

The Cutoff Structure of Top Trading Cycles in School Choice*

Jacob D. Leshno Irene Lo[†]

November 9, 2017

Abstract

The prominent Top Trading Cycles (TTC) mechanism has attractive properties for school choice, as it is strategy-proof, Pareto efficient, and allows school boards to guide the assignment by specifying priorities. However, the common combinatorial description of TTC does little to explain the relationship between student priorities and their eventual assignment. This creates difficulties in transparently communicating TTC to parents and in guiding policy choices of school boards.

We show that the TTC assignment can be described by n^2 admission thresholds, where n is the number of schools. These thresholds can be observed after mechanism is run, and can serve as non-personalized prices that allow parents to verify their assignment.

In a continuum model these thresholds can be computed directly from the distribution of preferences and priorities, providing a framework that can be used to evaluate policy choices. We provide closed form solutions for the assignment under a family of distributions, and derive comparative statics. As an application of the model we solve for the welfare maximizing investment in school quality, and find that a more egalitarian investment can be more efficient because it promotes more efficient sorting by students.

*This paper was previously circulated as “The Simple Structure of Top Trading Cycles in School Choice”. A current version of this paper and supplementary online appendix are available at [http://www.columbia.edu/~iyl2104/Leshno,Lo-TTC%20\(without%20online%20appendix\).pdf](http://www.columbia.edu/~iyl2104/Leshno,Lo-TTC%20(without%20online%20appendix).pdf) and [http://www.columbia.edu/~iyl2104/Leshno,Lo-TTC%20\(online%20appendix\).pdf](http://www.columbia.edu/~iyl2104/Leshno,Lo-TTC%20(online%20appendix).pdf) respectively. We thank (in random order) Cyndia Yu, John Tsitsiklis, Yeon-Koo Che, Larry Samuelson, Eduardo Azevedo, Nikhil Agarwal, Jay Sethuraman, Chiara Margaria, Itai Ashlagi, Laura Doval, Parag Pathak, Jonah Rockoff, Yusuke Narita, Fuhito Kojima, Alvin Roth and seminar participants in Berkeley, Columbia, Stanford, Yale, University of Toronto, Duke, NYU, Matching in Practice workshop, and Cirq Montreal Microeconomics meeting for their helpful comments.

[†]Columbia University, email: jleshno@columbia.edu and irene.lo@columbia.edu

1 Introduction

Since Abdulkadiroğlu and Sönmez (2003) formulated school choice as an assignment problem, many school districts have redesigned their school assignment mechanisms to give students more choice over their school assignment. Abdulkadiroğlu and Sönmez propose two mechanisms that are strategyproof for students and allow the school district to set a priority structure: the Deferred Acceptance (DA) mechanism (Gale and Shapley, 1962); and the Top Trading Cycles (TTC) mechanism (Shapley and Scarf (1974), attributed to David Gale). The mechanisms differ in that DA is stable but not necessarily efficient, and TTC is Pareto efficient for students but not necessarily stable. In many districts schools do not screen students and there is no need to sacrifice efficiency to guarantee stability¹. Despite this, almost every district that redesigned their school choice mechanism chose to implement the DA mechanism (Pathak and Sönmez, 2013; IIPSC, 2017), instead of the more efficient TTC.²

One of the main barriers to using TTC in practice is the difficulty designers had in communicating it to parents and school boards.³ The standard explanation of TTC is an algorithmic description in terms of sequentially clearing trade cycles, from which it is not directly apparent how priorities and preferences determine assignment. This makes it difficult for school boards to evaluate the effects of policy decisions on the TTC assignment and resulting welfare. It is also difficult for students to verify they were correctly assigned and be convinced that the mechanism is strategyproof.

¹There is no strategic concern that blocking pairs will match to each other outside the mechanism. This is because of two differences between school choice and two-sided settings like the medical match. First, in Boston and other districts schools cannot directly admit students without approval. Second, priorities are often determined by school zone, sibling status and lotteries, and are not controlled by the school, and so schools do not necessarily prefer students with higher priority. A notable exception is the NYC high school admissions system, see Abdulkadiroğlu, Pathak, and Roth (2009).

²To the authors' knowledge, the only instances of implementation of TTC in school choice systems are in the San Francisco school district (Abdulkadiroğlu et al. 2017) and previously in the New Orleans Recovery School District (Abdulkadiroğlu et al. 2017).

³Pathak (2016) writes that:

“I believe that the difficulty of explaining TTC, together with the precedent set by New York and Boston's choice of DA, are more likely explanations for why TTC is not used in more districts, rather than the fact that it allows for justified envy, while DA does not.”

In addition, Boston and NYC were early school redesigns that set a precedent in favor of DA. More details can be found in the discussion in Abdulkadiroğlu, Pathak, Roth, and Sönmez (2005) and in Pathak (2016), Abdulkadiroğlu, Che, Pathak, Roth, and Tercieux (2017).

We explain the TTC assignment using n^2 cutoffs $\{p_b^c\}$. These cutoffs parallel prices in competitive equilibrium, where students' priorities serve the role of endowments. Students can use priority at a school b to gain assignment for school c if their priority at school b is above the threshold p_b^c . Publicly publishing the cutoffs $\{p_b^c\}$ lets each student verify that she is assigned to her most preferred school for which she has sufficient priority.⁴ To help convey to students that TTC is strategyproof, we present cutoffs $\{p_b^c\}$ that do not change with the preferences of any individual student.

To understand the mapping from the economy to the cutoffs and TTC assignment, we formulate TTC in a continuum model. We show that the TTC assignment in the continuum can be directly calculated by solving a system of equations, and give closed form solutions for parameterized economies. This framework allows us to derive comparative statics, and evaluate welfare under various policy choices. We validate our continuum TTC model by demonstrating that it produces cutoffs for discrete economies that describe the discrete TTC assignment.

The continuum TTC model allows us to calculate welfare. For example, the cutoff representation yields budget sets for each student, providing a tractable expression for welfare under random utility models. The model also provides comparative statics for assignment and welfare with respect to changes in school attributes and student preferences. For example, consider investment in a school's quality that increases the utility for students who attend that school. Such an investment would change student preferences and therefore the TTC assignment. We decompose the marginal effect of such investment on welfare into the direct increase in utility of students assigned to the school, and the indirect effect that arises from changes in the assignment.

As an application of our model, we study optimal investment in school quality when all students prefer higher quality school but have unobserved taste shocks. We solve for the optimal investment under TTC for a parametric setting, and find that the optimal investment is equitable in the sense of making all schools equally over-demanded. A more equitable investment is more efficient because it allows students more choice, yielding better sorting on horizontal dimensions and therefore

⁴This threshold representation allows us to give the following non-combinatorial description of TTC. For each school b , each student receives b -tokens according to their priority at school b , where students with higher b -priority receive more b -tokens. The TTC algorithm publishes prices $\{p_b^c\}$. Students can purchase a single school using a single kind of token, and the required number of b -tokens to purchase school c is p_b^c . Note that $\{p_b^c\}$ can be obtained by running TTC and setting p_b^c to be the number of b -tokens of the lowest b -priority student that traded into c using priority at b . We thank Chiara Margaria, Laura Doval and Larry Samuelson for suggesting this explanation.

higher welfare. This holds even if the schools are not symmetric, as the benefits from more efficient sorting can outweigh benefits from targeting schools with more efficient investment opportunities.

Second, we explore the design of priorities for TTC. The priority structure under TTC is “bossy” in the sense that a change in the relative priority among top priority students can change the cutoffs and the assignment of low priority students, without changing the assignment of any high priority student. Such changes to the relative priority among top students cannot be detected by supply-demand equations as in Azevedo and Leshno (2016), and therefore it is not possible to determine the TTC thresholds directly through a supply-demand equation. We characterize the range of possible assignments generated by TTC after changes to relative priority of high-priority students, and show that a small change to the priorities will only change the assignment of a few students.

A third application of our model is to provide comparisons between mechanisms in terms of assignments and welfare. We solve for welfare under TTC and DA in a particular setting and quantify how much welfare is sacrificed to guarantee stability. We compare TTC and DA across different school choice environments and corroborate a conjecture by Pathak (2016) that the difference between the mechanisms is smaller when students have a preference for and priority at their neighborhood school. We also compare TTC to the Clinch and Trade mechanism by Morrill (2015b) in large economies and find that it is possible for TTC to produce fewer blocking pairs than the Clinch and Trade mechanism.

A key idea that allows us to define TTC in the continuum is that the TTC algorithm can be characterized by its aggregate behavior over many cycles. Any collection of cycles must maintain *trade balance*, that is, the number of students assigned to each school is equal to the number of students who claimed or traded a seat at that school. In the continuum this necessary condition yields a system of equations that fully characterizes TTC. These equations also provide a recipe for calculating the TTC assignment.

A few technical aspects of the analysis may be of interest. First, we note that the trade balance equations circumvent many of the measure theoretic complications in defining TTC in the continuum. Second, a connection to Markov chain theory allows us to show that a solution to the marginal trade balance equations always exists, and to characterize the possible trades.

1.1 Related Literature

Abdulkadiroğlu and Sönmez (2003) first introduced school choice as a mechanism design problem and suggested the TTC mechanism as a solution with several desirable properties. Since then, TTC has been considered for use in a number of school choice systems. Abdulkadiroğlu, Pathak, Roth, and Sönmez (2005) discuss how the city of Boston debated between using DA or TTC for their school choice systems and ultimately chose DA. Abdulkadiroğlu, Pathak, and Roth (2009) compare the outcomes of DA and TTC for the NYC public school system, and show that TTC gives higher student welfare. Kesten (2006) also studies the relationship between DA and TTC, and shows that they are equivalent mechanisms if and only if the priority structure is acyclic.

Our work contributes to the literature on threshold representations, which has been instrumental for empirical work on DA and variants of DA. Abdulkadiroğlu, Angrist, Narita, and Pathak (Forthcoming) use admission thresholds to construct propensity score estimates. Agarwal and Somaini (2014); Kapor, Neilson, and Zimmerman (2016) structurally estimate preferences from rank lists submitted to non-strategyproof variants of DA. Both build on the threshold representation of Azevedo and Leshno (2016). We hope that our threshold representation of TTC will be similarly useful for future empirical work on TTC.

Our finding that the TTC assignment can be represented in terms of cutoffs parallels the role of prices in competitive markets. Dur and Morrill (2016) show that the outcome of TTC can be expressed as the outcome of a competitive market where there is a price for each priority position at each school, and agents may buy and sell exactly one priority position. Their characterization provides a connection between TTC and competitive markets, but requires a price for each rank at each school and does not provide a method for directly calculating these prices without running TTC. He, Miralles, Pycia, Yan, et al. (2015) propose an alternative pseudo-market approach for discrete assignment problems that extends Hylland and Zeckhauser (1979) and also uses admission thresholds. Miralles and Pycia (2014) show a second welfare theorem for discrete goods, namely that any Pareto efficient assignment of discrete goods without transfers can be decentralized through prices and endowments, but requires an arbitrary endowment structure.

Our approach builds on existing axiomatic characterizations of TTC, and may extend to the class of Pareto efficient and strategyproofness mechanisms. Abdulka-

diroglu, Che, Pathak, Roth, and Tercieux (2017) show that TTC minimizes the number of blocking pairs subject to strategyproofness and Pareto efficiency. Additional axiomatic characterizations of TTC were given by Dur (2012) and Morrill (2013, 2015a). These characterizations explore the properties of TTC, but do not provide another method for calculating the TTC outcome or evaluating welfare. Ma (1994), Pápai (2000) and Pycia and Ünver (2015) give characterizations of more general classes of Pareto efficient and strategy-proof mechanisms in terms of clearing trade cycles. While our analysis focuses on the TTC mechanism, we believe that our trade balance approach will be useful in analyzing these general classes of mechanisms.

Several variants of TTC have been suggested in the literature. Morrill (2015b) introduces the Clinch and Trade mechanism, which differs from TTC in that it identifies students who are guaranteed admission to their first choice and assigns them immediately without implementing a trade. Hakimov and Kesten (2014) introduce Equitable TTC, a variation on TTC that aims to reduce inequity. In Section 4.2 we use our model to analyze such variants of TTC and compare their assignments. Other variants of TTC can also arise from the choice of tie-breaking rules. Ehlers (2014) shows that any fixed tie-breaking rule satisfies weak efficiency, and Alcalde-Unzu and Molis (2011), Jaramillo and Manjunath (2012) and Saban and Sethuraman (2013) give specific variants of TTC that are strategy-proof and efficient. The continuum model allows us to characterize the possible outcomes from different tie-breaking rules.

Several papers also study TTC in large markets. Che and Tercieux (2015a,b) study the properties of TTC in a large market where the number of items grows as the market gets large. Hatfield, Kojima, and Narita (2016) study the incentives for schools to improve their quality under TTC and find that a school may be assigned some less preferred students when it improves its quality.

This paper contributes to a growing literature that uses continuum models in market design (Avery and Levin, 2010; Abdulkadiroğlu, Che, and Yasuda, 2015; Ashlagi and Shi, 2015; Che, Kim, and Kojima, 2013; Azevedo and Hatfield, 2015). Our description of the continuum economy uses the setup of Azevedo and Leshno (2016), who characterize stable matchings in terms of cutoffs that satisfy a supply and demand equation. Our results from Section 4.2 imply that the TTC cutoffs depend on the entire distribution and cannot be computed from simple supply and demand equations.

1.2 Organization of the Paper

Section 2 provides the standard description of the TTC mechanism in the discrete model and presents our cutoff characterization. Section 3 presents the continuum TTC model and provides our main results that allow for direct calculation of the TTC cutoffs. Section 4 explores several applications: quantifying the effects of improving school quality under school choice and solving for optimal investment, showing the “bossiness” of the TTC priorities, and comparing TTC with other mechanisms. Section A provides the intuition for the continuum TTC model. Omitted proofs can be found in the appendix.

2 TTC in School Choice

2.1 The Discrete TTC Model

In this section, we describe the standard model for the TTC mechanism in the school choice literature, and outline some of the properties of TTC in this setting.

Let \mathcal{S} be a finite set of students, and let \mathcal{C} be a finite set of schools. Each school $c \in \mathcal{C}$ has a finite capacity $q_c > 0$. Each student $s \in \mathcal{S}$ has a strict preference ordering \succ^s over schools. Let $Ch^s(C) = \arg \max_{\succ^s} \{C\}$ denote s 's most preferred school out of the set C . Each school $c \in \mathcal{C}$ has a strict priority ordering \succ_c over students. To simplify notation, we assume that all students and schools are acceptable, and that there are more students than available seats at schools.⁵ It will be convenient to represent the priority of student s at school c by the student's percentile rank $r_c^s = |\{s' \mid s \succ_c s'\}| / |\mathcal{S}|$ in the school's priority ordering. Note that for any two students s, s' and school c we have that $s \succ_c s' \iff r_c^s > r_c^{s'}$ and that $0 \leq r_c^s < 1$.

A *feasible* assignment is $\mu : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$ where $|\mu^{-1}(c)| \leq q_c$ for every $c \in \mathcal{C}$. If $\mu(s) = c$ we say that s is assigned to c , and we use $\mu(s) = \emptyset$ to denote that the student s is unassigned. As there is no ambiguity, we let $\mu(c)$ denote the set $\mu^{-1}(c)$ for $c \in \mathcal{C} \cup \{\emptyset\}$. A *discrete economy* is $E = (\mathcal{C}, \mathcal{S}, \succ^{\mathcal{S}}, \succ_{\mathcal{C}}, q)$, where \mathcal{C} is the set of schools, \mathcal{S} is the set of students, $q = \{q_c\}_{c \in \mathcal{C}}$ is the capacity of each school, and $\succ^{\mathcal{S}} = \{\succ^s\}_{s \in \mathcal{S}}$, $\succ_{\mathcal{C}} = \{\succ_c\}_{c \in \mathcal{C}}$.

⁵This is without loss of generality, as we can introduce auxiliary students and schools that represent being unmatched.

The discrete Top Trading Cycles algorithm (TTC) calculates an assignment μ_{dTTC} . The algorithm runs in discrete steps as follows.

Algorithm 1 (Top Trading Cycles). *Initialize unassigned students $S = \mathcal{S}$, available schools $C = \mathcal{C}$, capacities $\{q_c\}_{c \in C}$.*

While there are still unassigned students and available schools:

- *Each available school $c \in C$ tentatively offers a seat by pointing to its highest priority remaining student.*
- *Each student $s \in S$ that was tentatively offered a seat points to her most preferred remaining school.*
- *Select at least one trading cycle, that is, a list of students s_1, \dots, s_ℓ , $s_{\ell+1} = s_1$ such that s_i points to the school pointing to s_{i+1} for all i , or equivalently s_{i+1} was offered a seat at s_i 's most preferred school. Assign all students in the cycles to the school they point to.⁶*
- *Remove the assigned students from S , reduce the capacity of the schools they are assigned to by 1, and remove schools with no remaining capacity from C .*

TTC satisfies a number of desirable properties. An assignment μ is *Pareto efficient for students* if no group of students can improve by swapping their allocations, and no individual student can improve by swapping her assignment for an unassigned object. A mechanism is *Pareto efficient for students* if it always produces an assignment that is Pareto efficient for students. A mechanism is *strategy-proof for students* if reporting preferences truthfully is a dominant strategy. It is well known that TTC, as used in the school choice setting, is both Pareto efficient and strategy-proof for students (Abdulkadiroğlu and Sönmez, 2003). Moreover, when type-specific quotas must be imposed, TTC can be easily modified to meet quotas while still maintaining constrained Pareto efficiency and strategy-proofness (Abdulkadiroğlu and Sönmez, 2003).

2.2 Cutoff Characterization

Our first main contribution is that the TTC assignment can be described in terms of n^2 cutoffs $\{p_b^c\}$, one for each pair of schools.

⁶Such a trading cycle must exist, since every vertex in the pointing graph with vertex set $S \cup C$ has out-degree 1.

Theorem 1. *The TTC assignment is given by*

$$\mu_{TTC}(s) = \max_{\succ^s} \{c \mid r_b^s \geq p_b^c \text{ for some } b\},$$

where p_b^c is the percentile in school b 's ranking of the worst ranked student at school b that traded a seat at school b for a seat at school c . If no such student exists, $p_b^c = 1$.

Cutoffs serve a parallel role to prices, with each student's priority at each school serving the role of endowments. The cutoffs $\mathbf{p} = \{p_b^c\}_{b,c}$ combine with each student s 's priorities r^s to give s a *budget set* $B(s, \mathbf{p}) = \{c \mid r_b^s \geq p_b^c \text{ for some } b\}$ of schools she can attend. TTC assigns each student to her favorite school in her budget set.

Theorem 1 provides an intuitive way for students to verify that they were correctly assigned by the TTC algorithm. The cutoff p_b^c can be easily identified after the mechanism has been run. Instead of only communicating the assignment of each student, the mechanism can make the cutoffs publicly known. Students can calculate their budget set from their privately known priorities and the publicly given cutoffs, allowing them to verify that they were indeed assigned to their most preferred school in their budget set. In particular, if a student does not receive a seat at a desired school c , it is because she does not have sufficiently high priority at *any* school, and so c is not in her budget set. We illustrate these ideas in Example 1.

Example 1. Consider a simple economy where there are two schools each with capacity $q = 120$, and a total of 300 students, 2/3 of whom prefer school 1. Figure 1(a) illustrates the preferences and priorities of each of the students.⁷ A student's priority determines a location in the square, with the horizontal axis indicating priority at school 1 and the vertical axis indicating priority at school 2.

The cutoffs \mathbf{p} and resulting budget sets $B(s, \mathbf{p})$ for each student are illustrated in Figure 1(b); the shaded areas show budget sets as a function of student priority. For example, a student has the budget set $\{1, 2\}$ if she has sufficiently high priority at either school 1 or school 2. Note that students' preferences are not indicated in Figure 1(b) as each student's budget set is independent of her preferences. Figure 1(c) depicts the students' assignments, which depend on their preferences. The plot has two squares: the left square gives the assignment of students who prefer school

⁷Student priorities were selected such that there is little correlation between student priority at either school and between student priorities and preferences.

1; and the right square gives the assignment of students who prefer school 2. The assignment of each student is simply her favorite school in her budget set.


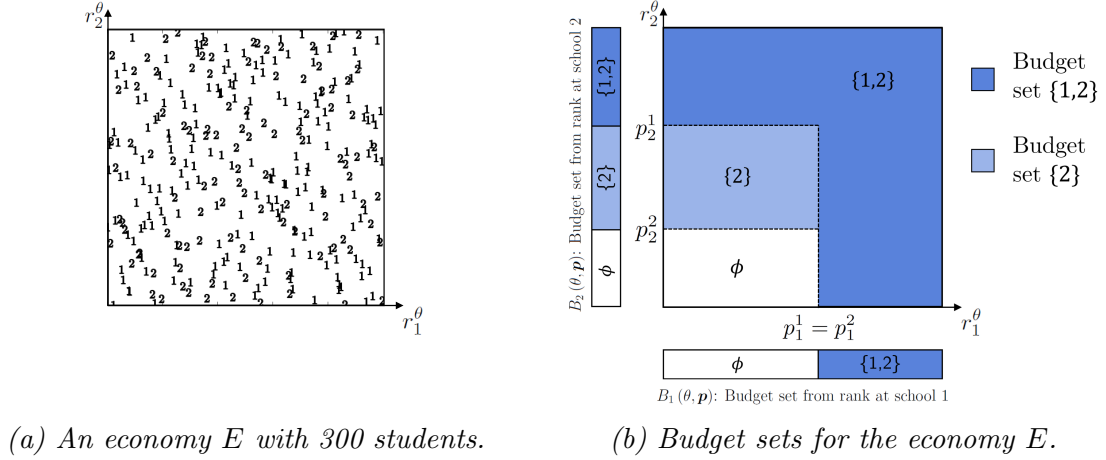


Figure 1: The TTC budget sets and assignment for the economy E in Example 1. All students find both schools acceptable, and students are labeled by their preferred school. Students' priority ranks at school 1 are given by the horizontal axis and priority ranks at school 2 by the vertical axis. The budget sets along the axis for school c list $B_c(s, \mathbf{p})$, the schools that enter a student's budget set because of her rank at school c .

Figure 1 shows the role of priorities in determining the TTC assignment in Example 1. Students with higher priority have a larger budget set of school from which they can choose. A student can choose her desired school if her priority for some school is sufficiently high. Priority for each school is considered separately, and priority from multiple schools cannot be combined. For example, a student who has top priority for one school and bottom priority at the other school is assigned to her top

choice, but a student who has the median priority at both schools will not be assigned to school 1.

Remark. This example also shows that the TTC assignment cannot be expressed in terms of one cutoff for each school, as the assignment in Example 1 cannot be described by fewer than 3 cutoffs.

2.3 The Structure of TTC Budget Sets

The cutoff structure for TTC allows us to provide some insight into the structure of the assignment. For each student s , let $B_b(s, \mathbf{p}) = \{c \mid r_b^s \geq p_b^c\}$ denote the set of schools that enter student s 's budget set because of her priority at school b . $B_1(s, \mathbf{p})$ and $B_2(s, \mathbf{p})$ are depicted along the x and y axes of Figure 1. Note that $B_b(s, \mathbf{p})$ depends only on the n cutoffs $\mathbf{p}_b = \{p_b^c\}_{c \in \mathcal{C}}$. A student's budget set is the union $B(s, \mathbf{p}) = \cup_b B_b(s, \mathbf{p})$. The following proposition shows that budget sets $B_b(s, \mathbf{p})$ can be given by cutoffs \mathbf{p}_b that share the same ordering over schools for every b . We let $\mathcal{C}^{(c)} = \{c, c+1, \dots, n\}$ denote the set of schools that have a higher index than c .

Proposition 1. *There exist cutoffs $\{p_b^c\}$ and a relabeling of school indices such that the cutoffs \mathbf{p} describe the TTC assignment,*

$$\mu_{\text{TTC}}(s) = \max_{\succ_s} \{c \mid r_b^s \geq p_b^c \text{ for some } b\},$$

and for any school b the cutoffs are ordered,⁸

$$p_b^1 \geq p_b^2 \geq \dots \geq p_b^b = p_b^{b+1} = \dots = p_b^n.$$

Therefore, the set of schools $B_b(s, \mathbf{p})$ student s can afford via her priority at school b is either the empty set ϕ or

$$B_b(s, \mathbf{p}) = \mathcal{C}^{(c)} = \{c, c+1, \dots, n\}$$

for some $c \leq b$. Moreover, each student's budget set $B(s, \mathbf{p}) = \cup_b B_b(s, \mathbf{p})$ is either $B(s, \mathbf{p}) = \phi$ or $B(s, \mathbf{p}) = \mathcal{C}^{(c)}$ for some c .

⁸The cutoffs \mathbf{p} defined in Theorem 1 do not necessarily satisfy this condition. However, the run of TTC produces the following relabeling of schools and cutoffs $\tilde{\mathbf{p}}$ that give the same assignment and satisfy the condition: the schools are relabeled in the order in which they reach capacity under TTC, and the cutoffs $\tilde{\mathbf{p}}$ are given by $\tilde{p}_b^c = \min_{a \leq c} p_b^a$.

The cutoff ordering proved in Proposition 1 implies that budget sets of different students are nested, and therefore that the TTC assignment is Pareto efficient. The cutoff ordering is a stronger property than Pareto efficiency, and is not implied by the Pareto efficiency of TTC. For example, serial dictatorship with a randomly drawn ordering will give a Pareto efficient assignment, but there is no relationship between a student's priorities and her assignment.

Proposition 1 allows us to give a simple illustration for the TTC assignment when there are $n \geq 3$ schools. For each school b , we can illustrate the set of schools $B_b(s, \mathbf{p})$ that enter a student's budget set because of her priority at school b as in Figure 2 (under the assumption that schools are labeled in the correct order). This generalizes the illustration along each axis in Figure 1, and can be used for any number of schools. It is possible that $p_b^c = 1$, meaning that students cannot use their priority at school b to trade into school c .




Figure 2: The schools $B_b(s, \mathbf{p})$ that enter a student's budget set because of her rank at school b . Students may use their rank at school b to obtain a budget set of $C^{(1)}, C^{(2)}, \dots, C^{(b)}, \emptyset$. The cutoffs p_b^c are weakly decreasing in c , and are equal for all $c \geq b$ (i.e. $p_b^b = p_b^{b+1} = \dots = p_b^n$).

2.4 Limitations

Although the cutoff structure is helpful in understanding the structure of the TTC assignment, there are several limitations to the cutoffs computed in Theorem 1 and Proposition 1. First, while the cutoffs can be determined by running the TTC algorithm, Theorem 1 does not provide a direct method for calculating the cutoffs from the economy primitives. In particular, it does not explain how the TTC assignment changes with changes in school priorities or student preferences. Second, the budget set $B(s, \mathbf{p})$ given by the cutoffs derived in Theorem 1 does not correspond to the set of schools that student s can be assigned to by reporting some preferences.^{9,10}

⁹More precisely, given economy E and student s , let economy E' be generated by changing the preferences ordering of s from \succ^s to \succ' . Let $\mu_{dTTC}(s | E)$ and $\mu_{dTTC}(s | E')$ be the assignment of s under the two economies, and let \mathbf{p} be the cutoffs derived by Theorem 1 for economy E . Theorem 1 shows that $\mu_{dTTC}(s | E) = \max_{\succ^s} B(s, \mathbf{p})$ but it may be $\mu_{dTTC}(s | E') \neq \max_{\succ'} B(s, \mathbf{p})$.

¹⁰For example, let E be an economy with three schools $\mathcal{C} = \{1, 2, 3\}$, each with capacity 1. There are three students s_1, s_2, s_3 such that the top preference of s_1, s_2 is school 1, the top preference of

We therefore introduce the continuum model for TTC which allows us to directly calculate the cutoffs, allowing for comparative statics. These cutoffs will also correspond to refined budget sets which provide the sets of schools that students could be assigned to by unilaterally changing their preferences.

3 Continuum Model and Main Results

3.1 Model

We model the school choice problem with a continuum of students and finitely many schools, as in Azevedo and Leshno (2016). There is a finite set of schools denoted by $\mathcal{C} = \{1, \dots, n\}$, and each school $c \in \mathcal{C}$ has the capacity to admit a mass $q_c > 0$ of students. A student $\theta \in \Theta$ is given by $\theta = (\succ^\theta, r^\theta)$. We let \succ^θ denote the student's strict preferences over schools, and let $Ch^\theta(C) = \max_{\succ^\theta}(C)$ denote θ 's most preferred school out of the set C . The priorities of schools over students are captured by the vector $r^\theta \in [0, 1]^C$. We say that r_b^θ is the *rank* of student θ at school b . Schools prefer students with higher ranks, that is $\theta \succ_b \theta'$ if and only if $r_b^\theta > r_b^{\theta'}$.

Definition 1. A *continuum economy* is given by $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ where $q = \{q_c\}_{c \in \mathcal{C}}$ is the vector of capacities of each school, and η is a measure over Θ .

We make some assumptions for the sake of tractability. First, we assume that all students and schools are acceptable. Second, we assume there is an excess of students, that is, $\sum_{c \in \mathcal{C}} q_c < \eta(\Theta)$. Finally, we make the following technical assumption that ensures that the run of TTC in the continuum economy is sufficiently smooth and allows us to avoid some measurability issues.¹¹

s_3 is school 3, and student s_i has top priority at school i . Theorem 1 gives the budget set $\{1\}$ for student s_1 , as $\mathbf{p}^1 = (\frac{2}{3}, 1, 1)$, $\mathbf{p}^2 = (1, \frac{2}{3}, 1)$ and $\mathbf{p}^3 = (1, 1, \frac{2}{3})$, since the only trades are of seats at c for seats at the same school c . However, if s_1 reports the preference $2 \succ 1 \succ 3$ she will be assigned to school 2, so an appropriate definition of budget sets should include school 2 in the budget set for student s_1 . Also note that no matter what preference student s_1 reports, she will not be assigned to school 3, so an appropriate definition of budget sets should not include school 3 in the budget set for student s_1 .

¹¹We can incorporate an economy where two schools have perfectly aligned priorities by considering them as a combined single school in the trade balance equations. The capacity constraints still consider the capacity of each school separately.

Assumption 1. *The measure η admits a density ν . That is for any measurable subset of students $A \subseteq \Theta$*

$$\eta(A) = \int_A \nu(\theta) d\theta.$$

Furthermore, ν is piecewise Lipschitz continuous everywhere except on a finite grid,¹² bounded from above, and bounded from below away from zero on its support.¹³

Assumption 1 is general enough to allow embeddings of discrete economies, and is satisfied by all the economies considered throughout the paper. However, it is not without loss of generality, e.g. it is violated when all schools share the same priorities over students.

An immediate consequence of Assumption 1 is that a school's indifference curves are of η -measure 0. That is, for any $c \in \mathcal{C}$, $x \in [0, 1]$ we have that $\eta(\{\theta \mid r_b^\theta = x\}) = 0$. This is analogous to schools having strict preferences in the standard discrete model. Given this assumption, as r_b^θ carries only ordinal information, we may assume each student's rank is normalized to be equal to her percentile rank in the school's preferences, i.e. for any $b \in \mathcal{C}$, $x \in [0, 1]$ we have that $\eta(\{\theta \mid r_b^\theta \leq x\}) = x$.

It is convenient to describe the distribution η by the corresponding marginal distributions. Specifically, for each point $\mathbf{x} \in [0, 1]^n$ and subset of schools $C \subseteq \mathcal{C}$, let $H_b^{c|C}(\mathbf{x})$ be the marginal density of students who are top ranked at school b among all students whose rank at every school a is no better than x_a , and whose top choice among the set of schools C is c .¹⁴ We omit the dependence on C when the relevant set of schools is clear from context, and write $H_b^c(\mathbf{x})$. The marginal densities $H_b^{c|C}(\mathbf{x})$ uniquely determine the distribution η .

As in the discrete model, an assignment is a mapping $\mu : \Theta \rightarrow \mathcal{C} \cup \{\emptyset\}$ specifying the assignment of each student. With slight abuse of notation, we let $\mu(c) = \{\theta \mid \mu(\theta) = c\}$ denote the set of student assigned to school c . An assignment μ is *feasible* if it respects capacities, i.e. for each school $c \in \mathcal{C}$ we have $\eta(\mu(c)) \leq q_c$. Two allocations μ and μ' are *equivalent* if they differ only on a set of students of zero measure, i.e. $\eta(\{\theta \mid \mu(\theta) \neq \mu'(\theta)\}) = 0$.

¹²A grid $G \subset \Theta$ is given by $G = \{\theta \mid \exists c \text{ s.t. } r_c^\theta \in D\}$, where $D = \{d_1, \dots, d_L\} \subset [0, 1]$ is a finite set of grid points. Equivalently, ν is Lipschitz continuous on the union of open hypercubes $\Theta \setminus G$.

¹³That is, there exists $M > m > 0$ such that for every $\theta \in \Theta$ either $\nu(\theta) = 0$ or $m \leq \nu(\theta) \leq M$.

¹⁴Formally $H_b^{c|C}(\mathbf{x}) \stackrel{\text{def}}{=} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \eta(\{\theta \in \Theta \mid r^\theta \in [\mathbf{x} - \varepsilon \cdot \mathbf{e}^b, \mathbf{x}) \text{ and } Ch^\theta(C) = c\})$, where \mathbf{e}^b is the unit vector in the direction of coordinate b .

Remark 1. In school choice, it is common for schools to have coarse priorities, and to refine these using a tie-breaking rule. Our economy \mathcal{E} captures the strict priority structure that results after applying the tie-breaking rule.

3.2 Main Results

Our main result establishes that in the continuum model the TTC assignment can be directly calculated from trade balance and capacity equations. This allows us to explain how the TTC assignment changes with changes in the underlying economy. It also allows us to derive cutoffs that are independent of a student's reported preferences, giving another proof that TTC is strategyproof.

We remark that directly translating the TTC algorithm to the continuum setting by considering individual trading cycles is challenging, as a direct adaptation of the algorithm would require the clearing of cycles of zero measure. We circumvent the technical issues raised by such an approach by formally defining the continuum TTC assignment in terms of trade balance and capacity equations, which characterize the TTC algorithm in terms of its aggregate behavior over multiple steps. To verify the validity of our definition, we show in Subsection 3.3 that continuum TTC can be used to calculate the discrete TTC outcome. We provide further intuition in Appendix A.

We begin with some definitions. A function $\gamma(t) : [0, \infty) \rightarrow [0, 1]^C$ is a *TTC path* if γ is continuous and piecewise smooth, $\gamma_c(t)$ is weakly decreasing for all c , and the initial condition $\gamma(0) = \mathbf{1}$ holds. A function $\tilde{\gamma}(t) : [t_0, \infty) \rightarrow [0, 1]^C$ is a *residual TTC path* if it satisfies all the properties of a TTC path except the initial condition, and is defined only for $t \geq t_0 > 0$. For a set $\{t^{(c)}\}_{c \in C} \subseteq \mathbb{R}_{\geq}^C$ of times we let $t^{(c^*)} \stackrel{\text{def}}{=} \min_c [t^{(c)}]$ denote the minimal time. For a point $x \in [0, 1]^C$, let

$$D^c(x) \stackrel{\text{def}}{=} \eta(\{\theta \mid r^\theta \not\prec x, Ch^\theta(\mathcal{C}) = c\})$$

denote the mass of students who want c among all students with rank better than x . We will refer to $D^c(x)$ as the demand for c . Recall that $H_b^c(x)$ is the marginal density of students who want c who are top ranked at school b among all students with rank no better than x . Note that $D^c(x)$ and $H_b^c(x)$ depend implicitly on the set of available schools \mathcal{C} .

Definition 2. Let $E = (\mathcal{C}, \Theta, \eta, q)$ be an economy. We say that the (residual) TTC

path $\gamma(t)$ and positive stopping times $\{t^{(c)}\}_{c \in \mathcal{C}} \in \mathbb{R}_{\geq}^{\mathcal{C}}$ satisfy trade balance and capacity for the economy E if the following hold.

1. $\gamma(\cdot)$ satisfies the *marginal trade balance equations* given by

$$\sum_{a \in \mathcal{C}} \gamma'_a(t) H_a^c(\gamma(t)) = \sum_{c \in \mathcal{C}} \gamma'_c(t) H_c^a(\gamma(t)) \quad (1)$$

for all $t \leq t^{(c^*)} = \min_c [t^{(c)}]$ for which the derivatives exist.

2. The minimal stopping time $t^{(c^*)}$ solves the *capacity equations*

$$\begin{aligned} D^{c^*}(\gamma(t^{(c^*)})) &= q_{c^*} \\ D^a(\gamma(t^{(c^*)})) &\leq q_a \quad \forall a \in \mathcal{C} \end{aligned} \quad (2)$$

and $\gamma_{c^*}(t)$ is constant for all $t \geq t^{(c^*)}$.

3. Define the residual economy $\tilde{E} = (\tilde{\mathcal{C}}, \Theta, \tilde{\eta}, \tilde{q})$ by $\tilde{\mathcal{C}} = \mathcal{C} \setminus \{c^*\}$, $\tilde{q}_c = q_c - D^c(\gamma(t^{(c^*)}))$ and $\tilde{\eta}(A) = \eta(A \cap \{\theta : r^\theta \leq \gamma(t^{(c^*)})\})$. Define the residual TTC path $\tilde{\gamma}(\cdot)$ by restricting $\gamma(\cdot)$ to $t \geq t^{(c^*)}$ and coordinates within $\tilde{\mathcal{C}}$. Then $\tilde{\gamma}$ and the stopping times $\{t^{(c)}\}_{c \in \tilde{\mathcal{C}}}$ satisfy trade balance and capacity for \tilde{E} .

Theorem 2. *Let $E = (\mathcal{C}, \Theta, \eta, q)$ be an economy. There exist a TTC path $\gamma(\cdot)$ and stopping times $\{t^{(c)}\}_{c \in \mathcal{C}}$ that satisfy trade balance and capacity for E , and the TTC assignment is given by*

$$\mu_{cTTC}(\theta) = \max_{\succ_{\theta}} \{c : r_b^\theta \geq p_b^c \text{ for some } b\},$$

where the n^2 TTC cutoffs $\{p_b^c\}$ are given by

$$p_b^c = \gamma_b(t^{(c)}) \quad \forall b, c.$$

Moreover, any $\gamma(\cdot), \{t^{(c)}\}_{c \in \mathcal{C}}$ that satisfy trade balance and capacity yield the same assignment μ_{cTTC} .

In other words, Theorem 2 provides the following a recipe for calculating the TTC assignment. First, find $\hat{\gamma}(\cdot)$ that solves equation 1 for all t . Second, calculate $t^{(c^*)}$ from the capacity equations (2) for $\hat{\gamma}(\cdot)$. Set $\gamma(t) = \hat{\gamma}(t)$ for $t \leq t^{(c^*)}$. To determine

the remainder of $\gamma(\cdot)$, apply the same steps to the residual economy \tilde{E} which has one less school.¹⁵ This recipe is illustrated in Example 2.

Theorem 2 shows that the cutoffs can be directly calculated from the primitives of the economy. In contrast to the cutoff characterization in the standard model (Theorem 1), this allows us to understand how the TTC assignment changes with changes in capacities, preferences or priorities. We remark that the existence of a smooth curve γ follows from our assumption that η has a density that is piecewise Lipschitz and bounded, and the existence of $t^{(c)}$ satisfying the capacity equations (2) follows from our assumptions that there are more students than seats and all students find all schools acceptable.

The following immediate corollary of Theorem 2 shows that in contrast with the cutoffs given by the discrete model, the cutoffs given by Theorem 2 always satisfy the cutoff ordering.

Corollary 1. *Let the schools be labeled such that $t^{(1)} \leq t^{(2)} \leq \dots \leq t^{(n)}$. Then the cutoff ordering in Proposition 1 holds, namely*

$$p_b^1 \geq p_b^2 \geq \dots \geq p_b^b = p_b^{b+1} = \dots = p_b^{|C|} \text{ for all } b.$$

To illustrate how Theorem 2 can be used to calculate the TTC assignment and understand how it depends on the parameters of the economy, we consider the following simple economy. This parameterized economy yields a tractable closed form solution for the TTC assignment. For other economies the equations may not necessarily yield tractable expressions, but the same calculations can be used to numerically solve for cutoffs for any economy satisfying our smoothness requirements.

Example 2. We demonstrate how to use Theorem 2 to calculate the TTC assignment for a simple parameterized continuum economy. The economy \mathcal{E} has two schools 1, 2 with capacities $q_1 = q_2 = q$ with $q < 1/2$. A fraction $p > 1/2$ of students prefer school 1, and student priorities are uniformly distributed on $[0, 1]$ independently for each school and independently of preferences. This economy is described by

$$H(x_1, x_2) = \begin{bmatrix} px_2 & (1-p)x_2 \\ px_1 & (1-p)x_1 \end{bmatrix},$$

¹⁵Continuity of the TTC path provides an initial condition for $\tilde{\gamma}$, namely that $\tilde{\gamma}_c(t^{(c^*)}) = \gamma_c(t^{(c^*)})$ for all c .

where $H_b^c(\mathbf{x})$ is given by the b -row and c -column of the matrix. A particular instance of this economy with $q = 4/10$ and $p = 2/3$ can be viewed as a smoothed continuum version of the economy in Example 1.

We start by solving for γ from the trade balance equations (1), which simplify to the differential equation¹⁶

$$\frac{\gamma_2'(t)}{\gamma_1'(t)} = \frac{1-p}{p} \frac{\gamma_2(t)}{\gamma_1(t)}.$$

Since $\gamma(0) = 1$, this is equivalent to $\gamma_2(t) = (\gamma_1(t))^{\frac{1}{p}-1}$. Hence for $0 \leq t \leq \min\{t^{(1)}, t^{(2)}\}$ we set

$$\gamma(t) = \left(1 - t, (1 - t)^{\frac{1}{p}-1}\right).$$

We next compute $t^{(c^*)} = \min\{t^{(1)}, t^{(2)}\}$. Observe that because $p > 1/2$ it must be that $t^{(1)} < t^{(2)}$. Otherwise, we have that $t^{(2)} = \min\{t^{(1)}, t^{(2)}\}$ and $D^1(\gamma(t^{(2)})) \leq q$, implying that $D^2(\gamma(t^{(2)})) = \frac{1-p}{p} D^1(\gamma(t^{(2)})) < q$. Therefore, we solve $D^1(\gamma(t^{(1)})) = q$ to get that $t^{(1)} = 1 - \left(\frac{p-q}{p}\right)^p$ and that

$$p_1^1 = \gamma_1(t^{(1)}) = \left(1 - \frac{q}{p}\right)^p, \quad p_2^1 = \gamma_2(t^{(1)}) = \left(1 - \frac{q}{p}\right)^{1-p}.$$

For the remaining cutoffs, we eliminate school 1 and reiterate the same steps for the residual economy where $\mathcal{C}' = \{2\}$ and $q'_2 = q_2 - D^2(\gamma(t^{(1)})) = q(2 - 1/p)$.

For the residual economy equations (1) are trivial, and we set

$$\gamma(t) = (1 - t^{(1)}, p_2^1 - (t - t^{(1)}))$$

for $t^{(1)} \leq t \leq t^{(2)}$. Solving the capacity equation (2) for $t^{(2)}$ yields that

$$p_1^2 = \gamma_1(t^{(2)}) = \left(1 - \frac{q}{p}\right)^p, \quad p_2^2 = \gamma_2(t^{(2)}) = (1 - 2q) \left(1 - \frac{q}{p}\right)^{-p}.$$

¹⁶The original trade balance equations are

$$\begin{aligned} \gamma_1'(t) p \gamma_2(t) + \gamma_2'(t) p \gamma_1(t) &= \gamma_1'(t) p \gamma_2(t) + \gamma_1'(t) (1-p) \gamma_2(t), \\ \gamma_1'(t) (1-p) \gamma_2(t) + \gamma_2'(t) (1-p) \gamma_1(t) &= \gamma_2'(t) p \gamma_1(t) + \gamma_2'(t) (1-p) \gamma_1(t). \end{aligned}$$

For instance, if we plug in $q = 2/5$ and $p = 2/3$ to match the economy in Example 1, the calculation yields the cutoffs $p_1^1 = p_1^2 \approx .54$, $p_2^1 \approx .73$ and $p_2^2 \approx .37$, which are approximately the same cutoffs as those for the discrete economy in Example 1.

Example 2 illustrates how the TTC cutoffs can be directly calculated from the trade balance equations and capacity equations, without running the TTC algorithm. Example 2 also shows that it is not possible to solve for the TTC cutoffs only from supply-demand equations. In particular, the following equations are equivalent to the condition that for given cutoffs $\{p_b^c\}_{b,c \in \{1,2\}}$, the demand for each school c is equal to the available supply q_c given by the school's capacity:

$$\begin{aligned} p \cdot (1 - p_1^1 \cdot p_2^1) &= q_1 = q \\ (1 - p) \cdot (1 - p_1^1 \cdot p_2^1) + p_1^1 (p_2^1 - p_2^2) &= q_2 = q. \end{aligned}$$

Any cutoffs $p_1^1 = p_1^2 = x$, $p_2^1 = (1 - q/p)/x$, $p_2^2 = (1 - 2q)x$ with $x \in [1 - q/p, 1]$ solve these equations, but if $x \neq \left(1 - \frac{q}{p}\right)^p$ then the corresponding assignment is different from the TTC assignment. Section 4.2 provides further details as to how the TTC assignment depends on features of the economy that cannot be observed from supply and demand alone. In particular, the TTC cutoffs depend on the relative priority among top-priority students, and not all cutoffs that satisfy supply-demand conditions produce the TTC assignment.

3.3 Consistency with the Discrete TTC Model

In this section we first show that any discrete economy can be translated into a continuum economy, and that the cutoffs obtained using Theorem 2 on this continuum economy give the same assignment as discrete TTC. This demonstrates that the continuum TTC model generalizes the standard discrete TTC model. We then show that the TTC assignment changes smoothly with changes in the underlying economy.

To represent a discrete economy $E = (\mathcal{C}, \mathcal{S}, \succ_c, \succ^S, q)$ by a continuum economy $\Phi(E) = (\mathcal{C}, \Theta, \eta, \frac{q}{N})$, we construct a measure η over Θ by placing a mass at (\succ^s, r^s) for each student s . To ensure the measure has a bounded density, we spread the mass of each student s over a small region $I^s = \{\theta \in \Theta \mid \succ^\theta = \succ^s, r^\theta \in [r_c^s, r_c^s + \frac{1}{N}) \ \forall c \in \mathcal{C}\}$ and identify any point $\theta^s \in I^s$ with student s . The following proposition shows that the continuum TTC assigns all $\theta^s \in I^s$ to the same school, which is the assignment of

student s in the discrete model. Moreover, we can directly use the continuum cutoffs for the discrete economy. Further details and a formal definition of the map Φ are in online Appendix D.4.

Proposition 2. *Let $E = (\mathcal{C}, \mathcal{S}, \succ_{\mathcal{C}}, \succ^{\mathcal{S}}, q)$ be a discrete economy, and let $\Phi(E) = (\mathcal{C}, \Theta, \eta, \frac{q}{N})$ be the corresponding continuum economy. Let \mathbf{p} be the cutoffs produced by Theorem 2 for economy $\Phi(E)$. Then the cutoffs \mathbf{p} give the TTC assignment for the discrete economy E , namely,*

$$\mu_{dTTC}(s | E) = \max_{\succ^s} \{c \mid r_b^s \geq p_b^c \text{ for some } b\},$$

and for every $\theta^s \in I^s$ we have that

$$\mu_{dTTC}(s | E) = \mu_{cTTC}(\theta^s | \Phi(E)).$$

In other words, Φ embeds each discrete economy into a continuum economy that represents it, and the TTC cutoffs in the continuum embedding give the same assignment as TTC in the discrete model. This shows that the TTC assignment defined in Theorem 2 provides a strict generalization of the discrete TTC assignment to a larger class of economies. We provide an example of an embedding of a discrete economy in Appendix B.

Next, we show that the continuum economy can also be used to approximate sufficiently similar economies. Formally, we show that the TTC allocations for strongly convergent sequences of economies are also convergent.

Theorem 3. *Consider two continuum economies $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ and $\tilde{\mathcal{E}} = (\mathcal{C}, \Theta, \tilde{\eta}, q)$, where the measures η and $\tilde{\eta}$ have total variation distance ε . Suppose also that both measures have full support. Then the TTC allocations in these two economies differ on a set of students of measure $O(\varepsilon|C|^2)$.*

In Section 4.2, we show that changes to the priorities of a set of high priority students can affect the final assignment of other students in a non-trivial manner. This raises the question of what the magnitude of these effects are, and whether the TTC mechanism is robust to small perturbations in student preferences or school priorities. Our convergence result implies that the effects of perturbations are proportional to the total variation distance of the two economies, and suggests that the TTC mechanism is fairly robust to small perturbations in preferences.

3.4 Proper budget sets

The standard definition for a student's budget set is the set of schools she can be assigned to by reporting some preference to the mechanism. Specifically, let $[E_{-s}; \succ']$ denote the economy where student s changes her report from \succ^s to \succ' , and let

$$B^*(s \mid E) \stackrel{def}{=} \bigcup_{\succ'} \mu_{dTTC}(s \mid [E_{-s}; \succ'])$$

denote the set of schools that student s can attain by some reported preference, holding fixed the reports of other students. Note that s cannot misreport her priority.

We observed in Section 2.4 that in the discrete model the budget set $B(s, \mathbf{p})$ produced by cutoffs $\mathbf{p} = \mathbf{p}(E)$ generated by Theorem 1 do not necessarily correspond to the set $B^*(s \mid E)$. The analysis in this section can be used to show that the budget sets $B^*(s \mid E)$ correspond to the budget sets $B(s, \mathbf{p}^*)$ for appropriate cutoffs \mathbf{p}^* .

Proposition 3. *Let $E = (\mathcal{C}, \mathcal{S}, \succ^{\mathcal{S}}, \succ^{\mathcal{C}}, q)$ be a discrete economy, and let*

$$\mathcal{P}(E) = \left\{ \mathbf{p} \mid p_b^c = \gamma_b \left(t^{(c)} \right) \text{ where } \gamma(\cdot), t^{(c)} \text{ satisfy trade balance and capacity for } \Phi(E) \right\}$$

be the set of all cutoffs that can be generated by some TTC path $\gamma(\cdot)$ and stopping times $\{t^{(c)}\}_{c \in \mathcal{C}}$. Then

$$B^*(s \mid E) = \bigcap_{\mathbf{p} \in \mathcal{P}(E)} B(s; \mathbf{p}).$$

Moreover, there exists $\mathbf{p}^ \in \mathcal{P}(E)$ such that for every student s*

$$B^*(s \mid E) = B(s; \mathbf{p}^*).$$

Proposition 3 allows us to construct proper budget sets for each agent that determine not only their assignment given their current preferences, but also their assignment given any other submitted preferences. This particular budget set representation of TTC makes it clear that it is strategyproof. In the appendix we prove Proposition 3 and constructively find \mathbf{p}^* .

4 Applications

4.1 Optimal Investment in School Quality

In this section, we explore how to invest in school quality when students are assigned through the TTC mechanism. School financing has been subject to major reforms, and empirical evidence suggests that increased financing has substantial impact on school quality (Hoxby, 2001; Cellini, Ferreira, and Rothstein, 2010; Jackson, Johnson, and Persico, 2016; Lafortune, Rothstein, and Schanzenbach, 2016). Under school choice, changes in school quality will affect student preferences over schools, and therefore change the assignment of students to schools. This may have a negative welfare effect, as schools that become popular will be excluded from some students' budget sets. Under heterogeneous preferences (Hastings, Kane, and Staiger, 2009; Abdulkadiroğlu, Agarwal, and Pathak, 2015) welfare depends on whether students can choose a school for which they have an idiosyncratically high preference. Observing students' budget sets allows us to track the welfare generated by student choices along horizontal dimensions.

We first provide more general comparative statics demonstrating how an increase in school quality affects the TTC assignment. We then examine the question of optimal investment in school quality under a stylized model. Omitted proofs and derivations can be found in online Appendix E.1.

Model with quality dependent preferences and comparative statics

We first enrich our model from Section 3 to allow student preferences to depend on school quality investments. An economy with quality dependent preferences is given by $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$, where $\mathcal{C} = \{1, 2, \dots, n\}$ is the set of schools and \mathcal{S} is the set of student types. A student $s \in \mathcal{S}$ is given by $s = (u^s(\cdot | \cdot), r^s)$, where $u^s(c | \delta)$ is the utility of student s for school c given the quality of each school $\delta = \{\delta_c\}$ and r^s_c is the student's rank at school c . We assume $u^s(c | \cdot)$ is differentiable, increasing in δ_c and non-increasing in δ_b for any $b \neq c$. The measure η over \mathcal{S} specifies the distribution of student types. School capacities are $q = \{q_c\}$, where $\sum q_c < 1$.

For a fixed quality δ , let η_δ be the induced distribution over Θ , and let $\mathcal{E}_\delta = (\mathcal{C}, \Theta, \eta_\delta, q)$ denote the induced economy.¹⁷ We assume for all δ that η_δ has a Lipschitz

¹⁷To make student preferences strict we arbitrarily break ties in favor of schools with lower indices.

continuous non-negative density ν_δ that is bounded below on its support and depends smoothly on δ . For a given δ , let μ_δ and $\{p_b^c(\delta)\}_{c \in \mathcal{C}}$ denote the TTC assignment and associated cutoffs. We omit the dependence on δ when it is clear from context.

When there are two schools, we can specify the direction of change of the TTC cutoffs when we slightly increase δ_ℓ for some $\ell \in \{1, 2\}$. We consider changes that do not change the strict order of school run-out times and without loss of generality assume that schools are numbered in order of their run-out times.

Proposition 4. *Suppose $\mathcal{E} = (\mathcal{C} = \{1, 2\}, \mathcal{S}, \eta, q)$ and δ induces an economy \mathcal{E}_δ such that the TTC cutoffs have a strict runout order $p_2^1(\delta) > p_2^2(\delta)$. Suppose $\hat{\delta}$ has higher school 2 quality $\hat{\delta}_2 \geq \delta_2$, the same quality $\delta_1 = \hat{\delta}_1$, and $\mathcal{E}_{\hat{\delta}}$ has the same runout order, i.e. $p_2^1(\hat{\delta}) \geq p_2^2(\hat{\delta})$. When we change from δ to $\hat{\delta}$ the cutoffs $p_b^c(\cdot)$ change as follows:*

- p_1^1 and p_2^1 both decrease, i.e., it becomes easier to trade into school 1; and
- p_2^2 increases, i.e. higher 2-priority is required to get into school $\ell = 2$.

Figure 3 illustrates the effect of improving the quality of school $\ell = 2$ when $\mathcal{C} = \{1, 2\}$. As in Hatfield, Kojima, and Narita (2016), there may be low 2-rank students who will gain assignment to school 2 after the quality change because of the decrease in p_1^2 . Notice also that small changes in the cutoffs can result in individual students' budget sets growing or shrinking by more than one school. In general, if the TTC cutoffs change slightly then for every pair of schools $b \neq c$ there will be students whose budget sets switch between $\mathcal{C}^{(b)}$ and $\mathcal{C}^{(c)}$.

In general, when $n \geq 3$, increasing the quality of a school ℓ can have non-monotone effects on the cutoffs, and it is not possible to specify the direction of change of the cutoffs p_c^b . However, with additional structure we can give more descriptive comparative statics. For example, consider the logit economy where students' utilities for each school c are randomly distributed as a logit with mean δ_c , independently of priorities and utilities for other schools. That is, utility for school c is given by $u^s(c | \delta) = \delta_c + \varepsilon_{cs}$ with η chosen so ε_{cs} are i.i.d. EV shifted to have mean 0 (McFadden, 1973). Schools have uncorrelated uniform priorities over the students. This model allows us to capture a fixed utility term δ_c that can be impacted by investment together with heterogeneous idiosyncratic taste shocks. The following proposition shows that under the logit economy we have closed form expressions for the TTC cutoffs. This allows us to describe the comparative statics.

We assume the utility of being unassigned is $-\infty$, so all students find all schools acceptable.




Figure 3: The effect of an increase in the quality of school 2 on TTC cutoffs and budget sets. Dashed lines indicate initial TTC cutoffs, and dotted lines indicate TTC cutoffs given increased school 2 quality. The cutoffs $p_1^1 = p_1^2$ and p_2^1 decrease and the cutoff p_2^2 increases. Students in the colored sections receive different budget sets after the increase. Students in dark blue improve to a budget set of $\{1, 2\}$ from \emptyset , students in light blue improve to $\{1, 2\}$ from $\{2\}$, and students in red have an empty budget set \emptyset after the change and $\{2\}$ before.

Proposition 5. Under the logit economy with fixed qualities δ the TTC cutoffs p_b^c for $b \geq c$ are given by¹⁸

$$p_b^c = \left(\frac{\prod_{c' < c} p_{c'}^{c-1} - \rho_c \pi_c}{\prod_{c' < c} p_{c'}^c} \right)^{\pi^{b|c}} \quad (3)$$

where $\pi^{b|c}$ is the probability that a student chooses school b given budget set $\mathcal{C}^{(c)}$, $\rho_c = \frac{q_c}{e^{\delta_c}} - \frac{q_{c-1}}{e^{\delta_{c-1}}}$ is the relative residual capacity for school c , $\pi_c = \sum_{c' \geq c} e^{\delta_{c'}}$ normalizes ρ_c for when the set of available schools is $\mathcal{C}^{(c)}$, and the schools are indexed in the run-out order $\frac{q_1}{e^{\delta_1}} \leq \frac{q_2}{e^{\delta_2}} \leq \dots \leq \frac{q_n}{e^{\delta_n}}$. Moreover, p_b^c is decreasing in δ_ℓ for $c < \ell$ and increasing in δ_ℓ for $b > c = \ell$.

Figure 4 illustrates how the TTC cutoffs change with an increase in the quality of school ℓ . Using equation (3), we derive closed form expressions for $\frac{dp_b^c}{d\delta_\ell}$, which can be found in online Appendix E.1.

¹⁸To simplify notation, when $c = 1$ we let $\prod_{c' < c} p_{c'}^{c-1} = 1$.




Figure 4: The effects of changing the quality δ_ℓ of school ℓ on the TTC cutoffs p_b^c under the logit economy. If $c < \ell$ then $\frac{dp_b^c}{d\delta_\ell} < 0$ for all $b \geq c$, so it becomes easier to get into the more popular schools. If $c > \ell$ then $\frac{dp_b^c}{d\delta_\ell} = 0$. If $c = \ell$ then $\frac{dp_b^c}{d\delta_\ell} = \frac{dp_b^c}{d\delta_\ell} > 0$ for all $b > \ell$, and p_ℓ^ℓ may increase or decrease depending on the specific problem parameters. Note that although p_b^c and p_ℓ^c look aligned in the picture, in general it does not hold that $p_b^c = p_\ell^c$ for all b .

Optimal investment in school quality

Consider a social planner who selects quality levels δ for schools in economy \mathcal{E} . We suppose that the social planner wishes to assign students to schools at which they attain high utility, and for simplicity also assume the social planner wishes to maximize social welfare. For a given assignment μ , the social welfare is given by

$$U(\delta) = \int_{s \in S} u^s(\mu(s) | \delta) d\eta.$$

To illustrate the effects of choice when using TTC, we first consider investment under neighborhood assignment μ_{NH} , which assigns each student to a fixed school regardless of quality and preferences. We assume this assignment fills the capacity of each school. Social welfare for the logit economy is

$$U_{NH}(\delta) = \sum_c q_c \cdot \delta_c,$$

because $\mathbb{E}[\varepsilon_{\mu(s)s}] = 0$ under neighborhood assignment. Under neighborhood assignment, the marginal welfare gain from increasing δ_ℓ is $\frac{dU_{NH}}{d\delta_\ell} = q_\ell$, as an increase in the school quality benefits each of the q_ℓ students assigned to school ℓ .

When the assignment is determined by TTC we need to use the budget set formulation of TTC to capture student welfare.¹⁹ Assume the schools are indexed according to the run-out order given by some fixed δ . A student who is offered the budget set $\mathcal{C}^{(c)} = \{c, \dots, n\}$ is assigned to the school $\ell = \arg \max_{b \in \mathcal{C}^{(c)}} \{\delta_b + \varepsilon_{bs}\}$, and the logit distribution implies that their utility is $U^c = \ln \left(\sum_{b \geq c} e^{\delta_b} \right)$ (Small and Rosen, 1981). Let N^c be the mass of agents with budget set $\mathcal{C}^{(c)}$. Then social welfare under the TTC assignment given quality δ simplifies to

$$U_{TTC}(\delta) = \sum_c N^c \cdot U^c.$$

This expression for welfare also allows for a simple expression for the marginal welfare gain from increasing δ_ℓ under TTC.

Proposition 6. *For the logit economy, the increase in social welfare $U_{TTC}(\delta)$ under TTC from a marginal increase in δ_ℓ is given by*

$$\frac{dU_{TTC}}{d\delta_\ell} = q_\ell + \sum_{c \leq \ell+1} \frac{dN^c}{d\delta_\ell} \cdot U^c.$$

Under neighborhood assignment $\frac{dU_{NH}}{d\delta_\ell} = q_\ell$.

Proposition 6 shows that a marginal increase $d\delta_\ell$ in the quality of school ℓ will have two effects. It will change the utility of the q_ℓ students assigned to ℓ by $d\delta_\ell$, which is the same effect as under neighborhood assignment. In addition, the quality increase changes student preferences, and therefore changes the assignment. The second term captures the additional welfare effect of changes in the assignment by looking at the change in the number of students offered each budget set. The additional term can be negative so that $\frac{dU_{TTC}}{d\delta_\ell} < q_\ell = \frac{dU_{NH}}{d\delta_\ell}$, because an increase in the quality of a school can lead to less efficient sorting of students to schools.

In particular, when there are two schools $\mathcal{C} = \{1, 2\}$ with $q_1 = q_2$ and $\delta_1 \geq \delta_2$ we

¹⁹The expected utility of student s assigned to school $\mu(s)$ depends on the student's budget $B(s, \mathbf{p})$ because of selection on taste shocks. Namely, $\mathbb{E}[u^s(\mu(s) | \delta)] = \delta_{\mu(s)} + \mathbb{E}[\varepsilon_{\mu(s)s} | \delta_{\mu(s)} + \varepsilon_{\mu(s)s} \geq \delta_c + \varepsilon_{cs} \ \forall c \in B(s, \mathbf{p})]$

have that

$$\begin{aligned}\frac{dU_{TTC}}{d\delta_1} &= q_1 + \frac{dN^1}{d\delta_1} \cdot U^1 + \frac{dN^2}{d\delta_1} \cdot U^2 = q_1 + \frac{dN^1}{d\delta_1} \cdot (U^1 - U^2) \\ &= q_1 - (q_1 \cdot e^{\delta_2 - \delta_1}) (\ln(e^{\delta_1} + e^{\delta_2}) - \delta_2) < q_1,\end{aligned}$$

where we use that $N^c = \left(\frac{q_c}{e^{\delta_c}} - \frac{q_{c-1}}{e^{\delta_{c-1}}} \right) (\sum_{b \geq c} e^{\delta_b})$.

An increase in the quality of the higher quality school 1 gives higher utility for students assigned to 1, which is captured by the first term. Additionally, it causes some students to switch their preferences to $1 \succ 2$, making school 1 run out earlier in the TTC algorithm, and removing school 1 from the budget set of some students. Students whose budget set did not change and who switched to $1 \succ 2$ are almost indifferent between the schools and hence unaffected. Students who lost school 1 from their budget set can prefer school 1 by a large margin, and incur significant loss. Thus there is a total negative effect from the changes in the assignment, which is captured by the second term, and the derivative is smaller than under neighborhood assignment. Improving the quality of school 1 when $\delta_1 \leq \delta_2$ will have the opposite effect, as it enlarges student budget sets. Specifically, $\frac{dU_{TTC}}{d\delta_1} = q_1 + q_2 \cdot e^{\delta_1 - \delta_2} (\ln(e^{\delta_1} + e^{\delta_2}) - \delta_1) > q_1$ which is larger than under neighborhood assignment. Note that, holding δ_2 fixed, the function $U_{TTC}(\delta_1)$ has a kink at $\delta_1 = \delta_2$.

We now provide an illustrative example of optimal investment with quality constraints under DA, TTC and neighborhood assignment.

Example 3. Consider a logit economy with two schools and $q_1 = q_2 = \frac{3}{8}$, and let $Q = q_1 + q_2$ denote the total capacity. The planner is constrained to choose quality levels δ such that $\delta_1 + \delta_2 = 2$ and $\delta_1, \delta_2 \geq 0$.

Under neighborhood assignment $U_{NH}/Q = 1$ for any choice of δ_1, δ_2 . Under TTC the unique optimal quality is $\delta_1 = \delta_2 = 1$, yielding $U_{TTC}/Q = 1 + \mathbb{E}[\max(\varepsilon_{1s}, \varepsilon_{2s})] = 1 + \ln(2) \approx 1.69$. This is because any assigned student has the budget set $B = \{1, 2\}$ and is assigned to the school for which he has higher idiosyncratic taste. Welfare is lower when $\delta_1 \neq \delta_2$, because fewer students choose the school for which they have higher idiosyncratic taste. For instance, given $\delta_1 = 2, \delta_2 = 0$ welfare is $U_{TTC}/Q = \frac{1}{2} (1 + e^{-2}) \log(1 + e^2) \approx 1.20$. Under Deferred Acceptance (DA) the unique optimal quality is also $\delta_1 = \delta_2 = 1$, yielding $U_{DA}/Q = 1 + \frac{1}{3} \ln(2) \approx 1.23$. This is strictly lower than the welfare under TTC because under DA only students that have sufficiently

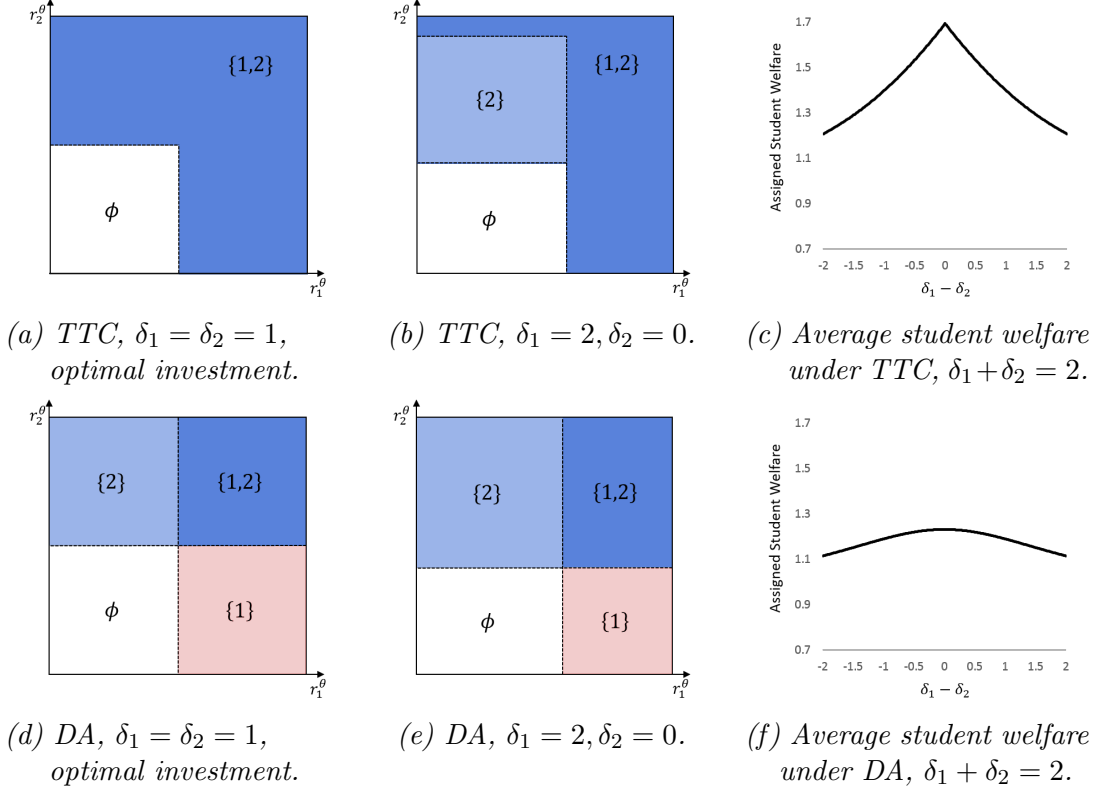


Figure 5: Illustration for Example 3. Figures (a) and (b) show the budget sets under TTC for different quality levels, and Figure (c) shows the average welfare of assigned students under TTC for quality levels $\delta_1 + \delta_2 = 2$ for different values of $\delta_1 - \delta_2$. Figures (d) and (e) show the budget sets under DA, and Figure (f) shows the average welfare of assigned students under DA.

high priority for both schools have the budget set $B = \{1, 2\}$. The remaining assigned students have a budget set $B = \{1\}$ or $B = \{2\}$, corresponding to the single school for which they have sufficient priority. If $\delta_1 = 2, \delta_2 = 0$ welfare under DA is $U_{DA}/Q \approx 1.11$.

TTC yields higher student welfare by providing all assigned students with a full budget set, thus maximizing each assigned student's contribution to welfare from horizontal taste shocks. However, the assignment it produces is not stable. In fact, both schools admit students whom they rank at the bottom, and thus virtually all unassigned students can potentially block with either school.²⁰ Example 3 shows that requiring a stable assignment will constrain many students from efficiently sorting on horizontal taste shocks.

²⁰Note that this may not be a concern in a school choice setting where assignments must be authorized by the department of education and blocking pairs cannot deviate and be assigned outside of the mechanism.

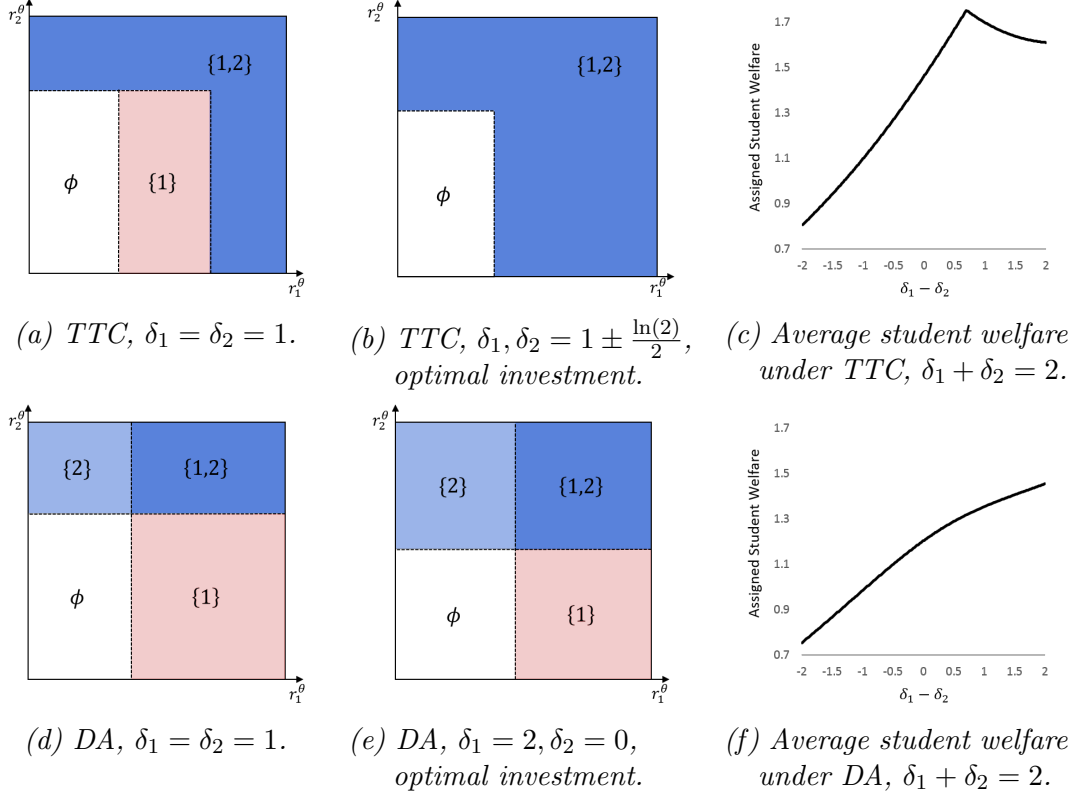


Figure 6: Illustration for Example 4. Figures (a) and (b) show the budget sets under TTC for different quality levels, and Figure (c) shows the average welfare of assigned students under TTC for quality levels $\delta_1 + \delta_2 = 2$ for different values of $\delta_1 - \delta_2$. Note that $\delta_1 = \delta_2 = 1$ is no longer optimal. Figures (d) and (e) show the budget sets under DA , and Figure (f) shows the average welfare of assigned students under DA .

We next provide an example where one school has larger capacity. Investment in the larger school yields more direct benefit as it effects more students, but balancing investments in both schools can yield larger budget sets for more students, leading to more welfare from horizontal taste shocks.

Example 4. Consider a logit economy with two schools and $q_1 = 1/2, q_2 = 1/4$, and let $Q = q_1 + q_2$ denote the total capacity. The planner is constrained to choose quality levels δ such that $\delta_1 + \delta_2 = 2$ and $\delta_1, \delta_2 \geq 0$.

Under neighborhood assignment the welfare optimal quality is $\delta_1 = 2, \delta_2 = 0$, yielding $U_{NH}/Q = 4/3 \approx 1.33$. Under TTC assignment the unique optimal quality is $\delta_1 = 1 + \frac{1}{2} \ln(2), \delta_2 = 1 - \frac{1}{2} \ln(2)$, yielding $U_{TTC}/Q = \ln\left(\frac{3e}{\sqrt{2}}\right) \approx 1.75$. Under these quality levels any assigned student has the budget set $B = \{1, 2\}$. Given $\delta_1 = 2, \delta_2 = 0$ welfare is $U_{TTC}/Q \approx 1.61$. The quality levels that are optimal in Example 3, namely

$\delta_1 = 1, \delta_2 = 1$, give welfare $U_{TTC}/Q \approx 1.46$. Under DA assignment the unique optimal quality is $\delta_1 = 2, \delta_2 = 0$, yielding $U_{DA}/Q \approx 1.45$. Given $\delta_1 = 1, \delta_2 = 1$ welfare under DA is $U_{DA}/Q \approx 1.20$.

Again we find that the optimal quality under TTC provides all assigned students with a full budget set, while the optimal qualities under neighborhood assignment and DA do not. The optimal quality levels under TTC in Example 4 imply that there is a $2/3$ chance a student prefers school 1, and therefore both schools run out at the same time and all assigned students are offered a choice between both schools. Increasing δ_1 further (and decreasing δ_2) would increase welfare holding the assignment fixed, but would result in worse sorting of students to schools on the horizontal taste shocks.

Finally, consider a central school board with a fixed amount of capital K to invest in the n schools. The cost of quality δ_c is the convex function $\kappa_c(\delta_c) = e^{\delta_c}$.²¹ Using Proposition 6 we solve for optimal investment in school quality. Social welfare is maximized when all assigned students have a full budget set, which occurs when the amount invested in each school is proportional to the number of seats at the school.

Proposition 7. *Social welfare is uniquely maximized when the amount κ_c invested in school c is proportional to the capacity q_c , that is,*

$$\kappa_c(\delta_c) = \frac{q_c}{\sum_b q_b} K$$

and all assigned students θ receive a full budget set, i.e., $B(\theta, \mathbf{p}) = \{1, 2, \dots, n\}$ for all assigned students θ .

Under optimal investment, the resulting TTC assignment is such that every assigned student receives a full budget set and is able to attend their top choice school. More is invested in higher capacity schools, as they provide more efficient investment opportunities, but the investment is balanced across schools to prevent any school from being over-demanded. This allows the TTC mechanism to offer assigned students a choice between all schools.

²¹Note that κ_c is the total school funding. This is equivalent to setting student utility of school c to be $\log(\kappa_c) = \log(\kappa_c/q_c) + \log(q_c)$, which is the log of the per-student funding plus a fixed school utility that is larger for bigger schools. Thus, schools with higher capacity also provide more efficient investment opportunities.

4.2 Design of TTC Priorities

To better understand the role of priorities in the TTC mechanism, we examine how the TTC assignment changes with changes in the priority structure. Notice that any student θ whose favorite school is c and who is within the q_c highest ranked students at c is guaranteed admission to c . In the following example, we consider changes to the relative priority of such highly ranked students and find that these changes can have an impact on the assignment of other students, without changing the assignment of any student whose priority changed.

Example 5. The economy \mathcal{E} has two schools 1, 2 with capacities $q_1 = q_2 = q$, students are equally likely to prefer each school, and student priorities are uniformly distributed on $[0, 1]$ independently for each school and independently of preferences. The TTC algorithm ends after a single round, and the resulting assignment is given by $p_1^1 = p_1^2 = p_2^1 = p_2^2 = \sqrt{1 - 2q}$. The derivation can be found in Appendix E.2.

Consider the set of students $\{\theta \mid r_c^\theta \geq m \ \forall c\}$ for some $m > 1 - q$. Any student in this set is assigned to his top choice, regardless of his rank. Suppose we construct an economy \mathcal{E}' by arbitrarily changing the rank of students within the set, subject to the restriction that their ranks must remain in $[m, 1]$.²² The range of possible TTC cutoffs for \mathcal{E}' is given by $p_1^1 = p_1^2, p_2^1 = p_2^2$ where

$$p_1^1 \in [p, \bar{p}], \quad p_2^2 = \frac{1}{p_1^1} (1 - 2q)$$

for $p = \sqrt{(1 - 2m + 2m^2)(1 - 2q)}$ and $\bar{p} = \sqrt{\frac{1 - 2q}{1 - 2m + 2m^2}}$. Figure 7 illustrates the range of possible TTC cutoffs for \mathcal{E}' and the economy \mathcal{E} for which TTC obtains one set of extreme cutoffs.

Example 5 has several implications. First, it shows that it is not possible to directly compute TTC cutoffs from student demand. The set of cutoffs such that student demand is equal to school capacity (depicted by the grey curve in Figure 7) are the cutoffs that satisfy $p_1^1 = p_1^2, p_2^1 = p_2^2$ and $p_1^1 p_2^2 = 1 - 2q$. Under any of these cutoffs the students in $\{\theta \mid r_c^\theta \geq m \ \forall c\}$ have the same demand, but the resulting TTC outcomes are different. It follows that the mechanism requires more information to determine the assignment. However, Theorem 3 implies that the changes in TTC outcomes are small if $1 - m$ is small.

²²The remaining students still have ranks distributed uniformly on the complement of $[1 - r, 1]^2$.




Figure 7: The range of possible TTC cutoffs in example 5 with $q = 0.455$ and $m = 0.6$. The points depict the TTC cutoffs for the original economy and the extremal cutoffs for the set of possible economies \mathcal{E}' , with the range of possible TTC cutoffs for \mathcal{E}' given by the bold curve. The dashed line is the TTC path for the original economy. The shaded squares depict the changes to priorities that generate the economy $\bar{\mathcal{E}}$ which has extremal cutoffs. In $\bar{\mathcal{E}}$ the priority of all top ranked students is uniformly distributed within the smaller square. The dotted line depicts the TTC path for $\bar{\mathcal{E}}$, which results in cutoffs $p_1^1 = \sqrt{\frac{1-2q}{1-2m+2m^2}} \approx 0.42$ and $p_2^2 = \sqrt{(1-2q)(1-2m+2m^2)} \approx 0.22$.

A second implication is that the TTC priorities can be ‘bossy’ in the sense that changes in the relative priority of high priority students can affect the assignment of other students, even when all high priority students receive the same assignment. Notice that in all the economies considered in Example 5, we only changed the relative priority within the set $\{\theta \mid r_c^\theta \geq m \exists c\}$, and all these students were always assigned to their top choice. However, these changes resulted in a different assignment for low priority students. For example, if $q = 0.455$ and $m = 0.4$, a student θ with priority $r_1^\theta = 0.35, r_2^\theta = 0.1$ could possibly receive his first choice or be unassigned depending on the choice of \mathcal{E}' . Such changes to priorities may naturally arise when there are many indifferences in student priorities, and tie-breaking is used. Since priorities are bossy, the choice of tie-breaking between high-priority students can have indirect effects on the assignment of low priority students.

4.3 Comparing Mechanisms

Both TTC and Deferred Acceptance (DA) (Gale and Shapley, 1962) are strategyproof, but differ in that TTC is efficient whereas DA is stable. In theory, the choice between the mechanisms requires a trade-off between efficiency and stability (this trade-off is evident in Example 3). Kesten (2006); Ehlers and Erdil (2010) show

the two mechanisms are equivalent only under strong conditions that are unlikely to hold in practice. However, Pathak (2016) evaluates the two mechanisms on application data from school choice in New Orleans and Boston, and reports that the two mechanisms produced similar outcomes. In Section 4.1 we compare DA and TTC in terms of welfare and assignment and find potentially large differences. Pathak (2016) conjectures that the neighborhood priority used in New Orleans and Boston led to correlation between student preferences and school priorities that may explain the similarity between the TTC and DA allocations.

To study this conjecture, we consider a simple model with neighborhood priority.²³ There are n neighborhoods, each with one school and a mass q of students. Schools have capacities $q_1 \leq \dots \leq q_n = q$, and each school gives priority to students in their neighborhood. For each student, the neighborhood school is their top ranked choice with probability α ; otherwise the student ranks the neighborhood school in position k drawn uniformly at random from $\{2, 3, \dots, n\}$. Student preference orderings over non-neighborhood schools are drawn uniformly at random.

We find that the proportion of students whose assignments are the same under both mechanisms scales linearly with the probability of preference for the neighborhood school α , supporting the conjecture of Pathak (2016).

Proposition 8. *The proportion of students who have the same assignments under TTC and DA is given by*

$$\alpha \frac{\sum_i q_i}{nq}.$$

Proof. We use the methodologies developed in Section 3.2 and in Azevedo and Leshno (2016) to find the TTC and DA allocations respectively. Students with priority are given a lottery number uniformly at random in $[\frac{1}{2}, 1]$, and students without priority are given a lottery number uniformly at random in $[0, \frac{1}{2}]$, where lottery numbers at different schools are independent. For all values of α , the TTC cutoffs are given by $p_j^i = p_i^j = 1 - \frac{q_i}{2q}$ for all $i \leq j$, and the DA cutoffs are given by $p_i = 1 - \frac{q_i}{2q}$. The derivations of the cutoffs can be found in Appendix E.3.

The students who have the same assignments under TTC and DA are precisely the students at neighborhood i whose ranks at school i are above $1 - \frac{q_i}{2q}$, and whose

²³Che and Tercieux (2015b) also show that when there are a large number of schools with a single seat per school and preferences are random both DA and TTC are asymptotically efficient and stable and give asymptotically equivalent allocations. As Example 3 shows, these results do not hold when there are many students and a few large schools.

first choice school is their neighborhood school. This set of students comprises an $\alpha \frac{\sum_i q_i}{nq}$ fraction of the entire student population, which scales proportionally with the correlation between student preferences and school priorities. \square

We can also compare TTC with the Clinch and Trade (C&T) mechanism introduced by Morrill (2015b). The C&T mechanism identifies students who are guaranteed admission to their favorite school c by having priority $r_c^\theta \geq 1 - q$ and assigns them to c by ‘clinching’ without trade. Morrill (2015b) gives an example where the C&T assignment has fewer blocking pairs than the TTC assignment. The fact that allowing students to clinch can change the assignment can be interpreted as another example for the bossiness of priorities under TTC: we can equivalently implement C&T by running TTC on a changed priority structure where students who clinched at school c have higher rank at c than any other student.²⁴ The following proposition builds on Example 5 and shows that C&T may produce more blocking pairs than TTC.

Proposition 9. *The Clinch and Trade mechanism can produce more, fewer or an equal number of blocking pairs compared to TTC.*

Proof. Morrill (2015b) provides an example where C&T produces fewer blocking pairs than TTC. Both mechanism give the same assignment for the symmetric economy in the beginning of Example 5. It remains to construct an economy \mathcal{E}_1 for which C&T produces more blocking pairs than TTC. Economy \mathcal{E}_1 is the same as $\bar{\mathcal{E}}$, except that school 2 rank is redistributed among students with $r_2^\theta \leq \underline{p}$ so that students with $r_1^\theta \geq \bar{p}$ have higher school 2 rank.²⁵ The C&T assignment for \mathcal{E}_1 is given by $p_1^1 = p_2^2 = 0.3$, while TTC gives $p_1^1 = \bar{p}$ and $p_2^2 = \underline{p}$ (and under both $p_1^1 = p_1^2$, $p_2^1 = p_2^2$). Under TTC unmatched students will form blocking pairs only with school 2, while under C&T all unmatched students will form a blocking pair with either school. See Figure 13 for an illustration. \square

²⁴For brevity, we abstract away from certain details of C&T mechanism that are important when not all schools run out at the same round.

²⁵Specifically, select $\ell_1 < \ell_2$. Among students with $r_2^\theta \leq \underline{p}$ and $r_1^\theta \geq \bar{p}$ the school 2 rank is distributed uniformly in the range $[\ell_2, \underline{p}]$. Among students with $r_2^\theta \leq \underline{p}$ and $r_1^\theta < \bar{p}$ the school 2 rank is distributed uniformly in the range $[0, \ell_1]$. Within each range r_1^θ and r_2^θ are still independent. See Figure 13 for an illustration.

5 Discussion

We can simplify how the TTC outcome is communicated to students and their families by using the cutoff characterization. The cutoffs $\{p_b^c\}$ are calculated in the course of running the TTC algorithm. The cutoffs can be published to allow parents to verify their assignment, or the budget set structure can be communicated using the language of tokens (see footnote 4). We hope that these methods of communicating TTC will make the mechanism more palatable to students and their parents, and facilitate a more informed comparison with the Deferred Acceptance mechanism, which also has a cutoff structure.

Examples provided in the paper utilized functional form assumptions to gain tractability. The methodology can be used more generally with numerical solvers. This provides a useful alternative to simulation methods that can be more efficient for large economies, or calculating an average outcome for large random economies. For example, most school districts uses tie-breaking rules, and current simulation methods perform many draws of the random tie-breaking lottery to calculate the expected outcomes. Our methodology directly calculates the expected outcome from the distribution. In Section 4.2 we characterize all the possible TTC outcomes for a class of tie-breaking rules, and find that the choice of tie-breaking rule can have significant effect on the assignment. We leave the problem of determining the optimal choice of tie-breaking lottery for future research.

The model assumes for simplicity that all students and schools are acceptable. It can be naturally extended to allow for unacceptable students or schools by erasing from student preferences any school that they find unacceptable or that finds them unacceptable. Type-specific quotas can be incorporated, as in Abdulkadiroğlu and Sönmez (2003), by adding type-specific capacity equations and erasing from the preference list of each type all the schools which do not have remaining capacity for their type.

References

ABDULKADIROĞLU, A., N. AGARWAL, AND P. A. PATHAK (2015): “The welfare effects of coordinated assignment: Evidence from the NYC HS match,” Discussion paper, National Bureau of Economic Research.

- ABDULKADIROĞLU, A., J. D. ANGRIST, Y. NARITA, AND P. A. PATHAK (Forthcoming): “Research design meets market design: Using centralized assignment for impact evaluation,” *Econometrica*.
- ABDULKADIROĞLU, A., Y.-K. CHE, P. A. PATHAK, A. E. ROTH, AND O. TERCIEUX (2017): “Minimizing Justified Envy in School Choice: The Design of New Orleans’ OneApp,” Discussion paper, National Bureau of Economic Research.
- ABDULKADIROĞLU, A., Y.-K. CHE, AND Y. YASUDA (2015): “Expanding choice in school choice,” *American Economic Journal: Microeconomics*, 7(1), 1–42.
- ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2009): “Strategy-proofness versus efficiency in matching with indifference: Redesigning the NYC high school match,” *The American Economic Review*, pp. 1954–1978.
- ABDULKADIROĞLU, A., P. A. PATHAK, A. E. ROTH, AND T. SÖNMEZ (2005): “The Boston public school match,” *American Economic Review*, pp. 368–371.
- ABDULKADIROĞLU, A., AND T. SÖNMEZ (2003): “School choice: A mechanism design approach,” *The American Economic Review*, 93(3), 729–747.
- AGARWAL, N., AND P. SOMAINI (2014): “Demand Analysis using Strategic Reports: An application to a school choice mechanism,” Discussion paper, National Bureau of Economic Research.
- ALCALDE-UNZU, J., AND E. MOLIS (2011): “Exchange of indivisible goods and indifference: The top trading absorbing sets mechanisms,” *Games and Economic Behavior*, 73(1), 1–16.
- ASHLAGI, I., AND P. SHI (2015): “Optimal allocation without money: An engineering approach,” *Management Science*, 62(4), 1078–1097.
- AVERY, C., AND J. LEVIN (2010): “Early Admissions at Selective Colleges,” *American Economic Review*, 100(5), 2125–56.
- AZEVEDO, E. M., AND J. W. HATFIELD (2015): “Existence of equilibrium in large matching markets with complementarities,” .
- AZEVEDO, E. M., AND J. D. LESHNO (2016): “A supply and demand framework for two-sided matching markets,” *Journal of Political Economy*.

- CELLINI, S. R., F. FERREIRA, AND J. ROTHSTEIN (2010): “The value of school facility investments: Evidence from a dynamic regression discontinuity design,” *The Quarterly Journal of Economics*, 125(1), 215–261.
- CHE, Y.-K., J. KIM, AND F. KOJIMA (2013): “Stable matching in large economies,” Discussion paper, mimeo.
- CHE, Y.-K., AND O. TERCIEUX (2015a): “An Analysis of Top Trading Cycles in Two-Sided Matching Markets,” .
- (2015b): “Efficiency and stability in large matching markets,” .
- DUR, U., AND T. MORRILL (2016): “Competitive equilibria in school assignment,” .
- DUR, U. M. (2012): “A characterization of the top trading cycles mechanism for the school choice problem,” *Available at SSRN 2147449*.
- EHLERS, L. (2014): “Top trading with fixed tie-breaking in markets with indivisible goods,” *Journal of Economic Theory*, 151, 64–87.
- EHLERS, L., AND A. ERDIL (2010): “Efficient assignment respecting priorities,” *Journal of Economic Theory*, 145(3), 1269–1282.
- GALE, D., AND L. S. SHAPLEY (1962): “College admissions and the stability of marriage,” *The American Mathematical Monthly*, 69(1), 9–15.
- HAKIMOV, R., AND O. KESTEN (2014): “The equitable top trading cycles mechanism for school choice,” Discussion paper, WZB discussion paper.
- HASTINGS, J., T. J. KANE, AND D. STAIGER (2009): “Heterogeneous preferences and the efficacy of public school choice,” *NBER Working Paper*, 2145.
- HATFIELD, J. W., F. KOJIMA, AND Y. NARITA (2016): “Improving Schools Through School Choice: A Market Design Approach,” *Journal of Economic Theory*.
- HE, Y., A. MIRALLES, M. PYCIA, J. YAN, ET AL. (2015): “A Pseudo-Market Approach to Allocation with Priorities,” Discussion paper, Toulouse School of Economics (TSE).

- HOXBY, C. M. (2001): “All School Finance Equalizations Are Not Created Equal,” *Quarterly Journal of Economics*, pp. 1189–1231.
- HYLLAND, A., AND R. ZECKHAUSER (1979): “The efficient allocation of individuals to positions,” *The Journal of Political Economy*, pp. 293–314.
- IIPSC (2017): “Innovation in Public School Choice,” <http://iipsc.org/projects/>, Accessed: 2017-01-11.
- JACKSON, C. K., R. C. JOHNSON, AND C. PERSICO (2016): “The effects of school spending on educational and economic outcomes: Evidence from school finance reforms,” *The Quarterly Journal of Economics*, 131(1), 157–218.
- JARAMILLO, P., AND V. MANJUNATH (2012): “The difference indifference makes in strategy-proof allocation of objects,” *Journal of Economic Theory*, 147(5), 1913–1946.
- KAPOR, A., C. NEILSON, AND S. ZIMMERMAN (2016): “Heterogeneous Beliefs and School Choice,” .
- KARLIN, S., AND H. M. TAYLOR (1975): “A first course in stochastic processes,” *Academic, San Diego*.
- KESTEN, O. (2006): “On two competing mechanisms for priority-based allocation problems,” *Journal of Economic Theory*, 127(1), 155–171.
- LAFORTUNE, J., J. ROTHSTEIN, AND D. W. SCHANZENBACH (2016): “School finance reform and the distribution of student achievement,” Discussion paper, National Bureau of Economic Research.
- MA, J. (1994): “Strategy-proofness and the strict core in a market with indivisibilities,” *International Journal of Game Theory*, 23(1), 75–83.
- McFADDEN, D. L. (1973): “Conditional Logit Analysis of Qualitative Choice Behavior,” *Frontiers in Econometrics*, pp. 105–142.
- MIRALLES, A., AND M. PYCIA (2014): “Prices and Efficient Assignments Without Transfers,” .

- MORRILL, T. (2013): “An alternative characterization of top trading cycles,” *Economic Theory*, 54(1), 181–197.
- (2015a): “Making just school assignments,” *Games and Economic Behavior*, 92, 18–27.
- (2015b): “Two simple variations of top trading cycles,” *Economic Theory*, 60(1), 123–140.
- PÁPAI, S. (2000): “Strategyproof assignment by hierarchical exchange,” *Econometrica*, 68(6), 1403–1433.
- PATHAK, P. A. (2016): “What Really Matters in Designing School Choice Mechanisms,” .
- PATHAK, P. A., AND T. SÖNMEZ (2013): “School admissions reform in Chicago and England: Comparing mechanisms by their vulnerability to manipulation,” *The American Economic Review*, 103(1), 80–106.
- PYCIA, M., AND M. U. ÜNVER (2015): “Incentive compatible allocation and exchange of discrete resources,” *Available at SSRN 1079505*.
- SABAN, D., AND J. SETHURAMAN (2013): “House allocation with indifference: a generalization and a unified view,” in *Proceedings of the fourteenth ACM conference on Electronic commerce*, pp. 803–820. ACM.
- SHAPLEY, L., AND H. SCARF (1974): “On cores and indivisibility,” *Journal of mathematical economics*, 1(1), 23–37.
- SMALL, K. A., AND H. S. ROSEN (1981): “Applied Welfare Economics with Discrete Choice Models,” *Econometrica*, 49(1), 105–130.

A Intuition for the Continuum TTC Model

In this section, we provide some intuition for our main results by considering a more direct adaptation of the TTC algorithm to continuum economies. Informally speaking, consider a continuum TTC algorithm in which schools offer seats to their highest priority remaining students, and students are assigned through clearing of trading cycles. This process differs from the discrete TTC algorithm as there is now a *set* of zero measure of highest priority students at each school, and the resulting trading cycles are also within sets of students of zero measure.

There are few challenges in turning this informal algorithm description into a precise definition. First, each cycle is of zero measure, but the algorithm needs to appropriately reduce school capacities as students are assigned. Second, a school will generally offer seats to multiple types of students at once. This implies each school may be involved in multiple cycles at a given point, a type of multiplicity that leads to non-unique TTC allocations in the discrete setting.

To circumvent the challenges above, we define the algorithm in terms of its aggregate behavior over many cycles. Instead of tracing each cleared cycle, we track the state of the algorithm by looking at the fraction of each school's priority list that has been cleared. Instead of progressing by selecting one cycle at a time, we determine the progression of the algorithm by conditions that must be satisfied by any aggregation of cleared cycles. These yield equations (1) and (2), which determine the characterization given in Theorem 2.

A.1 Tracking the State of the Algorithm through the TTC Path γ

Consider some point in time during the run of the discrete TTC algorithm before any school has filled its capacity. While the history of the algorithm up to this point includes all previously cleared trading cycles, in order to run the algorithm, it is sufficient to record only the top priority remaining student at each school. This is because knowing the top remaining student at each school allows us to know exactly which students were previously assigned, and which students remain unassigned. Assigned students are relevant for the remainder of the algorithm only insofar as they reduce the number of seats available. Because all schools have remaining capacity, all assigned students are assigned to their top choice, and we can calculate the remaining

capacity at each school.

To formalize this notion, let τ be some point during the run of the TTC algorithm before any school has filled its capacity. For each school c , let $\gamma_c(\tau) \in [0, 1]$ be the percentile rank of the remaining student with highest c -priority. That is, at time τ in the algorithm each school c is offering a seat to students s for whom $r_c^s = \gamma_c(\tau)$. Let $\gamma(\tau)$ be the vector $(\gamma_c(\tau))_{c \in \mathcal{C}}$. The set of students that have already been assigned at time τ is $\{s \mid r^s \not\leq \gamma(\tau)\}$, because any student s where $r_c^s > \gamma_c(\tau)$ for some c must have already been assigned. Likewise, the set of remaining unassigned students is $\{s \mid r^s \leq \gamma(\tau)\}$. See Figure 8 for an illustration. Since all assigned students were assigned to their top choice, the remaining capacity at school $c \in \mathcal{C}$ is $q_c - |\{s \mid r^s \not\leq \gamma(\tau) \text{ and } Ch^s(\mathcal{C}) = c\}|$. Thus, $\gamma(\tau)$ captures all the information needed for the remainder of the algorithm.




Figure 8: The set of students assigned at time τ is described by the point $\gamma(\tau)$ on the TTC path. Students in the grey region with rank better than $\gamma(\tau)$ are assigned, and students in the white region with rank worse than $\gamma(\tau)$ are unassigned.

This representation can be readily generalized to continuum economies. In the continuum, the algorithm progresses in continuous time. The state of the algorithm at time $\tau \in \mathbb{R}_{\geq}$ is given by $\gamma(\tau) \in [0, 1]^{\mathcal{C}}$, where $\gamma_c(\tau) \in [0, 1]$ is the percentile rank of the remaining students with highest c -priority. By tracking the progression of the algorithm through $\gamma(\cdot)$ we avoid looking at individual trade cycles, and instead track how many students were already assigned from each school's priority list.

A.2 Determining the Algorithm Progression through Trade Balance

The discrete TTC algorithm progresses by finding and clearing a trade cycle. This cycle assigns a set of discrete students; for each involved school c the top student is cleared and $\gamma_c(\cdot)$ is reduced. In the continuum each cycle is infinitesimal, and any change in $\gamma(\cdot)$ must involve many trade cycles. Therefore, we seek to determine the progression of the algorithm by looking at the effects of clearing many cycles.

Suppose at time τ_1 the TTC algorithm has reached the state $x = \gamma(\tau_1)$, where $\gamma(\cdot)$ is differentiable at τ_1 and $\mathbf{d} = -\gamma'(\tau_1) \geq 0$. Let $\varepsilon > 0$ be a small step size, and assume that by sequentially clearing trade cycles the algorithm reaches the state $\gamma(\tau_2)$ at time $\tau_2 = \tau_1 + \varepsilon$. Consider the sets of students offered seats and assigned seats during this time step from time τ_1 to time τ_2 . Let $c \in \mathcal{C}$ be some school. For each cycle, the measure of students assigned to school c is equal to the measure of seats offered²⁶ by school c . Therefore, if students are assigned between time τ_1 and τ_2 through clearing a collection of cycles, then the set of students assigned to school c has the same measure as the set of seats offered by school c . If $\gamma(\cdot)$ and η are sufficiently smooth, the measures of both of these sets can be approximately expressed in terms of $\varepsilon \cdot \mathbf{d}$ and the marginal densities $\{H_b^c(x)\}_{b,c \in \mathcal{C}}$, yielding an equation that determines \mathbf{d} . For the sake of clarity, we omit technical details in the ensuing discussion. A rigorous derivation can be found in online Appendix F.

We first identify the measure of students who were offered a seat at a school b or assigned to a school c during the step from time τ_1 to time τ_2 . If $\mathbf{d} = -\gamma'(\tau_1)$ and ε is sufficiently small, we have that for every school b

$$|\gamma_b(\tau_2) - \gamma_b(\tau_1)| \approx \varepsilon d_b,$$

that is, during the step from time τ_1 to time τ_2 the algorithm clears students with b -ranks between x and $x - \varepsilon d_b$. To capture this set of students, let

²⁶Strictly speaking, the measure of students assigned to each school is equal to the measure of seats at that school which were claimed (not the measure of seats offered). A seat can be offered but not claimed in one of two ways. The first occurs when the seat is offered at time τ but not yet claimed. The second is when a student is offered two or more seats at the same time, and claims only one of them. Both of these sets are of η -measure 0 under our assumptions, and thus the measure of seats claimed is equal to the measure of seats offered.

$$T_b(x, \varepsilon d_b) \stackrel{def}{=} \{\theta \in \Theta \mid r^\theta \leq x, r_b^\theta > x - \varepsilon d_b\}$$

denote the set of students with ranks in this range. For all ε , $T_b(x, \varepsilon d_b)$ is the set of top remaining students at b , and when ε is small, $T_b(x, \varepsilon d_b)$ is approximately the set of students who were offered a seat at school b during the step.²⁷

To capture the set of students that are assigned to a school c during the step, partition the set $T_b(x, \varepsilon d_b)$ according to the top choice of students. Namely, let

$$T_b^c(x, \varepsilon d_b) \stackrel{def}{=} \{\theta \in T_b(x, \varepsilon d_b) \mid Ch^\theta(\mathcal{C}) = c\},$$

denote the top remaining students on b 's priority list whose top choice is school c . Then the set of students assigned to school c during the step is $\cup_a T_a^c(x, \varepsilon d_a)$, the set of students that got an offer from some school $a \in \mathcal{C}$ and whose top choice is c .

We want to equate the measure of the set $\cup_a T_a^c(x, \varepsilon d_a)$ of students who were assigned to c with the measure of the set of students who are offered²⁸ a seat at c , which is approximately the set $T_c(x, \varepsilon d_c)$. By smoothness of the density of η , for sufficiently small δ we have that

$$\eta(T_b^c(x, \delta)) \approx \delta \cdot H_b^c(x).$$

Therefore, we have that²⁹

$$\begin{aligned} \eta(\cup_a T_a^c(x, \varepsilon d_a)) &\approx \sum_{a \in \mathcal{C}} \eta(T_a^c(x, \varepsilon d_a)) \approx \sum_{a \in \mathcal{C}} \varepsilon d_a \cdot H_a^c(x), \\ \eta(T_c(x, \varepsilon d_c)) &= \eta(\cup_a T_c^a(x, \varepsilon d_c)) \approx \sum_{a \in \mathcal{C}} \varepsilon d_c \cdot H_c^a(x). \end{aligned}$$

In sum, if the students assigned during the step from time τ_1 to time τ_2 are cleared via a collection of cycles, we must have the following condition on the gradient $\mathbf{d} = \gamma'(\tau_1)$

²⁷The students in the set $T_b(x, \varepsilon d_b) \cap T_a(x, \varepsilon d_a)$ could have been offered a seat at school a and assigned before getting an offer from school b . However, for small ε the intersection is of measure $O(\varepsilon^2)$ and therefore negligible.

²⁸In the continuum model the set of seats offered but not claimed (or traded) is of η -measure 0.

²⁹These approximations make use of the fact that $\eta(T_b(x, \varepsilon d_b) \cap T_a(x, \varepsilon d_a)) = O(\varepsilon^2)$ for small ε .

of the TTC path,

$$\sum_{a \in \mathcal{C}} \varepsilon d_a \cdot H_a^c(x) \approx \sum_{a \in \mathcal{C}} \varepsilon d_c \cdot H_c^a(x).$$

Formalizing this argument yields the marginal trade balance equations at $x = \gamma(\tau_1)$,

$$\sum_{a \in \mathcal{C}} \gamma'_a(\tau_1) \cdot H_a^c(x) = \sum_{a \in \mathcal{C}} \gamma'_c(\tau_1) \cdot H_c^a(x).$$

A.3 Interpretation of the Trade Balance Equations

The previous subsection showed that any small step clearing a collection of cycles must correspond to a gradient γ' that satisfies the trade balance equations. We next characterize the set of solutions to the trade balance equations and explain why any solution corresponds to clearing a collection of cycles.

Let $\gamma(\tau) = x$, and consider the set of *valid gradients* $\mathbf{d} = -\gamma'(\tau) \geq 0$ that solve the trade balance equations for x

$$\sum_{a \in \mathcal{C}} d_a \cdot H_a^c(x) = \sum_{a \in \mathcal{C}} d_c \cdot H_c^a(x).$$

Consider the following equivalent representation. Construct a graph with a node for each school. Let the weight of node b be d_b , and let the flow from node b to node c be $f_{b \rightarrow c} = d_b \cdot H_b^c(x)$. The flow $f_{b \rightarrow c}$ represents the flow of students who are offered a seat at b and wish to trade it for school c when the algorithm progress down school b 's priority list at rate d_b . Figure 9 illustrates such a graph for $\mathcal{C} = \{1, 2, 3, 4\}$. The node weights \mathbf{d} solve the trade balance equations if and only if the total flow into a node is equal to the total flow out of a node, i.e. if and only if f is a zero-sum flow. Standard arguments from network flow theory show that any zero-sum flow can be decomposed into a collection of cycles. In other words, the algorithm can find a collection of cycles that clears each school c 's priority list at rate d_c if and only if and only if \mathbf{d} is a solution to the trade balance equations.

To characterize the set of solutions we draw on a connection to Markov chains. Consider a continuous time Markov chain over the states \mathcal{C} , and transition rates from state b to state c equal to $H_b^c(x)$. The stationary distributions of the Markov chain are characterized by the balance equations, which state that the total probability flow out of state c is equal to the total probability flow into state c . Mathematically,




Figure 9: Example of a graph representation for the trade balance equations at x . There is an edge from b to c if $H_b^c(x) > 0$. The two communication classes are framed.

these are exactly the trade balance equations. Hence \mathbf{d} is a solution to the trade balance equations if and only if $\mathbf{d}/\|\mathbf{d}\|_1$ is a stationary distribution of the Markov chain.

This connection allows us to fully characterize the set of solutions to the trade balance equations through well known results about Markov chains. We restate them here for completeness. Given a transition matrix P , a recurrent communication class is a subset $K \subseteq \mathcal{C}$, such that the restriction of P to rows and columns with coordinates in K is an irreducible matrix, and $P_c^b = 0$ for every $c \in K$ and $b \notin K$. See Figure 9 for an example. There exists at least one recurrent communication class, and two different communication classes have an empty intersection. Let the set of communicating classes be $\{K_1, \dots, K_\ell\}$. For each communicating class K_i there is a unique vector \mathbf{d}^{K_i} that is a stationary distribution and $\mathbf{d}_c^{K_i} = 0$ for any $c \notin K_i$. The set of stationary distributions of the Markov chain is given by convex combinations of $\{\mathbf{d}^{K_1}, \dots, \mathbf{d}^{K_\ell}\}$.

An immediate implication is that a solution to the trade balance equations always exists. Moreover, if η has full support³⁰ then the TTC path γ is unique (up to rescaling of the time parameter). This is because full support of η implies that the matrix $H(x)$ is irreducible for every x , i.e. there is a single communicating class. Therefore there is a unique (up to normalization) solution $\mathbf{d} = -\gamma'(\tau)$ to the trade balance equations at $x = \gamma(\tau)$ for every x and the path is unique.

Lemma 1. *Let $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ be a continuum economy where η has full support. Then there exists a TTC path γ that is unique up to rescaling of the time parameter t . For $\tau \leq \min_{c \in \mathcal{C}} \{t^{(c)}\}$ we have that $\gamma(\cdot)$ is given by*

$$\frac{d\gamma(t)}{dt} = \mathbf{d}(\gamma(t))$$

³⁰ η has full support if for every open set $A \subset \Theta$ we have $\eta(A) > 0$.

where $\mathbf{d}(x)$ is the solution to the trade balance equations at x , and $\mathbf{d}(x)$ is unique up to normalization.

In general, there can be multiple solutions to the trade balance equations at x , and therefore multiple TTC paths. The Markov chain and recurrent communication class structure give intuition as to why the TTC assignment is still unique. Each solution \mathbf{d}^{K_i} corresponds to the clearing of cycles involving only schools within the set K . The discrete TTC algorithm may encounter multiple disjoint trade cycles, and the outcome of the algorithm is invariant to the order in which these cycles are cleared (when preferences are strict). Similarly here, the algorithm may encounter mutually exclusive combinations of trade cycles $\{\mathbf{d}^{K_1}, \dots, \mathbf{d}^{K_\ell}\}$, which can be cleared sequentially or simultaneously at arbitrary relative rates. Theorem 2 shows that just like the outcome of the discrete TTC algorithm does not depend on the cycle clearing order, the outcome of the continuum TTC algorithm does not depend on the order in which $\{\mathbf{d}^{K_1}, \dots, \mathbf{d}^{K_\ell}\}$ are cleared.

As an illustration, consider the unique solution \mathbf{d}^K for the communicating class $K = \{1, 2\}$, as illustrated in Figure 9. Suppose that at some point x we have $H_1^1(x) = 1/2$, $H_1^2(x) = 1/2$ and $H_2^1(x) = 1$. That is, the marginal mass of top ranked students at either school is 1, all the top marginal students of school 2 prefer school 1, and half of the top marginal students of school 1 prefer school 1 and half prefer school 2. The algorithm offers seats and goes down the school's priority lists, assigning students through a combination of two kinds of cycles: the cycle $1 \circlearrowleft$ where a student is offered a seat at 1 and is assigned to 1, and a cycle $1 \rightleftharpoons 2$ where a student who was offered a seat at 1 trades her seat with a student who was offered a seat at 2. Given the relative mass of students, the cycle $1 \rightleftharpoons 2$ should be twice as frequent as the cycles $1 \circlearrowleft$. Therefore, clearing cycles leads the mechanism to go down school 1's priority list at twice the speed it goes down school 2's list, or $d_1 = 2 \cdot d_2$, which is the unique solution to the trade balance equations at x (up to normalization).

Figure 10 illustrates the path $\gamma(\cdot)$ and the solution $\mathbf{d}(x)$ to the trade balance equations at x . Note that for every x we can calculate $\mathbf{d}(x)$ from $H(x)$. When there are multiple solutions to the trade balance equations at some x , we may select a solution $\mathbf{d}(x)$ for every x such that $\mathbf{d}(\cdot)$ is a sufficiently smooth gradient field. The TTC path $\gamma(\cdot)$ can be generated by starting from $\gamma(0) = \mathbf{1}$ and following the gradient field.




Figure 10: Illustration of the gradient field $\mathbf{d}(\cdot)$ and path $\gamma(\cdot)$ (ignoring the capacity equations).

A.4 When a School Fills its Capacity

So far we have described the progression of the algorithm while all schools have remaining capacity. To complete our description of the algorithm we need to describe how the algorithm detects that a school has exhausted all its capacity, and how the algorithm continues after a school is full.

As long as there is still some remaining capacity, the trade balance equations determine the progression of the algorithm along the TTC path $\gamma(\cdot)$. The mass of students assigned to school c at time τ is

$$D^c(\gamma(\tau)) = \eta(\{\theta \mid r^\theta \not\prec \gamma(\tau), Ch^\theta(\mathcal{C}) = c\}).$$

Because $\gamma(\cdot)$ is continuous and monotonically decreasing in each coordinate, $D^c(\gamma(\tau))$ is a continuous increasing function of τ . Therefore, the first time during the run of the continuum TTC algorithm at which any reached its capacity is given by $t^{(c^*)}$ that solves the capacity equations

$$\begin{aligned} D^{c^*}(\gamma(t^{(c^*)})) &= q_{c^*} \\ D^a(\gamma(t^{(c^*)})) &\leq q_a \quad \forall a \in \mathcal{C} \end{aligned}$$

where c^* is the first school to reach its capacity.

Once a school has filled up its capacity, we can eliminate that school and apply the algorithm to the residual economy. Note that the remainder of the run of the algorithm depends only on the remaining students, their preferences over the remaining schools, and remaining capacity at each school. After eliminating assigned students and schools that have reached their capacity we are left with a residual economy that has strictly fewer schools. To continue the run of the continuum TTC algorithm, we may recursively apply the same steps to the residual economy. Namely, to continue the algorithm after time $t^{(c^*)}$ start the path from $\gamma(t^{(c^*)})$ and continue the path using a gradient that solves the trade balance equations for the residual economy. The algorithm follows this path until one of the remaining schools fills its capacity, and another school is removed.

A.5 Comparison between Discrete TTC and Continuum TTC

Table 1 summarizes the relationship between the discrete and continuum TTC algorithms, and provides a summary of this section. It parallels the objects that define the continuum TTC algorithm with their counterparts in the discrete TTC algorithm. For example, running the continuum TTC algorithm on the embedding $\Phi(E)$ of a discrete economy E performs the same assignments as the discrete TTC algorithm, except that the continuum TTC algorithm performs these assignments continuously and in fractional amounts instead of in discrete steps.

Discrete TTC	→	Continuum TTC	Expression	Equation
Cycle	→	Valid gradient	$\mathbf{d}(x)$	trade balance equations
Algorithm progression	→	TTC path	$\gamma(\cdot)$	$\gamma'(\tau) = \mathbf{d}(\gamma(\tau))$
School removal	→	Stopping times	$t^{(c)}$	capacity equations

Table 1: The relationship between the discrete and continuum TTC processes.

Finally, we note that the main technical content of Theorem 2 is that there always exists a TTC path γ and stopping times $\{t^{(c)}\}$ that satisfy trade balance and capacity, and that these necessary conditions, together with the capacity equations (2), are sufficient to guarantee the uniqueness of the resulting assignment.

B Example: Embedding a discrete economy in the continuum model

Consider the discrete economy $E = (\mathcal{C}, \mathcal{S}, \succ^{\mathcal{S}}, \succ^{\mathcal{C}}, q)$ with two schools and six students, $\mathcal{C} = \{1, 2\}$, $\mathcal{S} = \{a, b, c, u, v, w\}$. School 1 has capacity $q_1 = 4$ and school 2 has capacity $q_2 = 2$. The school priorities and student preferences are given by


$$\begin{aligned} 1 & : a \succ u \succ b \succ c \succ v \succ w, & a, b, c : 1 \succ 2, \\ 2 & : a \succ b \succ u \succ v \succ c \succ w, & u, v, w : 2 \succ 1. \end{aligned}$$

In Figure 11, we display three TTC paths for the continuum embedding $\Phi(E)$ of the discrete economy E . The first path γ_{all} corresponds to clearing all students in recurrent communication classes, that is, all students in the maximal union of cycles in the pointing graph. The second path γ_1 corresponds to taking $K = \{1\}$ whenever possible. The third path γ_2 corresponds to taking $K = \{2\}$ whenever possible. We remark that the third path gives a different first round cutoff point p^1 , but all three paths give the same allocation.


B.1 Calculating the TTC paths

We first calculate the TTC path in the regions where the TTC paths are the same. In the following, we consider only solutions d to the trade balance equations that have been normalized so that $d \cdot \mathbf{1} = -1$. For brevity we call such solutions *valid directions*.


At every point (x_1, x_2) with $\frac{5}{6} < x_1 \leq x_2 \leq 1$ the H matrix is $\begin{bmatrix} x_2 - \frac{5}{6} & 0 \\ x_1 - \frac{5}{6} & 0 \end{bmatrix}$, so $d = [-1, 0]$ is the unique valid direction and the TTC path is defined uniquely for $t \in [0, \frac{1}{6}]$ by $\gamma(t) = (1 - t, 1)$. This section of the TTC path starts at $(1, 1)$ and ends at $(\frac{5}{6}, 1)$. At every point $(\frac{5}{6}, x_2)$ with $\frac{5}{6} < x_2 \leq 1$ the H matrix is $\begin{bmatrix} 0 & \frac{1}{6} \\ 0 & 0 \end{bmatrix}$, so $d = [0, -1]$ is the unique valid direction, and the TTC path is defined uniquely for $t \in [\frac{1}{6}, \frac{1}{3}]$ by $\gamma(t) = (\frac{5}{6}, \frac{7}{6} - t)$. This section of the TTC path starts at $(\frac{5}{6}, 1)$ and ends at $(\frac{5}{6}, \frac{5}{6})$.



TTC path γ_{all} clears all students in recurrent communication classes.



TTC path γ_1 clears all students who want school 1 before students who want school 2.



TTC path γ_2 clears all students who want school 2 before students who want school 1.

Figure 11: Three TTC paths and their cutoffs and allocations for the discrete economy in example B. In each set of two squares, students in the left square prefer school 1 and students in the right square prefer school 2. The first round TTC paths are solid, and the second round TTC paths are dotted. The cutoff points p^1 and p^2 are marked by filled circles. Students shaded light blue are assigned to school 1 and students shaded dark blue are assigned to school 2.

At every point (x_1, x_2) with $\frac{2}{3} < x_1, x_2 \leq \frac{5}{6}$ the H matrix is $\begin{bmatrix} 0 & \frac{1}{6} \\ \frac{1}{6} & 0 \end{bmatrix}$, and so $d = [-\frac{1}{2}, -\frac{1}{2}]$ is the unique valid direction, the TTC path is defined uniquely to lie

on the diagonal $\gamma_1(t) = \gamma_2(t)$, and this section of the TTC path starts at $(\frac{5}{6}, \frac{5}{6})$ and ends at $(\frac{2}{3}, \frac{2}{3})$. At every point $x = (\frac{1}{3}, x_2)$ with $\frac{1}{3} < x_2 \leq \frac{2}{3}$ the H matrix is $\begin{bmatrix} 0 & 6x_2 - 2 \\ 0 & 0 \end{bmatrix}$, and so $d = [0, -1]$ is the unique valid direction, and the TTC path is parallel to the y axis. Finally, at every point $(x_1, \frac{1}{3})$ with $0 < x_1 \leq \frac{2}{3}$, the measure of students assigned to school c_1 is at most 3, and the measure of students assigned to school c_2 is 2, so c_2 is unavailable. Hence, from any point $(x_1, \frac{1}{3})$ the TTC path moves parallel to the x_1 axis.




Figure 12: The valid directions $d(x)$ for the continuum embedding $\Phi(E)$. Valid directions $d(x)$ are indicated for points x in the grey squares (including the upper and right boundaries but excluding the lower and left boundaries), as well as for points x on the black lines. Any vector $d(x)$ is a valid direction in the lower left grey square. The borders of the squares corresponding to the students are drawn using dashed grey lines.

We now calculate the various TTC paths where they diverge.

At every point $x = (x_1, x_2)$ with $\frac{1}{2} < x_1, x_2 \leq \frac{2}{3}$ the H matrix is $\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$ (i.e. there are no marginal students), and so $\tilde{H} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Moreover, at every point $x = (x_1, x_2)$ with $\frac{1}{3} < x_1, x_2 \leq \frac{1}{2}$ the H matrix is $\begin{bmatrix} \frac{1}{6} & 0 \\ 0 & \frac{1}{6} \end{bmatrix}$, and so $\tilde{H} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Also, at every point $x = (x_1, x_2)$ with $\frac{1}{3} < x_1 \leq \frac{1}{2}$ and $\frac{1}{2} < x_2 \leq \frac{2}{3}$, the H matrix is $\begin{bmatrix} \frac{1}{6} & 0 \\ 0 & 0 \end{bmatrix}$ so again $\tilde{H} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. The same argument with the coordinates swapped gives that

$\tilde{H} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ when $\frac{1}{2} < x_1 \leq \frac{2}{3}$ and $\frac{1}{3} < x_2 \leq \frac{1}{2}$. Hence in all these regions, both schools are in their own recurrent communication class, and any vector d is a valid direction.

The first path corresponds to taking $d = [-\frac{1}{2}, -\frac{1}{2}]$, the second path corresponds to taking $d = [-1, 0]$ and the third path corresponds to taking $d = [0, -1]$. The first path starts at $(\frac{2}{3}, \frac{2}{3})$ and ends at $(\frac{1}{3}, \frac{1}{3})$ where school 2 fills. The third path starts at $(\frac{2}{3}, \frac{2}{3})$ and ends at $(\frac{2}{3}, \frac{1}{3})$ where school 2 fills. Finally, when $x = (\frac{1}{3}, x_2)$ with $\frac{1}{3} < x_2 \leq \frac{1}{2}$, the H matrix is $\begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$ and so $d = [0, -1]$ is the unique valid direction, and the second TTC path starts at $(\frac{1}{3}, \frac{1}{2})$ and ends at $(\frac{1}{3}, \frac{1}{3})$ where school 2 fills. All three paths continue until $(0, \frac{1}{3})$, where school 1 fills.

Note that all three paths result in the same TTC allocation, which assigns students a, b, c, w to school 1 and u, v to school 2. All three paths assign the students assigned before p^1 (students a, u, b, c for paths 1 and 2 and a, u, b for path 3) to their top choice school. All three paths assign all remaining students to school 1.

C Proofs for Section 2

C.1 Proof of Theorem 1

For each student s let $B(s, \mathbf{p}) = \{c \mid r_b^s \geq p_b^c \text{ for some } b\}$. It suffices to show that for each student s it holds that $\mu_{dTTC}(s) \in B(s, \mathbf{p})$, and that if $c \in B(s, \mathbf{p})$ then s prefers $\mu_{dTTC}(s)$ to c , i.e. $\mu(s) \succeq^s c$. The former is simple to show, since if we let b be the school such that s traded a seat at school b for a seat at school $\mu_{dTTC}(s)$, then by definition $p_b^{\mu_{dTTC}(s)} \leq r_b^s$ and $\mu_{dTTC}(s) \in B(s, \mathbf{p})$.

Now suppose for the sake of contradiction that $c \in B(s, \mathbf{p})$ and student s strictly prefers c to $\mu_{dTTC}(s)$, i.e. $c \succ^s \mu_{dTTC}(s)$. As $c \in B(s, \mathbf{p})$ there exists a school b' such that $r_{b'}^s \geq p_{b'}^c$. Let s' be the student with rank $r_{b'}^{s'} = p_{b'}^c$ at school b' . (Such a student exists since $p_{b'}^c \leq r_{b'}^s < 1$.) Then by definition student s' traded a seat at school b' , so since $r_{b'}^s \geq p_{b'}^c = r_{b'}^{s'}$ student s is assigned weakly before student s' . Additionally, since $c \succ^s \mu_{dTTC}(s)$ school c must reach capacity before student s is assigned, and so since student s' was assigned to school c student s was assigned strictly before student s . This provides the required contradiction.

C.2 Proof of Proposition 1

Let the schools be indexed such that they reach capacity in the order $1, 2, \dots, |\mathcal{C}|$. If a student s was assigned (strictly) after school $\ell - 1$ reached capacity and (weakly) before school ℓ reached capacity, we say that the student s was assigned in round ℓ .

We define new cutoffs $\{\tilde{p}_b^c\}$ by setting $\tilde{p}_b^c = \min_{c' \leq c} p_{b'}^{c'}$, so that it evidently holds that $\tilde{p}_b^1 \geq \tilde{p}_b^2 \geq \dots \geq \tilde{p}_b^b = \tilde{p}_b^{b+1} = \dots = \tilde{p}_b^n$ for all b . We show that the cutoffs $\{\tilde{p}_b^c\}$ give the same allocation as the cutoffs $\{p_b^c\}$, i.e. for each student s it holds that

$$\max_{\succsim^s} \{c \mid r_b^s \geq \tilde{p}_b^c \text{ for some } b\} = \mu_{dTTC}(s) = \max_{\succsim^s} \{c \mid r_b^s \geq p_b^c \text{ for some } b\}.$$

For each student s let $B(s, \tilde{\mathbf{p}}) = \{c \mid r_b^s \geq \tilde{p}_b^c \text{ for some } b\}$. It suffices to show that for each student s it holds that $\mu_{dTTC}(s) \in B(s, \tilde{\mathbf{p}})$, and that if $c \in B(s, \tilde{\mathbf{p}})$ then s prefers $\mu_{dTTC}(s)$ to c , i.e. $\mu(s) \succeq^s c$. The former is simple to show, since clearly $\tilde{\mathbf{p}} \leq \mathbf{p}$ and so $B(s, \tilde{\mathbf{p}}) \supseteq B(s, \mathbf{p}) \ni \mu_{dTTC}(s)$ (by Theorem 1).

The rest of the proof can be completed in much the same way as the proof of Theorem 1. Suppose for the sake of contradiction that $c \in B(s, \tilde{\mathbf{p}})$ and student s strictly prefers c to $\mu_{dTTC}(s)$, i.e. $c \succ^s \mu_{dTTC}(s)$. As $c \in B(s, \tilde{\mathbf{p}})$ there exists a school b' such that $r_{b'}^s \geq \tilde{p}_{b'}^c$. Let s' be the student with rank $r_{b'}^{s'} = \tilde{p}_{b'}^c$ at school b' . (Such a student exists since $\tilde{p}_{b'}^c \leq r_{b'}^s < 1$.) Then by definition student s' traded a seat at school b' , so since $r_{b'}^s \geq \tilde{p}_{b'}^c = r_{b'}^{s'}$ student s is assigned weakly before student s' . Additionally, since $c \succ^s \mu_{dTTC}(s)$ school c must reach capacity before student s is assigned. Finally, by definition there exists some $c' \leq c$ such that $\tilde{p}_{b'}^c = p_{b'}^{c'}$ and student s' was assigned to school c' , and so student s was assigned weakly before school c reached capacity, and hence strictly before student s . This provides the required contradiction.