

The Cutoff Structure of Top Trading Cycles in School Choice*

Jacob D. Leshno[†] Irene Lo[‡]

June 24, 2020

Abstract

This paper develops a tractable theoretical framework for the Top Trading Cycles (TTC) mechanism for school choice that allows quantifying welfare and optimizing policy decisions. We compute welfare for TTC and Deferred Acceptance (DA) under different priority structures, and find that the choice of priorities can have larger welfare implications than the choice of mechanism. We solve for the welfare-maximizing distributions of school quality for parametrized economies, and find that optimal investment decisions can be very different under TTC and DA.

Our framework relies on a novel characterization of the TTC assignment in terms of a cutoff for each pair of schools. These cutoffs parallel prices in competitive equilibrium, with students' priorities serving the role of endowments. We show that these cutoffs can be computed directly from the distribution of preferences and priorities in a continuum model, and derive closed-form solutions and comparative statics for parameterized settings. The TTC cutoffs clarify the role of priorities in determining the TTC assignment, but also demonstrate that TTC is more complicated than DA.

*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, Canice Prendergast, Utku Unver, Laura Doval, Parag Pathak, Jonah Rockoff, Eric Budish, Yusuke Narita, Fuhito Kojima, Alvin Roth, two anonymous referees, the editor, and seminar participants in Berkeley, Columbia, Stanford, Yale, University of Toronto, Duke, NYU, Market Design NBER, Matching in Practice workshop, and Cirq Montreal Microeconomics meeting for their helpful comments.

[†]University of Chicago Booth School of Business, email: jacob.leshno@ChicagoBooth.edu

[‡]Stanford University, email: ilo@stanford.edu

1 Introduction

School choice mechanisms are commonly used to determine school admission based on student preferences and school priorities. In their seminal paper, Abdulkadiroğlu & Sönmez (2003) proposed two candidate mechanisms for school choice, the Deferred Acceptance (DA) mechanism and the Top Trading Cycles (TTC) mechanism (Gale & Shapley 1962, Shapley & Scarf 1974). These mechanisms have been considered and implemented by school districts across the US (Abdulkadiroğlu & Sönmez 2003, Abdulkadiroğlu et al. 2005, Pathak & Sönmez 2013, IIPSC 2017).

In choosing between TTC and DA, school districts face a non-trivial trade-off between efficiency and transparency, among other considerations. Boston public schools commended TTC for its efficiency, but were concerned about its lack of transparency, citing uncertainty as to whether priorities under TTC were serving their intended role.¹ Most school districts have favored DA, which produces a stable assignment in which a student’s priority plays an intuitive role in determining her assignment.² However, DA does not generally produce an efficient assignment.

This paper provides a theoretical framework for clarifying this trade-off. We develop a characterization of TTC that allows us to theoretically quantify the efficiency gains of TTC over other mechanisms. It relies on a novel cutoff representation of the TTC outcome that clarifies the role of priorities in determining each student’s assignment, but also reaffirms some of the aforementioned concerns. Such cutoff characterizations have also proved useful for price-theoretic analysis and empirical work (Abdulkadiroğlu, Angrist, Narita & Pathak 2017, Agarwal & Somaini 2018, Kapor et al. 2020).

Empirical work suggests that in some settings TTC can lead to significantly higher welfare than DA (Abdulkadiroğlu, Agarwal & Pathak 2017, Calsamiglia et al. 2020), while in other examples TTC and DA produce very similar results (Abdulkadiroğlu et al. 2009, Pathak 2017). By developing a tractable framework with a continuum of students we are able to provide, to our knowledge, the first theoretical quantification of welfare under TTC. We theoretically show that TTC can lead to large efficiency gains over DA when preferences and priorities are uncorrelated. Our results contrast

¹Boston 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 [in TTC] makes the student assignment process less transparent.*”

²In the words of Boston Public School, “*students will receive their highest choice among their school choices for which they have high enough priority to be assigned*” (BPS 2005).

with the finding of Che & Tercieux (2019) that TTC and DA lead to similar welfare when the number of agents and items are both large and preferences are uncorrelated.

Our tools allow us to analyze the effects of correlation between priorities and preferences. Pycia (2019) shows that in the absence of such correlation, all strategyproof and efficient mechanisms produce similar outcome statistics. In many school choice settings priorities are set based on student characteristics that are correlated with their preferences, for example proximity to school. We evaluate TTC and DA in such a setting, and show that a priority design that correlates with student preferences can lead to higher welfare. Moreover, the choice of priority structure can have a larger welfare effect than the choice of mechanism. We also find that TTC and DA lead to similar outcomes when students have a preference and priority for their neighborhood school, corroborating a conjecture by Pathak (2017).

The tractability of the framework is enabled by a novel cutoff characterization of the TTC assignment. We show that the TTC assignment can be concisely described by cutoffs $\{p_b^c\}$, one for each pair of schools b, c . These cutoffs parallel prices in competitive equilibrium, with students' priorities serving the role of endowments. We show that these cutoffs can be tractably calculated in a continuum model. In addition, the cutoff representation yields for each student a budget set of schools at which she gained admission, and these budget sets allow for tractable expressions for welfare under random utility models. We provide closed form expressions for the cutoffs and for student welfare in parametric settings.

Our cutoff characterization clarifies the role of priorities in determining the TTC assignment. 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. Thus, priority at a school allows a student to gain admission at a range of schools.

While our cutoff characterization provides a simpler description of the TTC assignment, we show there is no simpler cutoff representation of TTC. Thus, the concern that TTC is 'less transparent' than DA remains valid, as DA can be described using a cutoff for each school, while TTC requires a cutoff for each pair of schools.

Using our cutoff characterization, we can leverage tools from price theory to understand how policy decisions affect the TTC assignment. We provide comparative statics with respect to changes in the popularity of schools and identify which students gain or lose. We evaluate policy decisions such as the design of priorities and investment in school quality and find significant implications for welfare.

Applying these tools to optimize a school district’s investment in schools reveals that optimal investment decisions under DA and TTC can be very different. We set up a model where students have a common preference for quality, which depends on school-specific investment,³ as well as horizontal taste shocks. Comparative statics reveal a tension between maximizing choice and investing in cost-effective schools where investment generates the most added quality. This is because a larger difference in the vertical quality of the schools leads to smaller budget sets for students and less efficient sorting on horizontal taste shocks.

We quantify this trade-off, and find that these effects have different magnitudes under DA and TTC. A simple parameterized economy shows these differences can lead to differences in optimal investment. Under TTC optimal investment makes all schools equally popular, maximizing student choice and welfare gains from horizontal sorting. Under DA the optimal investment is focused on the most cost-effective school, because DA does not allow as much choice and achieves less welfare gains from horizontal sorting.

Given that priorities serve a different role under DA and TTC, a school district that adapts TTC may want to accordingly change the priority structure. We analyze how changes to the priority structure affect the TTC assignment and find that TTC priorities are “bossy” in the sense that changes in priority among a set of students that do not alter their assignment can nonetheless alter the assignment of other students. This can be seen as another undesirable property of TTC, and suggests challenges in designing priority structures for TTC.

A few technical aspects of the analysis may be of interest. The tractability of our framework relies on a novel approach to analyzing TTC in terms of *trade balance equations*. 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 trade balance equations also circumvent many of the measure theoretic complications in

³Examples 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 2019), 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.

defining TTC in the continuum. Finally, we make use of a connection to Markov chain theory 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.

Cutoff representations have been instrumental for empirical work on DA and related mechanisms. Abdulkadiroğlu, Angrist, Narita & Pathak (2017) use admission cutoffs to construct propensity score estimates. Agarwal & Somaini (2018), Kapor et al. (2020) 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.

The seminal paper of (Shapley & Scarf 1974) introduced TTC (attributed to David Gale) as a method to find competitive prices when each agent has unit demand and endowment of a single indivisible good. 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 differs in that it requires a lower dimensional set of cutoffs that does not grow with the number of positions, and in that it provides a method for directly calculating these cutoffs which allows for welfare calculations and comparative statics.

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. 2019, 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 5.2 imply that the TTC cutoffs depend on the entire distribution and cannot be computed from simple supply and demand equations.

Several papers also study TTC in large markets. Hatfield et al. (2016) show that

even in a large market it is possible that a school is assigned less preferred students when it improves its quality. Our framework allows us to provide comparative statics that quantify the changes in the TTC assignment and identify the affected students. We examine a parameterized setting and find that newly assigned students include students of arbitrarily low priority. Che & Tercieux (2019, 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. Pycia (2019) shows that absent correlation between preferences and priorities, all strategyproof and efficient mechanisms produce similar outcome statistics. The results in Section 4 show that TTC has different properties in these different large markets.

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 methods 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 (2014) introduce Equitable TTC, a variation on TTC that aims to reduce inequity. In Section 5.2 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.

He et al. (2018) propose an alternative pseudo-market approach for discrete assignment problems that extends Hylland & Zeckhauser (1979) and also uses admission cutoffs. Miralles & Pycia (2020) 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.

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 technical contributions that allow for direct calculation of the TTC cutoffs, assignment and welfare. Section 4 applies our framework to calculate and compare welfare under DA and TTC, analyze allocation of resources to schools, and provide conditions under which TA and TTC give a similar assignment. Section 5 provides additional results on comparative statics, closed form expressions, and inspects effects of changes in priorities. Section 6 provides concluding remarks.

Appendix A provides the technical intuition for the continuum TTC model. Appendix B.1 shows that the continuum TTC model is consistent with the discrete TTC model and 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. It will be convenient to represent the priority of student s at school c by the student's percentile rank $r_c^s = |\{s' \mid s \succ_c s'\}| / |\mathcal{S}|$ in the school's priority ordering. Note that for any two students s, s' and school c we have that $s \succ_c s' \iff r_c^s > r_c^{s'}$ and that $0 \leq r_c^s < 1$.

A *feasible* assignment is $\mu : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$ where $|\mu^{-1}(c)| \leq q_c$ for every $c \in \mathcal{C}$. If $\mu(s) = c$ we say that s is assigned to c , and we use $\mu(s) = \emptyset$ to denote that the

student s is unassigned. As there is no ambiguity, we let $\mu(c)$ denote the set $\mu^{-1}(c)$ for $c \in \mathcal{C} \cup \{\emptyset\}$. A *discrete economy* is $E = (\mathcal{C}, \mathcal{S}, \succ^{\mathcal{S}}, \succ^{\mathcal{C}}, q)$, where \mathcal{C} is the set of schools, \mathcal{S} is the set of students, $q = \{q_c\}_{c \in \mathcal{C}}$ is the capacity of each school, and $\succ^{\mathcal{S}} = \{\succ^s\}_{s \in \mathcal{S}}$, $\succ^{\mathcal{C}} = \{\succ_c\}_{c \in \mathcal{C}}$.

Given an economy E , the discrete Top Trading Cycles algorithm (TTC) calculates an assignment $\mu_{dTTC}(\cdot | E) : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$. 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.⁴*
- *Remove the assigned students from S , reduce the capacity of the schools they are assigned to by 1, and remove schools with no remaining capacity from C .*

TTC satisfies a number of desirable properties. An assignment μ is *Pareto efficient for students* if no group of students can improve by swapping their allocations, and no individual student can improve by swapping her assignment for an unassigned object. A mechanism is *Pareto efficient for students* if it always produces an assignment that is Pareto efficient for students. A mechanism is *strategy-proof for students* if reporting preferences truthfully is a dominant strategy. It is well known that the TTC school choice mechanism 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).

⁴Such a trading cycle must exist, since every vertex in the pointing graph with vertex set $S \cup C$ has out-degree 1.

2.2 Cutoff Characterization and the Role of Priorities

Our first 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*

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

where p_b^c is the percentile in school b 's ranking of the worst ranked student at school b that traded a seat at school b for a seat at school c 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 a concise way of communicating the TTC assignment to students. 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. 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. 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. Numbers indicate students' preferred schools, and all students find both schools acceptable. Colors indicate students' assigned schools 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.




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.

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 $\mathbf{p}_b = \{p_b^c\}_{c \in C}$. 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 cut-

offs \mathbf{p}_b that share the same ordering over schools for every b .⁵ 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*

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

and for any school b the cutoffs are ordered,⁶

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

Therefore, under such relabeling the set of schools $B_b(s, \mathbf{p})$ student s can afford via her priority at school b is either the empty set ϕ or $B_b(s, \mathbf{p}) = \mathcal{C}^{(c)} = \{c, c + 1, \dots, n\}$ for some $c \leq b$. Moreover, each student's budget set $B(s, \mathbf{p}) = \cup_b B_b(s, \mathbf{p})$ is either $B(s, \mathbf{p}) = \phi$ or $B(s, \mathbf{p}) = \mathcal{C}^{(c)}$ for some c .

When there exist TTC cutoffs that satisfy inequality (1) we say that the schools are *labeled in order*, and for all pairs of schools $c \leq b$ we say that b is *more affordable* than c . 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 .

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 $\{p_b^c\}$

⁵Similar orderings have previously been identified by Liu & Pycia (2013) for asymptotically efficient mechanisms, and by Bogomolnaia & Moulin (2001) for the Probabilistic Serial mechanism.

⁶The cutoffs \mathbf{p} defined in Theorem 1 do not necessarily satisfy this condition. However, the following relabeling of schools and cutoffs $\tilde{\mathbf{p}}$ gives 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 $\tilde{p}_b^c = \min_{a \leq c} p_b^a$.




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 $\mathcal{C}^{(1)}, \mathcal{C}^{(2)}, \dots, \mathcal{C}^{(b)}, \emptyset$ 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 .

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 Transparency of the TTC Assignment

Our cutoff characterization provides a more transparent way to communicate the TTC outcome. A student can determine her assignment using just the cutoffs and her priorities, as she has a school c in her budget set if and only if her priority at some school b meets the cutoff p_b^c . Hence the school district can publish the TTC cutoffs and let students verify that they were assigned to their most preferred school in their budget set. This also suggests the following non-combinatorial description of the TTC assignment. For each school b , students receive b -tokens according to their priority at school b , and students with higher b -priority receive more b -tokens. The TTC algorithm publishes cutoffs $\{p_b^c\}$. Students can use one kind of token and purchase a single school, and the required number of b -tokens to purchase school c is p_b^c . Theorem 1 shows the cutoffs can be observed after running TTC.⁷

Thus our characterization makes the 'behind the scenes mechanized trading' (BPS 2005) in TTC more transparent. The cutoffs make it simple to see, ex post, which priorities can be traded to get into which school: priority at a school b can be traded for a seat at school c if and only if it meets the cutoff p_b^c .

Our characterization also highlights some aspects of TTC which may make it seem less transparent. Example 1 shows that the TTC assignment cannot be expressed in terms of a single cutoff for each school, as the assignment cannot be described by fewer than 3 cutoffs. One can extend this to construct economies with n schools where the assignment cannot be described by fewer than $\frac{1}{2}n^2$ cutoffs. Hence the

⁷We thank Chiara Margaria, Laura Doval and Larry Samuelson for suggesting this explanation.

cutoff structure under TTC is provably more complicated than the cutoff structure under DA.⁸

In addition, the characterization highlights that a priority at a school can add many schools to the student’s budget set, which may not be the intended purpose of the priority (e.g., sibling priority at a popular school).

Finally, we note that while in general n^2 cutoffs are required to describe the TTC assignment, the structure of the TTC budget sets makes it possible for each student to consider only n of these cutoffs when *verifying* their assignment. Specifically, Proposition 1 provides a *popularity ordering* of schools under TTC: if a student does not have a school c in their budget set, then they do not have any of the more popular schools in their budget set (i.e. schools b with $b < c$ when schools are labeled in order). Hence a student s can verify that her assigned school $\mu(s)$ is her best possible assignment by considering only the least popular school c that she prefers to $\mu(s)$ (i.e. the school $c = \max \{c' \mid c' \succ^s \mu(s)\}$ if the schools are labeled in order) and checking whether she has sufficiently high priority at some school b to be assigned to school c . However, we note that revealing a popularity ordering of schools may be seen as another undesirable property of TTC.

2.5 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

⁸This can be formalized in the language of VC theory (Vapnik & Chervonenkis 1971) as follows. Suppose we let H^{TTC} be the set family containing: (1) the sets of priorities for which students would have school i in their budget set under TTC for some economy; and (2) the sets of priorities for which students would not have school i in their budget set under TTC for some economy, for a given school i in a finite set of schools $\mathcal{C} = \{1, 2, \dots, n\}$. Suppose we similarly define H^{DA} . Then the VC dimension of H^{TTC} is $n + 1$, and the VC dimension of H^{DA} is 2.

her reported preferences.^{9,10} 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 3.5 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.

3 TTC in a Continuum Model

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 $\mathcal{C} = \{1, \dots, n\}$, and each school $c \in \mathcal{C}$ has the capacity to admit a mass $q_c > 0$ of students. A student $\theta \in \Theta$ is given by $\theta = (\succ^\theta, r^\theta)$. We let \succ^θ denote the student's strict preferences over schools, and let $Ch^\theta(C) = \max_{\succ^\theta}(C)$ denote θ 's most preferred school out of the set C . The priorities of schools over students are captured by the vector $r^\theta \in [0, 1]^C$. We say that r_b^θ is the *rank* of student θ at school b , or the b -rank of student θ . Schools prefer students with higher ranks, i.e. $\theta \succ_b \theta' \Leftrightarrow r_b^\theta > r_b^{\theta'}$.

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, $\sum_{c \in \mathcal{C}} q_c < \eta(\Theta)$.¹¹ Finally, we make the following technical assumption that

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

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

¹¹This assumption guarantees the existence of stopping times in Definition 2. If there is sufficient school capacity for all students, Definition 2 part (2) needs to consider the case where all students are assigned (the TTC path hits the boundary) before any school fills its capacity, i.e., there exists

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, there exists a grid G ,¹² such that for each preference order \succ and intervals $\{(x_c, y_c)\}_{c \in \mathcal{C}}$, if the rectangle $R = \{(\succ, r) \in \Theta \mid x_c < r_c < y_c\}$ does not intersect any gridlines, then the density ν is Lipschitz continuous on R and is either uniformly zero or strictly positive on R .¹³

Lipschitz continuity allows us to use the Picard-Lindelöf Theorem to show existence of a TTC path that satisfies trade balance equations. The additional condition on ν is used in the proof of uniqueness of the TTC outcome, where it is used to define contiguous selection of trade cycles.

Assumption 1 is satisfied by all economies considered throughout the paper, and it is general enough to allow embeddings of any discrete economy. While this assumption is not without loss of generality, e.g. it is violated when all schools share the same priorities over students, we consider Assumption 1 to be a technical restriction that requires the representation of the economy to be sufficiently smooth.¹⁴

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 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

t such that $\gamma_c(t) = 0$ for some $c \in C$ and $D^{c'}(\gamma(t)) \leq q_{c'}$ for all $c' \in C$. It is straightforward to verify that this happens if and only if it is feasible to assign all students to their most preferred remaining school. In this case, all students are assigned to their most preferred remaining school, and all remaining cutoffs are 0.

¹²A finite 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.

¹³That is, if $R \subset \Theta \setminus G$, then either $\nu(\theta) > 0 \forall \theta \in R$ or $\nu(\theta) = 0 \forall \theta \in R$.

¹⁴We can represent an economy where 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.

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 .¹⁵ 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 Calculating the TTC Assignment

In this section, we establish that in the continuum model the TTC assignment can be directly calculated from trade balance and capacity equations. This allows us to provide closed-form expressions for the TTC cutoffs, budget sets and welfare for parametric economies. It also allows us to explain how the TTC assignment changes with changes in the underlying economy.

We remark that our results are also of technical interest. Directly translating the TTC algorithm to the continuum setting by considering individual trading cycles is challenging, as a direct adaptation of the algorithm would require the clearing of cycles of zero measure. We circumvent the technical issues raised by such an approach by formally defining the continuum TTC assignment in terms of trade balance and capacity equations, which characterize the TTC algorithm in terms of its aggregate behavior over multiple steps. To verify the validity of our definition,

¹⁵Formally

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

where 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 .

we show in Subsection B 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 $t^{(c^*)} \stackrel{def}{=} \min_c [t^{(c)}]$ denote the minimal time. For a point $\mathbf{x} \in [0, 1]^{\mathcal{C}}$, let

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

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*

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

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 $\tilde{\gamma}(\cdot) : [t^{(c^*)}, \infty) \rightarrow [0, 1]^{\tilde{\mathcal{C}}}$ by restricting $\gamma(\cdot)$ to $t \geq t^{(c^*)}$ and coordinates within $\tilde{\mathcal{C}}$. Then $\tilde{\gamma}$ and the stopping times $\{t^{(c)}\}_{c \in \tilde{\mathcal{C}}}$ satisfy 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.

In other words, 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 Appendix B that our characterization is consistent with 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 $\hat{\gamma}(\cdot), \{\hat{t}^{(c)}\}_{c \in \mathcal{C}}$ that satisfy the trade balance and capacity equations yield the same assignment μ_{cTTC} , given by*

$$\mu_{cTTC}(\theta) = Ch^\theta(B(\theta, \mathbf{p})),$$

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

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

and the budget set for each student θ is given by $B(\theta, \mathbf{p}) = \{c : r_b^\theta \geq p_b^c \text{ for some } b\}$.

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.¹⁶ 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.¹⁷

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,*

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

3.3 Example: Calculating TTC Budget Sets and Assignment

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

Example 2. We demonstrate how to use Theorem 2 to calculate the TTC budget sets and 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

¹⁶Continuity of the TTC path provides an initial condition for $\tilde{\gamma}$, namely $\tilde{\gamma}_c(t^{(c^*)}) = \gamma_c(t^{(c^*)}) \forall c$.

¹⁷See footnote 11.

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

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 .

Calculating the TTC Cutoffs and Assignment 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 $0 \leq t \leq \min\{t^{(1)}, t^{(2)}\}$ 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 $t^{(1)} < t^{(2)}$.¹⁹ Therefore, we solve $D^1(\gamma(t^{(1)})) = q$ to get that $t^{(1)} = 1 - \left(\frac{p-q}{p}\right)^p$ and that

¹⁸The original trade balance equations are

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

¹⁹Otherwise, we have that $t^{(2)} = \min\{t^{(1)}, t^{(2)}\}$ and $D^1(\gamma(t^{(2)})) \leq q$, implying that $D^2(\gamma(t^{(2)})) = \frac{1-p}{p} D^1(\gamma(t^{(2)})) < q$.

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

For the remaining cutoffs, we eliminate school 1 and reiterate the same steps for the residual economy where $\tilde{\mathcal{C}} = \{2\}$ and $\tilde{q}_2 = q_2 - D^2(\gamma(t^{(1)})) = q(2 - 1/p)$.

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

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

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 $p_1^1 = p_1^2 \approx .54$, $p_2^1 \approx .73$ and $p_2^2 \approx .37$, which are approximately the same cutoffs as those for the discrete economy in Example 1.

Students in the upper L-shaped region $\{\theta \mid r_1^\theta \geq (1 - q/p)^p \text{ or } r_2^\theta \geq (1 - q/p)^{1-p}\}$ are assigned to their favorite school, students in the light blue rectangular region $\{\theta \mid r_1^\theta < (1 - q/p)^p \text{ and } r_2^\theta \in [(1 - 2q)(1 - q/p)^{-p}, (1 - q/p)^{1-p}]\}$ are assigned to school 2, and all other students are unassigned.

Necessity of the Trade Balance Equations 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 5.2 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.

3.4 Calculating TTC Welfare from Budget Sets

Model with Cardinal Utilities We are interested in quantifying the welfare of the TTC assignment. In order to do so, we augment the model from Section 3.1 to allow students to have cardinal utilities. A student $\tilde{\theta} \in \tilde{\Theta}$ is given by $\tilde{\theta} = (u^{\tilde{\theta}}, r^{\tilde{\theta}})$, where for each school c the quantity $u^{\tilde{\theta}}(c)$ denotes the cardinal utility for student $\tilde{\theta}$ of attending school c . Each cardinal-utility student type $\tilde{\theta} \in \tilde{\Theta}$ induces a student type $\theta(\tilde{\theta}) = (\succ, r^{\tilde{\theta}}) \in \Theta$ satisfying $b \succ c$ if and only if $u^{\tilde{\theta}}(b) > u^{\tilde{\theta}}(c)$.²⁰ The measure $\tilde{\eta}$ over $\tilde{\Theta}$ specifies the distribution of cardinal-utility student types, and induces a measure η , defined by $\eta(A) = \tilde{\eta}(\{\tilde{\theta} | \theta(\tilde{\theta}) \in A\})$ over Θ . We require that η satisfies Assumption 1. The TTC cutoffs and assignment depend only on ordinal preferences and can therefore be computed using the induced measure η .

Given an assignment μ , the welfare from the assignment is given by

$$W(\mu) = \int_{\tilde{\theta} \in \tilde{\Theta}, \mu(\tilde{\theta}) \neq \phi} u^{\tilde{\theta}}(\mu(\tilde{\theta})) d\tilde{\eta}.$$

We let $W_{TTC} = W(\mu_{TTC})$ denote the welfare from the TTC assignment.

Budget Sets and Welfare under TTC Given a set of cutoffs $\{p_b^c\}$, let $B(\tilde{\theta}, \mathbf{p}) = B(\theta(\tilde{\theta}), \mathbf{p})$ denote the budget set of a student of type $\tilde{\theta}$. The TTC assignment for a student $\tilde{\theta}$ can be defined in terms of the student's budget set $B(\tilde{\theta}, \mathbf{p})$ by

$$\mu_{cTTC}(\tilde{\theta}) = Ch^\theta(B(\tilde{\theta}, \mathbf{p})),$$

where \mathbf{p} are the TTC cutoffs and $\theta = \theta(\tilde{\theta})$ is the induced student type with ordinal preferences. These budget sets allow for a simple expression for welfare. For each set of schools $C \subset \mathcal{C}$, let

$$U^C(\tilde{\theta}) = \mathbb{E} \left[\max_{c \in C} \{u^{\tilde{\theta}}(c)\} \right]$$

denote the expected utility of a student $\tilde{\theta}$ with budget set C , and let

$$A_C = \{\tilde{\theta} | B(\tilde{\theta}, \mathbf{p}) = C\}$$

²⁰To guarantee strict preferences, if $u^{\tilde{\theta}}(b) = u^{\tilde{\theta}}(c)$ let $b \succ c$ iff $b > c$, where $>$ is an arbitrary fixed strict ranking of schools.

denote the set of students with budget set C under TTC, where \mathbf{p} are the TTC cutoffs. Then for each set of schools C and each student $\tilde{\theta} \in A_C$ the assignment of $\tilde{\theta}$ under TTC is given by $\mu_{cTTC}(\tilde{\theta}) = \operatorname{argmax}_{c \in C} \{u^{\tilde{\theta}}(c)\}$, and social welfare under the TTC assignment is equal to

$$W_{TTC} = \sum_{C \subset C} \int_{A_C} U^C(\tilde{\theta}) d\tilde{\eta}.$$

We illustrate how to calculate welfare under TTC in Example 2 for a cardinal utility model in Section 4.1.

3.5 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

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

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

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

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

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

for all students $s \in \mathcal{S}$. Moreover, there exists $\mathbf{p}^ \in \mathcal{P}(E)$ such that for every student $s \in \mathcal{S}$ the budget sets are given by $B^*(s | E) = B(s, \mathbf{p}^*)$.*

Proposition 2 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 2 and constructively find \mathbf{p}^* .

4 Applications: Comparing DA and TTC

The results of Section 3 provide us with tractable expressions for assignment and welfare under TTC. We apply these tools to analyze the design of policy levers under the TTC and DA assignment mechanisms. Section 4.1 applies our results to compare welfare across different priority structures under TTC and DA. The tractable expressions also allow us to analyze economies where preferences are endogenously determined by policy decisions, such as investment in school quality. Section 4.2 applies our results to evaluate investment in school quality that affects student utility and choice of assigned school, and compare optimal investment decisions under TTC and DA. Section 4.3 gives conditions under which TTC and DA produce a similar assignment.

4.1 Welfare under TTC vs. DA

We explore the effects of priority design in a stylized economy where the student's distance to a school is correlated with the student's preferences for the school.²¹ We compare DA and TTC under random priority as well as under a priority system that exploits observed distance. We calculate each of the resulting assignments and compare them in terms of the resulting welfare and total distance traveled to the assigned schools.

Formally, we consider an economy with a unit mass of students and two schools $\mathcal{C} = \{1, 2\}$ with capacities $q_1 = q_2 = 0.4$. The utility of a student is determined by the distances $d_1^{\tilde{\theta}}, d_2^{\tilde{\theta}}$ to each school and a random preference shock ε : $u^{\tilde{\theta}}(1) = 2 - d_1^{\tilde{\theta}} + \varepsilon^{\tilde{\theta}}$ for school 1 and $u^{\tilde{\theta}}(2) = 2 - d_2^{\tilde{\theta}}$ for school 2. For tractability, distances are distributed $(d_1^{\tilde{\theta}}, d_2^{\tilde{\theta}}) \sim U[0, 1]^2$ and $\varepsilon^{\tilde{\theta}}$ is distributed $\varepsilon^{\tilde{\theta}} \sim U[-1, 2]$ independently of $d_1^{\tilde{\theta}}, d_2^{\tilde{\theta}}$. The utility of being unassigned is normalized to 0, and all schools are acceptable to all

²¹The empirical matching literature commonly assumes students have a disutility from distance traveled to school, see for example Abdulkadiroğlu, Agarwal & Pathak (2017).

students.




Figure 4: Welfare and distance attained under the different assignments. As benchmarks, the plot includes the random assignment, the welfare-maximizing assignment (which can be obtained by a competitive market with monetary transfers), and the distance-minimizing assignment among all assignments that fill both schools.

Uncorrelated Priority

In this case, priority is distributed $(r_1^{\tilde{\theta}}, r_2^{\tilde{\theta}}) \sim U[0, 1]^2$ independently of $d_1^{\tilde{\theta}}, d_2^{\tilde{\theta}}, \varepsilon^{\tilde{\theta}}$. The probability that a randomly drawn student prefers school 1 is $2/3$, and therefore the joint distribution of ordinal preference and priorities is identical to that of the economy in Example 2. The TTC assignment depends only on ordinal preferences, and we can therefore use the calculation in Example 2 to get that the TTC cutoffs are $p_1^1 = p_1^2 \approx .54$, $p_2^1 \approx .73$ and $p_2^2 \approx .37$. Under these cutoffs a mass $N^{\{1,2\}} = 1 - p_1^1 \cdot p_2^1 = 0.6$ of students will have the budget set $\{1, 2\}$, another mass $N^{\{2\}} = p_1^1 (p_2^1 - p_2^2) = 0.2$ of students will have the budget set $\{2\}$, and the remaining 0.2 mass of students remain unassigned. Thus, welfare under TTC with uncorrelated priorities is

$$\begin{aligned}
 W_{TTC, uncorr} &= \sum_{CC} \int_{\tilde{\theta} \in A_C} u^{\tilde{\theta}}(C) d\tilde{\eta} = N^{\{1,2\}} \cdot \mathbb{E}[u^s(\{1, 2\})] + N^{\{2\}} \cdot \mathbb{E}[u^s(\{2\})] \\
 &= 0.6 \cdot 79/36 + 0.2 \cdot 3/2 = 1.617.
 \end{aligned}$$

The total distance traveled is $Dist_{TTC, uncorr} = 0.367$, which is lower than 0.4 because students who have the full budget set $\{1, 2\}$ are more likely to prefer and be assigned

to the closer school.

Under DA assignment only a mass 0.3 of students have the full budget set $\{1, 2\}$, and the remaining students are either unassigned or can choose a single school.²² Welfare is $W_{DA,uncorr} = 1.508$ and distance traveled is $Dist_{DA,uncorr} = 0.383$. Because fewer students are able to choose their favorite school, the DA assignment is worse than the TTC assignment in terms of both welfare and distance.

Figure 4 shows that welfare under both DA and TTC when priorities are uncorrelated with student preferences is comparable to the welfare given by the assignment that minimizes the total distance traveled, but results in almost double the average distance. In other words, school choice mechanisms with poorly designed priorities can perform poorly.

Proximity-Based Priority

School districts often wish to assign students to school that are close to their home, and use the observable distance to school to prioritize students who live closer to the school. We capture this by setting the priority of student s to be $(r_1^s, r_2^s) = (1 - d_1^s, 1 - d_2^s)$. This implies a correlation between having priority at a school and preference for a school. For example,

$$Prob(1 \succ 2 \mid r_1^s = x_1, r_2^s = x_2) = \frac{1}{3}(1 + x_2 - x_1).$$

Welfare under TTC is $W_{TTC,dist} = 1.7$ and distance traveled is $Dist_{TTC,dist} = 0.263$ (details of the calculations and Mathematica code are in Appendix E.1). Under DA welfare is $W_{DA,dist} = 1.679$ and distance traveled is $Dist_{DA,dist} = 0.215$. That is, TTC yields slightly higher welfare than DA but results in higher distance traveled. This gives support to the concerns about TTC originally expressed by BPS. If priorities are set by the district with the intent of minimizing distance traveled, then TTC increases welfare but shifts away from the intended goal of the priorities.

²²We find the DA assignment by solving for cutoffs (c_1, c_2) that solve the market clearing equations (Azevedo & Leshno 2016)

$$D_1(c_1, c_2) = (1 - c_1)(c_2 + 2/3(1 - c_2)) = q_1 \quad D_2(c_1, c_2) = (1 - c_2)(c_1 + 1/3(1 - c_1)) = q_2.$$

The unique solution is $c_1 = 0.5, c_2 = 0.4$.

Importance of the Priority Structure

Comparing across both priority structures, we see that the choice between TTC and DA has a smaller effect than switching from uncorrelated priorities to distance based priorities. In fact, both mechanisms with proximity-based priorities yield better welfare and distance than either mechanism with uncorrelated priorities. In other words, designing priorities that appropriately reflect student welfare can create larger benefits than choosing between mechanisms.

We also observe that TTC can yield higher utility than DA even when priorities are correlated with preferences. This is because the presence of taste shocks means that student preferences are not fully aligned with priorities. However, this may shift away from the intended goals of these priorities.

4.2 Optimal Investment

In this section, we consider a school district's problem of allocating resources to improve school, when 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 2019) 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 evaluate student welfare and solve for the optimal distribution of school quality under different mechanisms. In doing so, we quantify the trade-off between maximizing choice, and directing investment to the most cost-effective schools (which generate highest direct returns from investment). We find that this trade-off can result in very different optimal investment strategies under DA and TTC: we provide an example where the optimal distribution of quality under TTC is equitable, in that it makes all schools equally over-demanded; whereas the optimal distribution of quality under DA is targeted at the most cost-effective school.

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. We assume for all δ that η_δ has a Lipschitz continuous non-negative density ν_δ that satisfies Assumption 1 and depends smoothly on δ . For a given δ , let μ_δ and $\{p_b^c(\delta)\}_{c \in \mathcal{C}}$ denote the TTC assignment and associated cutoffs for the economy \mathcal{E}_δ . We omit the dependence on δ when it is clear from context.

We focus on the logit economy, which is a parametrized economy that incorporates quality-dependent preferences. It does so by combining heterogeneous idiosyncratic taste shocks with a common preference modifier δ_c .

Definition 3. A *logit economy* is an economy $\mathcal{E} = (\mathcal{C}, \Upsilon, v, q)$ with quality-dependent preferences, 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, i.e. $u^s(c | \delta) = \delta_c + \varepsilon_{cs}$ where ε_{cs} are i.i.d. extreme value variables shifted to have a mean of 0 (McFadden 1973). Schools' priorities are uncorrelated and uniformly distributed, and all students prefer any school to being unassigned.²³

Optimal investment under DA, TTC, and neighborhood assignment

We first provide two illustrative examples showing the welfare-optimal quality distribution for the logit economy under DA and TTC, as well as under neighborhood assignment, where students are assigned to a fixed school regardless of their preferences. These examples allow us to compare welfare across mechanisms, and illustrate the differences in optimal investment under these mechanisms. In the examples below we rely on closed-form expressions for cutoffs, budget sets and welfare in the logit

²³Formally, $u^s(\phi | \delta) = -\infty$. For welfare calculations we only consider assigned students.




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

economy, which we provide in Section 5.1. We also fix the school labels and consider various δ . For some values of δ the schools may not be labeled in 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. The district invests in quality levels δ , subject to the constraints that $\delta_1 + \delta_2 = 2$ and $\delta_1, \delta_2 \geq 0$. Budget sets and welfare under different investment levels can be seen in Figure 5.

Under neighborhood assignment $W_{NH}/Q = 1$ for any choice of δ_1, δ_2 , because the expected utility of a random student assigned to c is δ_c . Under TTC the unique optimal investment is $\delta_1 = \delta_2 = 1$. Any assigned student has the budget set $B = \{1, 2\}$ and is assigned to the school for which she has higher idiosyncratic taste. This yields welfare $W_{TTC}/Q = 1 + \mathbb{E}[\max(\varepsilon_{1s}, \varepsilon_{2s})] = 1 + \ln(2) \approx 1.69$. Welfare is lower when $\delta_1 \neq \delta_2$, because fewer students choose the school for which they have higher idiosyncratic taste; e.g. given $\delta_1 = 2, \delta_2 = 0$ welfare is $W_{TTC}/Q = \frac{1}{2}(1 + e^{-2}) \log(1 + e^2) \approx 1.20$.

Under DA the unique optimal investment is also $\delta_1 = \delta_2 = 1$, yielding $W_{DA}/Q = 1 + \frac{1}{3} \ln(2) \approx 1.23$. This is strictly lower than W_{TTC}/Q because under DA only students

that have sufficiently 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.

In Example 3, both TTC and DA optimize welfare from investment by investing equally ($\delta_1 = \delta_2 = 1$). 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.²⁴ 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 $q_1 > q_2$. To make investment in school 1 more cost-effective, we assume that (despite having more students) school 1 requires the same amount of resources to increase its quality for all its students. We also keep the constraint that $\delta_1 + \delta_2 = 2$. It is straightforward to see that under neighborhood assignment the welfare optimal investment is $\delta_1 = 2, \delta_2 = 0$, and we find that this investment is also welfare-optimal under DA. In contrast, we find the welfare-optimal investment under TTC is equitable, in the sense that it makes all schools equally overdemanded and gives all students full budget sets.

Example 4. Consider a logit economy with two schools and $q_1 = 1/2, q_2 = 1/4$, and let $Q = q_1 + q_2$ denote the total capacity. The district invests in quality levels δ , subject to the constraints that $\delta_1 + \delta_2 = 2$ and $\delta_1, \delta_2 \geq 0$. Budget sets and welfare under different investment levels can be seen in Figure 6.

Under neighborhood assignment the welfare optimal investment is $\delta_1 = 2, \delta_2 = 0$, yielding $W_{NH}/Q = 4/3 \approx 1.33$. Under TTC the unique optimal investment is $\delta_1 = 1 + \frac{1}{2} \ln(2), \delta_2 = 1 - \frac{1}{2} \ln(2)$, yielding $W_{TTC}/Q = \ln(3e/\sqrt{2}) \approx 1.75$. Given these quality levels any assigned student has the budget set $B = \{1, 2\}$. Under DA the unique optimal quality is $\delta_1 = 2, \delta_2 = 0$, yielding $U_{DA}/Q \approx 1.45$.

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 overdemanded. In contrast, the optimal investment under neighborhood assignment and

²⁴Note that this is not a concern in school choice settings where blocking pairs cannot be assigned outside of the mechanism.




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

DA allocates all resources to the school with higher capacity that yields the highest direct returns from investment. A unit of investment to the higher capacity school increases the utility of more students and hence gives larger direct effects on welfare. But a more equitable investment leads to more welfare gains from sorting on horizontal tastes. Under neighborhood assignment and DA the direct effects exceed the benefits from improved sorting. TTC allows for more choice, and so the benefits from improved sorting under TTC are greater than under DA or neighborhood assignment, and exceed the direct effects.²⁵

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.²⁶ Despite the

²⁵Note that the direct effects and the benefits from improved sorting under either mechanism depend on the parameters of the economy. For general parameters the welfare gain from sorting can be lower or higher than the welfare gains from directing all investment to the more efficient school.

²⁶Note that κ_c is the total school funding. This is equivalent to setting the student utility of

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

4.3 When do DA and TTC give the same assignment?

Kesten (2006), Ehlers & Erdil (2010) show that DA and TTC produce identical assignments only under strong conditions that are unlikely to hold in practice. However, Pathak (2017) reports that the two mechanisms produce a similar rank distribution (i.e., number of students who received their first choice, second choice, etc.) in New Orleans and Boston. Pathak (2017) conjectures that correlation between student preferences and school priorities may explain the similarity in the rank distribution.

To study the effect of correlation between preferences and priorities, 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 high correlation between preferences and priorities (α close to 1) implies that DA and TTC give a similar allocation in a stronger sense:

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

the proportion of students whose assignments are the same under both mechanisms scales linearly with the probability of preference for the neighborhood school α . The proof can be found in Appendix E.2.

Proposition 4. *The proportion of students who have the same assignments under TTC and DA is given by*

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

5 Additional Results

5.1 The Impact of a Marginal Increase in School Quality

This section considers the model with quality dependent preferences presented in Section 4.2. We describe how allocation and welfare change with a marginal increase in the quality of some school. In addition, we provide closed form expressions for the logit economy.

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 5. *Consider $\mathcal{E} = (\mathcal{C} = \{1, 2\}, \Upsilon, v, q)$ and $\delta, \hat{\delta}$ that induce economies $\mathcal{E}_\delta, \mathcal{E}_{\hat{\delta}}$ such that $\eta_\delta, \eta_{\hat{\delta}}$ have full support.²⁷ Suppose that $\hat{\delta}$ increases the quality of school 2, i.e. $\hat{\delta}_2 \geq \delta_2$ and $\hat{\delta}_1 = \delta_1$. Then a change from δ to $\hat{\delta}$ changes the cutoffs 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 5 is illustrated in Figure 7. 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.

²⁷We assume full support to preclude multiplicity of the TTC cutoffs. The theorem extends to general economies with appropriately chosen TTC cutoffs (see Appendix D.4 for more details).




Figure 7: 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.

Closed form expressions for the logit economy

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 quantify the effect and provide precise comparative statics that mirror the intuition from Proposition 5. Proposition 6 gives the TTC assignment in closed form for the logit economy.

Proposition 6. Consider a logit economy (Definition 3) where schools are indexed so that $\frac{q_1}{e^{\delta_1}} \leq \frac{q_2}{e^{\delta_2}} \leq \dots \leq \frac{q_n}{e^{\delta_n}}$. Then the TTC cutoffs p_b^c for $b \geq c$ are given by²⁸

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

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, with $R^0 = 1$.

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

Figure 8 illustrates how the TTC cutoffs change with an increase in the quality

²⁸To simplify notation, when $c = 1$ we let $\prod_{c' < c} p_{c'}^{c-1} = 1$ and set $\rho_1 = q_1/e^{\delta_1}$.

of school ℓ . Using equation (7), we derive closed form expressions for $\frac{dp_b^c}{d\delta_\ell}$, which can be found in online Appendix E.3.




Figure 8: 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_ℓ^ℓ may increase or decrease depending on the specific problem parameters. Note that although p_b^c and p_ℓ^c look aligned in the picture, in general it does not hold that $p_b^c = p_\ell^c$ for all b .

Proposition 6 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 6.

Comparative statics of student welfare

We consider a social planner who can affect quality levels δ of schools in logit economy \mathcal{E} . Consider first the 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 is

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

because $\mathbb{E}[\varepsilon_{\mu(s)s}] = 0$ under neighborhood assignment. Under neighborhood assignment, the marginal welfare gain from increasing δ_ℓ is $\frac{dW_{NH}}{d\delta_\ell} = q_\ell$, 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.²⁹ A student who is offered the budget set $\mathcal{C}^{(c)} = \{c, \dots, n\}$ is assigned to the school $\ell = \arg \max_{b \in \mathcal{C}^{(c)}} \{\delta_b + \varepsilon_{bs}\}$, and her expected utility is $U^c = \ln(\sum_{b \geq c} e^{\delta_b})$ (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

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

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

Proposition 7. *For the logit economy, the change in social welfare $W_{TTC}(\delta)$ under TTC from a marginal increase in δ_ℓ is given by*

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

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

Proposition 7 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 is³⁰

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

An increase in the quality of school 1 gives higher utility for students assigned to 1,

²⁹Under TTC the expected utility of student s assigned to school $\mu(s)$ depends on the student's budget set $B(s, \mathbf{p})$ because of the dependency of $\mu(s)$ on student preferences. Namely, $\mathbb{E}[u^s(\mu(s) | \delta)] = \delta_{\mu(s)} + \mathbb{E}[\varepsilon_{\mu(s)s} | \delta_{\mu(s)} + \varepsilon_{\mu(s)s} \geq \delta_c + \varepsilon_{cs} \forall c \in B(s, \mathbf{p})]$

³⁰Recall that we assume that schools are labeled in order, and thus school 1 is the more selective school. We use that $N^1 = q_1 + q_1 e^{\delta_1 - \delta_2}$, $N^2 = q_2 - q_1 e^{\delta_2 - \delta_1}$.

which is captured by the first term. Additionally, it causes some students to switch their preferences to $1 \succ 2$, making school 1 run out earlier in the TTC algorithm, and removing school 1 from the budget set of some students. Students whose budget set did not change and who switched to $1 \succ 2$ are almost indifferent between the schools and hence 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,

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

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 $N^2 = 0$), 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.³¹

5.2 Design and Bossyness of TTC Priorities

We turn to investigate the design of priorities under TTC. Section 2.2 shows that priorities serve a different role under TTC and DA. Section 4.1 highlights the importance of the priority design in determining welfare. The following analysis provides some initial insight about how changes in priorities change the TTC assignment.

In the following example, we consider changes to the priority among highly ranked students (we provide a similar example in a discrete economy in Appendix E.6). 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 . We find that changes to the priority of these students can have an impact on the assignment of other students, without changing the assignment of any student whose priority changed. This implies that the TTC priorities are “bossy”, and that the question of designing priorities for TTC may be non-trivial.

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

³¹That is, if $\delta_1 = \delta_2$ then $\frac{dW_{TTC}}{d\delta_1} < q_1$ and $\frac{dW_{TTC}}{d\delta_2} < q_2$. If we fix $\delta_1 + \delta_2$ and consider $W_{TTC}(\Delta)$ as a function of $\Delta = \delta_1 - \delta_2$ the function $W_{TTC}(\Delta)$ will have a kink at $\Delta = 0$ (see Figure 5c).

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.4.

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 $[m, 1]^2$.³² 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 9 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.




Figure 9: The range of possible TTC cutoffs in example 5 with $q = 0.455$ and $m = 0.6$. The points depict the TTC cutoffs for the original economy and the extremal cutoffs for the set of possible economies \mathcal{E}' , with the range of possible TTC cutoffs for \mathcal{E}' given by the bold curve. The dashed line is the TTC path for the original economy. The shaded squares depict the changes to priorities that generate the economy $\bar{\mathcal{E}}$ which has extremal cutoffs. In $\bar{\mathcal{E}}$ the priority of all top ranked students is uniformly distributed within the smaller square. The dotted line depicts the TTC path for $\bar{\mathcal{E}}$, which results in cutoffs $p_1^1 = \sqrt{(1 - 2q) \frac{1 - 2m + 2m^2}{m^2}} \approx 0.36$ and $p_2^2 = \sqrt{(1 - 2q) \frac{m^2}{1 - 2m + 2m^2}} \approx 0.25$.

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 9)

³²The remaining students still have ranks distributed uniformly on the complement of $[m, 1]^2$.

are the cutoffs that satisfy $p_1^1 = p_1^2$, $p_2^1 = p_2^2$ and $p_1^1 p_2^2 = 1 - 2q$. Under any of these cutoffs the students in $\{\theta \mid r_c^\theta \geq m \forall c\}$ have the same demand, but the resulting TTC outcomes are different. It follows that the mechanism requires more information to determine the assignment. However, Theorem 3 in Appendix A 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. This is because the changes to priorities changed the trading cycles selected under TTC, and thus changed the population of students who seek to participate in trade. When there are fewer students seeking to trade into school 1, priority at school 1 becomes less useful. 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.

The Clinch and Trade Mechanism

We can also use Example 5 to compare TTC with the Clinch and Trade (C&T) mechanism introduced by Morrill (2015*b*). 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 (2015*b*) 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. We can calculate the C&T assignment by observing that 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.³³ The following proposition builds on Example 5 and shows that C&T may produce more blocking pairs than TTC.

³³For brevity, we abstract away from certain details of C&T mechanism that are important when not all schools run out at the same round.

Proposition 8. *The Clinch and Trade mechanism can produce more, fewer or an equal number of blocking pairs compared to TTC.*

6 Discussion

The cutoff characterization developed in the paper provides a more transparent description of the TTC assignment. We hope that this characterization will help school districts in their evaluation of TTC by providing tools to quantify the welfare gains and the role of priorities under TTC. The cutoff characterization can simplify how the TTC outcome is communicated to students and their families, but also reaffirms concerns about the complexity of TTC.

The tractability of the framework allows us to investigate other design decisions by the school district, such as the allocation of resources to schools or comparison between different priority structures. These decisions can have substantial implications for welfare, and our analysis shows that the resulting differences in welfare can be larger than the welfare difference between DA and TTC. We also demonstrate that the optimal decision itself can depend on the choice of school choice mechanism.

A number of examples provided in the paper utilized functional form assumptions to gain tractability. We also demonstrated that our 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 assignment from the distribution. We leave the problem of determining the optimal choice of tie-breaking lottery for future research.

Cutoff characterizations have been instrumental for empirical work on DA and related mechanisms (Abdulkadiroğlu, Angrist, Narita & Pathak 2017, Agarwal & Somaini 2018, Kapor et al. 2020). We hope that the cutoff characterization of TTC will be similarly useful.

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

- Abdulkadirođlu, A., Agarwal, N. & Pathak, P. A. (2017), ‘The welfare effects of coordinated assignment: Evidence from the new york city high school match’, *American Economic Review* **107**(12), 3635–89.
- Abdulkadirođlu, 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.
- Abdulkadirođlu, 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.
- Abdulkadirođlu, A., Che, Y.-K. & Yasuda, Y. (2015), ‘Expanding choice in school choice’, *American Economic Journal: Microeconomics* **7**(1), 1–42.
- Abdulkadirođlu, A., Pathak, P. A., Roth, A. E. & Sönmez, T. (2005), ‘The boston public school match’, *American Economic Review* pp. 368–371.
- Abdulkadirođlu, A., Pathak, P. & Roth, A. (2009), ‘Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match’, *American Economic Review* **99**(5), 1954–1978.
- Abdulkadirođlu, A. & Sönmez, T. (2003), ‘School choice: A mechanism design approach’, *American Economic Review* **93**(3), 729.
- Agarwal, N. & Somaini, P. (2018), ‘Demand analysis using strategic reports: An application to a school choice mechanism’, *Econometrica* **86**(2), 391–444.
- 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), 2125–2156.
- 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* **124**(5), 1235–1268.
- Bogomolnaia, A. & Moulin, H. (2001), ‘A new solution to the random assignment problem’, *Journal of Economic Theory* **100**(2), 295–328.
- BPS (2005), Recommendation to implement a new bps assignment algorithm, Technical report.
- Calsamiglia, C., Fu, C. & Güell, M. (2020), ‘Structural estimation of a model of school choices: The boston mechanism versus its alternatives’, *Journal of Political Economy* **128**(2), 642–680.
- 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. (2019), ‘Stable matching in large economies’, *Econometrica* **87**(1), 65–110.
- Che, Y.-K. & Tercieux, O. (2018), ‘Payoff equivalence of efficient mechanisms in large matching markets’, *Theoretical Economics* **13**(1), 239–271.
- Che, Y.-K. & Tercieux, O. (2019), ‘Efficiency and stability in large matching markets’, *Journal of Political Economy* **127**(5), 2301–2342.
- 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. (1962), ‘College admissions and the stability of marriage’, *American Mathematical Monthly* **69**(1), 9–15.
- Hakimov, R. & Kesten, O. (2014), ‘The equitable top trading cycles mechanism for school choice’, *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. (2018), ‘A pseudo-market approach to allocation with priorities’, *American Economic Journal: Microeconomics* **10**(3), 272–314.
- 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. (2019), ‘Reducing inequality through dynamic complementarity: Evidence from head start and public school spending’, *American Economic Journal: Economic Policy* **11**(4), 310–49.

- Kapor, A. J., Neilson, C. A. & Zimmerman, S. D. (2020), ‘Heterogeneous beliefs and school choice mechanisms’, *American Economic Review* **110**(5), 1274–1315.
- Karlin, S. & Taylor, H. M. (1975), *A first course in stochastic processes*, Academic press.
- 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.
- Liu, Q. & Pycia, M. (2013), ‘Ordinal efficiency, fairness, and incentives in large markets’, *Mimeo, UCLA* .
- 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. (2020), ‘Foundations of pseudomarkets: Walrasian equilibria for discrete resources’, *working paper* .
- 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. (2017), ‘What really matters in designing school choice mechanisms’, *Advances in Economics and Econometrics* **1**, 176–214.

- 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. (2019), ‘Evaluating with statistics: Which outcome measures differentiate among matching mechanisms?’, *working paper* .
- 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.
- Vapnik, V. & Chervonenkis, A. Y. (1971), ‘On the uniform convergence of relative frequencies of events to their probabilities’, *Theory of Probability and its Applications* **16**(2), 264.

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 10 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.

This representation can be readily generalized to continuum economies. In the




Figure 10: 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.

continuum, the algorithm progresses in continuous time. The state of the algorithm at time $\tau \in \mathbb{R}_{\geq}$ is given by $\gamma(\tau) \in [0, 1]^c$, 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 offered³⁴ by school c . Therefore, if students are assigned between time τ_1 and τ_2

³⁴Strictly 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

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 11. For the sake of clarity, we omit technical details in the ensuing discussion. A rigorous derivation can be found in online Appendix F.

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

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

that is, during the step from time τ_1 to time τ_2 the algorithm clears students with b -ranks between $\gamma_b(\tau_1) = x$ 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.³⁵

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

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.

³⁵The 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.

of students that got an offer from some school $a \in \mathcal{C}$ and whose top choice is c .




Figure 11: 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 , 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 that³⁶

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

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 $x =$

³⁶These 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 .

$\gamma(\tau_1)$,

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

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 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 12 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 \mathbf{d} is a solution to the trade balance equations.




Figure 12: 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 12 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.³⁷ 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.

³⁷ η has full support if for every open set $A \subset \Theta$ we have $\eta(A) > 0$.

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.

As an illustration, consider the unique solution \mathbf{d}^K for the communicating class $K = \{1, 2\}$, as illustrated in Figure 12. 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 13 illustrates the path $\gamma(\cdot)$ and the solution $\mathbf{d}(x)$ to the trade balance equations at x . Note that for every x we can calculate $\mathbf{d}(x)$ from $H(x)$. When there are multiple solutions to the trade balance equations at some x , we may select a solution $\mathbf{d}(x)$ for every x such that $\mathbf{d}(\cdot)$ is a sufficiently smooth gradient field. The TTC path $\gamma(\cdot)$ can be generated by starting from $\gamma(0) = \mathbf{1}$ and following the gradient field.




Figure 13: Illustration of the gradient field $\mathbf{d}(\cdot)$ and path $\gamma(\cdot)$ (ignoring the capacity equations).

A.4 When a School Fills its Capacity

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

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

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

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	\rightarrow	Continuum TTC	Expression	Equation
Cycle	\rightarrow	Valid gradient	$\mathbf{d}(x)$	trade balance equations
Algorithm progression	\rightarrow	TTC path	$\gamma(\cdot)$	$\gamma'(\tau) = \mathbf{d}(\gamma(\tau))$
School removal	\rightarrow	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 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.

Specifically, for a discrete economy $E = (\mathcal{C}, \mathcal{S}, \succ_{\mathcal{C}}, \succ^{\mathcal{S}}, q)$ with $N = |\mathcal{S}|$ students, we define the continuum economy $\Phi(E) = (\mathcal{C}, \Theta, \eta, \frac{q}{N})$ as follows. For each student $s \in \mathcal{S}$ and school $c \in \mathcal{C}$, recall that $r_c^s = |\{s' \mid s \succ_c s'\}| / |\mathcal{S}|$ is the percentile rank of s at c . We identify each student $s \in \mathcal{S}$ with the N -dimensional cube $I^s = \succ^s \times \prod_{c \in \mathcal{C}} [r_c^s, r_c^s + \frac{1}{N}]$ of student types with preferences \succ^s , and define η to have constant density $\frac{1}{N} \cdot N^N$ on $\cup_s I^s$ and 0 everywhere else.

Proposition 9. *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. The intuition behind this result is that TTC is essentially performing the same assignments in both models, with discrete

TTC assigning students to schools in discrete steps, and continuum TTC assigning students to schools continuously, in fractional amounts. By considering the progression of continuum TTC at the discrete time steps when individual students are fully assigned, we obtain the same outcome as discrete TTC.

Proof of Proposition 9. We show that given a discrete economy, the cutoffs of TTC in a continuum embedding Φ give the same assignment as TTC on the discrete model,

$$\mu_{dTTC}(s | E) = \max_{\succ_s} \{c : r_b^s \geq p_b^c \text{ for some } b\} = \mu_{cTTC}(\theta^s | \Phi(E)) \quad \forall \theta^s \in I^s.$$

Fix a discrete cycle selection rule ψ . We construct a TTC path γ such that TTC on the discrete economy E with cycle selection rule ψ gives the same allocation as $TTC(\gamma | \Phi(E))$. Since the assignment of discrete TTC is unique (Shapley & Scarf 1974), and the assignment in the continuum model is unique (Proposition 2), this proves the proposition.

Consider a point during the run of discrete TTC when all schools are still available. At this point, denote by x_c the c -rank of the student pointed to by school c for all $c \in \mathcal{C}$, and denote by $S(x)$ the set of assigned students. By construction, $x \in X = \{0, \frac{1}{N}, \frac{2}{N}, \dots, 1\}^{\mathcal{C}}$. In the next step the discrete TTC clears a cycle and schools point to their favorite remaining student. Let K be the set of schools in the cycle, and let $d_c = \mathbf{1}_{\{c \in K\}}$. Denote by y_c the c -rank of the student pointed to by school c after the cycle is cleared for all $c \in \mathcal{C}$, and denote by $S(y)$ the set of assigned students after the cycle is cleared. Note that $\mathbf{x} - \mathbf{y} = \frac{1}{N}\mathbf{d}$.

Suppose that in continuum TTC there is a TTC path such that $\gamma(t_1) = \mathbf{x} + \mathbf{1} \cdot \frac{1}{N} \in X$. First, notice that by time t_1 the continuum TTC has assigned $\theta \in I^s$ if and only if $s \in S(x)$. Second, we will show that $\gamma(t) = \mathbf{x} - (t - t_1) \frac{1}{N}\mathbf{d} + \frac{1}{N}$ for $t \in [t_1, t_1 + 1)$ satisfies the trade balance equations, and thus the continuum TTC can progress to $\gamma(t_1 + 1) = \mathbf{y} + \mathbf{1} \cdot \frac{1}{N} \in X$. To see that, observe that $H_b^c(\mathbf{x} + \mathbf{1} \cdot \frac{1}{N}) = 1$ if in the discrete TTC school c is the favorite school of the student with b -rank x_b , and $H_b^c(\mathbf{x} + \mathbf{1} \cdot \frac{1}{N}) = 0$ otherwise. On the path $\gamma(t)$ we have that for every $b, c \in K$

$$H_b^c(\gamma(t)) = H_b^c\left(\mathbf{x} + \mathbf{1} \cdot \frac{1}{N}\right) \cdot (1 - (t - t_1))$$

and if $b \in K$ and $c \notin K$ then $H_b^c(\gamma(t)) = 0$.

Therefore for any $c \in K$

$$\sum_{a \in \mathcal{C}} d_a H_a^c(\gamma(t)) = (1 - (t - t_1)) = \sum_{a \in \mathcal{C}} d_c H_c^a(\gamma(t)),$$

and for any $c \notin K$

$$\sum_{a \in \mathcal{C}} d_a H_a^c(\gamma(t)) = 0 = \sum_{a \in \mathcal{C}} d_c H_c^a(\gamma(t)).$$

Thus, the trade balance equations hold for $t \in [t_1, t_1 + 1)$, and there is a continuum TTC path such that $\gamma(t_1) = \mathbf{x}$, $\gamma(t_2) = \mathbf{y}$.

The claim follows by induction on the number of cycles cleared so far in discrete TTC. \square

Proposition 9 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.1.

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 with density bounded below away from 0. Then the TTC allocations in these two economies differ on a set of students of measure $O(\varepsilon|\mathcal{C}|^2)$.*


In Section 5.2, 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 no more than proportional to the total variation distance of the two economies, and suggests that the TTC mechanism is fairly robust to small perturbations in preferences.

B.1 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, $\mathcal{C} = \{1, 2\}$, $\mathcal{S} = \{s_1, s_2, s_3, s_4, s_5, s_6\}$. School 1 has capacity $q_1 = 4$ and school 2 has capacity $q_2 = 2$. The school priorities and student preferences are given by

$$\begin{aligned} 1 & : s_1 \succ s_2 \succ s_3 \succ s_4 \succ s_5 \succ s_6, & s_1, s_3, s_4 : 1 \succ 2, \\ 2 & : s_1 \succ s_3 \succ s_2 \succ s_5 \succ s_4 \succ s_6, & s_2, s_5, s_6 : 2 \succ 1. \end{aligned}$$


In Figure 14, 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.



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 14: Three TTC paths and their cutoffs and allocations for the discrete economy in example B.1. Students labeled in white prefer $1 \succ 2$, and in black prefer $2 \succ 1$. The color of the square indicates the student's assignment, with green indicating school 1, blue indicating school 2. The first round TTC paths are solid, and the second round TTC paths are dotted.

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 15.

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 $(\frac{5}{6}, \frac{5}{6})$.

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 1 is at most 3, and the measure of students assigned to school 2 is 2, so school 2 is unavailable. Hence, from any point $(x_1, \frac{1}{3})$ the TTC path moves parallel to the x_1 axis.




Figure 15: 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 s_1, s_3, s_4, s_6 to school 1 and s_2, s_5 to school 2. All three paths assign the students assigned before p^1 (students s_1, s_2, s_3, s_4, s_5 for paths 1 and 2 and s_1, s_2, s_3, s_5 for path 3) to their top choice school. All three paths assign all remaining students to school 1.