

The Cutoff Structure of Top Trading Cycles in School Choice*

Jacob D. Leshno Irene Lo[†]

April 10, 2018

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 not help explain the relationship between student priorities and their eventual assignment.

We show that the TTC assignment can be described by admission cutoffs for each pair of schools. These cutoffs parallel prices in competitive equilibrium, with students' priorities serving the role of endowments. In a continuum model these cutoffs 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 distribution of school quality, and find that a more egalitarian distribution 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”. We thank (in random order) Cyndia Yu, John Tsitsiklis, Yeon-Koo Che, Larry Samuelson, Dilip Ravindran, Eduardo Azevedo, Nikhil Agarwal, Jay Sethuraman, Chiara Margaria, Itai Ashlagi, Utku Unver, 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, Market Design NBER, Matching in Practice workshop, and Cireq Montreal Microeconomics meeting for their helpful comments.

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

1 Introduction

School choice mechanisms are commonly used to determine school admission based on student preferences and school priorities. Informed by advances in school choice theory, many school districts redesigned their school choice mechanisms (Abdulkadiroğlu & Sönmez 2003, Abdulkadiroğlu et al. 2005, Pathak & Sönmez 2013, IIPSC 2017), with most choosing to implement the Deferred Acceptance (DA) mechanism (Gale & Shapley 1962).¹ However, DA is inefficient, in that it may produce assignments that are Pareto dominated for students. For example, students are commonly given priority for their neighborhood schools, and DA may assign two students to their respective neighborhood schools even if both students would prefer to swap their assignments.²

Such inefficiencies are addressed by the Top Trading Cycles (TTC) mechanism (Shapley & Scarf (1974), attributed to David Gale), which was proposed for school choice by Abdulkadiroğlu & Sönmez (2003). In the TTC mechanism each school offers a seat to its highest priority student, and students can trade seats with other students. The mechanism sequentially assigns students by offering seats and finding trading cycles.

While the theoretical literature promotes TTC for being Pareto efficient and strategy-proof for students, there is no guarantee that TTC Pareto improves upon DA. Moreover, while the sequential clearing of trade cycles is simple to state, it results in an opaque mapping between a student’s priorities and their assignment. In particular, a major drawback of the algorithmic description of TTC is that it makes it difficult to discern how a student’s priorities determine their assignment under TTC. This is exacerbated by the fact that priority at a school has different implications under DA and TTC; under TTC (in contrast to DA) it is possible for a student to gain admission to one school by having priority at another school. This means that

¹The only instances where TTC was implemented for school choice are in the San Francisco school district and previously in the New Orleans Recovery School District (Abdulkadiroğlu, Che, Pathak, Roth & Tercieux 2017).

²Such a swap will not harm any other students, but can lead to an assignment that is unstable with respect to the priority structure. While this may allow strategic agents to form blocking pairs in other contexts (such as the NRMP), this is not a concern for many school districts (such as the Boston Public Schools system) because of two attributes of school choice. First, priority for a school is often determined by school zone, sibling status and lotteries. Thus, schools do not prefer higher priority students. Second, schools cannot enroll students without the districts approval. (The NYC high school admissions system is a notable exception, see Abdulkadiroğlu et al. (2009)).

school boards could potentially redesign their priority structures to obtain their goals under TTC, but in general the appropriate priority structures under DA and TTC will be different, and current theory provides almost no guidance as to how to design such priority structures under TTC.³

This lack of transparency has hindered the adoption of TTC.⁴ It caused difficulty in communicating it to parents and school boards, raised concerns about whether priorities would achieve their intended goals under TTC, and made it difficult to convince students and parents that the mechanism was strategy-proof.⁵ It also made it difficult for students to verify they were correctly assigned.

In this paper we develop a characterization of TTC that explains the role of priorities in determining the TTC assignment. We show that the TTC assignment can be concisely described by cutoffs $\{p_b^c\}$ for each pair of schools b, c . These cutoffs parallel prices in competitive equilibrium, with students' priorities serving the role of endowments. Thus, the role of priorities under TTC is as follows. Students can use priority at school b to gain admission to school c if their priority at school b is above the cutoff p_b^c . Each student is assigned to her most preferred school for which she gained admission. If students privately know their priorities, publicly publishing the cutoffs $\{p_b^c\}$ allows each student to determine their assignment.⁶ We show that there is a labeling of schools $\{1, \dots, n\}$ such that for any b the cutoffs are ordered $p_b^1 \geq p_b^2 \geq \dots p_b^b = \dots = p_b^n$. Additionally, to help convey to students that TTC is strategy-proof, we derive cutoffs that are independent of the reported preferences of a given student.

To facilitate tractable analysis of TTC, we formulate a continuum model of TTC

³Comparisons between DA and TTC rely on simulations, and typically use the same priority structure for both mechanisms instead of optimizing the priority structures used for each mechanism, see for example Abdulkadiroğlu et al. (2009) and Pathak (2016).

⁴Boston Public Schools considered both TTC and DA when redesigning its school choice in 2005, and decided in favor of using DA, stating (BPS 2005): *“The behind the scenes mechanized trading makes the student assignment process less transparent.”* Pathak (2016) writes: *“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.”*

⁵Boston’s final school committee report states (BPS 2005): *“The Top Trading Cycles Algorithm allows students to trade their priority for a seat at a school with another student. This trading shifts the emphasis onto the priorities and away from the goals BPS is trying to achieve by granting these priorities in the first place.”* Additionally, the report suggests that *“this trading of priorities could lead families to believe they can still benefit from strategizing, as they may be encouraged to rank schools to which they have priority, even if they would not have put it on the form if the opportunity for trading did not exist.”*

in which the TTC assignment can be directly calculated from the distribution of preferences and priorities by solving a system of equations. We present closed form solutions for parameterized economies. The discrete TTC model is a particular case of the continuum framework; for discrete problems the continuum TTC model calculates cutoffs that give the discrete TTC assignment. We show that the TTC assignment changes smoothly with changes in the underlying economy, implying that the continuum economy can also be used to approximate sufficiently similar economies.

The tractability of our framework relies on a novel approach to analyzing TTC. 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. For smooth continuum economies we reformulate the trade balance equations into a system of equations that fully characterizes TTC. These equations provide a recipe for calculating the TTC assignment.

The tractable continuum framework allows us to analyze the performance of TTC. We provide comparative statics, calculate assignment probabilities under lotteries and evaluate welfare. In particular, when priorities are partly determined by random lottery, the probability that a student gains admission to a school can be directly derived as the probability her random priority is above the required cutoffs. The cutoff representation also yields for each student a budget set of schools at which she gained admission, and these budget sets allow tractable expressions for welfare under random utility models.

As an illustration of the framework, we apply it to study the effects of making a school more desirable. As a shorthand, we refer to such changes as an increase in the quality of the school.⁷ To evaluate the effects of increasing the quality of a school

⁶This cutoff 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 cutoffs $\{p_b^c\}$. 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 . Theorem 1 shows the cutoffs can be observed after the run of TTC. We thank Chiara Margaria, Laura Doval and Larry Samuelson for suggesting this explanation.

⁷Examples of such changes include increases in school infrastructure spending (Cellini et al. 2010), increases in school district funding (Hoxby 2001, Jackson et al. 2016, Johnson & Jackson 2017), reduction in class size (Krueger 1999, Chetty et al. 2011) and changes in an individual school's funding (Dinerstein et al. 2014), but our theoretical model is not specific to any of these examples.

it is necessary to account for changes in the assignment due to changes in student preferences. First, we derive comparative statics that show how the assignment and student welfare change with changes in a school’s quality. We decompose the marginal change in student welfare into the direct increase in utility of students assigned to the more desirable school and the indirect effect that arises from changes in the assignment. A marginal increase in the quality of a popular school can have a negative indirect effect on welfare: as some students switch into the school and gain a marginal utility increase, other students are denied admission and can suffer substantial losses. We quantify these effects in a parametric setting, showing that increasing the quality of a popular school can decrease the welfare gains from sorting on idiosyncratic preferences.

Second, we consider a school district’s problem of allocating resources to improve schools, taking utilitarian welfare as a proxy for the school district’s objective. The framework allows us to solve for the optimal distribution of school quality under TTC for a parametric setting. We find that the optimal distribution of quality is equitable, in the sense that it makes all schools equally over-demanded. An equitable distribution of quality is efficient because it allows students more choice, yielding better sorting on idiosyncratic preferences and therefore higher welfare. This can hold even if some schools are more efficient at utilizing resources, as the benefits from more efficient sorting can outweigh benefits from targeting more efficient schools.

As another application, we explore the design of priorities for TTC and find that it is “bossy” in the sense that a change in the priority of a student that does not alter her assignment can nonetheless alter the assignment of other students. This implies that it is not possible to determine the TTC cutoffs directly through a supply-demand equation as in Azevedo & Leshno (2016). We characterize the range of possible assignments generated by TTC after changes to the 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 provides comparisons between mechanisms in terms of assignments and welfare. We solve for welfare under TTC and DA in a parametric setting and quantify how much welfare is sacrificed due to stability. A comparison between TTC and DA across different school choice environments corroborates a conjecture by Pathak (2016) that the difference between the mechanisms becomes smaller with increased alignment between student preferences and school pri-

orities. We also compare TTC to the Clinch and Trade mechanism (Morrill 2015*b*) in large economies and find that it is possible for TTC to produce fewer blocking pairs than the Clinch and Trade mechanism.

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

Cutoff representations have been instrumental for empirical work on DA and variants of DA. Abdulkadiroğlu, Angrist, Narita & Pathak (2017) use admission cutoffs to construct propensity score estimates. Agarwal & Somaini (Forthcoming), Kapor et al. (2016) structurally estimate preferences from rank lists submitted to non-strategy-proof variants of DA. Both build on the cutoff representation of Azevedo & Leshno (2016). We hope that our cutoff representation of TTC will be similarly useful for empirical work on TTC.

Dur & Morrill (2017) 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. Our characterization also provides a connection between TTC and competitive markets, but requires a lower dimensional set of cutoffs and provides a method for directly calculating these cutoffs. He et al. (Forthcoming) propose an alternative pseudo-market approach for discrete assignment problems that extends Hylland & Zeckhauser (1979) and also

uses admission cutoffs. Miralles & 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 require an arbitrary endowment structure.

This paper contributes to a growing literature that uses continuum models in market design (Avery & Levin 2010, Abdulkadiroğlu et al. 2015, Ashlagi & Shi 2015, Che et al. 2017, Azevedo & Hatfield 2015). Our description of the continuum economy uses the setup of Azevedo & 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.

Ma (1994), Pápai (2000) and Pycia & Ünver (2017) 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. Abdulkadiroğlu, Che, Pathak, Roth & Tercieux (2017) show that TTC minimizes the number of blocking pairs subject to strategy-proofness and Pareto efficiency. Additional axiomatic characterizations of TTC were given by Dur (2012) and Morrill (2013, 2015*a*). These characterizations explore the properties of TTC, but do not provide another method for calculating the TTC outcome or evaluating welfare.

Several variants of TTC have been suggested in the literature. Morrill (2015*b*) 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 & Kesten (Forthcoming) introduce Equitable TTC, a variation on TTC that aims to reduce inequity. In Section 4.2 we show how our model can be used 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 & Molis (2011), Jaramillo & Manjunath (2012) and Saban & 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. Hatfield et al. (2016) study the

incentives for schools to improve their quality under TTC and find that even in a large market a school may be assigned less preferred students when it improves its quality. Our results in Section 4.1 quantify these effects. Che & Tercieux (2015, 2018) study the properties of TTC in a large market where the heterogeneity of items grows as the market gets large, whereas our setting considers a large population of agents and a fixed number of item types. The results in Section 4 show that TTC has different properties in these different large markets.

1.2 Organization of the Paper

Section 2 presents our cutoff characterization under the standard discrete TTC model. Section 3 presents the continuum TTC model and provides our main results that allow for direct calculation of the TTC cutoffs. In this section we also demonstrate how our model can be used to calculate the TTC assignment, see Example 2. Section 4 explores several applications: quantifying the effects of improving school quality under school choice and solving for the optimal distribution of quality, showing the “bossiness” of the TTC priorities, and comparing TTC with other mechanisms. Appendix A provides the intuition for the continuum TTC model. Appendix B provides an example of a computation of the discrete TTC allocation through the continuum framework. Omitted proofs can be found in the online 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 $\mathcal{C} = \{1, \dots, n\}$ 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.⁸ It

⁸This is without loss of generality, as we can introduce auxiliary students and schools that represent being unmatched.

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 $E = (\mathcal{C}, \mathcal{S}, \succ^{\mathcal{S}}, \succ_c, q)$, 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 $\succ^{\mathcal{S}} = \{\succ^s\}_{s \in \mathcal{S}}$, $\succ_c = \{\succ_c\}_{c \in \mathcal{C}}$.

Given an economy E , the discrete Top Trading Cycles algorithm (TTC) calculates an assignment $\mu_{dTTC}(\cdot \mid E) : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$. We omit the dependence on E when it is clear from context. 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 $\{q_c\}_{c \in \mathcal{C}}$.*

While there are still unassigned students and available schools:

- *Each available school $c \in C$ offers a seat by pointing to its highest priority remaining student.*
- *Each student $s \in S$ who was offered a seat points to her most preferred remaining school.*
- *Select at least one trading cycle, that is, a list of students $s_1, \dots, s_\ell, s_{\ell+1} = s_1$ such that 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.⁹*
- *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*

⁹Such a trading cycle must exist, since every vertex in the pointing graph with vertex set $S \cup C$ has out-degree 1.

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 & 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 & 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.

Theorem 1. *Let E be an economy. The TTC assignment is given by*

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

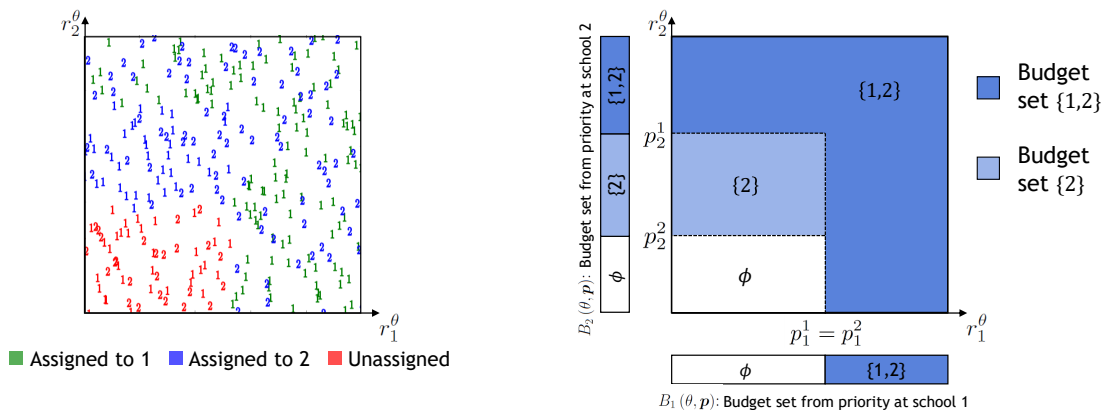
where p_b^c is the percentile in school b 's ranking of the worst ranked student at school b that traded a seat at school b for a seat at school c during the run of the TTC algorithm on E . If no such student exists, $p_b^c = 1$.

Cutoffs serve a parallel role to prices in Competitive Equilibrium, and each student's vector of priorities at each school serves as her endowment. For each student s , the cutoffs $\mathbf{p} = \{p_b^c\}_{b,c}$ combine with 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.

The cutoffs p_b^c in Theorem 1 can be easily identified after the mechanism has been run. Hence Theorem 1 provides an intuitive way for students to verify that they were correctly assigned by the TTC algorithm. 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. Student priorities were selected such that there is little correlation between student priority at either school and between student priorities and preferences. Figure 1a illustrates the preferences and priorities of each of the students. Each colored number represents a student. The location of the student in the square indicates their priority, with the horizontal axis indicating priority at school 1 and the vertical axis indicating priority at school 2. The number indicates the student's preferred school, and all students find both schools acceptable. The color indicates the student's assignment under TTC.

The cutoffs \mathbf{p} and resulting budget sets $B(s, \mathbf{p})$ for each student are illustrated in Figure 1b. The colors in the body of the figure indicate the budget sets given to students as a function of their priority at both schools. The colors along each axis indicate the schools that enter a student's budget set because of her priority at the school whose priority is indicated by that axis. 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 1b as each student's budget set is independent of her preferences. The assignment of each student is her favorite school in her budget set.



(a) The economy E and the TTC assignment.

(b) Budget sets for the economy E .

Figure 1: The economy and TTC budget sets for Example 1.

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 schools 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 . 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 $B(s, \mathbf{p}) = \cup_b B_b(s, \mathbf{p})$. Figure 1(b) depicts $B_1(s, \mathbf{p})$ and $B_2(s, \mathbf{p})$ for the economy of Example 1 along the x and y axes respectively.

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 exists a relabeling of school indices such that there exist cutoffs $\mathbf{p} = \{p_b^c\}$ that describe the TTC assignment*

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

and for any school b the cutoffs are ordered,¹⁰

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

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

¹⁰The 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 $\tilde{p}_b^c = \min_{a \leq c} p_b^a$.

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 .

When there exist TTC cutoffs that satisfy inequality (1) we say that the schools are *labeled in order*. 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 order). This generalizes the illustration along each axis in Figure 1(b), 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 priority at school b . The cutoffs p_b^c are weakly decreasing in c , and are equal for all $c \geq b$ (i.e. $p_b^b = p_b^{b+1} = \dots = p_b^n$). That is, a student's priority at b can add one of the sets $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \dots, \mathcal{C}^{(b)}, \phi$ to her budget set. If any school enters a student's budget because of her priority at b , then school b must also enter her budget set because of her priority at b .

Dur & Morrill (2017) provide a characterization of TTC as a competitive equilibrium where a priority value function $v(r, b)$ specifies the price of priority r at school b and students are allowed to buy and sell one priority. Given TTC cutoffs $\{p_b^c\}$ where schools are labeled in order, the TTC assignment and priority value function $v(r, b) = n - \min \{c \mid r \geq p_b^c\}$ constitute a competitive equilibrium. We introduce a framework in Section 3 that allows a direct calculation of this competitive equilibrium

as a solution to a set of equations.

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 possible school assignments that student s can achieve by unilaterally changing her reported preferences.^{11,12} We therefore introduce the continuum model for TTC which allows us to directly calculate the cutoffs, allowing for comparative statics. Using the continuum model, we present in Section 3.4 cutoffs that yield refined budget sets which provide for each student the set of schools that she could be assigned to by unilaterally changing her preferences. Thus the appropriate cutoff structure also makes it clear that TTC is strategy-proof.

3 Continuum Model and Main Results

3.1 Model

We consider the school choice problem with a continuum of students and finitely many schools, as in Azevedo & Leshno (2016). There is a finite set of schools denoted by

¹¹More 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 $\mu_{dTTC}(s | E') \neq \max_{\succ'} B(s, \mathbf{p})$.

¹²For 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 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 .

$\mathcal{C} = \{1, \dots, n\}$, 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 $r^\theta \in [0, 1]^{\mathcal{C}}$. We say that r_b^θ is the *rank* of student θ at school b , or the *b-rank* of student θ . 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.

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,¹³ bounded from above, and bounded from below away from zero on its support.¹⁴

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.¹⁵

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

¹³A 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 $\Theta \setminus G$.

¹⁴That is, there exists $M > m > 0$ such that for every $\theta \in \Theta$ either $\nu(\theta) = 0$ or $m \leq \nu(\theta) \leq M$.

¹⁵We can incorporate an economy where two schools have perfectly aligned priorities by considering them as a combined single school in the trade balance equations, as defined in Definition 2. The capacity constraints still consider the capacity of each school separately.

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 following induced marginal distributions. 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 .¹⁶ 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 students 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. $\eta(\{\theta \mid \mu(\theta) \neq \mu'(\theta)\}) = 0$.

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 strategy-proof.

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

¹⁶Formally

$$\begin{aligned} H_b^{c|C}(\mathbf{x}) &\stackrel{def}{=} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \eta(\{\theta \in \Theta \mid r^\theta \in [(x_b - \varepsilon) \cdot \mathbf{e}^b, \mathbf{x}] \text{ and } Ch^\theta(C) = c\}) \\ &= \int_{\{\theta \mid r^\theta \in [x_b \cdot \mathbf{e}^b, \mathbf{x}] \text{ and } Ch^\theta(C) = c\}} \nu(\theta) d\theta, \end{aligned}$$

where \mathbf{e}^b is the unit vector in the direction of coordinate b . In other words, $H_b^{c|C}(\mathbf{x})$ is the density of students θ with priority $r_b^\theta = x_b$ and $r_a^\theta \leq x_a$ for all $a \in C$ whose most preferred school in C is c .

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 $\gamma(t) : [0, \infty) \rightarrow [0, 1]^{\mathcal{C}}$ 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 $\tilde{\gamma}(t) : [t_0, \infty) \rightarrow [0, 1]^{\tilde{\mathcal{C}}}$ is a *residual TTC path* if it satisfies all the properties of a TTC path except the initial condition, and $\tilde{\gamma}_c(t)$ is defined only for $t \geq t_0 > 0$ and $c \in \tilde{\mathcal{C}} \subset \mathcal{C}$. For a set $\{t^{(c)}\}_{c \in \mathcal{C}} \in \mathbb{R}_{\geq}^{\mathcal{C}}$ of times we let $t^{(c^*)} \stackrel{\text{def}}{=} \min_c [t^{(c)}]$ denote the minimal time. For a point $\mathbf{x} \in [0, 1]^{\mathcal{C}}$, let

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

denote the mass of students whose rank at some school b is better than x_b and their first choice is school c . We will refer to $D^c(\mathbf{x})$ as the demand for c . Recall that $H_b^c(\mathbf{x})$ 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 $D^c(\mathbf{x})$ and $H_b^c(\mathbf{x})$ depend implicitly on the set of available schools \mathcal{C} , as well as on the economy \mathcal{E} .

A TTC path γ can capture the progression of a continuous time TTC algorithm, with the interpretation that $\gamma_c(t)$ is the highest c -priority of any student who remains unassigned by time t . The stopping times $\{t^{(c)}\}_{c \in \mathcal{C}}$ indicate when each school fills its capacity. To verify whether γ and $\{t^{(c)}\}_{c \in \mathcal{C}}$ can correspond to a run of TTC we introduce trade balance conditions and capacity constraints as defined below.

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 the trade balance and capacity equations for the economy E if the following hold.

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

$$\sum_{a \in \mathcal{C}} \gamma'_a(t) H_a^c(\gamma(t)) = \sum_{a \in \mathcal{C}} \gamma'_c(t) H_c^a(\gamma(t)) \quad (2)$$

for all $c \in \mathcal{C}$ and 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*

$$\begin{aligned} D^{c^*} (\gamma (t^{(c^*)})) &= q_{c^*} \\ D^a (\gamma (t^{(c^*)})) &\leq q_a \quad \forall a \in \mathcal{C} \end{aligned} \tag{3}$$

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

3. If $\mathcal{C} \setminus \{c^*\} \neq \emptyset$, 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 $\gamma(\cdot) : [t^{(c^*)}, \infty) \rightarrow [0, 1]^{\tilde{\mathcal{C}}}$ 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 the trade balance and capacity equations for \tilde{E} .

A brief motivation for the definition is as follows. TTC progresses by clearing trading cycles, and in each trading cycle the number of seats offered by a school is equal to the number of students assigned to that school. Equation (2) states that over every small time increment the mass of students assigned to a school must be equal to the mass of offers made by the school. While all schools have remaining capacity, every assigned student is assigned to his first choice, and thus $D^c (\gamma (t))$ gives the mass of students assigned to school c at time $t \leq t^{(c^*)}$ in the algorithm. The time $t^{(c^*)}$ when school c^* fills its capacity can be calculated as a solution to Equation (3). Once a school exhausts its capacity we can eliminate that school and recursively calculate the TTC assignment on the remaining problem with $n - 1$ schools, which is stated as condition (3). We provide more comprehensive intuition for the definition and the results in Appendix A.

Our main result is that the trade balance and capacity equations fully characterize and provide a way to directly calculate the TTC assignment from the problem primitives. We show in Section 3.3 that this characterization is consistent with the discrete TTC.

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

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

In other words, Theorem 2 provides the following a recipe for calculating the TTC assignment. First, find $\hat{\gamma}(\cdot)$ that solves the marginal trade balance equations (2) for all t . Second, calculate $t^{(c^*)}$ from the capacity equations (3) 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.¹⁷ This recipe is illustrated in Example 2. The TTC path used in this recipe may not be the unique TTC path, but all TTC paths yield the same TTC assignment.

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 (3) 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 schools are labeled in order, that is,*

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

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.

¹⁷Continuity 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 .

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

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

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 $p = 2/3$ is illustrated in Figure 3. This economy can be viewed as a smoothed continuum version of the economy in Example 1.




Figure 3: The TTC path, cutoffs, and budget sets for a particular instance of the economy \mathcal{E} in Example 2. Students in the dark blue region have a budget set of $\{1, 2\}$, students in the light blue region have a budget set of $\{2\}$, and students in the white region have a budget set of ϕ .

We start by solving for γ from the trade balance equations (2), which simplify to the differential equation¹⁸

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

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

¹⁸The original trade balance equations are

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

$\min \{t^{(1)}, t^{(2)}\}$ 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 the marginal trade balance equations (2) are trivial, and we define the residual TTC path by

$$\gamma(t) = (p_1^1, p_2^1 - (t - t^{(1)}))$$

for $t^{(1)} \leq t \leq t^{(2)}$. Solving the capacity equation (3) for $t^{(2)}$ yields that

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

For instance, if we plug in $q = 4/10$ and $p = 2/3$ to match the economy in Example 1, the calculation yields the cutoffs $p_1^1 = p_2^2 \approx .54$, $p_2^1 \approx .73$ 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 can also be used to show 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 $E = (\mathcal{C}, \mathcal{S}, \succ_{\mathcal{C}}, \succ^{\mathcal{S}}, q)$ with $N = |\mathcal{S}|$ students 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.5.

Proposition 2. *Let $E = (\mathcal{C}, \mathcal{S}, \succ_{\mathcal{C}}, \succ^{\mathcal{S}}, q)$ be a discrete economy with $N = |\mathcal{S}|$ students, 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,*

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

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

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

In other words, Φ embeds a 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 $O(\varepsilon|\mathcal{C}|^2)$.*

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 discrete economy where student s changes her report from \succ^s to \succ' (holding others' reported preferences fixed), and let

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

denote the set of possible school assignments that student s can achieve by unilaterally changing her reported preferences. 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 $B^*(s | E)$. The analysis in this section can be used to show that the budget sets $B^*(s | E)$ 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

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

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

$$B^*(s | E) = B(s, \mathbf{p}^*).$$

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 strategy-proof. In the appendix we prove Proposition 3 and constructively find \mathbf{p}^* .

4 Applications

4.1 Effects of Changes in the Distribution of School Quality

We apply our model to analyze economies where preferences for schools are endogenously determined by the allocation of resources to schools. Empirical evidence suggests that increased financing affects student achievements (Jackson et al. 2016, Lafortune et al. 2016, Johnson & Jackson 2017) as well as demand for housing (Hoxby 2001, Cellini et al. 2010), which indicate increased demand for schools. Similarly, Krueger (1999) finds that smaller classes have a positive impact on student performance, and Dinerstein et al. (2014) finds that increased funding for public schools increases enrollment in public schools and reduces demand for private schools.

Under school choice, such resource allocation decisions can change the desirability of schools and therefore change the assignment of students to schools. We explore the implication of such changes in a stylized model. As a shorthand, we refer to an

increase in the desirability of a school as an increase in the quality of the school. We explore comparative statics of the allocation and evaluate student welfare. Omitted proofs and derivations can be found in the online Appendix E.1.

Model with quality dependent preferences

We enrich the model from Section 3 to allow student preferences to depend on school quality $\delta = \{\delta_c\}_{c \in \mathcal{C}}$, where the desirability of school c is increasing in δ_c . An economy with quality dependent preferences is given by $\mathcal{E} = (\mathcal{C}, \Upsilon, v, q)$, where $\mathcal{C} = \{1, 2, \dots, n\}$ is the set of schools and Υ is the set of student types. A student $s \in \Upsilon$ 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 $\delta = \{\delta_c\}_{c \in \mathcal{C}}$ and r_c^s 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 v over Υ specifies the distribution of student types. School capacities are $q = \{q_c\}$, where $\sum q_c < 1$. We will refer to δ_c as the quality of c .

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.¹⁹ We assume for all δ that η_δ has a Lipschitz 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 $\{p_b^\delta(\delta)\}_{c \in \mathcal{C}}$ denote the TTC assignment and associated cutoffs for the economy \mathcal{E}_δ . We omit the dependence on δ when it is clear from context.

Comparative statics of the allocation

The following proposition gives the direction of change of the TTC cutoffs when there are two schools and δ_ℓ increases for some $\ell \in \{1, 2\}$. Throughout this subsection, when considering a fixed δ we assume that schools are labeled in order, unless stated otherwise.

Proposition 4. *Suppose $\mathcal{E} = (\mathcal{C} = \{1, 2\}, \Upsilon, v, q)$ and δ induces an economy \mathcal{E}_δ such that the TTC path γ that, if possible, assigns seats at school 1 before seats at school 2, yields $p_2^1(\delta) > p_2^2(\delta)$.²⁰ Consider $\hat{\delta}$ that increases the quality of school 2, i.e. $\hat{\delta}_2 \geq \delta_2$*

¹⁹To make student preferences strict we arbitrarily break ties in favor of schools with lower indices. We assume the utility of being unassigned is $-\infty$, so all students find all schools acceptable.

²⁰Formally, γ is defined by requiring that for all t it holds that $\gamma'(t)$ is the valid direction at $\gamma(t)$ with support that is minimal under the order $\{1\} < \{1, 2\} < \{2\}$.

and $\delta_1 = \hat{\delta}_1$, and which induces $\mathcal{E}_{\hat{\delta}}$ with TTC path $\hat{\gamma}$ that also assigns seats at 1 before 2 when possible and yields $p_2^1(\hat{\delta}) \geq p_2^2(\hat{\delta})$.

Then a change from δ to $\hat{\delta}$ induces the cutoffs $p_b^c(\cdot)$ to 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 2.

Proposition 4 is illustrated in Figure 4. As first shown in Hatfield et al. (2016), an increase in the desirability of school 2 can cause low 2-rank students to be assigned to school 2. Note that individual students' budget sets can grow or shrink by more than one school.

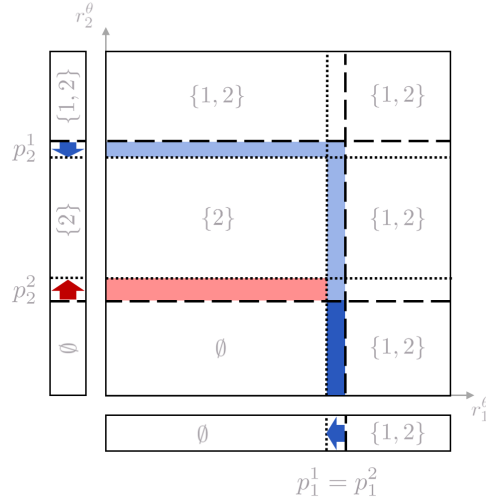


Figure 4: The effect of an increase in the quality of school 2 on TTC cutoffs and budget sets. Dashed lines indicate initial TTC cutoffs, and dotted lines indicate TTC cutoffs given increased school 2 quality. The cutoffs $p_1^1 = p_1^2$ and p_2^1 decrease and the cutoff p_2^2 increases. Students in 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.

When there are $n \geq 3$ schools, it is possible to show that an increase in the quality of a school ℓ can either increase or decrease any cutoff. With additional structure we can provide precise comparative statics that mirror the intuition from Proposition (4).

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 ε_{cs} distributed

as i.i.d. extreme value shifted to have a mean of 0 (McFadden 1973). We assume that the total measure of students is normalized to 1, that there are more students than school seats, i.e. $\sum_c q_c < 1$, and that all students prefer any school to being unassigned²¹. Schools' priorities are uncorrelated and uniformly distributed. This model combines heterogeneous idiosyncratic taste shocks with a common preferences modifier δ_c . Proposition 5 gives the TTC assignment in closed form for the logit economy.

Proposition 5. *Under the logit economy schools are labeled in order if $\frac{q_1}{e^{\delta_1}} \leq \frac{q_2}{e^{\delta_2}} \leq \dots \leq \frac{q_n}{e^{\delta_n}}$, and in such cases the TTC cutoffs p_b^c for $b \geq c$ are given by²²*

$$p_b^c = (R^c)^{\frac{e^{\delta_b}}{\pi_c}} \prod_{a < c} (R^a)^{\frac{e^{\delta_b}}{\pi_a} - \frac{e^{\delta_b}}{\pi_{a+1}}} \quad (4)$$

where $\pi_c = \sum_{c' \geq c} e^{\delta_{c'}}$ is the normalization term for schools in $\mathcal{C}^{(c)}$, for all $c \geq 1$ the quantity $R^c = 1 - \sum_{c' < c} q_{c'} - \frac{\pi_c}{e^{\delta_c}} q_c$ is the measure of unassigned, or remaining, students after the c th round, and $R^0 = 1$.

Moreover, p_b^c is decreasing in δ_ℓ for $c < \ell$ and increasing in δ_ℓ for $b > c = \ell$.

Figure 5 illustrates how the TTC cutoffs change with an increase in the quality of school ℓ . Using equation (13), we derive closed form expressions for $\frac{dp_b^c}{d\delta_\ell}$, which can be found in online Appendix E.1.

Remark 2. Proposition 5 can be used to calculate admission probability under multiple tie-breaking as follows. Consider an economy where priorities are determined by a multiple tie-breaking rule where the priority of each student at each school is generated by an independent $U[0, 1]$ lottery draw. As a result, students priorities will be uniformly distributed over $[0, 1]^C$ and uncorrelated with student preferences. If in addition student preferences are given by the MNL model, this is a logit economy. In the logit economy the ex-ante probability that a student will gain admission to school c is given by

$$1 - \prod_{b \in \mathcal{C}} p_b^c$$

with p_b^c given by Proposition 5.

²¹Formally, $u^s(\phi | \delta) = -\infty$. For welfare calculations we only consider assigned students.

²²To simplify notation, when $c = 1$ we let $\prod_{c' < c} p_{c'}^{c-1} = 1$ and set $\rho_1 = q_1/e^{\delta_1}$.




Figure 5: 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 $b \geq c$, i.e., 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 $\frac{dp_b^c}{d\delta_\ell} = \frac{dp_b^\ell}{d\delta_\ell} > 0$ for all $b > \ell$, and p_b^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 .

Comparative statics of student welfare

We consider a social planner who can affect quality levels δ of 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 the sake of simplicity consider students' social welfare as a proxy for the social planner's objective. Given assignment μ , the social welfare is given by

$$U(\delta) = \int_{s \in \Upsilon, \mu(s) \neq \phi} u^s(\mu(s) | \delta) dv.$$

As a benchmark, we first consider neighborhood assignment μ_{NH} which assigns each student to a fixed school regardless of her 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 .

The budget set formulation of TTC allows us to tractably capture student welfare under TTC.²³ A student who is offered the budget set $\mathcal{C}^{(c)} = \{c, \dots, n\}$ is assigned

²³Under TTC the expected utility of student s assigned to school $\mu(s)$ depends on the stu-

to the school $\ell = \arg \max_{b \in \mathcal{C}^{(c)}} \{\delta_b + \varepsilon_{bs}\}$, and her expected utility is $U^c = \ln \left(\sum_{b \geq c} e^{\delta_b} \right)$ (Small & Rosen 1981). Let N^c be the mass of agents with budget set $\mathcal{C}^{(c)}$. Then social welfare under the TTC assignment given 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 change 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 under TTC a marginal increase in the quality of school ℓ will have two effects. As under neighborhood assignment, it will increase the utility of the q_ℓ students assigned to ℓ by $d\delta_\ell$. In addition, the quality increase changes student preferences, and therefore changes the assignment. The second term captures the indirect effect on welfare due to changes in the assignment. This effect is captured by changes in the number of students offered each budget set.

The indirect effect can be negative. In particular, when there are two schools $\mathcal{C} = \{1, 2\}$ the welfare effect of a quality increase to school 1 is²⁴

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

An increase in the quality of 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

dent's budget set $B(s, \mathbf{p})$ because of the dependency of $\mu(s)$ on student preferences. Namely, $\mathbb{E}[u^s(\mu(s) | \delta)] = \delta_{\mu(s)} + \mathbb{E}[\varepsilon_{\mu(s)s} | \delta_{\mu(s)} + \varepsilon_{\mu(s)s} \geq \delta_c + \varepsilon_{cs} \forall c \in B(s, \mathbf{p})]$

²⁴Recall that we assume that schools are labeled in order, and thus school 1 is the more selective school. We use that $N^1 = q_1 + q_1 e^{\delta_1 - \delta_2}$, $N^2 = q_2 - q_1 e^{\delta_2 - \delta_1}$.

set did not change and who switched to $1 \succ 2$ are almost indifferent between the schools and hence almost unaffected. Students who lost school 1 from their budget set may prefer school 1 by a large margin, and hence incur significant loss. Thus, there is a total negative effect from changes in the assignment, which is captured by the second term.

If a positive mass of students receive the budget set $\{2\}$ (that is, $N^2 > 0$), improving the quality of school 2 will have the opposite indirect effect. Specifically,

$$\frac{dU_{TTC}}{d\delta_2} = q_2 + q_1 \cdot e^{\delta_2 - \delta_1} (\ln(e^{\delta_1} + e^{\delta_2}) - \delta_2) > q_2$$

which is larger than the marginal effect under neighborhood assignment.

If admission cutoffs into both schools are equal (that is, $p_2^1 = p_2^2$ and $N^2 = 0$) we say that both schools are equally over-demanded. In such a case, a marginal increase in the quality of either school will have a negative indirect effect on welfare.²⁵

Selecting the quality distribution to maximize student welfare

We now provide an illustrative example showing the welfare optimal quality distribution under DA, TTC and neighborhood assignment. This example also allows us to compare welfare across mechanisms. In the examples below we fix the school labels and consider various δ . For some values of δ the schools may be labeled out of order.

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. Quality levels δ are constrained by $\delta_1 + \delta_2 = 2$ 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

²⁵That is, if $\delta_1 = \delta_2$ then $\frac{dU_{TTC}}{d\delta_1} < q_1$ and $\frac{dU_{TTC}}{d\delta_2} < q_2$. If we fix $\delta_1 + \delta_2$ and consider $U_{TTC}(\Delta)$ as a function of $\Delta = \delta_1 - \delta_2$ the function $U_{TTC}(\Delta)$ will have a kink at $\Delta = 0$ (see Figure 6c).

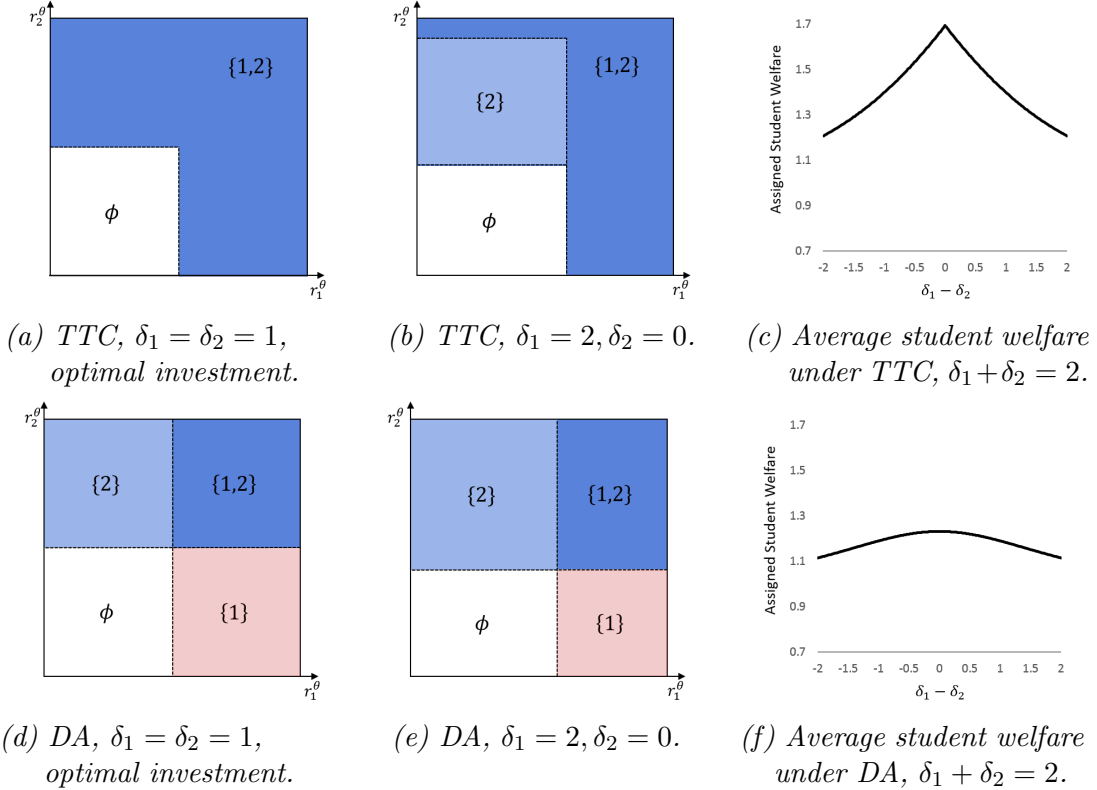


Figure 6: Illustration for Example 3. Figures (a) and (b) show the budget sets under TTC for different quality levels, and Figure (c) shows the average welfare of assigned students under TTC for quality levels $\delta_1 + \delta_2 = 2$ for different values of $\delta_1 - \delta_2$. Figures (d) and (e) 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 $B = \{1, 2\}$. Two thirds of 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 $U_{DA}/Q \approx 1.11$.

In Example 3, 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.²⁶ Requiring a stable assignment will constrain two thirds of the assigned students from efficiently sorting on horizontal taste shocks.

We next provide an example where the two schools have different capacity, with

²⁶Note that this is not a concern in school choice settings where blocking pairs cannot be assigned outside of the mechanism.

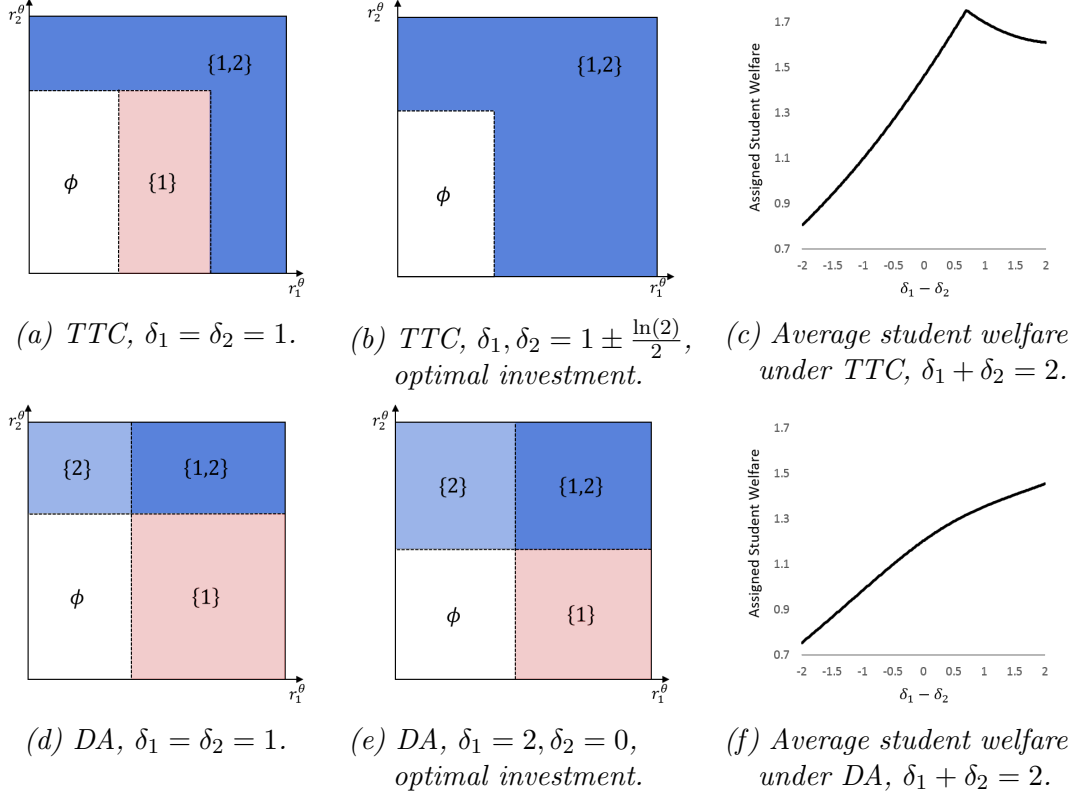


Figure 7: Illustration for Example 4. Figures (a) and (b) show the budget sets under TTC for different quality levels, and Figure (c) shows the average welfare of assigned students under TTC for quality levels $\delta_1 + \delta_2 = 2$ for different values of $\delta_1 - \delta_2$. Note that $\delta_1 = \delta_2 = 1$ 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.

$q_1 > q_2$. To make investment in school 1 more efficient, we assume that (despite having more students) school 1 requires the same amount of resources to increase its quality for all its students. Thus, we keep the constraint that $\delta_1 + \delta_2 = 2$. It is straightforward to see that under neighborhood assignment the welfare optimal distribution of quality is $\delta_1 = 2, \delta_2 = 0$. In contrast, we find the welfare optimal distribution under TTC can be closer to egalitarian.

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. Quality levels δ are constrained by $\delta_1 + \delta_2 = 2$ 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 $\delta_1 = 1, \delta_2 = 1$, yield $U_{TTC}/Q \approx 1.46$. 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 in Example 4 we find that the optimal quality distribution under TTC provides all assigned students with a full budget set, making all schools equally over-demanded. The optimal quality distribution under neighborhood assignment and DA allocates all resources to the more efficient school. While quality directed to the larger school affects more students and yields more direct benefit, under TTC an egalitarian distribution leads to more welfare gains from sorting on horizontal tastes. For general parameters the welfare gain from sorting can be lower or higher than the welfare gains from directing all resources to the more efficient school.

Finally, consider a central school board with a fixed amount of resources K to be allocated to the n schools. We assume that the cost of quality δ_c is the convex function $\kappa_c(\delta_c) = e^{\delta_c}$. This specification makes bigger schools more efficient.²⁷ Using Proposition 6 we solve for the optimal distribution of school quality. Despite the heterogeneity among schools, social welfare is maximized when all assigned students have a full budget set, which occurs when the amount allocated to each school is proportional to the number of seats at the school.

Proposition 7. *Consider a logit economy with cost function $\kappa_c(\delta_c) = e^{\delta_c} \forall c$ and resource constraint $\sum_c \kappa_c(\delta_c) \leq K$. Social welfare is uniquely maximized when the resources κ_c allocated to school c are 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.

²⁷Note that κ_c is the total school funding. This is equivalent to setting the student utility of school c to be $u^s(c | \kappa_c) = \log(\kappa_c) + \varepsilon_{cs} = \log(\kappa_c/q_c) + \log(q_c) + \varepsilon_{cs}$, which is the log of the per-student funding plus a fixed school utility that is larger for bigger schools.

More is invested in higher capacity schools, as they provide more efficient investment opportunities, but the investment is balanced across schools.

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. 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 $[m, 1]^2$.²⁸ 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

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

for $\underline{p} = \sqrt{(1 - 2q) \frac{m^2}{1 - 2m + 2m^2}}$ and $\bar{p} = \sqrt{(1 - 2q) \frac{1 - 2m + 2m^2}{m^2}}$. Figure 8 illustrates the range of possible TTC cutoffs for \mathcal{E}' and the economy $\bar{\mathcal{E}}$ 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 8) 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

²⁸The remaining students still have ranks distributed uniformly on the complement of $[m, 1]^2$.




Figure 8: 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 $p_1 = \sqrt{(1-2q)\frac{1-2m+2m^2}{m^2}} \approx 0.36$ and $p_2 = \sqrt{(1-2q)\frac{m^2}{1-2m+2m^2}} \approx 0.25$.

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.

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 $\{\theta \mid \exists c \text{ s.t. } r_c^\theta \geq m\}$, 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. 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

In Section 4.1 we compared the welfare effects of changes in school resource allocation under various school choice mechanisms. Our formulation of TTC also allows us to compare TTC with other school choice mechanisms. In this section, we provide a theoretical explanation for observed similarities between assignments under TTC and Deferred Acceptance (DA), as well as a comparison of the number of blocking pairs induced by TTC and the closely related Clinch and Trade mechanism.

Both TTC and Deferred Acceptance (DA) (Gale & Shapley 1962) are strategy-proof, but differ in that TTC is efficient whereas DA is stable. Kesten (2006), Ehlers & 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 produce similar outcomes. In Section 4.1 we compared DA and TTC in terms of welfare and assignment and found that large differences were possible.²⁹ 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 in these cities.

To study this conjecture, we consider a simple model with neighborhood priority. 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*

²⁹Che & Tercieux (2015) 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.

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 & Leshno (2016) to find the TTC and DA allocations respectively. For each school, students with priority are given a lottery number uniformly at random in $[\frac{n-1}{n}, 1]$, and students without priority are given a lottery number uniformly at random in $[0, \frac{n-1}{n}]$, where lottery numbers at different schools are independent. For all values of α , the TTC cutoffs are given by $p_j^i = p_i^j = 1 - \frac{q_i}{nq}$ for all $i \leq j$, and the DA cutoffs are given by $p_i = 1 - \frac{q_i}{nq}$. 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 $1 - \frac{q_i}{nq}$, and whose 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) suggests that this design choice is desirable because it can reduce the number of blocking pairs induced by the assignment, and 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 of 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.³⁰ 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.*

³⁰For brevity, we abstract away from certain details of C&T mechanism that are important when not all schools run out at the same round.

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 6). 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. The differences between the cutoff structures of these mechanisms can help clarify the different role of priorities under these mechanisms.

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 for 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 & 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., Angrist, J. D., Narita, Y. & Pathak, P. A. (2017), ‘Research design meets market design: Using centralized assignment for impact evaluation’,

Econometrica **85**(5), 1373–1432.

Abdulkadiroğlu, A., Che, Y.-K., Pathak, P. A., Roth, A. E. & Tercieux, O. (2017), Minimizing justified envy in school choice: The design of new orleans' oneapp, Technical report, National Bureau of Economic Research.

Abdulkadiroğlu, A., Che, Y.-K. & Yasuda, Y. (2015), 'Expanding choice in school choice', *American Economic Journal: Microeconomics* **7**(1), 1–42.

Abdulkadiroğlu, A., Pathak, P. A. & Roth, A. E. (2009), 'Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match', *The American Economic Review* pp. 1954–1978.

Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E. & Sönmez, T. (2005), 'The boston public school match', *American Economic Review* pp. 368–371.

Abdulkadiroğlu, A. & Sönmez, T. (2003), 'School choice: A mechanism design approach', *The American Economic Review* **93**(3), 729–747.

Agarwal, N. & Somaini, P. (Forthcoming), 'Demand analysis using strategic reports: An application to a school choice mechanism', *Econometrica* .

Alcalde-Unzu, J. & Molis, E. (2011), 'Exchange of indivisible goods and indifferences: The top trading absorbing sets mechanisms', *Games and Economic Behavior* **73**(1), 1–16.

Ashlagi, I. & Shi, P. (2015), 'Optimal allocation without money: An engineering approach', *Management Science* **62**(4), 1078–1097.

Avery, C. & Levin, J. (2010), 'Early admissions at selective colleges', *American Economic Review* **100**(5), 2125–56.

Azevedo, E. M. & Hatfield, J. W. (2015), 'Existence of equilibrium in large matching markets with complementarities'.

Azevedo, E. M. & Leshno, J. D. (2016), 'A supply and demand framework for two-sided matching markets', *Journal of Political Economy* .

BPS (2005), Recommendation to implement a new bps assignment algorithm, Technical report.

- Cellini, S. R., Ferreira, F. & Rothstein, J. (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., Kim, J. & Kojima, F. (2017), Stable matching in large economies, Technical report.
- Che, Y.-K. & Tercieux, O. (2015), ‘Efficiency and stability in large matching markets’.
- Che, Y.-K. & Tercieux, O. (2018), ‘Payoff equivalence of efficient mechanisms in large matching markets’, *Theoretical Economics* **13**(1), 239–271.
- Chetty, R., Friedman, J. N., Hilger, N., Saez, E., Schanzenbach, D. W. & Yagan, D. (2011), ‘How does your kindergarten classroom affect your earnings? evidence from project star’, *The Quarterly Journal of Economics* **126**(4), 1593–1660.
- Dinerstein, M., Smith, T. et al. (2014), ‘Quantifying the supply response of private schools to public policies’.
- Dur, U. M. (2012), ‘A characterization of the top trading cycles mechanism for the school choice problem’, *Available at SSRN 2147449* .
- Dur, U. & Morrill, T. (2017), ‘Competitive equilibria in school assignment’, *Games and Economic Behavior* .
- Ehlers, L. (2014), ‘Top trading with fixed tie-breaking in markets with indivisible goods’, *Journal of Economic Theory* **151**, 64–87.
- Ehlers, L. & Erdil, A. (2010), ‘Efficient assignment respecting priorities’, *Journal of Economic Theory* **145**(3), 1269–1282.
- Gale, D. & Shapley, L. S. (1962), ‘College admissions and the stability of marriage’, *The American Mathematical Monthly* **69**(1), 9–15.
- Hakimov, R. & Kesten, O. (Forthcoming), ‘Equitable top trading cycles mechanisms: Theory and experiments’, *International Economic Review* .
- Hatfield, J. W., Kojima, F. & Narita, Y. (2016), ‘Improving schools through school choice: A market design approach’, *Journal of Economic Theory* .

- He, Y., Miralles, A., Pycia, M. & Yan, J. (Forthcoming), ‘A pseudo-market approach to allocation with priorities’, *American Economic Journal: Microeconomics* .
- Hoxby, C. M. (2001), ‘All school finance equalizations are not created equal’, *Quarterly Journal of Economics* pp. 1189–1231.
- Hylland, A. & Zeckhauser, R. (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., Johnson, R. C. & Persico, C. (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. & Manjunath, V. (2012), ‘The difference indifference makes in strategy-proof allocation of objects’, *Journal of Economic Theory* **147**(5), 1913–1946.
- Johnson, R. C. & Jackson, C. K. (2017), Reducing inequality through dynamic complementarity: Evidence from head start and public school spending, Technical report, National Bureau of Economic Research.
- Kapor, A., Neilson, C. & Zimmerman, S. (2016), ‘Heterogeneous beliefs and school choice’.
- Karlin, S. & Taylor, H. M. (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.
- Krueger, A. B. (1999), ‘Experimental estimates of education production functions’, *The quarterly journal of economics* **114**(2), 497–532.
- Lafortune, J., Rothstein, J. & Schanzenbach, D. W. (2016), School finance reform and the distribution of student achievement, Technical report, 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. & Pycia, M. (2014), ‘Prices and efficient assignments without transfers’.
- Morrill, T. (2013), ‘An alternative characterization of top trading cycles’, *Economic Theory* **54**(1), 181–197.
- Morrill, T. (2015a), ‘Making just school assignments’, *Games and Economic Behavior* **92**, 18–27.
- Morrill, T. (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. & Sönmez, T. (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. & Ünver, M. U. (2017), ‘Incentive compatible allocation and exchange of discrete resources’, *Theoretical Economics* **12**(1), 287–329.
- Saban, D. & Sethuraman, J. (2013), House allocation with indifferences: a generalization and a unified view, *in* ‘Proceedings of the fourteenth ACM conference on Electronic commerce’, ACM, pp. 803–820.
- Shapley, L. & Scarf, H. (1974), ‘On cores and indivisibility’, *Journal of mathematical economics* **1**(1), 23–37.
- Small, K. A. & Rosen, H. S. (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 a 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 (2) and (3), 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, 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 time 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 9 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 9: 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 offered³¹ 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} . We provide an illustrative example with two schools in Figure 10. 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

³¹Strictly speaking, the measure of students assigned to each school during the time step is equal to the measure of seats at that school which were claimed by the student offered the seat or traded by the student offered the seat during the time step (not the measure of seats offered). A seat can be offered but not claimed or traded in one of two ways. The first occurs when the seat is offered at time τ but not yet claimed or traded. The second is when a student is offered two or more seats at the same time, and trades only one of them. Both of these sets are of η -measure 0 under our assumptions, and thus the measure of seats claimed or traded is equal to the measure of seats offered.

b -ranks between $\gamma_b(\tau_1) = x$ and $\gamma_b(\tau_2) = x - \varepsilon d_b$. To capture this set of students, let

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

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.³²

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 .




Figure 10: The set of students that are assigned during a small time step between τ_1 and τ_2 . The dot indicates $\gamma(\tau_1) = x$. The highlighted areas indicate the students $T_b^c(x, \varepsilon d_b)$ who are offered a seat during this step. Student in the blue (red) region receive an offer from school 1 (school 2). The pattern indicates whether a student received an offer from his preferred school. Trade balance is satisfied when there is an equal mass of students in the checkered regions.

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 offered a seat at c ,

³²The 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.

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 that³³

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

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 $\mathbf{d} = \gamma'(\tau_1)$ 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 Solutions to 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

³³These 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 .

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 progresses down school b 's priority list at rate d_b . Figure 11 illustrates such a graph for $\mathcal{C} = \{1, 2, 3, 4\}$. Given a collection of cycles let d_b be the number of cycles containing node b . It is straightforward that any node weights \mathbf{d} obtained in this way give a zero-sum flow, i.e. total flow into each node is equal to the total flow out of the node. Standard arguments from network flow theory show that the opposite also holds, that is, 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.




Figure 11: 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.

To characterize the set of solutions to the trade balance equations 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, 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 11 for an example. There exists at least one recurrent communication class, and two different communication classes have 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. As an illustrative example, we provide the following result for when η has full support.³⁴ In this case, 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*

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

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

On the Multiplicity of TTC Paths

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_i . 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.

³⁴ η has full support if for every open set $A \subset \Theta$ we have $\eta(A) > 0$.

As an illustration, consider the unique solution \mathbf{d}^K for the communicating class $K = \{1, 2\}$, as illustrated in Figure 11. 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 schools' 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).

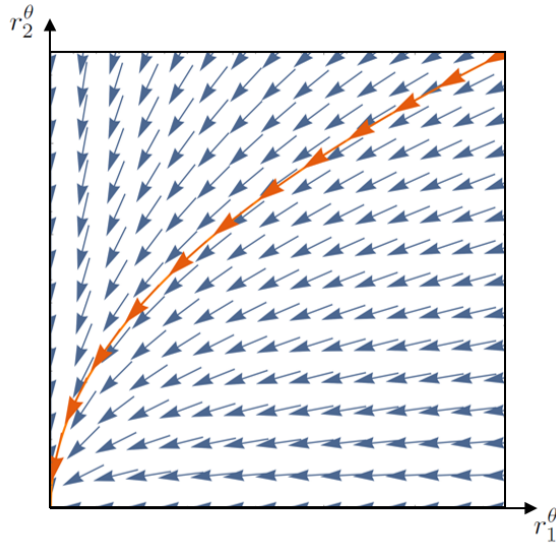


Figure 12: Illustration of the gradient field $\mathbf{d}(\cdot)$ and path $\gamma(\cdot)$ (ignoring the capacity equations).

Figure 12 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.

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

$$D^c(\gamma(\tau)) = \eta(\{\theta \mid r^\theta \notin \gamma(\tau), Ch^\theta(\mathcal{C}) = c\}).$$

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 school 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 presents 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 the trade balance and capacity equations, and that these necessary conditions, together with the capacity equations (3), 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

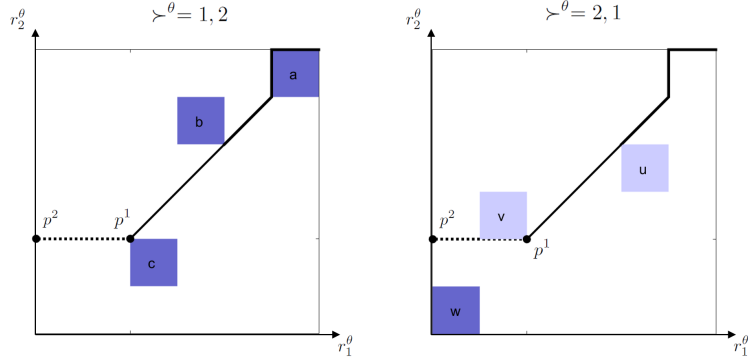
$$\begin{aligned}
 1 & : a \succ u \succ b \succ c \succ v \succ w, & a, b, c & : 1 \succ 2, \\
 2 & : a \succ b \succ u \succ v \succ c \succ w, & u, v, w & : 2 \succ 1.
 \end{aligned}$$

In Figure 13, 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

In this section, we calculate the TTC paths γ_{all} , γ_1 and γ_2 . We consider only solutions d to the trade balance equations (2) that have been normalized so that $d \cdot \mathbf{1} = -1$. For brevity we call such solutions *valid directions*. The relevant valid directions are shown in Figure 14.


We first calculate the TTC path in the regions where the TTC paths are the same. 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 13: 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 box prefer school 1 and students in the right box 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 dark blue are assigned to school 1 and students shaded dark light 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 14: 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 gray square. The borders of the squares corresponding to the students are drawn using dashed gray 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). 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 $\begin{bmatrix} \frac{1}{6} & 0 \\ 0 & \frac{1}{6} \end{bmatrix}$. 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 $\begin{bmatrix} \frac{1}{6} & 0 \\ 0 & 0 \end{bmatrix}$. The same argument with the coordinates swapped gives that $H = \begin{bmatrix} 0 & 0 \\ 0 & \frac{1}{6} \end{bmatrix}$ 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.