into Case B.1). Thus, TTC^* yields the following matching

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & i_2 & s_1 \end{pmatrix}.$$

We first claim that under $\varphi(P)$, i_2 must remain unmatched. Indeed, if i_2 is matched under $\varphi(P)$, then consider the new preference profile where $s_2 P'_{i_2} s_1 P'_{i_2} i_2$. Note that $(P'_{i_2}, P_{-\{i_2\}})$ falls into the cases considered in Claim 1. Hence, by Claim 1, we know that $\varphi(P'_{i_2}, P_{-\{i_2\}})(i_3) = \text{TTC}^*(P'_{i_2}, P_{-\{i_2\}})(i_3) = s_1$ and $\varphi(P'_{i_2}, P_{-\{i_2\}})(i_2) = \text{TTC}^*(P'_{i_2}, P_{-\{i_2\}})(i_2) = i_2$. Thus, from profile $(P'_{i_2}, P_{-\{i_2\}})$, i_2 can misreport her preference profile as P_{i_2} . In turn, she gets matched and is strictly better-off, which contradicts the strategy-proofness of φ . Hence, under $\varphi(P)$, i_2 must be unmatched. Next, we claim that i_3 is assigned s_1 under $\varphi(P)$. Indeed, if i_3 is not assigned s_1 under $\varphi(P)$, then i_1 must be assigned s_1 since it is acceptable to her (and we already know that i_2 must be unmatched). But then (i_3, s_1) would block $\varphi(P)$ but does not block $\text{TTC}^*(P)$, which contradicts our assumption that φ has less justified envy than TTC^* . To conclude, under $\varphi(P)$, i_3 gets s_1 , i_2 is unmatched, and so, since s_2 is acceptable to i_1 , i_1 gets matched to s_2 . Thus, $\varphi(P) = \text{TTC}^*(P)$.

Case 2. $s_2 P_{i_3} s_1$. TTC* yields the following matching

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ i_1 & s_1 & s_2 \end{pmatrix}.$$

We first claim that $\varphi(P)$ matches i_3 to her top choice s_2 . Indeed, if i_3 is not matched to s_2 under $\varphi(P)$ then in order for (i_3, s_2) not to block $\varphi(P)$, i_1 must match to s_2 . But then, in order for (i_2, s_1) not to block $\varphi(P)$, which is necessary since it does not block $\text{TTC}^*(P)$, i_2 must also match s_1 . So if i_3 is not matched to s_2 under $\varphi(P)$ the only candidate for $\varphi(P)$ is

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & i_3 \end{pmatrix}.$$

Now, let us assume that i_3 ranks s_1 acceptable under P_{i_3} . Next, consider the preference profile $(P'_{i_3}, P_{-\{i_3\}})$ where $s_1 P'_{i_3} s_2 P'_{i_3} i_3$. Note that $(P'_{i_3}, P_{-\{i_3\}})$ falls into Case 1 considered just above. Hence, we know that $\varphi(P'_{i_3}, P_{-\{i_3\}}) = \text{TTC}^*(P'_{i_3}, P_{-\{i_3\}})$ and so i_3 is matched to s_1 under $\varphi(P'_{i_3}, P_{-\{i_3\}})$. Since i_3 is unmatched under $\varphi(P)$, because we assumed that i_3 ranks s_1 acceptable under P_{i_3} , we found a profitable deviation for i_3 , a contradiction with the strategy-proofness of φ . Thus, provided that P_{i_3} ranks s_1 as acceptable, we obtained $\varphi(P_{i_3}, P_{-\{i_3\}})(i_3) = \text{TTC}^*(P_{i_3}, P_{-\{i_3\}})(i_3) = s_2$.

Let us now assume that i_3 ranks s_1 unacceptable under P_{i_3} . Consider a deviation of i_3 to P'_{i_3} satisfying $s_2 P'_{i_3} s_1 P'_{i_3} i_3$, i.e., where s_1 is ranked as acceptable. We just saw that, in such a case, $\varphi(P'_{i_3}, P_{-\{i_3\}})(i_3) = \text{TTC}^*(P'_{i_3}, P_{-\{i_3\}})(i_3) = s_2$ and so i_3 gets matched to s_2 under $\varphi(P'_{i_3}, P_{-\{i_3\}})$. Here again, we find a profitable deviation for i_3 , which contradicts the strategy-proofness of φ .

We conclude that $\varphi(P)$ matches i_3 to her top choice s_2 . Now, $\text{TTC}^*(P)$ matches i_2 with s_1 , and, in order not have the blocking pair (i_2, s_1) under $\varphi(P)$, i_2 and s_1 must also be matched together under $\varphi(P)$. We conclude that $\varphi(P) = \text{TTC}^*(P)$. \square

Claim 3. Assume that $s_2 P_{i_1} s_1$ and $s_2 P_{i_2} s_1$.

$$\varphi(P) = \text{TTC}^*(P).$$

Proof. The proof is similar to that of Claim 2. \square

Claim 4. Assume that $s_2 P_{i_1} s_1$ and $s_1 P_{i_2} s_2$.

$$\varphi(P) = \text{TTC}^*(P).$$

Proof. Without loss of generality, assume that $s_2 P_{i_3} s_1$ (the same argument applies when $s_1 P_{i_3} s_2$). TTC* yields the following matching

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ i_1 & s_1 & s_2 \end{pmatrix}.$$

If, under $\varphi(P)$, i_3 is not matched to her top choice s_2 , then in order for (i_3, s_2) not to block $\varphi(P)$ (which is necessary, since it does not block $\text{TTC}^*(P)$), i_1 must be matched to s_2 under $\varphi(P)$. But then for (i_2, s_1) not to block $\varphi(P)$ (which is necessary since it does not block $\text{TTC}^*(P)$), i_2 must be matched to s_1 . Thus, if i_3 is not matched to her top choice s_2 , the only candidate for $\varphi(P)$ is

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & i_3 \end{pmatrix}.$$

Now, consider P'_{i_1} such that $s_1 P'_{i_1} s_2 P'_{i_1} i_1$. Since i_1 ranks s_2 acceptable under P'_{i_1} , $(P'_{i_1}, P_{-\{i_1\}})$ falls in to the profile of preferences considered in Claim 2. Hence, $\varphi(P'_{i_1}, P_{-\{i_1\}})$ and $\text{TTC}^*(P'_{i_1}, P_{-\{i_1\}})$ both yield the same matching given by

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ i_1 & s_1 & s_2 \end{pmatrix}.$$

Now, if the true preference profile is $(P'_{i_1}, P_{-\{i_1\}})$ and i_1 misreports to P_{i_1} , then i_1 gets matched to s_2 under $\varphi(P_{i_1}, P_{-\{i_1\}})$. Hence, the misreport P_{i_1} is profitable to i_1 , which contradicts the strategy-proofness of φ .

We conclude that under $\varphi(P)$, i_3 must be matched to her top choice s_2 . But now, if i_2 is not matched to s_1 under $\varphi(P)$ then (i_2, s_1) blocks $\varphi(P)$ but does not block $\text{TTC}^*(P)$, which is a contradiction. Hence, i_2 must be matched to s_1 , and we conclude that $\varphi(P) = \text{TTC}^*(P)$. \square

These four claims together establish the proposition. \blacksquare

References

- Abdulkadiroglu, Atila and Yeon-Koo Che**, “The Role of Priorities in Assigning Indivisible Objects: A Characterization of Top Trading Cycles,” 2010. Working Paper, Columbia University.
- , **Parag Pathak, and Christopher Walters**, “Free to Choose: Can School Choice Reduce Student Achievement?,” 2016. Forthcoming, *American Economic Journal: Applied Economics*.
- Abdulkadiroglu, Atila and Tayfun Sönmez**, “Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems,” *Econometrica*, 1998, *66*(3), 689–701.
- and — , “School Choice: A Mechanism Design Approach,” *American Economic Review*, 2003, *93*, 729–747.
- , **Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez**, “The Boston Public School Match,” *American Economic Review, Papers and Proceedings*, 2005, *95*, 368–371.
- , — , and — , “The New York City High School Match,” *American Economic Review, Papers and Proceedings*, 2005, *95*, 364–367.
- , — , and — , “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the New York City High School Match,” *American Economic Review*, 2009, *99*(5), 1954–1978.
- Alcalde, José and Salvador Barberá**, “Top dominance and the possibility of strategy-proof stable solutions to matching problems,” *Economic Theory*, 1994, *4*, 417–435.
- Anderson, Ross, Itai Ashlagi, David Gamarnik, and Alvin E. Roth**, “Finding long chains in kidney exchange using the traveling salesmen problem,” *Proceedings of the National Academy of Sciences of the United States of America*, 2015, *112*(3), 663–668.
- Andersson, Tommy, Lars Ehlers, and Lars-Gunnar Svensson**, “Budget-Balance, Fairness and Minimal Manipulability,” *Theoretical Economics*, 2014, *9*, 753–777.
- Arribillaga, R. Pablo and Jordi Massó**, “Comparing generalized median voter schemes according to their manipulability,” *Theoretical Economics*, 2016, pp. 547–586.
- and — , “Comparing Voting by Committees According to their Manipulability,” 2016. , *American Economic Journal: Microeconomics*.
- Bade, Sophie**, “Random Serial Dictatorship: The One and Only,” *Working Paper, Department of Economics, Royal Holloway College*, 2016.

- Balinski, Michel and Tayfun Sönmez**, “A Tale of Two Mechanisms: Student Placement,” *Journal of Economic Theory*, 1999, *84*, 73–94.
- BPS**, “Recommendation to Implement a New BPS Assignment Algorithm,” 2005. Presentation to the Boston School Committee by Carleton Jones, May 11.
- , “Introducing the Boston Public Schools 2012,” 2012. Boston Public Schools, Available at: <http://www.bostonpublicschools.org/register>.
- Calsamiglia, C., Guillaume Haeringer, and Flip Klijn**, “Constrained School Choice: An Experimental Study,” *American Economic Review*, 2010, *100*, 1860–1874.
- Che, Yeon-Koo and Olivier Tercieux**, “An Analysis of Top Trading Cycles in Two-Sided Matching Markets,” 2016. Columbia University and PSE, Unpublished mimeo.
- and —, “Efficiency and Stability in Large Matching Markets,” 2016. Columbia University and PSE, Unpublished mimeo.
- and —, “Payoff Equivalence of Efficient Mechanisms in Large Markets,” 2017. Columbia University and PSE, Unpublished mimeo.
- Chen, Yan and Onur Kesten**, “Chinese College Admissions and School Choice Reforms: A Theoretical Analysis,” 2016. Forthcoming, *Journal of Political Economy*.
- Dubins, Lester E. and David A. Freedman**, “Machiavelli and the Gale-Shapley algorithm,” *American Mathematical Monthly*, 1981, *88*, 485–494.
- Dur, Umut**, “A Characterization of the Top Trading Cycles Mechanism for the School Choice Problem,” 2013. Working paper, NC State.
- , **Scott Kominers, Parag Pathak, and Tayfun Sönmez**, “Reserve Design: Unintended Consequences and The Demise of Walk Zones in Boston,” 2016. Working paper, MIT.
- Ergin, Haluk**, “Efficient Resource Allocation on the Basis of Priorities,” *Econometrica*, 2002, *70*, 2489–2498.
- Gale, David and Lloyd S. Shapley**, “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 1962, *69*, 9–15.
- Haeringer, Guillaume and Flip Klijn**, “Constrained School Choice,” *Journal of Economic Theory*, 2009, *144*, 1921–1947.
- Hakimov, Rustamdjan and Onur Kesten**, “The Equitable Top Trading Cycles Mechanism for School Choice,” 2014. WZB Discussion Paper 2014-210.

- Harris, Douglas, Jon Valant, and Betheny Gross**, “The New Orleans OneApp,” *Education Next*, 2015, 15(4).
- Kesten, Onur**, “On two competing mechanisms for priority-based allocation problems,” *Journal of Economic Theory*, 2006, 127, 155–171.
- , “School Choice with Consent,” *Quarterly Journal of Economics*, 2010, 125(3), 1297–1348.
- Kojima, Fuhito and Mihai Manea**, “Axioms for Deferred Acceptance,” *Econometrica*, 2010, 78, 633–653.
- Landsmark, Ted, Helen Dajer, and Lisa Gonsalves**, “Report and Recommendations of the Boston Public Schools Student Assignment Task Force,” 2014. Report to the Boston School Committee, September 22.
- Leshno, Jacob and Irene Lo**, “The Simple Structure of Top Trading Cycles in School Choice: A Continuum Model,” 2017. Columbia University.
- Ma, Jinpeng**, “Strategy-Proofness and the Strict Core in a Market with Indivisibilities,” *International Journal of Game Theory*, 1994, 23, 75–83.
- Maskin, Eric**, “Elections and Strategic Voting: Condorcet and Borda,” 2016. NBER Conference on Mathematical Economics, Johns Hopkins, <http://www.econ2.jhu.edu/conferences/mathecon2016/papers/Maskin.pdf>.
- Morrill, Thayer**, “Making Just School Assignments,” *Games and Economic Behavior*, 2015, 92, 18–27.
- , “Two Simple Variations of Top Trading Cycles,” *Economic Theory*, 2015, pp. 123–140.
- Pápai, Szilazi**, “Strategyproof Assignment by Hierarchical Exchange,” *Econometrica*, 2000, 68, 1403–1433.
- Pathak, Parag A.**, “The Mechanism Design Approach to Student Assignment,” *Annual Reviews*, 2011, 3, 513–536.
- , “What Really Matters in Designing School Choice Mechanisms,” 2016. Advances in Economics and Econometrics, 11th World Congress of the Econometric Society, eds. Larry Samuelson.
- **and Tayfun Sönmez**, “Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism,” *American Economic Review*, 2008, 98(4), 1636–1652.
- **and** —, “School Admissions Reform in Chicago and England: Comparing Mechanisms by their Vulnerability to Manipulation,” *American Economic Review*, 2013, 103(1), 80–106.

- Pycia, Marek and Utku Ünver**, “Trading Cycles for School Choice,” 2011. Working paper.
- and — , “Incentive Compatible Allocation and Exchange of Discrete Resources,” 2016. Forthcoming, *Theoretical Economics*.
- Roth, Alvin E.**, “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research*, 1982, 7, 617–628.
- , “Incentive Compatibility in a Market with Indivisible Goods,” *Economics Letters*, 1982, 9(2), 127–132.
- , “The College Admission Problem is not Equivalent to the Marriage Problem,” *Journal of Economic Theory*, 1985, 36, 277–288.
- and **Andrew Postlewaite**, “Weak Versus Strong Domination in a Market with Indivisible Goods,” *Journal of Mathematical Economics*, 1977, 4, 131–139.
- and **Elliott Peranson**, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 1999, 89, 748–780.
- , **Tayfun Sönmez**, and **Utku Ünver**, “Kidney Exchange,” *Quarterly Journal of Economics*, 2004, 119, 457–488.
- RSD**, “Overview of the New Public School Enrollment Process,” 2012. Recovery School District Presentation to Public.
- Shapley, Lloyd and Herbert Scarf**, “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1974, 1, 23–28.
- Svensson, Lars-Gunnar**, “Strategy-proof Allocation of Indivisible Goods,” *Social Choice and Welfare*, 1999, 16.
- Vanacore, Andrew**, “Recovery School District chief plans central enrollment system, technical training, more,” 2011. September 6, New Orleans Times-Picayune.
- , “Centralized enrollment in Recovery School District gets first tryout,” 2012. April 16, New Orleans Times-Picayune.