

Designing and Optimizing Matching Markets

Irene Lo

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2018

ProQuest Number: 10931329

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10931329

Published by ProQuest LLC (2018). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

© 2018

Irene Lo

All Rights Reserved

ABSTRACT

Designing and Optimizing Matching Markets

Irene Lo

Matching market design studies the fundamental problem of how to allocate scarce resources to individuals with varied needs. In recent years, the theoretical study of matching markets such as medical residency, public housing and school choice has greatly informed and improved the design of such markets in practice. Impactful work in matching market design frequently makes use of techniques from computer science, economics and operations research to provide end-to-end solutions that address design questions holistically. In this dissertation, I develop tools for optimization in market design by studying matching mechanisms for school choice, an important societal problem that exemplifies many of the challenges in effective marketplace design.

In the first part of this work I develop frameworks for optimization in school choice that allow us to address operational problems in the assignment process. In the school choice market, where scarce public school seats are assigned to students, a key operational issue is how to reassign seats that are vacated after an initial round of centralized assignment. We propose a class of reassignment mechanisms, the Permuted Lottery Deferred Acceptance (PLDA) mechanisms, which generalize the commonly used Deferred Acceptance school choice mechanism and retain its desirable incentive and efficiency properties. We find that under natural conditions on demand all PLDA mechanisms achieve equivalent allocative welfare, and the PLDA based on reversing the tie-breaking lottery during the reassignment round minimizes reassignment. Empirical investigations on data from NYC high school admissions support our theoretical findings. In this part, we also provide a framework for optimization when using the prominent Top Trading Cycles (TTC) mechanism. We show that the TTC assignment can be described by admission cutoffs, which explain the role of priorities in determining the TTC assignment and can be used to tractably analyze TTC. In a large-scale continuum model we show how to compute these cutoffs directly from the distribution of preferences and priorities, providing a framework for evaluating policy choices. As an application of the model we solve for optimal investment in school quality under choice and find that an egalitarian

distribution can be more efficient as it allows students to choose schools based on idiosyncracies in their preferences.

In the second part of this work, I consider the role of a marketplace as an information provider and explore how mechanisms affect information acquisition by agents in matching markets. I provide a tractable “Pandora’s box” model where students hold a prior over their value for each school and can pay an inspection cost to learn their realized value. The model captures how students’ decisions to acquire information depend on priors and market information, and can rationalize a student’s choice to remain partially uninformed. In such a model students need market information in order to optimally acquire their personal preferences, and students benefit from waiting for the market to resolve before acquiring information. We extend the definition of stability to this partial information setting and define regret-free stable outcomes, where the matching is stable and each student has acquired the same information as if they had waited for the market to resolve. We show that regret-free stable outcomes have a cutoff characterization, and the set of regret-free stable outcomes is a non-empty lattice. However, there is no mechanism that always produces a regret-free stable matching, as there can be information deadlocks where every student finds it suboptimal to be the first to acquire information. In settings with sufficient information about the distribution of preferences, we provide mechanisms that exploit the cutoff structure to break the deadlock and approximately implement a regret-free stable matching.

Contents

List of Figures	v
List of Tables	vii
1 Introduction	1
1.1 Simplifying Frameworks for Optimization in School Choice	3
1.2 Mechanisms for Matching with Incomplete Information	5
2 School Choice	6
2.1 The School Choice Model	7
2.2 Policy Goals and Design Objectives	9
2.3 Mechanisms	11
2.4 Optimization and Operations in School Choice	15
I Simplifying Frameworks for Optimization in School Choice	17
3 Dynamic Matching in School Choice: Efficient Seat Reassignment after Late Cancellations	18
3.1 Reassignment in Matching Markets	22
3.1.1 Current Reassignment Systems in School Choice	22
3.1.2 Related Literature	23
3.2 Model	25

3.2.1	Mechanisms	28
3.3	Main Results	35
3.3.1	Axiomatic Justification of PLDA Mechanisms	39
3.4	Intuition for Main Results	42
3.5	PLDA for a Discrete Set of Students	49
3.5.1	Discrete Model	49
3.5.2	PLDA Mechanisms & Their Properties	49
3.6	Empirical Analysis of PLDA Mechanisms	55
3.6.1	Data	56
3.6.2	Simulations	56
3.6.3	Results	58
3.6.4	Strategy-proofness of PLDA	62
3.7	Proposals & Discussion	63
4	The Cutoff Structure of Top Trading Cycles in School Choice	66
4.1	Prior Work on TTC	69
4.1.1	Practical Implications	70
4.2	TTC in School Choice	72
4.2.1	The Discrete TTC Model	72
4.2.2	Cutoff Characterization	74
4.2.3	The Structure of TTC Budget Sets	76
4.2.4	Limitations	79
4.3	Continuum Model and Main Results	80
4.3.1	Model	80
4.3.2	Main Results	82
4.3.3	Consistency with the Discrete TTC Model	88
4.3.4	Proper budget sets	91
4.4	Intuition for the Continuum TTC Model	94

4.4.1	Tracking the State of the Algorithm through the TTC Path γ	94
4.4.2	Determining the Algorithm Progression through Trade Balance	96
4.4.3	Interpretation of Solutions to the Trade Balance Equations	98
4.4.4	When a School Fills its Capacity	102
4.4.5	Comparison between Discrete TTC and Continuum TTC	103
4.5	Applications	107
4.5.1	Optimal Investment in School Quality	107
4.5.2	Design of TTC Priorities	117
4.5.3	Comparing Mechanisms	119
4.6	Discussion	124
II	Mechanisms for School Choice with Incomplete Information	127
5	The Information Acquisition Costs of Matching Markets	128
5.1	Prior Work	131
5.2	Model	133
5.2.1	Stability with Costly Information Acquisition	135
5.2.2	Regret-Free Stable Outcomes	136
5.3	The Structure of Regret-Free Stable Outcomes	139
5.4	Mechanisms	142
5.4.1	Direct Mechanisms	143
5.4.2	Choice-Based Mechanisms	144
5.5	Implementing Regret-Free Stable Matchings	145
5.5.1	All You Need are Cutoffs	146
5.5.2	Regret-Free Choice-Based Mechanisms and Information Deadlock	148
5.5.3	Regret-Free Learning	154
5.6	Discussion	159

Bibliography	163
Appendices	172
A Appendix for Chapter 3	172
A.1 Definitions and Notation	173
A.2 Proof of Theorem 3.1	173
A.3 Proof of Theorem 3.3	179
A.4 Proof of Theorem 3.4	198
B Appendix for Chapter 4	203
B.1 Omitted Proofs for Section 4.3	204
B.2 Omitted Proofs for Section 4.4	230
B.2.1 Derivation of Marginal Trade Balance Equations	230
B.3 Omitted Proofs for Applications (Section 4.5)	233
C Appendix for Chapter 5	246
C.1 Regret-Free Mechanisms in Extensive Form	247
C.1.1 General Mechanisms	247
C.1.2 Choice-Based Mechanisms	248
C.1.3 Accept-Waitlist-Reject Mechanisms	249
C.2 Examples demonstrating Impossibility of Regret-Free Stable Mechanisms	251
C.2.1 Direct One-Shot Mechanisms	251
C.2.2 Impossibility of Student-Optimal Regret-Free Stable Mechanisms	252
C.3 Estimating Regret-Free Stable Cutoffs	254
C.3.1 Continuity and Convergence of Market-Clearing Cutoffs	254
C.3.2 Omitted Proofs for Section 5.5	255
C.3.2.1 Example 5.2	255
C.3.2.2 Example 5.3	256

List of Figures

3.1	Running DA with a reversed lottery eliminates the cascade of reassignments.	21
3.2	Timeline of the two-round mechanism design problem	27
3.3	An example of type equivalence and the order condition.	40
3.4	Number of reassigned students for various PLDA mechanisms.	59
4.1	An example economy and its TTC budget sets.	76
4.2	The schools that enter a student's budget set because of her priority at a single school.	79
4.3	The TTC path, cutoffs, and budget sets for an example economy.	86
4.4	Intuition for the TTC path.	95
4.5	Visualizing assigned and unassigned students.	97
4.6	Example of a graph representation for the trade balance equations.	99
4.7	Illustration of the gradient field and TTC path for an example economy.	101
4.8	Three TTC paths and their cutoffs and allocations for an example discrete economy.	105
4.9	The valid directions for the continuum embedding of the example discrete economy.	106
4.10	The effect of an increase in the desirability of a school for two schools.	109
4.11	The effects of changing the desirability of a school for the logit economy.	111
4.12	Budget sets and welfare for investment in schools with equal capacity.	114
4.13	Budget sets and welfare for investment in schools with unequal capacity.	116
4.14	The range of possible TTC cutoffs in an example economy.	119
4.15	Illustration of economy showing the difference between the TTC and the Clinch and Trade mechanisms.	123

A.1 Constructing the permutation P for $n = 2$ schools. 185

B.1 Illustration of sets of students offered seats and assigned seats. 206

List of Tables

3.1	Centralized PLDA simulation results: 2004-2005 NYC high school admissions.	59
3.2	Decentralized PLDA simulation results: 2004-2005 NYC high school admissions. . .	61
4.1	The relationship between the discrete and continuum TTC processes.	103

Acknowledgements

First and foremost, I would like to acknowledge the numerous coauthors and collaborators involved in the work in this dissertation. Chapter 3 is based on the paper “*Dynamic Matching in School Choice: Efficient Seat Reassignment after Late Cancellations*”, which is joint work with Itai Feigenbaum, Yash Kanoria and Jay Sethuraman. Chapter 4 is based on “*The Cutoff Structure of Top Trading Cycles in School Choice*”, which is joint work with Jacob Leshno. Chapter 5 is based on “*The Information Acquisition Costs of Matching Markets*”, which is joint work with Nicole Immorlica, Jacob Leshno, and Brendan Lucier. This dissertation would not have been possible without all of their contributions.

Special acknowledgements and gratitude are due to my advisors Jay Sethuraman and Jacob Leshno, who took me under their wings and taught me how to grow and mature as a researcher. To Jay, for being an eternal fountain of simplifying insight, sage advice and moral support. To Jacob, for the countless hours spent mentoring and challenging me, and for surreptitiously trying to turn me into an economist. You have both been instrumental in helping me shape my research agenda and I couldn’t have asked for better mentors through my academic journey.

I am greatly indebted to the other members of my thesis committee, Garud Iyengar, Yash Kanoria and Yuri Faenza, for their feedback, and to others who have provided invaluable advice and mentorship throughout my years as a Ph.D. student. To Yash, for being an incredibly efficient collaborator, encouraging mentor and proactive supporter, and for essentially being an advisor to me. To Maria, for encouraging me to pursue a Ph.D. at Columbia and showing me that one can brandish technical prowess with humility and grace. To Nicole and Brendan, for being the best mentors at the best internship and ushering me into the algorithmic game theory community with

open arms. To Vineet, Adam, Omar, Carri, Costis, Yeon Koo for their encouragement and patient guidance during my job market year. To Ward, Don, Karl and Cliff, for their efforts in educating us and for introducing me to the world of operations research. To all the wonderful IEOR staff, especially Jaya, Krupa, Liz, Kristen and Jenny, for their tireless behind-the-scenes support and warm friendship.

I owe much of the beauty, joy and laughter in my Ph.D. experience to my fellow Ph.D. students at Columbia and beyond. I am thankful to Itai, for contributing so much to my experience at Columbia as a collaborator, officemate, and friend. To Allen, for bringing the light of community and friendship to our department and for always believing in me. To Jonathan, for the years of jokes and friendship and for always helping me keep my mathematical mind alive. To Diana, Ha Eun and Jay, for being not only fellow Ph.Ds but also warriors in prayer and precious sisters in Christ. To Chun, Antoine and Gonzalo, for blazing the trail ahead of me. To Daniel, Will, Joann and Vashist, for turning a potentially stressful year into a fun and rewarding adventure and being the best job market buddies I could have asked for. To Kira and Rediet, for shaping my academic experience so powerfully through MD4SG and for being inspirational leaders and magnificent friends. To Francois, Octavio, Mike and all the other students in Columbia IEOR. I look forward to being peers for years to come!

I am deeply grateful to my family and friends, and everyone who has been a part of my wonderful support network. To Alexa, Lauren, Yiyi and all my brothers and sisters in Christ at EPC and Bridges, for your years of loving accountability and patient prayer. To George, for your transcontinental journeys to support me in person, for being my rock in a year of uncertainty and transition, and for teaching me sacrificial love through putting my needs and wants above your own. To my parents, for supporting me in my move halfway across the world, for giving me space to make mistakes and pursue my passions, and for always believing in me, unconditionally loving me and challenging me to be the best version of myself.

Finally, I am thankful to my Lord and God, from whom comes all knowledge and truth, by whom there will be no more want and through whom all things are being made anew. To God be the glory.

Soli Deo Gloria

Chapter 1

Introduction

The allocation of scarce resources is a fundamental societal problem. Whether we are assigning tenants to public housing, medical supplies and donations to recipients, or children to schools and foster care, many of the important problems in modern-day society involve distributing a limited supply of goods to human beings with varied and competing needs. As such, the effective design of assignment processes requires not only operational and algorithmic tools, but also a sensitivity to economic incentives and societal influences. Indeed, in many of these applications we wish to assign goods to strategic users who have private information about their preferences, while accommodating societal aversions to using money to redistribute goods and incentivize behavior. The field of matching market design harnesses rigorous techniques from computer science, economics, and operations research to address such issues in a holistic manner.

In this dissertation, I develop tools for optimization in market design by studying matching mechanisms for school choice. In the school choice problem, cities seek to incorporate student preferences when assigning students to seats in public schools. As many school districts have fewer seats at desirable schools than students wishing to attend those schools, school choice is a socially important problem that exemplifies many of the challenges in matching market design. It is also a setting where theoretical solutions can be validated using data, and insights from theory have had direct and substantial impact. Since a market design approach to school choice was first proposed by Abdulkadirođlu and Sönmez (2003), academics and practitioners have together

developed algorithms for centralized school seat assignment with better incentive properties and efficiency than their prior counterparts (Abdulkadirođlu et al., 2017a). This literature proposed two incentive compatible mechanisms, Deferred Acceptance (DA) and Top Trading Cycles (TTC), and characterized the trade-off between their guarantees for efficiency and equity. The work in this dissertation joins a growing operations literature that quantifies operational tradeoffs in the school choice problem and provides simple policy levers for optimizing school district objectives (Arnosti, 2015; Ashlagi and Nikzad, 2016; Ashlagi and Shi, 2014, 2015; Shi, 2015).

I believe that an operations research perspective can bring two significant contributions to the sphere of market design. The first is practical: as a discipline that trains its students both in mathematical theory and solving real-world problems, we can furnish simplifying mathematical frameworks for market structures that are of practical relevance and help policymakers discern how decisions affect outcomes. The second is tactical: weaving together tools from combinatorial optimization and game theory, we can propose algorithms that incentivize strategic users to provide useful information for achieving desirable outcomes. Accordingly, in my dissertation I have sought to further our theoretical understanding of both DA and TTC in ways that provide clear operational insights, and have proposed new algorithms for school choice with desirable properties. In doing so, I address operational problems in the implementation of school choice, such as reducing congestion when reassigning students to vacated seats after late cancellations, deciding how to invest in school quality, and designing market mechanisms to minimize the cost that students must incur to learn their preferences. In order to provide solutions with practical relevance, I also corroborate my findings with data from school systems, as well as with personal interactions with parents, principals, and school district administrators.

1.1 Simplifying Frameworks for Optimization in School Choice

As market design theory increasingly shapes the design and operations of real-life marketplaces, it is important for designers to provide simple policy levers that practitioners can use to optimize platform objectives. In Chapter 3 we provide a framework for addressing a key operational issue in school choice – reassigning seats that are vacated after an initial round of centralized assignment. Every year around 10% of students assigned a seat in the NYC public high school system eventually do not attend a public school in the following year, opting instead to attend private or charter schools, and their vacated seats can be reassigned. Practical solutions to the reassignment problem must be simple to implement, truthful and efficient. We propose and axiomatically justify a class of reassignment mechanisms, the Permuted Lottery Deferred Acceptance (PLDA) mechanisms, which generalize the commonly used Deferred Acceptance (DA) school choice mechanism to a two-round setting and retain its desirable incentive and efficiency properties. We also provide guidance to school districts as to how to choose the appropriate mechanism in this class for their setting.

Centralized admissions are typically conducted in a single round using Deferred Acceptance, with a lottery used to break ties in each school’s prioritization of students. Our Permuted Lottery Deferred Acceptance mechanisms reassign vacated seats using a second round of DA with a lottery based on a suitable *permutation* of the first-round lottery numbers. We demonstrate that under natural conditions on aggregate student demand for schools, the second-round tie-breaking lottery can be correlated arbitrarily with that of the first round without affecting allocative welfare, and hence the correlation between tie-breaking lotteries can be chosen to attain other objectives. Using this framework, we show how the identifying characteristic of PLDA mechanisms, their permutation, can be chosen to control reallocation. In most school choice systems across the United States, seats vacated after the initial round are reassigned using decentralized waitlists that create significant student movement after the start of the school year, which is costly for both students and schools. We show that *reversing* the lottery order between rounds minimizes reassignment among all PLDA mechanisms, allowing us to alleviate costly student movement between schools without affecting the efficiency of the final allocation. Empirical investigations based on data from NYC high school

admissions support our theoretical findings.

In Chapter 4, we provide a framework for optimization when using the Top Trading Cycles (TTC) school choice mechanism. While TTC has attractive properties for school choice and is often considered by school systems, the commonly used combinatorial description of the mechanism obfuscates many of these properties, and the mechanism has essentially not been adopted in practice. Moreover there is little guidance in the literature as to how to design the inputs to TTC, such as schools' priorities for students, to optimize a school district's objectives. Drawing on ideas from general equilibrium theory, we show that the TTC assignment can be described by admission cutoffs. These cutoffs parallel prices in competitive equilibrium, with students' priorities serving the role of endowments. In a continuum model these cutoffs can be computed directly from the distribution of preferences and priorities, providing a framework that can be used to evaluate policy choices.

We characterize the TTC assignment in terms of cutoffs p_j^i for every *pair* of schools (i, j) . A student is able to attend a school i if for any school j her priority at j meets the cutoff p_j^i . We use a novel formulation of TTC in terms of *trade balance equations* in order to provide a procedure for computing these cutoffs as solutions to a system of differential equations. Using this procedure, we provide closed form solutions for the TTC assignment under a family of distributions, and derive comparative statics. For example, we show that increasing the desirability of a school may result in admitting students with lower priority. Our formulation also gives an alternative to current simulation techniques for evaluating the impact of policy decisions on school assignment and student welfare. As an illustration, we use our framework to solve for optimal investment in school quality under the TTC assignment for a parametrized economy, and show that choice incentivizes a welfare-maximizing school district to invest more equitably in all schools instead of just in the best schools. Our formulation can be used to better design TTC priorities, optimize the use of TTC and empirically compare TTC with other assignment mechanisms, and we hope that it will inspire future work in all these directions.

1.2 Mechanisms for Matching with Incomplete Information

One of the essential roles of a marketplace is to communicate information about supply and demand. Market design affects not only how goods and services are allocated to recipients, but also the information that must be acquired in order to do so. In Chapter 5 we study how school choice mechanisms affect the information acquisition costs borne by students. In a matching model where students pay a cost to learn their preferences, we show that traditional school choice mechanisms that do not account for costly information acquisition can lead to information deadlocks, where it is strictly optimal for every agent to wait for other agents to provide additional market information before paying the cost to learn their own preferences. To overcome this problem, we propose mechanisms that learn sufficient information about aggregate student demand to approximate the optimal outcome.

In our “Pandora’s box” model, school priorities are known, and students have a prior over their cardinal utilities for each school and can pay a cost to see the actual values. We define *stability under incomplete information*, where an outcome, consisting of a matching and acquired information, is stable if any student who has higher priority at a given school than a student assigned to that school either (1) knows her value for that school and prefers her current assignment, or (2) does not know her value for that school and is not willing to pay the cost to learn it. In settings with costly information acquisition students need information about their possible matches to optimally acquire information, and may benefit from waiting for the market to resolve before acquiring information. We refine the set of stable outcomes to the set of *regret-free stable outcomes*, under which the information acquired by students is as if they acquired information optimally after knowing the preferences and information acquisition processes of all other students. We characterize the set of regret-free stable outcomes using cutoffs and show they have a lattice structure. However, without information about student priors, it can be impossible to compute a regret-free stable matching. In settings with sufficient information about student preferences, we propose mechanisms that use the cutoff structure to approximately implement a regret-free stable matchings.

Chapter 2

School Choice

The goal of every school district is to provide each student with admission to a desirable school. However, in most districts, seats at good schools are a scarce resource and a decision has to be made as to which students are able to secure the most coveted seats. Historically, most schools districts have adopted neighborhood schooling policies, where students attend schools in their residential area. This results in students from wealthier families attending the most desirable schools, as they are the ones who are able to afford property in the neighborhoods of these schools. Increasingly, in cities across the United States students are assigned to schools via choice-based systems, where students are able to express their preferences over and be assigned to schools beyond their residential area. Such systems provide more equitable access to opportunity for students from all socio-economic backgrounds, and can improve efficiency by matching students to the schools that provide the educational experience that is best for them. However they also require more sophisticated assignment processes in order to ensure that students' true preferences are elicited and that the assignment process is equitable and efficient.

The mechanism design approach to school choice was first formulated by Balinski and Sönmez (1999) for college admissions, and by Abdulkadiroğlu and Sönmez (2003) for K-12 admissions. These papers took the view that school choice is an assignment problem where the inputs, each student's preferences, are private information that can be strategically reported by agents in order to affect the assignment outcome, and proposed centralized *mechanisms*, assignment algorithms with

strategic input, for computing desirable assignments. This initiated a rich and growing theoretical literature on the design of suitable mechanisms for a single round of centralized school assignment, and also led to significant academic input in the redesign of school choice systems across the US. Academics have worked closely with school authorities to redesign school choice systems in New York City (2003), Boston (2005), New Orleans (2012), Denver (2012), and Washington DC (2013), implementing centralized mechanisms that appear to outperform the uncoordinated and ad hoc assignment systems that they replaced (see, e.g. Abdulkadiroğlu et al., 2017a). I provide brief summaries of the works most relevant to the optimization of quantitative objectives in school choice, and point the interested reader to recent surveys by Abdulkadiroğlu (2013) and Pathak (2011) for more comprehensive overviews of the literature.

2.1 The School Choice Model

The model for the school choice problem first introduced by Abdulkadiroğlu and Sönmez (2003) can be summarized as follows. A finite set of students \mathcal{S} must be assigned to seats at a finite set of schools \mathcal{C} . Each student $s \in \mathcal{S}$ has *preferences* \succ^s for the school they attend, which are strict ordinal rankings over the subset of schools that they find acceptable. Each school $i \in \mathcal{C}$ has *priorities*, a weak ordinal ranking \succeq_i over the set of eligible students. Some schools give priority to students according to strict preferences formed using test scores or admissions portfolios, and other schools give priority to students based on coarse criteria, such as priority for living in the neighborhood or for having a sibling attending the school. In the canonical school choice model, it is assumed that all students know their preferences and that these preferences do not depend on the preferences of other students or the schools attended by other students. The goal of the school choice problem is to design a *mechanism* M that takes as input school priorities \succeq and reported student preferences \succ^r and outputs an appropriate *assignment* $\mu : \mathcal{S} \rightarrow \mathcal{C}$ of students to schools.

In this dissertation, we will often consider the following *continuum* formulation of the school choice problem, which was proposed by Azevedo and Leshno (2016). A *set* \mathcal{S} of students must be assigned to seats at a finite set of schools $\mathcal{C} = \{1, 2, \dots, n\}$. Each student $s \in \mathcal{S}$ has a type

$\theta = (\succ^\theta, p^\theta)$ which encodes both the student's ordinal preferences \succ^θ over schools as well as the schools' priorities $p^\theta \in [0, 1]^C$ for the student. We may think of p_c^θ as the percentile rank of a student of type θ in the school's priority ordering. A *continuum economy* is given by $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ where $q = \{q_c\}_{c \in C}$ is the vector of capacities of each school, and η is a measure over the space of student types Θ . An *assignment* is a mapping $\mu : \Theta \rightarrow \mathcal{C} \cup \{\emptyset\}$ from students to schools. The goal of the school choice problem in the continuum setting is again to design a *mechanism* M that takes as input school priorities p and reported student preferences \succ^r and outputs an appropriate assignment of students to schools.

The continuum model deviates from the standard model in a way that captures two salient features of the school choice problem. While traditional matching problems are often classified as either two-sided (e.g. hospital residency matching (Roth and Peranson, 1999), which view both residents and hospitals as strategic) or one-sided (e.g. housing allocation (Abdulkadiroğlu and Sönmez, 1999), which view only the tenants as strategic), the school choice problem lies somewhere in the middle, as it takes input from agents on both sides of the market with the goal of providing assignments that are efficient and equitable from the *student* perspective. The continuum model visually emphasizes the focus on student outcomes by considering schools' priorities as properties of students, rather than of schools, and invites us to view school priorities as a design lever for achieving what the mechanism designer believes is a desirable assignment for students. In addition, the continuum model of Azevedo and Leshno (2016) captures a setting where the number of students is much larger than the number of schools.¹ In doing so, it implicitly assumes that students are non-atomic and that any single student misreporting their preferences does not affect outcomes for any other students, a simplifying assumption that allows us to provide clean intuition and high-level insights for design. There is now a growing literature that uses continuum models in market design (see, e.g. Avery and Levin, 2010; Abdulkadiroğlu et al., 2015; Ashlagi and Shi, 2015; Che et al., 2015; Azevedo and Hatfield, 2015) .

¹This is contrast with e.g. Che and Tercieux (2015, 2018), which study the properties of TTC in a large market where the heterogeneity of items grows as the market gets large.

2.2 Policy Goals and Design Objectives

Since the academic study and design of school choice systems was initiated by Balinski and Sönmez (1999); Abdulkadiroğlu and Sönmez (2003), there have been remarkable levels of collaboration between academics and practitioners in redesigning school choice systems. This has given us significant insight into the desirable qualities of school choice assignments and assignment mechanisms, a few of which I highlight below. More detailed discussions of the first-order issues in school choice in practice can be found in Pathak (2016), as well as in papers about the redesign processes in New York (Abdulkadiroğlu et al., 2005a), Boston (Abdulkadiroğlu et al., 2005b) and New Orleans (Abdulkadiroğlu et al., 2017c) .

Clear Incentives

One of the clear themes that has emerged in the redesign of school choice systems is the importance of providing the correct incentives for students to tell the truth. When redesigning the Boston school choice system, officials at Boston Public Schools (BPS) strongly advocated for selecting a non-manipulable mechanism on the grounds that it would equalize the opportunity to access schools BPS (2005):

“[A] strategy-proof algorithm ‘levels the playing field’ by diminishing the harm done to parents who do not strategize or do not strategize well. [...] [T]he need to strategize provides an advantage to families who have the time, resources, and knowledge to conduct the necessary research.”

When we presented our class of reassignment mechanisms (Chapter 3) to practitioners, one of the first questions we received from officials at BPS and the New York City Department of Education alike was whether the mechanisms were manipulable by students. Similarly Clark medalist Parag Pathak, who was involved in the redesign of dozens of school choice systems including those at New York City and Boston, echoes the sentiment that manipulability erodes equity and trust, saying *“the idea that a manipulable mechanism frustrates participants and creates inequities for sincere participants is a theme that I have seen in cities other than Boston”*(Pathak, 2016). Thus

motivated, most of the theoretical literature on school choice has focused on providing mechanisms that incentivize students to truthfully report their preferences.

Definition 2.1. A mechanism is **strategy-proof** if it is a dominant strategy for each student to report her true preferences, i.e. for all students s , for all preferences \succ^{-s} of other students and priorities p , student s (weakly) prefers her assignment under truthful reporting to her assignment under any misreport \succ^r .

Equity and Fairness

As suggested in the previous section, another key attribute of commonly-used school choice mechanisms is that they are perceived to be fair. While definitions of equity and fairness in market design are nascent, one partial notion of fairness that was proposed in Abdulkadiroğlu and Sönmez (2003) has proved to be enduring in practice. An assignment *eliminates justified envy* or *respects priorities* if, when two students both wish to attend a given school i , the student with higher priority at school i is given precedence. In the school choice literature, in line with prior work in two-sided matching markets, this has also been termed *stability*.

Definition 2.2. An assignment μ is **stable** or **respects priorities** if for every student $s \in \mathcal{S}$ and school $i \in \mathcal{C}$ and such that student s prefers i to her assignment $i \succ^s \mu(i)$, the school i is full of students with higher priority than s i.e. (i) $\eta(\mu(i)) = q_i$, and (ii) $p_i^{s'} \geq p_i^s$ for all $s' \in \mu^{-1}(i)$. A mechanism M is **stable** or **respects priorities** if it always produces a stable assignment.

In a one-sided assignment setting, stability provides a rationale for the assignment: students who are not assigned to a school have insufficient priority for being assigned to that school. In a two-sided matching setting, a stable mechanism incentivizes schools and students to adopt the matching proposed by the mechanism, as students and schools cannot Pareto improve their outcomes by assigning students outside of the system. Hence in school choice, in addition to providing some measure of fairness, a stable mechanism also eliminates justifiable challenges to the outcome of the mechanism and helps prevent unraveling of the market.

Efficiency

Finally, one of the primary aims of school choice is to increase student welfare by allowing heterogeneous students to choose the schools that best suit their educational needs. Hence from an optimization perspective one of the most natural criteria for a school choice mechanism is that it produces an efficient assignment. Two desirable notions of efficiency commonly put forth in the literature are as follows.

Definition 2.3. An assignment μ is **non-wasteful** if no student desires an unused seat over her own, i.e. if $i \succ^s \mu(i)$ then $\eta(\mu^{-1}(i)) = q_i$.

Definition 2.4. An assignment μ is **feasible** if the number of students assigned to a school does not exceed its capacity, $\eta(\mu^{-1}(i)) \leq q_i \forall i$. A feasible assignment μ is **Pareto efficient** (for students) if there is no feasible assignment μ' that weakly improves outcomes for all students and strictly improves for at least one, i.e. there is no feasible assignment μ' such that $\mu'(s) \succ^s \mu(s)$ for some $s \in \mathcal{S}$ and $\mu'(s') \succeq^{s'} \mu(s') \forall s' \in \mathcal{S}$.

While the notion of Pareto efficiency has many proponents in the theoretical literature, in practice school districts have tended to emphasize other notions of efficiency, such as the proportion of students being assigned to their (reported) first choice school. Further discussion on the importance of efficiency in practice can be found in Chapter 4, in Abdulkadiroğlu et al. (2017c), and in a number of empirical papers (see e.g. Abdulkadiroğlu et al. (2009)).

2.3 Mechanisms

An ideal school choice mechanism would embody all the desirable properties of strategy-proofness, stability and efficiency. However, it is well known that there is no mechanism that is both stable and Pareto efficient (see e.g. Gale and Shapley (1962); Roth (1982)). The two main mechanisms proposed by the theoretical literature, Deferred Acceptance (DA) and Top Trading Cycles (TTC), represent two extremes in the tradeoff between these properties: both mechanisms are strategy-proof, and DA is stable while TTC is Pareto efficient. Moreover, DA implements the student-optimal stable assignment (Gale and Shapley, 1962), and TTC minimizes the violations of priorities

required to achieve an efficient outcome (Abdulkadiroğlu et al., 2017c). I describe both mechanisms and their qualitative properties, and also provide some background as to how they can be tuned to optimize quantitative outcomes.

Deferred Acceptance (DA)

The Deferred Acceptance mechanism for school choice is based on the Gale-Shapley Deferred Acceptance algorithm for one-to-one stable matching Gale and Shapley (1962), and was one of the two mechanisms first proposed by Abdulkadiroğlu and Sönmez (2003) for school choice. It takes as input student preferences \succ^S and strict school priorities and runs via student proposals as follows.² In step 1, each student applies to their most-preferred school. A school with a capacity of q tentatively assigns a seat to each of the q highest-ranked eligible applicants and rejects any remaining applicants. In each subsequent step, students who are not tentatively assigned apply to their most-preferred school that has not yet rejected them. A school with a capacity of q tentatively assigns a seat to the q highest-ranked students who have applied to the school in *any* step and rejects any remaining applicants. The algorithm runs until there are no new student applications, at which point it terminates and assigns each student to her tentatively assigned school seat.

The Gale-Shapley Deferred Acceptance mechanism has many desirable properties for school choice. Dubins and Freedman (1981) and Roth (1982) proved that it is strategy-proof for students, and Gale and Shapley (1962) demonstrated that it always finds the *student-optimal stable matching* μ^* , which satisfies the property that every student s (weakly) prefers to their assignment $\mu^*(s)$ under μ^* to their assignment $\mu(s)$ under any stable matching μ . DA also admits the following description via admissions cutoffs:

Definition 2.5 (*Deferred Acceptance (Azevedo and Leshno, 2016)*). The Deferred Acceptance mechanism is a function $DA_\eta((\succ^s, p^s)_{s \in S})$ mapping student preferences and strict school priorities into an assignment μ , defined by a vector of cutoffs $\mathbf{C} \in \mathbb{R}_+^{\mathfrak{N}}$ as follows. Each student s is assigned to her most-preferred school as per her preferences, among those where her priority exceeds the cutoff:

²In practice, weak school priorities are turned into strict school priorities via random tie-breaking.

$$\mu(s) = \max_{\succ_s}(\{i \in \mathcal{S} : p_i^s \geq C_i\} \cup \{n + 1\}). \quad (2.1)$$

Moreover, \mathbf{C} is market-clearing, namely

$$\eta(\mu(i)) \leq q_i \text{ for all } i \in \mathcal{S}, \text{ with equality of } C_i > 0. \quad (2.2)$$

The cutoff characterization simplifies the description of the outcome of Deferred Acceptance and elucidates the role of priorities in determining the outcome. It has also been instrumental in providing methods for empirical and counterfactual analysis of school systems (see, e.g. Agarwal and Somaini, 2018; Abdulkadiroğlu et al., 2017b; Fack et al., 2015). In addition, the cutoff characterization has been used to provide frameworks for optimization when the mechanism of choice is Deferred Acceptance. Shi (2015) uses the cutoff description of Deferred Acceptance to provide a method for optimizing the design of priorities via a reduction to an assortment planning problem, and Ashlagi and Shi (2014) employs the cutoff characterization to design optimal correlations in tie-breaking lotteries when the goal is the increase community cohesion in school seat assignment. In Chapter 3 we define a class of reassignment mechanisms based on running two rounds of Deferred Acceptance, show that they have a similar cutoff characterization, and use it to show the effect of cross-round correlations in tie-breaking lotteries on allocative efficiency and reassignment.

Top Trading Cycles (TTC)

The Top Trading Cycles mechanism for school choice is based on the Top Trading Cycles algorithm for housing allocation, first proposed in Shapley and Scarf (1974) and attributed to David Gale. It is also one of the two mechanisms first proposed by Abdulkadiroğlu and Sönmez (2003) for school choice. It runs in discrete steps as follows. In step 1, each student points to their most-preferred school, and each school points to their top priority student. Some number of cycles in the pointing graph (there is at least one) are selected, students in the cycles are assigned to the school they are pointing to, and schools in the cycles decrease their remaining capacity by 1. In each subsequent step, each unassigned student points to their most-preferred school, each school

with positive residual capacity points to their top priority remaining student, and more cycles are selected and students assigned. The algorithm runs until either all students are assigned or there are no seats with remaining capacity.

TTC also has many desirable properties for school choice. It is simple to verify that it always results in a Pareto efficient assignment, and Abdulkadiroğlu and Sönmez (2003) and Roth and Postlewaite (1977) demonstrated that it is strategy-proof for students. As a result, TTC has been considered for use in a number of school choice systems (see, e.g. Abdulkadiroğlu et al., 2005b). However, TTC has not been widely adopted for school choice.³ Pathak (2016) hypothesizes that this may be due to a lack of a simple explanation for TTC that emphasizes its desirable properties, and a lack of understanding of the role of priorities in TTC. For example, a report from Boston Public Schools states BPS (2005):

“The Top Trading Cycles Algorithm allows students to trade their priority for a seat at a school with another student. This trading shifts the emphasis onto the priorities and away from the goals BPS is trying to achieve by granting these priorities in the first place.”

In Chapter 4 we provide a cutoff characterization of the TTC mechanism that emphasizes that it is strategy-proof and elucidates the role of priorities in the TTC allocation. This allows us to reframe the primary choice between DA and TTC as follows: DA sacrifices some student welfare but strictly respects priorities, and TTC uses priorities to guide the selection of an efficient allocation. Our cutoff characterization also lays the framework for designing TTC priorities to achieve desirable outcomes.

Other Mechanisms

Several variants of DA and TTC have also been suggested in the literature. Kesten (2006) studies the relationship between DA and TTC, and shows that they are equivalent if and only if the priority structure is acyclic. In subsequent work Kesten (2010) also proposes the Efficiency Adjusted

³The only instances where TTC was implemented for school choice are in the San Francisco school district and previously in the New Orleans Recovery School District (Abdulkadiroğlu et al. (2017c)).

Deferred Acceptance, which relaxes the justified envy condition to improve efficiency. Morrill (2015b) suggests the Clinch and Trade mechanism, which differs from TTC in that it identifies students who are guaranteed admission to their first choice and assigns them immediately without implementing a trade. Hakimov and Kesten (Forthcoming) introduce Equitable TTC, a variation on TTC that aims to reduce inequity. All these mechanisms provide intermediate alternatives to DA and TTC in the tradeoff between efficiency and stability

2.4 Optimization and Operations in School Choice

There is a growing operations literature on designing the school choice process to optimize quantitative objectives. In the Boston school choice system (which uses the Deferred Acceptance mechanism), Ashlagi and Shi (2014) consider how to improve community cohesion in school choice by correlating the lotteries of students in the same community, and Ashlagi and Shi (2015) show how to maximize welfare given busing cost constraints.

Several papers also explore how school districts can use rules for breaking ties in school priorities as policy levers. Most school choice systems turn weak school priorities into strict school priorities using the same tie-breaking lottery across all schools before running DA, resulting in a mechanism known as DA-STB. Arnosti (2015); Ashlagi and Nikzad (2016); Ashlagi et al. (2015) show that DA-STB assigns more students to one of their top k schools (for small k) compared to DA using independent lotteries at different schools, and Abdulkadiroğlu et al. (2009) empirically compare these tie-breaking rules. A concrete design recommendation of Ashlagi and Nikzad (2016) is that in order to improve the efficiency of the assignment “popular” schools should use single tie-breaking to break ties, which is the tie-breaking rule used in our work. Erdil and Ergin (2008) also exploit indifferences to improve allocative efficiency. Similarly, when using TTC in a setting with weak priorities, Ehlers (2014) shows that any fixed tie-breaking rule satisfies weak efficiency, and Alcalde-Unzu and Molis (2011), Jaramillo and Manjunath (2012) and Saban and Sethuraman (2013) give specific variants of TTC that are strategy-proof and efficient.

There are still many important operational questions in the assignment of students to schools.

Most of the existing literature on school choice is essentially static and considers a model where all students know their preferences, submit them to a centralized system and are assigned in a single round. In Chapter 3 we address the problem of designing reassignment mechanisms and aftermarkets, a clear operational issue where a systematic market design approach can have significant welfare benefits (see, e.g. Narita, 2016; Pathak, 2016). In Chapter 4 we consider how the choice of a school choice mechanism can also affect how policy decisions outside of the school choice system, such as investment in school quality, impact student welfare. In Chapter 5 we consider a setting where students do not know their preferences and must exert effort to learn them, and explore how sequential matching mechanisms can be designed to reduce unnecessary costly learning. These works illustrate the benefits of an operations research approach, which can consider the detailed constraints of the market and provide practical solutions and overarching insights.

Part I

Simplifying Frameworks for Optimization in School Choice

Chapter 3

Dynamic Matching in School Choice: Efficient Seat Reassignment after Late Cancellations

In public school systems throughout the United States, students submit preferences that are used to assign them to public schools. As this occurs fairly early in the school year, students typically do not know their options outside of the public school system when submitting their preferences. Consequently, a significant fraction of students who are allotted a seat in a public school eventually do not use it, leading to considerable inefficiency. In the NYC public high school system, over 80,000 students are assigned to public school each year in March, and about 10% of these students choose to not attend a public school in September, possibly opting to attend a private or charter school instead.¹ Schools find out about many of the vacated seats only after classes begin, when students do not show up to class; such seats are reassigned in an ad hoc manner by the schools using decentralized procedures that can run months into the school year. A well-designed reassignment process, run after students learn about their outside options, could lead to significant gains in overall welfare. Yet no systematic way of reassigning students to unused seats has been proposed

¹In the 2004–2005 school year, 9.22% of a total of 81,884 students dropped out of the public school system after the first round. Numbers for 2005–2006 and 2006–2007 are similar.

in the literature. Our goal is to design an explicit reassignment mechanism run at a late stage of the matching process that efficiently reassigns students to vacated seats.

During the past fifteen years, insights from matching theory have informed the design of school choice programs in cities around the world. Although there is a vast and growing literature that explores many aspects of school choice systems and informs how they are designed in practice, most models considered in this literature are essentially static. Incorporating dynamic considerations in designing assignment mechanisms, such as learning new information at an intermediate time, is an important aspect that has only recently started to be addressed. Our work provides some initial theoretical results in this area and suggests that simple adaptations of one-shot mechanisms can work well in a more general setting.

We consider a two-round model of school assignment with finitely many schools, where students learn their outside option after the first-round assignment, resulting in vacant seats which can be reassigned. In the first round, schools have weak priorities over students, and students submit strict ordinal preferences over schools. Students receive a first-round assignment based on these preferences via Deferred Acceptance with Single Tie-Breaking (DA-STB), a variant of the standard Deferred Acceptance mechanism (DA) where ties in school preferences are broken via a single lottery ordering across all schools. Afterwards, students may be presented with better outside options (such as admission to a private school), and may no longer be interested in the seat allotted to them. In the second round, students are invited to submit new ordinal preferences over schools, reflecting changes in their preferences induced by learning their outside options. The goal is to reassign students so that the resulting assignment is efficient and the two-round mechanism is strategy-proof and does not penalize students for participating in the second round. As a significant fraction of vacated seats are reassigned only after the start of the school year, a key additional goal is to ensure that the reassignment process minimizes the number of students who need to be reassigned.

We introduce a class of reassignment mechanisms with desirable properties: the *Permuted Lottery Deferred Acceptance (PLDA) mechanisms*. PLDA mechanisms compute a first-round assignment by running DA-STB, and then a second-round assignment by running DA-STB with a *permuted* lottery. In the second round, each school first prioritizes students who were assigned to it

in the first round, which guarantees each student a second-round assignment that she prefers to her first-round assignment, then prioritizes students according to their initial priorities at the school, and finally breaks ties at all schools via a permutation of the (first-round) lottery numbers. Our proposed PLDA mechanisms are based on school choice mechanisms currently used in the main round of assignment, and can be implemented either as *centralized PLDAs*, which run a centralized second round with updated preferences, or as *decentralized PLDAs*, which run a decentralized second round via a waitlist system that closely mirrors current reassignment processes.

Our key insight is that the mechanism designer can design the correlation between tie-breaking lotteries to achieve operational goals. In particular, *reversing* the lottery between rounds minimizes reassignment without sacrificing student welfare. Our main theoretical result is that under an intuitive *order condition*, all PLDAs produce the same distribution over the final assignment, and reversing tie-breaking lotteries between rounds implements the centralized Reverse Lottery DA (RLDA), which minimizes the number of reassigned students. We axiomatically justify PLDA mechanisms: absent school priorities, PLDAs are equivalent to the class of mechanisms that are two-round strategy-proof while satisfying natural efficiency requirements and symmetry properties.

In a setting where all students agree on a ranking of schools and there are no priorities our results are very intuitive. By reversing the lottery, we move a few students many schools up their preference list rather than many students a few schools up, thereby eliminating unnecessary cascades of reassignment (see Figure 3.1). Surprisingly, however, our theoretical result holds in a general setting with heterogeneous student preferences and arbitrary priorities at schools. The order condition can be interpreted as aggregate student preferences resulting in the same *order* of popularity of schools in the two rounds. Our results show that if student preferences and school priorities produce such agreement in aggregate demand across the two rounds, then reversing the lottery between rounds preserves ex ante allocative efficiency and minimizes reassignment.

We empirically assess the performance of RLDA using data from the New York City public high school system. We first investigate a class of centralized PLDAs that includes RLDA, rerunning DA using the original lottery order (termed Forward Lottery Deferred Acceptance or FLDA), and rerunning DA using an independent random lottery. We find all these mechanisms provide




Figure 3.1: Running DA with a reversed lottery eliminates the cascade of reassignments.

There are 6 students with identical preferences over schools, and 6 schools each with a single priority group. All students prefer schools in the order $1 \succ 2 \succ \dots \succ 6$. The student assigned to school 1 in the first round leaves after the first round; otherwise all students find all schools acceptable in both rounds. Running DA with the same tie-breaking lottery reassigns each student to the school one better on her preference list, whereas reversing the tie-breaking lottery reassigns only the student initially assigned to 6 from her least preferred to her most preferred school.

similar allocative efficiency, but RLDA reduces the number of reassigned students significantly. For instance, in the NYC data set from 2004–2005, we find that FLDA results in about 7,800 reassignments, and RLDA results in about 3,400 reassignments out of a total of about 74,000 students who remained in the public school system, i.e. less than half the number of reassignments under FLDA. The gains become even more marked if we compare with current practice: RLDA results in less than 40% of the 8,600 reassignments under decentralized FLDA with waitlists.²

To better evaluate the decentralized waitlist systems currently used in practice, we also empirically explore the performance of decentralized FLDA and RLDA as a function of the time available to clear the market. We find that the timing of information revelation can greatly impact the allocative efficiency and the congestion of decentralized waitlist systems. If, as in Figure 3.1, congestion is caused by students taking time to vacate previously assigned seats, then reversing the lottery increases allocative efficiency during the early stages of reassignment and decreases the congestion in the system. However, if congestion is caused by students taking time to decide on offers from the waitlist, then these findings are reversed. In both cases, for reasonable timescales the welfare gains from centralizing the reassignment are substantial.

²A decentralized version of FLDA is used in most cities and in NYC kindergarten admissions. It was also used in NYC high school admissions until a few years ago, when the system abandoned reassignments entirely, anecdotally due to the excessive logistical difficulties created by market congestion. See Section 3.1.1 for more details.

3.1 Reassignment in Matching Markets

3.1.1 Current Reassignment Systems in School Choice

Schools systems in cities across the US, including in New York, Boston, Washington DC, Denver, Seattle, New Orleans, and Chicago, use similar centralized processes for admissions to public schools. Students seeking admission to a school submit their preference lists over schools to a central authority by December through March, for admission starting the subsequent fall. Each school may have priority classes of students, such as priority for students who live in the neighborhood, priority for siblings of students who are already enrolled at that school, or priority for students whose families have low income. An allocation of seats to students is produced by using the student-proposing Gale-Shapley Deferred Acceptance algorithm with single-tiebreaking. Students must register in their assigned school by April or early May.

In March and April students are also admitted to private and charter school via processes run concurrently with the public school assignment process. This results in an attrition rate of about 8-10% of the seats assigned in the main round of public school admissions. Some schools account for this attrition by making “over offers” in the first round and accepting more students than they have seats for (see, e.g. Szuffita). However, such oversubscription of students is usually conservative, due to hard constraints on space and teacher capacity.³ As a result, most schools have unused seats at the end of the first round that can be reassigned to students who want those seats, and many schools find out about these vacant seats only after the start of the school year.

Reassignments in most school choice systems, such as in New York kindergarten, Boston, Washington DC, Denver, Seattle, New Orleans, and Chicago, are performed using a decentralized waitlist system. Students are put on waitlists for all schools that they ranked above their first-round assignment, in the order of the first-round priorities (after tie-breaking).⁴ Students who do not register by the deadline are presumed to be uninterested and their seats are offered to waitlisted students in sequence, with further seats becoming available over time as students receive new offers from

³Capacity constraints are binding in most schools. Most states impose maximum class sizes and fund schools based on enrollment after the first 2-3 weeks of classes, which incentivizes schools to enroll as many students as permissible.

⁴Late applicants are also included in these waitlists, typically with inferior lottery numbers.

outside the system. The seats previously occupied by students who were reassigned are also offered to students on the waitlists. Students offered seats by the waitlist system usually have several days to a week to make a decision, and are only bound by the final offer they choose to accept.⁵ Overall, this typically results in a “huge slow round robin” (Szuffita) of reassignment that continues all summer until classes begin, and in some cities (e.g. NYC, Boston, and Washington DC) up to several months after classes begin.

Our proposed class of mechanisms may be viewed as a generalization of these waitlist systems as follows. Waitlists are PLDA mechanisms where 1) the second round is implemented in a decentralized fashion as information about vacated seats propagates through the system, and 2) the tie-breaking lotteries used in the second round are the same as those used in the first round. We show that *permuting* the tie-breaking lottery numbers before creating waitlists provides a class of reassignment mechanisms that, given sufficient time to clear the market, result in similar eventual allocative efficiency while allowing the designer additional flexibility for optimizing other objectives.

3.1.2 Related Literature

There is a vast literature on dynamic matching and reassignments. The reassignment of donated organs has been extensively studied in work on kidney exchange (see, e.g. Roth et al., 2004; Anderson et al., 2015, 2017; Ashlagi et al., 2017). Reassignments due to cancellations also frequently arise in online assignment settings such as kidney transplantation (see, e.g. Zenios, 1999; Su and Zenios, 2006) and public housing allocation (see, e.g. Kaplan, 1987; Arnosti and Shi, 2017). An important difference is that these are *online* settings where agents and objects arrive over time and are matched on an ongoing basis. In such settings matches are typically irrevocable, and so optimal assignment policies account for typical cancellation and arrival statistics and optimize for agents arriving in the future (see, e.g. Dickerson and Sandholm, 2015). In our setting the matching for the entire system is coordinated in time, and we improve welfare by controlling both the initial

⁵Students who have accepted an offer off the waitlist of one school are allowed to accept offers off the waitlists of other schools. Since registration for one school automatically cancels the student’s previous registrations, such an action would automatically release the seat the student accepted from the first school, making that seat available to other students on the waitlist.

assignment and subsequent reassignment of objects among the same set of agents.

Another relevant strand in the reassignment literature is the work of Abdulkadiroğlu and Sönmez on house allocation models with existing tenants (or housing endowments) (Abdulkadiroğlu and Sönmez, 1999). Our second round can be thought of as school seat allocation where some agents already own a seat and we wish to reassign seats to reach an efficient assignment. There are also a growing number of papers that consider a dynamic model for school admissions (see, e.g. Compte and Jehiel, 2008; Combe et al., 2016). A critical distinction between these works and ours is that in our model, the initial endowment is determined endogenously by preferences, and so we propose reassignment mechanisms that are impervious to students manipulating their first-round endowment to improve their final assignment. We also exploit indifference to minimize reassignment, a direction which is not explored in these works.

A number of recent papers, such as Dur (2012a); Kadam and Kotowski (2014); Pereyra (2013), focus on the strategic issues in dynamic reassignment. and also propose using modified versions of DA in each round. These works develop solution concepts in finite markets with specific cross-period constraints and propose DA-like mechanisms that implement them. In recent complementary work Narita (2016) analyzes preference data from NYC school choice, observes that a significant fraction of preferences are permuted after the initial match, and proposes a modified version of DA with desirable properties in this setting. We similarly propose PLDA mechanisms for their desirable incentive and efficiency properties. In addition, our large market and consistency assumptions allow us to uncover considerable structure in the problem and provide conditions under which we can optimize over the entire class of PLDA mechanisms.

Our work also has some connections to the queueing literature. The class of mechanisms that emerges in our setting involves choosing a permutation of the initial lottery order, and we find that the reverse lottery minimizes reassignment within this class. This is similar to the choice of a service policy in a queueing system (e.g. FIFO, LIFO, SRPT etc.), whereby a particular policy is chosen in order to minimize cost functions such as expected waiting time (see, e.g. Lee and Srinivasan, 1989). “Work-conserving” service policies such as these result in identical throughput but different expected waiting times, and we similarly find that different PLDA mechanisms differ

in the number of reassignments even when they have identical allocative efficiency. Our continuum model parallels fluid limits and deterministic models employed in queueing (Whitt, 2002), revenue management (Talluri and Van Ryzin, 2006), and other contexts in operations management.

3.2 Model

We consider the problem of assigning a set of students \mathcal{S} to seats in a finite set of schools $\mathcal{C} = \{1, \dots, n\}$. Each student can attend at most one school. There is a *continuum*⁶ of students with an associated measure η : for any (measurable) subset $S \subseteq \mathcal{S}$, we let $\eta(S)$ denote the mass of students in S . The outside option is $n + 1 \notin \mathcal{C}$. The capacities of the schools are $q_1, \dots, q_n \in \mathbb{R}_+$, and $q_{n+1} = \infty$. A set of students of η -measure at most q_i can be assigned to school i .

Each student submits a strict preference ordering over her acceptable schools, and each school partitions eligible students into priority groups. Each student has a *type* $\theta = (\succ^\theta, \hat{\succ}^\theta, p^\theta)$ that encapsulates both her preferences and school priorities. The student's first- and second-round preferences, respectively \succ^θ and $\hat{\succ}^\theta$, are strict ordinal preferences over $\mathcal{C} \cup \{n + 1\}$, and schools before (after) $\{n + 1\}$ in the ordering are acceptable (unacceptable). The student's *priority class* p^θ encodes her priority p_i^θ at school i . Each school i has n_i priority groups. We assume that schools prefer higher priority groups, students ineligible for school i have priority $p_i = -1$, and that $p_i \in \{-1, 0, 1, \dots, n_i - 1\}$. Eligibility and priority groups are exogenously determined and publicly known.

Each student $s = (\theta^s, L(s)) \in \mathcal{S}$ also has a first-round lottery number $L(s) \in [0, 1]$. We sometimes use the notation $(\succ^s, \hat{\succ}^s, p^s)$ as a less cumbersome alternative to $(\succ^{\theta^s}, \hat{\succ}^{\theta^s}, p^{\theta^s})$. We let Θ be the set of all student types, so that $\mathcal{S} = \Theta \times [0, 1]$ denotes the set of students. For each $\theta \in \Theta$ we let $\zeta(\theta) = \eta(\{\theta\} \times [0, 1])$ be the measure of all students with type θ .

We assume that all students have consistent preferences, defined as follows.

⁶Our continuum model can be viewed as a two-round version of the model introduced by Azevedo and Leshno (2016). Continuum models have been used in a number of papers on school choice; see Agarwal and Somaini (2018); Ashlagi and Shi (2014); Azevedo and Leshno (2016). Intuitively, one could think of the continuum model as a reasonable approximation of the discrete model when the number of students is large, although we do not establish a formal relationship between the discrete and continuum models, as that is beyond the scope of this work.

Definition 3.1. Preferences $(\succ, \hat{\succ})$ are **consistent** if the second-round preferences $\hat{\succ}$ are obtained from the first-round preferences \succ via truncation, i.e.: (1) (a school does not become acceptable only in the second round) for every $i \in \mathcal{S}$, $i \hat{\succ} \{n+1\}$ implies $i \succ \{n+1\}$, and (2) (the relative ranking of schools is unchanged across rounds) for every $i, j \in \mathcal{S}$, if $i \hat{\succ} \{n+1\}$ and $i \hat{\succ} j$ then $i \succ j$.

We say that the type θ is **consistent** if the preferences $(\succ^\theta, \hat{\succ}^\theta)$ are consistent.

Assumption 3.1. *If $\zeta(\theta) > 0$ then the type θ is consistent.*

Assumption 3.2. *For all consistent types $\theta \in \Theta$ it holds that $\zeta(\theta) > 0$.*

The consistency assumption is required in order to meaningfully define strategy-proofness in our two-round setting, as we require truthful reporting in the first round to be optimal for both the student's first-round assignment as well as her second-round assignment. We use the full support assumption only to characterize our class of proposed mechanisms (Theorem 3.3) and do not need it for our positive results (Theorems 3.1 and 3.2 and supporting results) .

We also assume that the first-round lottery numbers are drawn independently and uniformly from $[0, 1]$ and do not depend on preferences. This means that for all $\theta \in \Theta$ and intervals (a, b) with $0 \leq a \leq b \leq 1$, $\eta(\{\theta\} \times (a, b)) = (b - a)\zeta(\theta)$.⁷

An *assignment* $\mu : \mathcal{S} \rightarrow \mathcal{C}$ specifies the school that each student is assigned to. For any assignment μ , we let $\mu(s)$ denote the school to which student s is assigned, and overloading notation we let $\mu(i)$ denote the *set* of students assigned to school i . We assume that $\mu(i)$ is η -measurable and that the assignment is *feasible*, i.e. $\eta(\mu(i)) \leq q_i$ for all $i \in \mathcal{C}$ and if $\mu(s) = i$ then $p_i^s \geq 0$. We let μ and $\hat{\mu}$ denote the first- and second-round assignments respectively.

Timeline. Students report first-round preferences \succ .⁸ The mechanism designer obtains a first-round assignment μ by running DA-STB with lottery L and announces μ and L . Students then observe their outside options and update their preferences accordingly to $\hat{\succ}$. Finally, students report their second-round preference $\hat{\succ}$, and the mechanism designer obtains a second-round assignment

⁷This can be justified via an axiomatization of the kind obtained by Al-Najjar (2004).

⁸Since we will be considering mechanisms that are strategy-proof in the large, we assume that students report truthfully; we do not distinguish between reported preferences and true preferences.




Figure 3.2: Timeline of the two-round mechanism design problem

$\hat{\mu}$ by running a reassignment mechanism M and announces $\hat{\mu}$. We illustrate the timeline in Figure 3.2.

Informational Assumptions. Eligibility and priorities are exogenously determined and publicly known. The mechanism is publicly announced before preferences are submitted. Before first-round reporting, each student knows her first-round preferences, and that her second-round preferences will be obtained from these preferences via truncation. Each student has imperfect information regarding her own second-round preferences (i.e., the point of truncation) at that stage, and believes with positive probability her preferences in both rounds will be identical.⁹ We assume students know the distribution η over student types and lotteries (an assumption we need only for our characterization result, Theorem 3.3). Each student is assumed to learn her lottery number after the first round, as in practice students are often permitted to inquire about their position in each school’s waitlist; however our results hold even if students do not learn their lottery numbers.

Definition 3.2. A student $s \in \mathcal{S}$ is a **reassigned student** if she is assigned to a different school in \mathcal{C} in the second round than in the first round. That is, s is a reassigned student under reassignment $\hat{\mu}$ if $\mu(s) \neq \hat{\mu}(s)$ and $\mu(s) \neq \{n + 1\}$, $\hat{\mu}(s) \neq n + 1$.¹⁰

The majority of reassignments happen around the start of the school year, a time when they are costly for schools and students alike. Hence, in addition to providing an efficient final assignment, we also want to reduce congestion by minimizing the number of reassigned students.

⁹This ensures that students will report their full first-round preferences in the first round, instead of truncating in the first round based on their beliefs about their second-round preferences.

¹⁰Several alternative definitions of reassigned students—such as counting students who are initially unassigned and end up at a school in \mathcal{C} , and/or counting initially assigned students who end up unassigned—could also be considered. We note that our results continue to hold for all these alternative definitions.

3.2.1 Mechanisms

A *mechanism* is a function that maps the realization of first-round lotteries L , school priorities p , and students' first-round preference reports \succ into an assignment μ . A *reassignment mechanism* is a function that maps the realization of first-round lotteries L , first-round assignment μ , school priorities p , and students' second-round reports $\hat{\succ}$ into a second-round assignment $\hat{\mu}$.¹¹ A *two-round mechanism* obtained from a reassignment mechanism M is a two-round mechanism where the first-round mechanism is DA-STB (see Definition 3.3), and the second-round mechanism is M .

In the first round, seats are assigned according to the student-optimal Deferred Acceptance (DA) algorithm with single tie-breaking (STB), which constructs an assignment as follows. A *single* lottery ordering of the students L is used to resolve ties in the priority groups at all schools, resulting in an instance of the two-sided matching problem with strict preferences and priorities. In each step, unassigned students apply to their most-preferred school that has not yet rejected them. A school with a capacity of q *tentatively accepts* the q highest-ranked eligible applicants (according to its priority ranking of the students after breaking all ties) and rejects any remaining applicants. The algorithm runs until there are no new student applications, at which point it terminates and assigns each student to her tentatively assigned school seat. The strict student preferences, weak school priorities, and the use of DA-STB mirror current practice in many school choice systems, such as those in New York City, Chicago, and Denver.

In the continuum model, Azevedo and Leshno (2016) have shown that DA may be formally defined in terms of admissions scores and cutoffs.

Definition 3.3 (Deferred Acceptance (Azevedo and Leshno, 2016)). The Deferred Acceptance mechanism with single tie-breaking (DA-STB) is a function $DA_\eta((\succ^s, p^s)_{s \in \mathcal{S}}, L)$ mapping student preferences, priorities and lottery numbers into an assignment μ , defined by a vector of cutoffs $\mathbf{C} \in \mathbb{R}_+^n$ as follows. Each student s is given a score $r_i^s = p_i^s + L(s)$ at school i and assigned to her most-preferred school as per her preferences, among those where her score exceeds the cutoff:

¹¹Here we restrict our design space to second-round assignments that depend on the first-round reports only indirectly, through the first-round assignment μ . We believe that this is a reasonable restriction, given that the second round occurs a significant period of time after the first round.

$$\mu(s) = \max_{\succ_s}(\{i \in \mathcal{C} : r_i^s \geq C_i\} \cup \{n+1\}). \quad (3.1)$$

Moreover, \mathbf{C} is market-clearing, namely

$$\eta(\mu(i)) \leq q_i \text{ for all } i \in \mathcal{C}, \text{ with equality if } C_i > 0. \quad (3.2)$$

Azevedo and Leshno (2016) showed that the set of assignments satisfying equations (3.1) and (3.2) forms a non-empty complete lattice, and typically consists of a single uniquely determined assignment.¹² This unique assignment in the continuum further corresponds to the limit of the set of stable matches obtained in finite markets as the number of students grows (with school capacities growing proportionally). Throughout this paper, in the (knife edge) case where there are multiple assignments satisfying Definition 3.3, we pick the student-optimal matching.

Given cutoffs $\{C_i\}_{i=1}^n$, we will also find it helpful to define for each priority class π the *cutoffs within the priority class* at each school $C_{\pi,i} \in [0, 1]$ by $C_{\pi,i} = 0$ if $C_i \leq \pi_i$, $C_{\pi,i} = 1$ if $C_i \geq \pi_i + 1$, and $C_{\pi,i} = C_i - \pi_i$ otherwise. Thus, $C_{\pi,i}$ is the lowest lottery number a student in the priority class π can have and still be able to attend school i .

We now turn to the mechanism design problem. We emphasize that we keep the first round consistent with currently used mechanisms and consider only two-round mechanisms whose first round mechanism is DA-STB, i.e. the only freedom afforded the planner is the design of the reassignment mechanism. We propose the following class of two-round mechanisms. Intuitively, these mechanisms run DA-STB twice, once in each round. They explicitly correlate the lotteries used in the two rounds via a permutation P , and give each student top priority in the school she was assigned to in the first round to guarantee that each student receives a (weakly) better assignment in the second round.

Definition 3.4 (Permuted Lottery Deferred Acceptance (PLDA) mechanisms). Let $P : [0, 1] \rightarrow [0, 1]$ be a measure-preserving bijection. Let L be the realization of first-round lottery numbers, and let μ be the first-round assignment obtained by running DA with lottery L .

¹²A sufficient condition for the assignment to be unique is when the demand $D(C)$ for schools given cutoffs C , defined by $D_i(P) = \eta(s \in \mathcal{S} \mid \max_{\succ_s} \{j \mid r_j^s \geq P_j\} = i)$, is continuously differentiable in the cutoffs. Moreover for almost all demand functions the resulting assignment is unique for all but a measure zero set of capacity vectors.

Define a new economy $\hat{\eta}$, where to each student $s \in \mathcal{S}$ with priority vector p^s , and first-round lottery and assignment $L(s) = \ell$, $\mu(s) = i$, we: (1) assign a lottery number $P(\ell)$; and (2) give top second-round priority $\hat{p}_i^s = n_i$ at their first-round assignment i and unchanged priority $\hat{p}_j^s = p_j^s$ at all other schools $j \neq i$. PLDA(P) is the two-round mechanism obtained using the reassignment mechanism $DA_{\hat{\eta}}((\hat{\succ}^s, \hat{p}^s)_{s \in \mathcal{S}}, P \circ L)$.

We use $\hat{C}_{\pi,i}^P$ to denote the second-round cutoff for priority class π in school i under PLDA(P).

We highlight two particular PLDA mechanisms. The RLDA (reverse lottery) mechanism uses the reverse permutation $R(x) = 1 - x$; and the FLDA (forward lottery) mechanism, which preserves the original lottery order, uses the identity permutation $F(x) = x$. By default, school districts often use a decentralized version of the FLDA mechanism, implemented via waitlists. In this paper, we provide evidence that supports using the *centralized RLDA mechanism* in a school system like that in NYC, where a large proportion of vacated seats are revealed close to or after the start of the school year, and where reassignments are costly for both students and the school administration.

The PLDA mechanisms are an attractive class of two-round assignment mechanisms for a number of reasons. They are intuitive to understand and simple to implement in systems already using DA. (A decentralized implementation would be even simpler to integrate with current practice; the currently used waitlist mechanism for reassignments can be retained with the simple modification of permuting the lottery numbers just before waitlists are constructed.) In addition, we will show that the PLDA mechanisms have desirable incentive and efficiency properties, which we now describe.

Any reassignment mechanism that takes away a student's initial assignment against her will be impractical. Thus, we require our mechanism to respect first-round guarantees:

Definition 3.5. A two-round mechanism (or a second-round assignment $\hat{\mu}$) **respects guarantees** if every student (weakly) prefers her second-round assignment to her first-round assignment, that is, $\hat{\mu}(s) \hat{\succ}^s \mu(s)$ for every $s \in \mathcal{S}$.

One of the reasons for the success of DA in practice is that it respects priorities: if a student is not assigned to a school she wants, it is because that school is full of students who have higher priority at that school. This leads to the following natural requirement in our two-round context:

Definition 3.6. A two-round mechanism (or a second-round assignment $\hat{\mu}$) **respects priorities (subject to guarantees)** if for every $i \in \mathcal{C}$ and eligible student $s \in \mathcal{S}$ such that $i \succ^s \hat{\mu}(s)$ and every student s' such that $\hat{\mu}(s') = i \neq \mu(s')$ it holds that s' is eligible for i and $p_i^{s'} \geq p_i^s$.

Thus, our definition of respecting priorities (subject to guarantees) requires that every student who was upgraded to a school i in the second-round must have a (weakly) higher priority at that school than every eligible student s who prefers i to her second-round assignment.

We now turn to incentive properties. In the school choice problem it is reasonable to assume that students will be strategic in how they interact with the mechanism at each stage. Hence, it is desirable that whenever a student (with consistent preferences) reports preferences, conditional on everything that has happened up to that point, it is a dominant strategy for her to report truthfully. To describe the properties formally, we start by fixing an arbitrary profile of first and second round preferences (\succ^{-s}, ζ^{-s}) for all the students other than student s . For any preference report of student s in the first round she will receive an assignment that is probabilistic because of the lottery used to break ties in the first round; then, after observing her first-round assignment and her updated outside option, she can submit a second-round preference report, based on which her final assignment is computed. This leads to two natural notions of strategy-proofness.

Definition 3.7. A two-round mechanism is **strongly strategy-proof** if for each student s (with consistent preferences) truthful reporting is a dominant strategy, i.e. for each realization of lottery numbers and profile of first- and second-round reported preferences of the students other than s , reporting her preferences truthfully in each of the two rounds is a best response for s .

Our definition of *strong strategy-proofness* is rather demanding: it requires that no student be able to manipulate the mechanism even if she has full knowledge of the first and second round preferences of *all* other students and the lottery numbers. We shall also consider a weaker version of strategy-proofness that applies when a manipulating student does not know the lottery number realizations when she submits her first-round preference report and learns all lottery numbers only after the end of the first round. In that case, each student views her first-round assignment as a probability vector; her second-round assignment is also random, but is a deterministic function of

the first-round outcome, the second-round reports, and the first-round lottery numbers. We make precise the notion of a successful manipulation in this setting as follows.

Definition 3.8. A two-round mechanism is **weakly strategy-proof** if the following conditions hold:

- Knowing the specific realization of first-round assignments (and the second round preferences of the students other than s), it is optimal for student s to submit her second-round preference truthfully, given what the other students do;
- For each student s (with consistent preferences), and for each profile of first- and second-round preferences of the students other than s , the probability that student s is assigned to one of her top k schools in the second round is maximized when she reports truthfully in the first round (assuming truthful reporting in the second round), for each $k = 1, 2, \dots, n$.

In other words, in each stage of the dynamic game, the second-round assignment from truthful reporting stochastically dominates the outcomes of all other strategies. We emphasize that the uncertainty in the assignment is solely due to the lottery numbers, which students initially do not know.

Note that a two-round mechanism that uses the first-round assignment as the initial endowment for a mechanism like top trading cycles in the second round will *not* be two-round strategy-proof, because students can benefit from manipulating their first-round reports to obtain a more popular initial assignment that they could use to their advantage in the second round.

Finally, we discuss some efficiency properties. To be efficient, clearly a mechanism should not leave unused any seats that are desired by students.

Definition 3.9. A two-round mechanism is **non-wasteful** if no student is assigned to a school she is eligible for that she prefers less than a school not at capacity; that is, for each student $s \in \mathcal{S}$ and schools i, j , if $\hat{\mu}(s) = i$ and $j \succ^s i$ and $p_j^s \geq 0$, then $\eta(\hat{\mu}(j)) = q_j$.

It is also desirable for a two-round mechanism to be Pareto efficient. We do not want any students to be able to improve their utility by swapping probability shares in second-round assignments. However, we also require that our reassignment mechanism respect guarantees and

priorities (see Definitions 3.5 and 3.6), which is incompatible with Pareto efficiency even in a static, one-round setting.¹³ This motivates the following definitions. Consider a second-round assignment $\hat{\mu}$. A *Pareto-improving cycle* is an ordered set of types $(\theta_1, \theta_2, \dots, \theta_m) \in \Theta^m$, sets of students $(\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_m) \in \mathcal{S}^m$, and schools $(1', 2', \dots, m') \in \mathcal{C}^m$, such that $\eta(\mathcal{S}_i) > 0$ and $(i+1)' \succ^{\theta_i} i'$ (where we define $(m+1)' = 1'$), for all i , and such that for each i , $\theta^s = \theta_i$ and $\hat{\mu}(s) = i'$ for all $s \in \mathcal{S}_i$.

Let \hat{p} be the second-round priorities obtained by giving each student s a top second-round priority $\hat{p}_i^s = n_i$ at their first-round assignment $\mu(s) = i$ (if $i \in S$) and unchanged priority $\hat{p}_j^s = p_j^s$ at all other schools $j \neq i$. We say that a Pareto-improving cycle (in a second-round assignment) *respects (second-round) priorities* if $\hat{p}_{(i+1)'}^{\theta_i} \geq \hat{p}_{(i+1)'}^{\theta_{i+1}}$ for all i (where we define $\theta_{m+1} = \theta_1$).

Definition 3.10. A two-round mechanism is **constrained Pareto efficient** if the second-round assignment has no Pareto-improving cycles that respect second-round priorities.

We remark that this is the same notion of efficiency that is satisfied by static, single-round DA-STB (Definition 3.3)—the resulting assignment has no Pareto-improving cycles that respect priorities. In other words, the constrained Pareto efficiency requirement is informally to be “as efficient as static DA”. We also note here that as a result of the requirement to respect second round priorities, Pareto improving cycles considered must include only reassigned students.

Finally, for equity purposes, it is desirable that a mechanism be anonymous.

Definition 3.11. A two-round mechanism is **anonymous** if students with the same first-round assignment and the same first- and second-round preference reports have the same distribution over second-round assignments.

We show that PLDA mechanisms satisfy all the aforementioned properties.

Proposition 3.1. *Suppose student preferences are consistent. Then PLDA mechanisms respect guarantees and priorities, and are strongly two-round strategy-proof, non-wasteful, constrained Pareto efficient, and anonymous.*

¹³When schools have strict preferences, an assignment respects priorities if and only if it is stable, and it is well known that in two-sided matching markets with strict preferences, there exist preference structures for which every stable assignment can be Pareto improved.

Proof. Fix a permutation P and some PLDA with permutation P . We show that this particular PLDA satisfies all the desired properties. Let η be a distribution of students, and let $\hat{\mathbf{C}}^P$ be the second-round cutoffs corresponding to the assignment given by the PLDA for this distribution of student types.

PLDA respects guarantees because fewer students are guaranteed at each school than the capacity of the school. PLDA is non-wasteful because the second round terminates with a stable matching where all schools find all students acceptable, which is non-wasteful.

We now show that the PLDA mechanism is strongly two-round strategy-proof. Since students are non-atomic, no student can change the cutoffs $\hat{\mathbf{C}}^P$ by changing her first- or second-round reports. Hence it is a dominant strategy for all students to report truthfully in the second round. Moreover, for any student s , the only difference between having a first-round guarantee at a school i and having no first-round guarantee is that in the former case, $\hat{r}_i^s = \hat{p}_i^s + L(s)$ increases from r_i^s by $n_i - p_i^s$. This means that having a guarantee at a school i changes the student's second-round assignment in the following way. She receives a seat in school i whereas without the guarantee she would have received a seat in some school j that she reported preferring less to school i , and her second-round assignment is unchanged otherwise.

Therefore, students want their first-round guarantee to be the best under their second-round preferences, and so it is a dominant strategy for students with consistent preferences to report truthfully in the first round.

PLDA is constrained Pareto efficient, since we use single tie-breaking and the output is the student-optimal stable matching with respect to the updated second-round priorities \hat{p} . This is easily seen via the cutoff characterization. Let the second-round cutoffs be \hat{P} , where overloading notation we let \hat{P}_i denote the cutoff for school \tilde{i} . Fix a Pareto-improving cycle $(\Theta^m, \mathcal{S}^m, \mathcal{C}^m)$. Without loss of generality we may assume that $\hat{p}_i^s + L(s) \geq P_i$ for all $s \in \mathcal{S}_i$, since the set of students for whom this is not true has measure 0. Moreover, since all students $s \in \mathcal{S}_i$ prefer school $(i+1)'$ to their assigned school $\hat{\mu}(s) = i'$, without loss of generality we may also assume that $\hat{p}_{(i+1)'}^s + L(s) < P_{i+1}$ for all $s \in \mathcal{S}_i$, since the set of students for whom this is not true has measure 0. This means that for all $s_i \in \mathcal{S}_i$ and $s_{i+1} \in \mathcal{S}_{i+1}$ it holds that $\hat{p}_{(i+1)'}^{s_i} + L(s_i) < P_{i+1} \leq \hat{p}_{(i+1)'}^{s_{i+1}} + L(s_{i+1})$,

and so $\hat{p}_{(i+1)'}^{s_i} \leq \hat{p}_{(i+1)'}^{s_{i+1}}$.

Suppose for the sake of contradiction that the cycle $(\Theta^m, \mathcal{S}^m, \mathcal{C}^m)$ respects second-round priorities. Then for each $s_i \in \mathcal{S}_i$ and $s_{i+1} \in \mathcal{S}_{i+1}$ it holds that $\hat{p}_{(i+1)'}^{s_i} \geq \hat{p}_{(i+1)'}^{s_{i+1}}$, and so $L(s_i) > L(s_{i+1})$. But since this holds for all i we obtain a cycle of lottery numbers $L(s_1) > L(s_2) > \dots > L(s_m) > L(s_1)$, which provides the necessary contradiction. \square

We will show in Section 3.3.1 that in a setting without priorities, the PLDA mechanisms are the only mechanisms that satisfy all these properties (and some additional technical requirements), even if we only require weak strategy-proofness (Theorem 3.3).

Finally, it is simple to show that the natural counterparts to PLDA mechanisms in a discrete setting (with a finite number of students) respect guarantees and priorities, and are non-wasteful, constrained Pareto efficient, and anonymous. We make these claims formal in Section 3.5 and also provide an informal argument that the discrete PLDA mechanisms are also approximately strategy-proof when the number of students is large.

3.3 Main Results

In this section, we will show that the defining characteristic of a PLDA mechanism—the permutation of lotteries between the two rounds—can be chosen to achieve desired operational goals. We first provide a simple and intuitive *order condition*, and show that under this condition, all PLDA mechanisms give the same ex ante allocative efficiency. Thus when the primitives of the market satisfy the order condition, it is possible to pursue secondary operational goals without sacrificing allocative efficiency. Next, in the context of reassigning school seats at the start of the school year, we consider the specific problem of minimizing reassignment, and show that when the order condition is satisfied, *reversing* the lottery minimizes reassignment among all centralized PLDA mechanisms. In Section 3.6, we empirically demonstrate using data from NYC public high schools that reversing the lottery minimizes reassignment (amongst a subclass of centralized PLDA mechanisms) and does not significantly affect allocative efficiency even when the order condition does not hold exactly. Our results suggest that centralized RLDA is a good choice of mechanism when the

primary goal is to minimize reassignments while providing a second-round assignment with high allocative efficiency. In Section 3.3.1 we provide an axiomatic justification for PLDA mechanisms, and later in Section 3.7 we discuss how the choice of lottery permutation can be used to achieve other operational goals, such as maximizing the number of students with improved assignments.

We begin by defining the order condition, which we will need to state our main results.

Definition 3.12. The **order condition** holds on a set of primitives $(\mathcal{C}, \mathcal{S}, \eta, q)$ if for every priority class π , the first- and second-round school cutoffs under RLDA within that priority class are in the same order, i.e., for all $i, j \in \mathcal{C}$,

$$C_{\pi,i} > C_{\pi,j} \Rightarrow \hat{C}_{\pi,i}^R \geq \hat{C}_{\pi,j}^R.$$

We emphasize that the order condition is a condition on the market primitives, namely, school capacities and priorities and student preferences (though checking whether it holds involves investigating the output of RLDA). We may interpret the order condition as an indication that the relative demand for the schools is consistent between the two rounds. Informally speaking, it means that the revelation of the outside options does not change the *order* in which schools are overdemanded. One important setting where the order condition holds is the case of uniform dropouts and a single priority type. In this setting, each student independently with probability ρ either remains in the system and retains her first-round preferences in the second round, or drops out of the system entirely; student first-round preferences and school capacities are arbitrary. We establish the order condition and provide direct proofs of several of our theoretical results for the setting with uniform dropouts in Section 3.4, in order to give a flavor of the arguments employed to establish our results in the general setting.

To compare the allocative efficiency of different mechanisms, we define *type-equivalence* of assignments. In words, two second-round assignments are type-equivalent if the masses of different student types θ assigned to each school are the same across the two assignments.

Definition 3.13. Two second-round assignments $\hat{\mu}$ and $\hat{\mu}'$ are said to be **type-equivalent** if¹⁴

¹⁴We remark that the type-equivalence condition is well defined in the space of interest. Specifically, although for

In our continuum model, if two two-round mechanisms produce type-equivalent second-round assignments we may equivalently interpret them as providing each individual student of type θ with the same ex ante distribution (before lottery numbers are assigned) over assignments.

We are now ready to state the main results of this section. The first is a surprising finding that under the order condition, all PLDAs are allocatively equivalent.

Theorem 3.1 (Order condition implies type-equivalence). *If the order condition (Definition 3.12) holds, all PLDA mechanisms produce type-equivalent second-round assignments.*

Thus, if the order condition holds, the measure of students of type $\theta \in \Theta$ assigned to each school in the second round is independent of the permutation P . The intuition behind this result is that if the cutoffs are in the same order under RLDA, then the cutoffs are in the same order under any PLDA, which can be used to show that aggregate final outcomes are equivalent across mechanisms. We remark that type equivalence does not imply an equal (or similar) amount of reassignment, as type-equivalence depends only on the second-round assignment, while reassignment (Definition 3.2) measures the difference between the first- and second-round assignments. For example, in the example in Figure 3.1 if schools have a single priority class then FLDA and RLDA each give each remaining student a one fifth chance of being assigned to each of schools 1 through 5, but FLDA performs 5 reassignments whilst RLDA performs only 1. This brings us to our second result.

Theorem 3.2 (Reverse lottery minimizes reassignment). *If all PLDA mechanisms produce type-equivalent second-round assignments, then RLDA minimizes the measure of reassigned students among PLDA mechanisms.*

Proof. Fix $\theta = (\succ^\theta, \hat{\succ}^\theta, p^\theta) \in \Theta$ and school $i \in \mathcal{C}$. We will show that, among all type equivalent mechanisms, RLDA minimizes the measure of reassigned students with type θ who were assigned to school i in the second round. The idea is that RLDA never reassigns a student of type θ into a school i if it has reassigned a student of type θ out of school i .

Formally, for every permutation P , let the measure of students with type θ leaving and entering school i in the second round under PLDA(P) be denoted by $\ell_P = \eta(\{s \in \mathcal{S} : \theta^s = \theta, \mu(s) =$

general random mechanisms these measures are random variables, in the case of PLDA mechanisms, these measures are a deterministic function of priorities and preferences, and the equality is well defined.

$i, \hat{\mu}^P(s) \neq i\}$ and $e_P = \eta(\{s \in \mathcal{S} : \theta^s = \theta, \mu(s) \neq i, \hat{\mu}^P(s) = i\})$ respectively. Due to type-equivalence, there is a constant c such that $\ell_P = e_P - c$ for all permutations P . We show that either $\ell_R = 0$ and $e_R = c$, or $e_R = 0$, implying that $e_R \leq e_P$ for all permutations P .

If both $e_R > 0$ and $\ell_R > 0$, then students of type θ who entered i in the second round of RLDA had worse first- and second-round lottery numbers than students who left i in the second round of RLDA, which contradicts the reversal of the lottery. Formally, suppose $e_R > 0$ and $\ell_R > 0$. If $n + 1 \succ^\theta i$ then $e_R = 0$, so we may assume $i \succ^\theta n + 1$. Since $e_R > 0$, there exists some student $s \in \mathcal{S}$ with type $\theta^s = \theta$ for whom $i = \hat{\mu}^R(s) \succ^\theta \mu(s)$. By consistency, we have $i \succ^\theta \mu(s)$, and therefore s could not afford (meet the cutoff for) i in the first round. Since $\ell_R > 0$, there exists some student $s' \in \mathcal{S}$ with type $\theta^{s'} = \theta$ for whom $j = \hat{\mu}^R(s') \succ^\theta \mu(s') = i$. By definition, s' could afford i in the first round and s could not, and hence $L(s') > L(s)$. Note that since $i \succ^\theta n + 1$, it follows that $j \succ^\theta i \succ^\theta n + 1$. Now, since s' received a better second-round assignment under RLDA than s and both s and s' were reassigned under RLDA, it follows that $R(L(s')) > R(L(s))$, which is a contradiction. Since $e_P = \ell_P + c \geq c$ and $e_P \geq 0$ this completes the proof. \square

Our results present a strong case for using the centralized RLDA mechanism when the main goals are to achieve allocative efficiency and minimize the number of reassigned students. Theorems 3.1 and 3.2 show that when the order condition holds, centralized RLDA is unequivocally optimal in the class of PLDA mechanisms, since all PLDA mechanisms give type-equivalent assignments¹⁵ and centralized RLDA minimizes the number of reassigned students. In addition, we remark that the order condition can be checked easily by running RLDA (e.g., on historical data).¹⁶

Next, we give examples of when the order condition holds and does not hold, and illustrate the resulting implications for type-equivalence. We illustrate these in Figure 3.3.

Example 3.1. There are $n = 2$ schools, each with a single priority group. School 1 has lower capacity and is initially more overdemand. Student preferences are such that when all students

¹⁵A reasonable utility model in the continuum would yield that type-equivalence implies welfare equivalence.

¹⁶We are not suggesting that the mechanism should involve checking the order condition and then using centralized RLDA only if this condition is satisfied (based on the guarantee in Theorems 3.1 and 3.2). However, one could check whether the order condition holds on historical data and accordingly decide whether to use the centralized RLDA mechanism or not.

who want only 2 drop out the order condition holds, and when all students who want only 1 drop out school 2 becomes more overdemanded under RLDA and the order condition does not hold.

School capacities are given by $q_1 = 2$, $q_2 = 5$. There is measure 4 of each of the four types of first-round student preferences. Let θ_i denote the student type that finds only school i acceptable, and let $\theta_{i,j}$ denote the type that finds both schools acceptable and prefers i to j . (We will define the second round preferences of each student type below; each type will either leave the system completely or keep the same preferences.) If we run DA-STB, the first-round cutoffs are $(C_1, C_2) = \left(\frac{3}{4}, \frac{1}{2}\right)$.

Suppose that all type θ_2 students leave the system, and all students of other types stay in the system and keep the same preferences as in the first round. This frees up 2 units at school 2. Under RLDA, the second-round cutoffs are $(\hat{C}_1^R, \hat{C}_2^R) = \left(1, \frac{3}{4}\right)$. In this case, the order condition holds and FLDA and RLDA are type-equivalent. It is simple to verify that both FLDA and RLDA assign the same measure $\hat{\mu}(i)$ of students of type $(\theta_1, \theta_{1,2}, \theta_{2,1})$ to school i , where

$$\hat{\mu}^F = \hat{\mu}^R = (\hat{\mu}(1), \hat{\mu}(2)) = ((1, 1, 0), (0, 2, 3)).$$

Suppose that all type θ_1 students leave the system, and all students of other types stay in the system and keep the same preferences as in the first round. This frees up 1 unit at school 1. Under RLDA, no new students are assigned to school 2, and the previously bottom-ranked (but now top-ranked) measure 1 of students who find school 1 acceptable are assigned to school 1. Hence the second-round cutoffs are $(\hat{C}_1^R, \hat{C}_2^R) = \left(\frac{7}{8}, 1\right)$. In this case, the order condition does not hold. Type equivalence also does not hold (and in fact FLDA and RLDA give different ex ante assignments to students of every remaining type), since the FLDA and RLDA assignments are

$$\hat{\mu}^F = \left((2, 0, 0), \left(\frac{1}{3}, \frac{7}{3}, \frac{7}{3}\right)\right), \quad \hat{\mu}^R = ((1.5, 0.5, 0), (1, 2, 2)).$$

3.3.1 Axiomatic Justification of PLDA Mechanisms

We have shown that PLDA mechanisms satisfy a number of desirable properties. Namely, PLDA mechanisms respect guarantees and priorities, and are two-round strategy-proof (in a strong sense), non-wasteful, constrained Pareto efficient, and anonymous. In this section, we show that in a setting

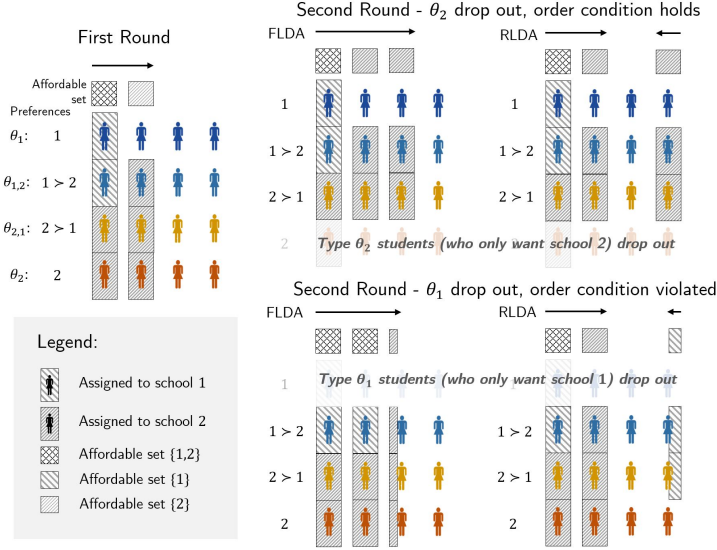


Figure 3.3: In Example 3.1, FLDA and RLDA are type equivalent when the order condition holds, and give different assignments to students of every type when the order condition does not hold.

The initial economy and first-round assignment are depicted on the top left. On the right, we show the second-round assignments under FLDA and RLDA when type θ_2 students (who want only school 2) drop out, and when type θ_1 students (who want only school 1) drop out. Students toward the left have larger first round lottery numbers. The patterned boxes above each column of students indicate the affordable sets for students in that column. When students who want only school 2 drop out, the order condition holds, and FLDA and RLDA are type-equivalent. When students who want only school 1 drop out, school 2 becomes more overdemanded in RLDA, and FLDA and RLDA give different ex ante assignments to students of every remaining type.

with a single priority class the PLDA mechanisms are the *only* mechanisms that satisfy both these properties as well as two mild technical conditions on the symmetry of the mechanism, even when we require only the weaker version of two-round strategy-proofness.

Definition 3.14. A two-round mechanism satisfies the **averaging** axiom if for every type θ and pair of schools (i, j) the randomization of the mechanism does not affect the measure of students with type θ assigned to (i, j) in the first and second rounds, respectively. That is, for all θ, i, j , there exists a constant $c_{\theta, i, j}$ such that $\eta(\{s \in \mathcal{S} : \theta^s = \theta, \mu(s) = i, \hat{\mu}(s) = j\}) = c_{\theta, i, j}$ w.p. 1.

Definition 3.15. A two-round mechanism is **non-atomic** if any single student changing her preferences has no effect on the assignment probabilities of other students.

Our characterization result is the following.

Theorem 3.3. *Suppose that student preferences are consistent and student types have full support (Assumptions 3.1 and 3.2). A non-atomic two-round assignment mechanism where the first round is DA-STB respects guarantees and is*

- *non-wasteful,*
- *(weakly) two-round strategy-proof,*
- *constrained Pareto efficient,*
- *anonymous, and*
- *averaging,*

if and only if the second-round assignment is given by PLDA. (Here the permutation P may depend on the measure of student preference types $\zeta(\cdot)$.)

We remark that we require two-round strategy-proofness only for students whose true preference type is consistent. This is because preference inconsistencies across rounds can lead to conflicts between the desired first-round assignment with respect to first-round preferences and the desired first-round guarantee with respect to second-round preferences, making it unclear how to even define a best response. Moreover, it may be reasonable to assume that students who are sophisticated enough to strategize about misreporting in the first round in order to affect the guarantee structure in the second round will also know their second-round preferences over schools in \mathcal{C} (i.e., everything except where they rank their outside option) at the beginning of the first round, and hence will have consistent preferences.¹⁷ We remark also that the ‘only if’ direction of this result is the only place where we require the full support assumption (Assumption 3.2).

The main focus of our result is the effect of cross-round constraints. By assumption, the first-round mechanism is DA-STB. It is relatively straightforward to deduce that the second-round mechanism also has to be DA-STB. Strategy-proofness in the second round, together with non-wastefulness, respecting priorities and guarantees, and anonymity, constrain the second round to

¹⁷One obvious objection is that students may also obtain extra utility from staying at a school between rounds, or, equivalently, they may have a disutility for moving, creating inconsistent preferences where the school they are assigned to in the first round becomes preferred to previously more desirable schools. We remark that Theorem 3.3 extends to the case of students whose preferences incorporate additional utility if they stay put, provided that the utility is the same at every school for a given student or satisfies a similar non-crossing property.

be DA, with each student given a guarantee at the school she was assigned to in the first round, and constrained Pareto efficiency forces the tiebreaking to be in the same order at all schools. The cross-round constraints are more complicated, but can be understood using *affordable sets*. A student's affordable set is the set of schools that she can choose to attend, i.e., the first-round affordable set is the set of schools for which she meets the first-round cutoff, and the affordable set is the set of schools for which she meets the first- or second-round cutoff. The set of possible affordable sets is uniquely determined by the order of cutoffs. By carefully using two-round strategy-proofness and anonymity, we show that a student's preference type does not affect the joint distribution over her first-round affordable set and affordable set, and hence her second-round lottery is a permutation of her first-round lottery that does not depend on her preference type.

Our result mirrors similar large market cutoff characterizations for single-round mechanisms by Liu and Pycia (2016) and Ashlagi and Shi (2014), which show, in settings with a single and multiple priority types respectively, that a mechanism is non-atomic, strategy-proof, symmetric, and efficient (in each priority class) if and only if it can be implemented by lottery-plus-cutoff mechanisms, which provide random lottery numbers to each student and admit them to their favorite school for which they meet the admission cutoff. We obtain such a characterization in a two-round setting using the fact that the mechanism respects guarantees and introducing an *affordable set* argument to isolate the second round from the first. This simplification allows us to employ arguments similar to those used in Liu and Pycia (2016) and Ashlagi and Shi (2014) to show that the first- and second-round mechanisms can be individually characterized using lottery-plus-cutoff mechanisms.

3.4 Intuition for Main Results

In this section, we provide some intuition for our main results. We also furnish full proofs for a special case of our model to give the interested reader a taste of the general proof techniques in a more transparent setting. This section may be skipped at a first reading without loss of continuity.

We begin with some definitions and intuition for our general results. A key conceptual insight is that we can simplify the analysis by shifting away from student assignments, which depend on

student preferences, and considering instead the options that a student is allowed to choose from, which are independent of preferences. Specifically, if we define the *affordable set* for each student as the set of schools for which she meets either the first- or second-round cutoffs, then each student is assigned to her favorite school in her affordable set at the end of the second round, and changing the student's preferences does not change her affordable set in our continuum model. Moreover, affordable sets and preferences uniquely determine demand.

The main technical idea that we use in establishing our main results is that the order condition is equivalent to the following seemingly much more powerful “global” order condition.

Definition 3.16. We say that PLDA(P) satisfies the **local order condition** on a set of primitives $(\mathcal{C}, \mathcal{S}, \eta, q)$ if, for every priority class π , the first- and second-round school cutoffs within that priority class are in the same order under PLDA(P). That is, for all $i, j \in \mathcal{C}$,

$$C_{\pi,i} > C_{\pi,j} \Rightarrow \hat{C}_{\pi,i}^P \geq \hat{C}_{\pi,j}^P.$$

We say that the **global order condition** holds on a set of primitives $(\mathcal{C}, \mathcal{S}, \eta, q)$ if:

- (*Consistency across rounds*) PLDA(P) satisfies the local order condition on $(\mathcal{C}, \mathcal{S}, \eta, q) \forall P$;
- (*Consistency across permutations*) For every priority class π , for all pairs of permutations P, P' and schools $i, j \in \mathcal{C} \cup \{n+1\}$, it holds that $\hat{C}_{\pi,i}^P > \hat{C}_{\pi,j}^P \Rightarrow \hat{C}_{\pi,i}^{P'} \geq \hat{C}_{\pi,j}^{P'}$.

In other words, the global order condition requires that all PLDA mechanisms result in the same order of school cutoffs in both rounds. Surprisingly, if the cutoffs are in the same order in both rounds under RLDA, then they are in the same order in both rounds under any PLDA.

Theorem 3.4. *The order condition (Definition 3.12) holds for a set of primitives $(\mathcal{C}, \mathcal{S}, \eta, q)$ if and only if the global order condition holds for $(\mathcal{C}, \mathcal{S}, \eta, q)$.*

We provide some intuition as to why Theorem 3.4 holds by using the affordable set framework. Under the reverse permutation, the sets of schools that enter a student's affordable set in the first and second rounds respectively are maximally misaligned. Hence, if the cutoff order is consistent

across both rounds under the reverse permutation, then the cutoff order should also be consistent across both rounds under any other permutation.

The affordable set framework also sheds some light on the power of the global order condition. Fix a mechanism and suppose that the first- and second-round cutoffs are in the same order. Then each student s 's affordable set is of the form $X_i = \{i, i + 1, \dots, n\}$ for some $i = i(s)$, where schools are indexed in decreasing order of their cutoffs for the relevant priority group $\pi = p^{\theta^s}$, and the probability that a student receives some affordable set is independent of her preferences. Moreover, since affordable sets are nested $X_1 \supseteq X_2 \supseteq \dots \supseteq X_n$, and since the lottery order is independent of student types, the demand for schools is uniquely identified by the proportion of students whose affordable set contains school i for each i . When the global order condition holds, this is true for every PLDA mechanism individually, which provides enough structure to induce type-equivalence.

We now introduce a special case of our model. For this special case, we will prove that the order condition holds, and show that all PLDA mechanisms give type-equivalent assignments.

Definition 3.17. (Informal) A market satisfies **uniform dropouts** if there is exactly one priority group at each school, students leave the system independently with some fixed probability ρ , (as formalized in Equation (3.3) below), and the students who remain in the system retain their preferences.

Before formalizing the definition and results for this setting, we provide some intuition for why the global order condition always holds under uniform dropouts. In the uniform dropouts model, each student drops out of the system with probability ρ , e.g. due to leaving the city after the first round for reasons that are independent of the school choice system. The second-round problem can thus be viewed as a rescaled version of the first-round problem; in particular, the measure of remaining students who were assigned to each school i in the first round is $(1 - \rho)q_i$, the measure of students of each type θ assigned to each school is scaled down by $1 - \rho$, the capacity of each school is still q_i , and the measure of students of each type θ who are still in the system is scaled down by $1 - \rho$. Thus schools fill in the same order regardless of the choice of permutation.

Let us now formalize our definitions and results. Throughout the rest of this section, since there are no priorities, we will let student types be defined either by $\theta = (\succ^\theta, \hat{\succ}^\theta, \mathbf{1})$ or simply by

$\theta = (\succ^\theta, \hat{\succ}^\theta)$. We define uniform dropouts with probability ρ by

$$\begin{aligned}\zeta(\{\theta \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = (n+1) \succ \dots\}) &= \rho \zeta(\{\theta \in \Theta : \succ^\theta = \succ\}), \\ \zeta(\{\theta \in \Theta : \succ^\theta = \succ, \hat{\succ}^\theta = \succ\}) &= (1 - \rho) \zeta(\{\theta \in \Theta : \succ^\theta = \succ\}),\end{aligned}\quad (3.3)$$

i.e. for every strict preference \succ over schools, students with first-round preferences \succ with probability ρ find the outside option $n+1$ the most attractive in the second round, and with probability $1 - \rho$ retain the same preferences in the second round.¹⁸

We show first that the global order condition (Definition 3.16) holds in the setting with uniform dropouts. The high level steps and algebraic tools used in this proof are similar to those used to show that the order condition is equivalent to the global order condition in our general framework (Theorem 3.4), although the analysis in each step is greatly simplified. We provide some intuition as to the differences in this section, and furnish the full proof of Theorem 3.4 in the Appendix.

Theorem 3.5. *In any market with uniform dropouts (Definition 3.17), the global order condition (Definition 3.16) holds.*

Proof. The main steps in the proof are as follows: (1) Assuming that every student's affordable set is X_i for some i , for every school j , guess the proportion of students who should receive an affordable set that contains j . (2) Calculate the corresponding second-round cutoffs \tilde{C}_j for school j . (3) Show that these cutoffs are in the same order as the first-round cutoffs. (4) Use the fact that the cutoffs are in the same order to verify that the cutoffs are market-clearing, and deduce that the constructed cutoffs are precisely the PLDA(P) cutoffs.

Throughout this proof, we amend the second-round score of a student s under PLDA(P) to be $\hat{r}_i^s = P(L(s)) + \mathbf{1}_{\{L(s) \geq C_i\}}$, meaning that we give each student a guarantee at any school for which she met the cutoff in the first round. By consistency of preferences, it is easily seen that this has no effect on the resulting assignment or cutoffs. Let the first-round cutoffs be C_1, C_2, \dots, C_n , where without loss of generality we index the schools such that $C_1 \geq C_2 \geq \dots \geq C_n$.

¹⁸We remark that there is a well-known technical measurability issue w.r.t. a continuum of random variables, and that this issue can be handled; see, for example, Al-Najjar (2004).

(1) In the setting with uniform dropouts, since the second-round problem is a rescaled version of the first-round problem (with a $(1 - \rho)$ fraction of the original students remaining), we guess that we want the proportion of students with an affordable set containing school j to be $\frac{1}{1-\rho}$ times the original proportion. (In the general setting, we no longer have a rescaled problem and so we instead guess that the proportion of students with each affordable set is the same as that under RLDA.)

(2) We translate this into cutoffs in the following way. Let $f_i^P(x) = |\{\ell : \ell \geq C_i \text{ or } P(\ell) \geq x\}|$ be the proportion of students who receive school i in their (second-round) affordable set with the amended second-round scores under permutation P if the first- and second-round cutoffs are C_i and x respectively. Notice that $f_i(x)$ is non-increasing for all i , $f_i(0) = 1$, $f_i(1) = 1 - C_i$, and if $i < j$ then $f_i(x) \leq f_j(x)$ for all $x \in [0, 1]$. Let the cutoff $\tilde{C}_i^P \in [0, 1]$ be the minimal cutoff satisfying the equation $f_i(\tilde{C}_i^P) = \frac{1}{1-\rho}(1 - C_i)$, and let $\tilde{C}_i^P = 0$ if $C_i < \rho$. (In the general setting the cutoffs are defined using the same functions $f_i^P(\cdot)$ with the proportions being equal to those that arise under RLDA, as mentioned in step (1) above.)

(3) We now show that the cutoffs $\tilde{\mathbf{C}}$ are in the right order. Suppose that $i < j$. If $\tilde{C}_i^P = 0$ then $C_j \leq C_i \leq \rho$ and so $\tilde{C}_j^P = 0 \leq \tilde{C}_i^P$ as required. Hence we may assume that $\tilde{C}_i^P, \tilde{C}_j^P > 0$. In this case, since $f_j(\cdot)$ is non-increasing and \tilde{C}_j^P is minimal, we can deduce that $\tilde{C}_j^P \leq \tilde{C}_i^P$ if $f_j(\tilde{C}_j^P) \geq f_j(\tilde{C}_i^P)$. It remains to establish the latter. Using the definition of f_j and f_i , we have

$$\begin{aligned} f_j(\tilde{C}_i^P) &= f_i(\tilde{C}_i^P) + \left| \left\{ \ell : \ell \in [C_j, C_i), P(\ell) < \tilde{C}_i^P \right\} \right| \\ &\leq \frac{1}{1-\rho}(1 - C_i) + (C_i - C_j) \leq \frac{1}{1-\rho}(1 - C_j) = f_j(\tilde{C}_j^P), \end{aligned}$$

where both inequalities hold since $C_j \leq C_i$. It follows that $\tilde{C}_i^P \geq \tilde{C}_j^P$, as required. (In the general setting, since we cannot give closed form expressions for the proportions $f_i(\tilde{C}_i^P)$ in terms of the cutoffs C_i , this step requires using the intermediate value theorem and an inductive argument.)

(4) We now show that $\tilde{\mathbf{C}}^P$ is the set of market-clearing DA cutoffs for the second round of PLDA(P). Note that $\gamma_i = C_{i-1} - C_i$ is the proportion of students whose first-round affordable set is X_i (where $C_0 = 1$). Since dropouts are uniform at random, this is the proportion of such students out of the total number of remaining students both before and after dropouts.

Consider first the case $\tilde{C}_i^P > 0$. Now $f_i(\tilde{C}_i^P)$ is the proportion of students whose second-round

affordable set contains i , and since $C_1 \geq C_2 \geq \dots \geq C_n$ and $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \dots \geq \tilde{C}_n^P$, it follows that the affordable sets are nested. Hence the proportion of students (of those remaining after students drop out) whose second-round affordable set is X_i is given by (where $f_0(\cdot) \equiv 1$)

$$\gamma_i^P = f_i(\tilde{C}_i^P) - f_{i-1}(\tilde{C}_{i-1}^P) = \frac{C_{i-1} - C_i}{1 - \rho} = \frac{\gamma_i}{1 - \rho}.$$

For each student type $\theta = (\succ, \succ)$ and set of schools X , let $D^\theta(X)$ be the maximal school in X under \succ , and let $\theta' = (\succ, \hat{\succ})$ be the student type consistent with θ that finds all schools unacceptable in the second round. Then a set of students of measure

$$\sum_{j \leq i} \sum_{\theta \in \Theta: D^\theta(X_j)=i} \gamma_j^P \zeta(\theta) = \sum_{j \leq i} \gamma_j \sum_{\theta \in \Theta: D^\theta(X_j)=i} \frac{\zeta(\theta)}{1 - \rho} = \sum_{j \leq i} \gamma_j \sum_{\theta \in \Theta: D^\theta(X_j)=i} \zeta(\theta) + \zeta(\theta')$$

choose to go to school i in the second round under the second-round cutoffs $\tilde{\mathbf{C}}^P$. We observe that the expression on the right gives the measure of the set of students who choose to go to school i in the first round under first-round cutoffs \mathbf{C} .

In the case where $\tilde{C}_i^P = 0$ the above expressions give upper bounds on the measure of the set of students who choose to go to school i in the second round under the second-round cutoffs $\tilde{\mathbf{C}}^P$. Since \mathbf{C} are market-clearing cutoffs, and $\tilde{C}_i^P > 0 \Rightarrow C_i^P > 0$, it follows that $\tilde{\mathbf{C}}^P$ are market-clearing cutoffs too. We have shown that in PLDA(P), the second-round cutoffs are exactly the constructed cutoffs $\tilde{\mathbf{C}}^P$ and they satisfy $\tilde{C}_1^P \geq \dots \geq \tilde{C}_n^P$, and so the global order condition holds. \square

The general proof of Theorem 3.4 uses the cutoffs for RLDA in steps (1) and (2) above to guess the proportion of students who receive an affordable set that contains school j , and requires that each student priority type be carefully accounted for. However, the general structure of the proof is similar, and the tools used are straightforward generalizations of those used in the proof above.

We next show that Theorem 3.1 holds with uniform dropouts. Specifically, we show that all PLDA mechanisms give type-equivalent assignments.

Proposition 3.2. *In any market with uniform dropouts (Definition 3.17), all PLDA mechanisms produce type-equivalent assignments.*

Proof. The proposition follows immediately from the fact that the proportion γ_i^P of students whose second-round affordable set is X_i does not depend on P . In more detail, consider first the case when all schools reach capacity in the second round of PLDA. We showed in the proof of Theorem 3.5 that for all i and all student types θ , the proportion of students of type θ with affordable set X_i in the second round under PLDA(P) is given by $\gamma_i^P = \frac{\gamma_i}{1-\rho}$, where γ_i is the proportion of students of type θ with affordable set X_i in the first round. It follows that all PLDAs are “type-equivalent” to each other because they are type-equivalent to the first-round assignment in the following sense. For each preference order \succ , let $\tilde{\succ}$ be the preferences obtained from \succ by making the outside option the most desirable, i.e., $n+1 \tilde{\succ} \dots$. Then

$$\begin{aligned} \eta(\{s \in \mathcal{S} : \theta^s = (\succ, \succ), \hat{\mu}^P(s) = i\}) &= \frac{1}{1-\rho} \eta(\{s \in \mathcal{S} : \theta^s = (\succ, \succ), \mu(s) = i\}) \\ &= \eta(\{s \in \mathcal{S} : \theta^s \in \{(\succ, \succ), (\succ, \tilde{\succ})\}, \mu(s) = i\}), \end{aligned}$$

where the second equality holds since students stay in the system uniformly-at-random with probability $1 - \rho$. Under uniform dropouts this holds for all student types that remain in the system, and so it follows that $\hat{\mu}^P$ is type-equivalent to $\hat{\mu}^{P'}$ for all permutations P, P' .

When some school does not reach capacity in the second round, we can show by induction on the number of such schools that all PLDAs are type-equivalent to RLDA. \square

Remark. Most of the results of this section extend to the following generalization of the uniform dropouts setting. A market satisfies *uniform dropouts with inertia* if there is exactly one priority group at each school, students leave the system independently with some fixed probability ρ , remain and wish to stay at their first round assignment with some fixed probability ρ' (have ‘inertia’), and otherwise remain and retain their first round preferences.¹⁹ It can be shown that in such a market, the global order condition always holds, and RLDA minimizes reassignment amongst all type-equivalent allocations. Moreover, if all students are assigned in the first round, it can also be shown that PLDA mechanisms produce type-equivalent allocations.

¹⁹This market is slightly beyond the scope of our general model, as the type of the student now also has to encode second-round preferences that depend on the first-round assignment, namely whether they have inertia.

3.5 PLDA for a Discrete Set of Students

Before verifying our theoretical results through simulations, we formally define and show how to implement PLDA mechanisms in a discrete setting with a finite number of students. We also prove that they retain almost all the desired incentive and efficiency properties discussed in Section 3.2.1.

3.5.1 Discrete Model

A finite set $\mathcal{S} = \{1, 2, \dots, N\}$ of students are to be assigned to a set $\mathcal{C} = \{1, \dots, n\}$ of schools. Each student can attend at most one school. As in the continuum model, for every school i , let $q_i \in \mathbb{N}_+$ be the *capacity* of school i , i.e., the number of students the school can accommodate. Let $n+1 \notin \mathcal{C}$ denote the outside option, and assume $q_{n+1} = \infty$. For each set of students $S \subseteq \mathcal{S}$ we let $\eta(S) = |S|$ be the number of students in the set. As in the continuum model, each student $s = (\theta^s, L(s)) \in \mathcal{S}$ has a type $\theta^s = (\succ^s, \hat{\succ}^s, p^s)$ and a first-round lottery number $L(s) \in [0, 1]$, which encode both student preferences and school priorities.

The first-round lottery numbers $L(s)$ are i.i.d. random variables drawn uniformly from $[0, 1]$ and do not depend on preferences. These random lottery numbers L generate a uniformly random permutation of the students based on the order of their lottery numbers.

An *assignment* $\mu : \mathcal{S} \rightarrow \mathcal{C}$ specifies the school that each student is assigned to. For an assignment μ , we let $\mu(s)$ denote the school to which student s is assigned, and in a slight abuse of notation, we let $\mu(i)$ denote the *set* of students assigned to school i . As in the continuum model, we say that a student $s \in \mathcal{S}$ is a *reassigned student* if she is assigned to a school in \mathcal{C} in the second round that is different to her first-round assignment.

3.5.2 PLDA Mechanisms & Their Properties

We now formally define PLDA mechanisms in a setting with a finite number of students. In order to do so, we use the algorithmic description of DA and extend it to a two-round setting. This also provides a clear way to implement PLDA mechanisms in practice. We first reproduce the widely deployed DA algorithm, and then proceed to define PLDAs.

Definition 3.18. The Deferred Acceptance algorithm with single tie-breaking is a function $DA((\succ^s, p^s)_{s \in \mathcal{S}}, L)$ mapping the student preferences in the first round, priorities and lottery numbers into an assignment μ constructed as follows. In each step, unassigned students apply to their most-preferred school that has not yet rejected them. A school with a capacity of q tentatively assigns a seat to each of its q highest-ranked applicants, ranked according to its priority ranking of the students with ties broken by giving preference to higher lottery numbers L (or tentatively assigns seats to all applicants, if fewer than q have applied), and rejects any remaining applicants, and the algorithm moves on to the next step. The algorithm runs until there are no new student applications, at which point it terminates and assigns each student to her tentatively assigned school seat.

Definition 3.19 (Permuted Lottery Deferred Acceptance (PLDA) mechanisms). Let P be a permutation of \mathcal{S} . Let L be the realization of first-round lottery numbers, and let μ be the first-round assignment obtained by running DA with lottery L . The permuted lottery deferred acceptance mechanism associated with P (PLDA(P)) is the mechanism that then computes a second-round assignment $\hat{\mu}^P$ by running DA on the same set of students \mathcal{S} but with student preferences $\hat{\succ}$, a modified lottery $P \circ L$, and modified priorities \hat{p} that give each student top priority at the school she was assigned to in the first round. Specifically, each school i 's priorities $\hat{\succ}_i$ are defined by lexicographically ordering the students first by whether they were assigned to i in the first round, and then according to p_i . PLDA(P) is the two-round mechanism obtained from using the reassignment mechanism $DA((\hat{\succ}^s, \hat{p}^s)_{s \in \mathcal{S}}, P \circ L)$.

We now formally define desirable properties from Section 3.2.1 in our discrete model. We remark that the definitions of respecting guarantees, strategy-proofness and anonymity do not reference school capacities and so carry over immediately. Similarly, the definitions for respecting priorities, non-wastefulness and constrained Pareto efficiency do not require non-atomicity and so our definition of η ensures that they also carry over. For completeness, we rewrite these properties without reference to η .

Definition 3.20. A two-round mechanism M respects priorities (subject to guarantees) if

(i) for every school $i \in \mathcal{C}$ and student $s \in \mathcal{S}$ who prefers i to her assigned school $i \succ^s \hat{\mu}(s)$, we have $|\hat{\mu}(i)| = q_i$, and (ii) for all students s' such that $\hat{\mu}(s') = i \neq \mu(s')$, we have $p_i^{s'} \geq p_i^s$.

Definition 3.21. A two-round mechanism is **non-wasteful** if no student is denied a seat at a school that has vacant seats; that is, for each student $s \in \mathcal{S}$ and schools i , if $i \succ^s \hat{\mu}(s)$, then $|\hat{\mu}(i)| = q_i$.

Let $\hat{\mu}$ be a second-round assignment. A *Pareto-improving cycle* is an ordered set of students $(s_1, s_2, \dots, s_m) \in \mathcal{S}^m$ and schools $(1', 2', \dots, m') \in \mathcal{C}^m$ such that $(i+1)' \succ^i i'$ (where \succ^i denotes the second-round preferences of student s_i , and we define $(m+1)' = 1'$), and $\hat{\mu}(s_i) = i'$ for all i .

Let \hat{p} be the second-round priorities obtained by giving each student s a top second-round priority $\hat{p}_i^s = n_i$ at their first-round assignment $\mu(s) = i$ (if $i \in \mathcal{C}$) and unchanged priority $\hat{p}_j^s = p_j^s$ at all other schools $j \neq i$. We say that a Pareto-improving cycle (in a second-round assignment) *respects (second-round) priorities* if $\hat{p}_{(i+1)'}^{s_i} \geq \hat{p}_{(i+1)'}^{s_{i+1}}$, for all i (where we define $s_{m+1} = s_1$).

Definition 3.22. A two-round mechanism is **constrained Pareto efficient** if the second-round assignment has no Pareto-improving cycles that respect second-round priorities.

In a setting with a finite number of students, PLDA mechanisms exactly satisfy all these properties except for strategy-proofness.

Proposition 3.3. *Suppose student preferences are consistent. Then PLDA mechanisms respect guarantees and priorities, and are non-wasteful, constrained Pareto efficient, and anonymous.*

Proof. The proofs of all these properties are almost identical to those in the continuum setting.

As an illustration, we prove that PLDA is constrained Pareto efficient in the discrete setting by using the fact that both rounds use single tie-breaking and the output is stable with respect to the second-round priorities \hat{p} .

Fix a Pareto-improving cycle C . Since s_i is assigned a seat at a school i when she prefers school $i+1 = \mu(s_{i+1})$, by the stability of DA she must either be in a strictly worse priority group than s_{i+1} at school $i+1$, or in the same priority group but have a worse lottery number. If C respects (second-round) priorities, then it must hold that for all i that students s_i and s_{i+1} are in the same priority group at school $i+1$ and s_i has a worse lottery number than s_{i+1} . But since this holds

for all i , single tie-breaking implies that we obtain a cycle of lottery numbers, which provides the necessary contradiction. \square

Proposition 3.3 states that in a setting with a finite number of students, PLDA mechanisms satisfy all our desired properties except for strategy-proofness. The following example illustrates that in a setting with a finite number of students, PLDA mechanisms may not satisfy two-round strategy-proofness. The intuition is that without non-atomicity, students are able to manipulate the first-round assignments of other students to change the guarantees, and hence change the second-round stability structure. In some cases in small markets, students are able to change the set of stable outcomes to benefit themselves.

Example 3.2 (PLDA with a finite number of students is not strategy-proof.). Consider a setting with $n = 2$ schools and $N = 4$ students, $\mathcal{S} = \{w, x, y, z\}$. Each school has capacity 1 and a single priority class. For readability, we let \emptyset denote the outside option. The students have the following preferences:

- $1 \succ_w \emptyset \succ_w 2$ and $\emptyset \hat{\succ}_w 1 \hat{\succ}_w 2$,
- $1 \succ_x 2 \succ_x \emptyset$, second-round preferences identical,
- $2 \succ_y 1 \succ_y \emptyset$, second-round preferences identical,
- $2 \succ_z \emptyset \succ_z 1$, second-round preferences identical.

We show that the RLDA mechanism is not strategy-proof. Consider the lottery that yields $L(w) > L(x) > L(y) > L(z)$. If the students report truthfully, the first-round assignment and second-round reassignment are

$$\begin{aligned} \mu(\mathcal{S}) &= (\mu(w), \mu(x), \mu(y), \mu(z)) = (1, 2, \emptyset, \emptyset), \text{ and} \\ \hat{\mu}(\mathcal{S}) &= (\hat{\mu}(w), \hat{\mu}(x), \hat{\mu}(y), \hat{\mu}(z)) = (\emptyset, 2, 1, \emptyset) \end{aligned}$$

respectively. However, consider what happens if student x says that only school 1 is acceptable to

her by reporting preferences $\succ^r = \hat{\succ}^r$ given by $1 \succ^r \emptyset \succ^r 2$ and $1 \hat{\succ}^r \emptyset \hat{\succ}^r 2$. Then

$$\mu(\mathcal{S}) = (1, \emptyset, 2, \emptyset), \quad \hat{\mu}(\mathcal{S}) = (\emptyset, 1, 2, \emptyset),$$

which is a strictly beneficial change for student x in the second round (and, in fact, weakly beneficial for all students).

Note that this reassignment was not stable in the second round when students reported truthfully, since, in that case, school 2 had second-round priorities $p_2^x > p_2^z > p_2^y > p_2^w$ and so school 2 and student z formed a blocking pair. In other words, for this particular realization of lottery numbers, student x is able to select a beneficial second-round assignment $\hat{\mu}$ that was previously unstable by changing student y 's first-round assignment so that student z cannot block $\hat{\mu}$.

In addition, the second-round outcome for student x under misreporting stochastically dominates her outcome from truthful reporting, when all other students report truthfully and the randomness is due the first-round lottery order. For if the lottery order is $L(w) > L(x) > L(y) > L(z)$ then student x can change her second-round assignment from 2 to 1 by reporting 2 as unacceptable, and this is the only lottery order for which student x receives a second-round assignment of 2 under truthful reporting. This is because if $L(s) > L(x)$ for $s \in \{y, z\}$ then student s is assigned to school 2 and stays there in both rounds, if $L(x) > L(w), L(y), L(z)$ then student x is assigned to school 1 and stays there in both rounds, and finally if $L(w) > L(x) > L(z) > L(y)$ then student x is assigned to school 1 in the second round. Moreover, for any lottery order where student x received 1 in the first or second round under truthful reporting, she also received school 1 in the same round by misreporting. This is because any stable matching in which student x is assigned to 1 remains stable after student x truncates. Indeed, student x is not part of any unstable pair, as she got her first choice, and any unstable pair not involving student x remains unstable under the true preferences, as only student x changes her preferences. Hence the second-round assignment student x receives by misreporting stochastically dominates the assignment she would have received under truthful reporting. This violates strategy-proofness.

This example shows that, as noted in Section 3.7, PLDA mechanisms are not two-round strategy-

proof in the finite setting. However, there are convergence results in the literature that suggest that PLDA mechanisms are almost two-round strategy-proof in large markets. Azevedo and Leshno (2016) have shown that if a sequence of (large) discrete economies converges to some limiting continuum economy with a unique stable matching (defined via cutoffs), then the stable matchings of the discrete economies converge to the stable matching of the continuum. Azevedo and Budish (2017) have shown that Deferred Acceptance is “strategy-proof in the large”. We conjecture that the proportion of students who are able to successfully manipulate PLDA mechanisms decreases polynomially in the size of the market. While a formal proof of such a result is beyond the scope of this paper, we provide a heuristic argument as follows. By definition, PLDA mechanisms satisfy the efficiency and anonymity requirements in finite markets as well. In the second round it is clearly a dominant strategy to be truthful, and, intuitively, for a student to benefit from a first-round manipulation, her report should affect the second-round cutoffs in a manner that gives her a second-round assignment she would not have received otherwise. If the market is large enough, the cutoffs will converge to their limiting values, and the probability that she could benefit from such a manipulation would be negligible.

Moreover, we believe that students will be unlikely to try to misreport under the PLDA mechanisms.²⁰ This is because, as Example 3.2 illustrates, successful manipulations require that students strategically change their first-round assignment and correctly anticipate that this changes the set of second-round stable assignments to their benefit. Such deviations are very difficult to plan and require sophisticated strategizing and detailed information about other students’ preferences.

The heuristic argument that PLDA mechanisms are strategy-proof in the large also suggests that our theoretical results should approximately hold for large discrete economies. A similar argument can be used to show that an approximate version of our characterization result (Theorem 3.3) holds for finite markets with no priorities, as PLDA mechanisms satisfy an approximate version of the averaging axiom in large finite markets. Our type-equivalence result (Theorem 3.1) and result showing that RLDA minimizes transfers (Theorem 3.2) should also be approximately valid in the large market limit. Specifically, consider a sequence of markets of increasing size. If the global order

²⁰as compared to the currently used DA mechanism.

condition holds in the continuum limit, this should lead to approximate type-equivalence under all PLDAs and to RLDA approximately minimizing transfers among PLDAs in the finite markets as market size grows. Moreover, if the order condition holds, then in large finite economies and for every permutation P , the set of students who violate a local order condition on $PLDA(P)$ will be small relative to the size of the market.

3.6 Empirical Analysis of PLDA Mechanisms

In this section, we use data from the New York City (NYC) high school choice system to simulate and evaluate the performance of centralized PLDA mechanisms under different permutations P . The simulations indicate that our theoretical results are real-world relevant. Different choices of P are found to yield similar allocative efficiency: the number of students assigned to their k -th choice for each rank k , as well as the number of students remaining unassigned, are similar for different permutations P . At the same time, the difference in the number of reassigned students is significant and is minimized under RLDA.

Motivated by current practice, we also simulate decentralized versions of FLDA and RLDA. In a version where students take time to vacate previously assigned seats, reversing the lottery increases allocative efficiency during the early stages of reassignment and decreases the number of reassignments at every stage. However, in a version where students take time to decide on offers from the waitlist, the efficiency comparisons are reversed.²¹ In both versions both FLDA and RLDA took tens of stages to converge. Our simulations suggest that decentralized waitlist mechanisms can achieve some of the efficiency gains of a centralized mechanism but incur significant congestion costs, and the effects of reversing the tie-breaking order before constructing waitlists will depend on the specific time and informational constraints of the market.

²¹This is due to a phenomenon that occurs when the second round is decentralized (not captured by our theoretical model), where under the reverse lottery the students with the worst lottery in the first round increase the waiting time for other students in the second round by considering multiple offers off the waitlist that they eventually decline.

3.6.1 Data

We use data from the high school admissions process in NYC for the academic years 2004–2005, 2005–2006, and 2006–2007, as follows.

First-round preferences. In our simulation, we take the first-round preferences \succ of every student to be the preferences they submitted in the main round of admissions. The algorithm used in practice is essentially strategy-proof (see Abdulkadiroğlu et al., 2005a), justifying our assumption that reported preferences are true preferences.²²

Second-round preferences. In our simulation, students either drop out of the system entirely in the second round or maintain the same preferences. Students are considered to drop out if the data does not record them as attending any public high school in NYC the following year (this was the case for about 9% of the students each year).²³

School capacities and priorities. Each school’s capacity is set to the number of students assigned to it in the first-round assignment in the data. This is a lower bound on the true capacity, but lets us compute the final assignment under PLDA with the true capacities, since the occupancy of each school with vacant seats decreases across rounds in our setting. School priorities over students are obtained directly from the data. (We obtain similar results in simulations with no school priorities.)

3.6.2 Simulations

We ran simulations using a centralized implementation of PLDA as well as two decentralized versions of PLDA.

Centralized PLDA. We first consider the following family of centralized PLDA mechanisms, parameterized by a single parameter α that smoothly interpolates between RLDA and FLDA. Each student s receives a uniform i.i.d. first-round lottery number $L(s)$ (a normal variable with

²²The algorithm is not completely strategy-proof, since students may rank no more than 12 schools. However, only a very small percentage of students rank 12 schools. Another issue is that there is some empirical evidence that students do not report their true preferences even in school choice systems with strategy-proof mechanisms; (see, e.g., Hassidim et al., 2015; Narita, 2016).

²³For a minority of the students (9.2%-10.45%), attendance in the following year could not be determined by our data, and hence we assume they drop out randomly at a rate equal to the dropout rate for the rest of the students (8.9%-9.2%).

mean 0 and variance 1), which generates a uniformly random lottery order.²⁴ The second-round ‘permuted lottery’ of s is given by $\alpha L(s) + \tilde{L}(s)$, where $\tilde{L}(s)$ is a new i.i.d. normal variable with mean 0 and variance 1, and α is identical for all the students. RLDA corresponds to $\alpha = -\infty$ and FLDA corresponds to $\alpha = \infty$. For a fixed real α , every realization of second-round scores corresponds to some permutation of first-round lottery numbers, with α roughly capturing the correlation of the second-round order with that of the first round. We quote averages across simulations.

Decentralized PLDA. In order to evaluate the performance of waitlist systems, we also ran simulations using two versions of decentralized PLDA with second rounds run in multiple “stages”:

Version 1. At stage ℓ , school i has residual capacity \tilde{q}_i^ℓ equal to the number of students previously assigned to school i who rejected school i in the previous stage (and \tilde{q}_i^1 is the number of students assigned to school i in the first round who dropped out of the system). Each school i proposes to the top \tilde{q}_i^ℓ students on their waitlist (including students who dropped out) and removes them from the waitlist, students who dropped out reject all offers, and all remaining students are (tentatively) assigned to their favorite school that offered them a seat in the first round or in the second round thus far and reject the rest. The stages of reassignment continue until there are no new proposals.

Version 2. At stage ℓ , school i has residual capacity \tilde{q}_i^ℓ equal to the number of students previously assigned to school i who rejected school i in the previous stage (and \tilde{q}_i^1 is the number of students assigned to school i in the first round who dropped out of the system). We run DA-STB on the residual economy where each school i has capacity \tilde{q}_i^ℓ and each student only finds schools strictly better than their current assignment acceptable.²⁵ This results in some students being reassigned and new residual capacities for stage $\ell + 1$, equal to the sum of the number of unfilled seats at the end of stage ℓ and the number of students who left the school due to an upgrade in stage ℓ . The stages of reassignment continue until there are no new proposals.

Version 1 of the decentralized PLDA mechanisms mirrors a decentralized process where students take time to make decisions. However, it does so in a rather naive fashion by assuming that students

²⁴School preferences are then generated by considering students in the lexicographical ordering first in terms of priority, then by lottery number. We may equivalently renormalize the set of realized lottery numbers to lie in the interval $[0, 1]$ before computing scores.

²⁵We provide results using school-proposing DA, as this more closely mirrors the structure of waitlist systems. Results using student-proposing DA were similar.

take the same amount of time to accept an offer, to reject an offer, or to inform a school that they were previously assigned to that they have been assigned to a different school. Version 2 captures a decentralized process where students also take time to both make and communicate decisions, but take much longer to tell schools that they were previously assigned to that they have been assigned to a different school. Accordingly the efficiency outcomes at a given stage of version 2 dominate those of version 1 at the same stage, as more information is communicated during each stage under version 2.

Version 2 simulations a setting where the main driver behind congestion is chains of student reassignment. Version 2 is more realistic in settings where schools are the primary drivers behind updated information, since a school is much more likely to ask for decisions from students who are undecided about an offer from the school rather than from students who have already accepted an offer from the school. In many school districts information about previously assigned students being reassigned to other schools is processed centrally, and it is also reasonable to assume that this would occur on a slower timescale than rejections of offers. In practice we expect that the dynamics of waitlist systems would lie somewhere on the spectrum between these two extreme versions of decentralized PLDA.

3.6.3 Results

The results of our centralized PLDA computational experiments based on 2004–2005 NYC high school admissions data appear in Table 3.1 and Figure 3.4. Results for 2005–2006 and 2006–2007 were similar. Figure 3.4 shows that the mean number of reassignments is minimized at $\alpha = -\infty$ (RLDA) and increases with α , which is consistent with our theoretical result in Theorem 3.2. The mean number of reassignments is as large as 7,800 under FLDA compared to just 3,400 under RLDA.

Allocative efficiency appears not to vary much across values of α : the number of students receiving at least their k -th choice for each $1 \leq k \leq 12$, as well as the number of unassigned students, vary by less than 1% of the total number of students. There is a slight trade-off between allocative efficiency due to reassignment and allocative efficiency from assigning previously unassigned students,




Figure 3.4: Number of reassigned students versus α . The number of reassigned students under the extreme values of α , namely, $\alpha = \infty$ (FLDA) and $\alpha = -\infty$ (RLDA), are shown via dotted lines.

α	Reassignments #	Unassigned %	$k = 1$ %	$k \leq 2$ %	$k \leq 3$ %
Round 1 (No Reassignment)	0	9.31	50.14	64.14	72.44
Round 2					
FLDA: ∞	7797	5.89	55.41	69.85	78.03
8.00	7606	5.90	55.40	69.85	78.02
6.00	7512	5.90	55.40	69.85	78.03
4.00	7325	5.89	55.38	69.84	78.02
2.00	6863	5.89	55.33	69.81	78.02
0.00	5220	5.87	54.96	69.65	77.97
-2.00	3686	5.81	54.52	69.37	77.82
-4.00	3480	5.79	54.47	69.33	77.78
-6.00	3433	5.79	54.46	69.32	77.77
-8.00	3416	5.79	54.45	69.31	77.77
RLDA: $-\infty$	3391	5.79	54.45	69.30	77.75

Table 3.1: Centralized PLDA simulation results: 2004–2005 NYC high school admissions.

We show the mean percentage of students remaining unassigned or getting at least their k th choice, averaged across 100 realizations for each value of α . All percentages are out of the total number of students remaining in the second round. The data contained 81,884 students, 74,366 students remaining in the second round, and 652 schools. The percentage of students who dropped out was 9.18%. The variation in the number of reassignments across realizations was only about 100 students.

with the percentage of unassigned students and percentage of students obtaining their top choice both decreasing in α by about 0.1% and 1% of students respectively.²⁶ We further find that for most students, the likelihoods of getting one of their top k choices under FLDA and under RLDA are very close to each other. (For instance, for 87% of students, these likelihoods differ by less than 3% for all k .) This is consistent with what we would expect based on our theoretical finding of type-equivalence (Theorem 3.1) of the final assignment under different PLDA mechanisms.

The results of our decentralized PLDA computation experiments appear in Table 3.2. When implementing PLDAs in a decentralized fashion, our measures of congestion can be more nuanced. We let a *reassignment* be a movement of a student from a school in \mathcal{C} to a different school in \mathcal{C} , possibly during an interim stage of the second round, and let a *temporary reassignment* be a movement of a student from a school in $\mathcal{C} \cup \{n + 1\}$ to a different school in \mathcal{C} that is not their final assignment. We will also be interested in the number of stages it takes to clear the market.

In the first version of decentralized PLDAs, FLDA reassigns more students than RLDA but far outperforms RLDA in terms of minimizing congestion and maximizing efficiency. FLDA takes on average 17 stages to converge, while RLDA requires 33. FLDA performs 780 temporary transfers while RLDA performs 2420, creating much more unnecessary congestion. FLDA takes 2 and 5 stages to achieve 50% and 90% respectively of the total increase in number of students assigned to their top school, whereas RLDA takes 3 and 9 stages respectively. FLDA also dominates RLDA in terms of the number of students assigned to one of their top k choices in the first ℓ stages, for all k and all ℓ , and the percentage of unassigned students in the first ℓ stages for almost all small ℓ .

In the second version of decentralized PLDAs, FLDA still reassigns more students and now achieves less allocative efficiency than RLDA during the initial stages of reassignment. RLDA has fewer unassigned students by stage ℓ than FLDA for all ℓ . RLDA also dominates FLDA in terms of the number of students assigned to one of their top k choices in the first 2 stages, and achieves most of its allocative efficiency by the second stage, improving the allocative efficiency by fewer than 100 students from that point onwards. In the limit FLDA is still slightly more efficient than RLDA,

²⁶Intuitively, prioritizing students with lower lotteries both decreases the number of unassigned students and decreases allocative efficiency by artificially increasing the constraints from providing first-round guarantees.

α	Reassignments # total (# temporary)	Unassigned %	$k = 1$ %	$k \leq 2$ %	$k \leq 3$ %
<i>Round 1 (No Reassignments)</i>	0	9.31	50.14	64.17	72.45
<i>Round 2 FLDA, Version 1</i>					
Stage 1	3461 (447)	7.89	52.68	66.62	74.47
Stage 2	2126 (206)	7.04	53.93	68.03	76.14
Stage 3	1258 (80)	6.55	54.60	68.83	76.96
Stage 4	727 (30)	6.27	54.97	69.28	77.42
Stage 5	425 (11)	6.11	55.18	69.53	77.68
Total (Stage ≈ 17)	8590 (780)	5.87	55.46	69.87	78.05
<i>Round 2 RLDA, Version 1</i>					
Stage 1	1004 (835)	7.85	51.38	65.70	74.09
Stage 2	1077 (577)	7.18	52.24	66.72	75.15
Stage 3	838 (369)	6.78	52.82	67.39	75.83
Stage 4	640 (234)	6.52	53.23	67.86	76.30
Stage 9	180 (24)	5.97	54.22	69.02	77.45
Total (Stage ≈ 33)	5818(2419)	5.79	54.51	69.37	77.80
<i>Round 2 FLDA, Version 2</i>					
Stage 1	4139 (457)	7.62	53.21	67.14	75.21
Stage 2	2333 (166)	6.69	54.50	68.66	76.75
Stage 3	1137 (42)	6.24	55.06	69.35	77.48
Stage 4	511 (9)	6.03	55.30	69.65	77.80
Total (Stage ≈ 12)	8503 (677)	5.89	55.47	69.87	78.04
<i>Round 2 RLDA, Version 2</i>					
Stage 1	2863 (199)	6.15	54.14	68.85	77.24
Stage 2	489 (17)	5.88	54.38	69.16	77.58
Stage 3	165 (2)	5.82	54.46	69.26	77.69
Stage 4	63 (0)	5.79	54.49	69.30	77.73
Total (Stage ≈ 9)	3624 (220)	5.79	54.51	69.33	77.76

Table 3.2: Decentralized PLDA simulation results: 2004–2005 NYC high school admissions.

We show the mean number of reassignments (number of movements of a student from a school in \mathcal{C} to a different school in \mathcal{C}) as well as the mean number of temporary reassignments (number of movements of a student from a school in $\mathcal{C} \cup \{n+1\}$ to a school in \mathcal{C} that is not their final assignment) in parentheses. We also show mean percentage of students remaining unassigned, or getting at least their k th choice. All figures are averaged across 100 realizations for each value of α , and all percentages are out of the total number of students remaining in the second round. The data contained 81,884 students, 74,366 students remaining in the second round, and 652 schools.

and so for large ℓ FLDA achieves higher allocative welfare than RLDA after ℓ stages. However FLDA also requires more stages to converge, taking on average 12 stages compared to 9 for RLDA.

Our empirical findings have mixed implications for implementing decentralized waitlists. Our clearest finding is the benefit of centralization in reducing congestion. In most school districts students are given up to a week to make decisions. If students take this long both to reject undesirable offers and to vacate previously assigned seats, our simulations on NYC data suggest that in the best case the market could take at least 4 months to clear. Even if students make quick decisions, if it takes them a week to vacate their previously assigned seats, our simulations suggest that the market would take at least 2 months to clear. In both cases the congestion costs are prohibitive. If, despite these congestion costs, a school district wishes to implement decentralized waitlists, our results suggest that the optimal permutation for the second-round lottery for constructing waitlists will depend on the informational constraints in the market.

3.6.4 Strategy-proofness of PLDA

One of the aspects of the DA mechanism that makes it successful in school choice in practice is that it is strategy-proof. While we have shown that PLDA mechanisms are two-round strategyproof in a continuum setting, it is natural to ask to what extent PLDA mechanisms are two-round strategy-proof in practice. We provide a numerical upper bound on the incentives to deviate from truthful reporting using computational experiments based on 2004–2005 NYC high school data, and find that on average a negligible proportion of students ($< 0.01\%$) could benefit from misreporting within their consideration set of programs. Specifically, 0.8% of sampled students could misreport in a potentially beneficial manner in at least one of 100 sampled lotteries, and no students could benefit in more than 3 of 100 sampled lotteries from misreporting. Moreover for 99.8% of lotteries the proportion of students who could successfully manipulate their report was at most 1%.

These upper bounds were computed as follows. Approximately 2700 students were sampled, and RLDA was run for each of these students using 100 different sampled lotteries. For a given student, let C be the set of schools that were a part of the student’s first round preferences in the data. We allowed the student to unilaterally misreport in the first round, reporting at most one

school from C in the first round instead of their true preferences. We then counted the number of such students who by doing so could either (1) change their first-round assignment (for the worse) but second-round assignment for the better, or (2) create a rejection cycle. This provides a provable upper bound on the number of students who can benefit from misreporting (and possibly reordering) a subset of C in the first round. We omit the formal details in the interest of space.

3.7 Proposals & Discussion

Summary of findings. We have proposed the PLDA mechanisms as a class of reassignment mechanisms with desirable incentive and efficiency properties. These mechanisms can be implemented with a centralized second round at the start of the school year, or with a decentralized second round via waitlists, and a suitable implementation can be chosen depending on the timing of information arrival and subsequent congestion in the market. Moreover, the key defining characteristic of the mechanisms in this class, the permutation used to correlate the tie-breaking lotteries between rounds, can be used to optimize various objectives. We propose implementing centralized RLDA at the start of the school year, as both in our theory and in simulations on data this allows us to maintain efficiency while eliminating the congestion caused by sequentially reassigning students, and minimizes the number of reassignments required to reach an efficient assignment.

RLDA is practical. We have shown that RLDA is an attractive choice when the objectives are to achieve allocative efficiency and minimize the number of reassigned students. In addition, RLDA has the nice property of being equitable in an intuitive manner, as students who receive a poor draw of the lottery in the first round are prioritized in the second round. This may make RLDA more palatable to students than other PLDA mechanisms. Indeed, Random Hall, an MIT dorm, uses a mechanism for assigning rooms that resembles the reverse lottery mechanism we have proposed. Freshmen rooms are assigned using serial dictatorship. At the end of the year (after seniors leave), students can claim the rooms vacated by the seniors using serial dictatorship where the initial lottery numbers (from their first match) are reversed.²⁷

²⁷The MIT Random Hall matching is more complicated, because sophomores and juniors can also claim the vacated rooms, but the lottery only gets reversed at the end of freshman year. Afterward, if a sophomore switches room, her

Optimizing other objectives. Our results suggest that PLDA mechanisms are an attractive class of mechanisms in more general settings, and the choice of mechanism within this class will vary with the policy goal. If, for instance, it were viewed as more equitable to allow more students to receive (possibly small) improvements to their first-round assignment, then the FLDA mechanism that simply runs DA again would optimize over this objective. Moreover, our type-equivalence result (Theorem 3.1) shows that when the relative overdemand for schools stays the same this choice can be made without sacrificing allocative efficiency.

Discussion of axiomatic characterization. We axiomatically justified the class of PLDA mechanisms in settings where schools do not have priorities (Theorem 3.3). In a model with priorities, we find that natural extensions of our axioms continue to describe PLDA mechanisms, but also include undesirable generalizations of PLDA mechanisms. Specifically, suppose that we add an axiom requiring that for each school i , the probability that a student who reports a top choice of i then receives it in the first or second round be independent of their priority at other schools. This new set of axioms describes a class of mechanisms that includes the PLDA mechanisms. However, there also exists an example market and a mechanism satisfying this new set of axioms such that the joint distribution over the two rounds of assignments does not match any PLDA. Characterizing the class of mechanisms satisfying these axioms in the richer setting with school priorities remains an open question. It may also be possible to characterize PLDA mechanisms in a setting with priorities using a different set of axioms. We leave both questions for future research.

Inconsistent preferences. Another natural question is how to deal with inconsistent student preferences. Narita (2016) observed that in the current reapplication process in the NYC public school system, although only about 7% of students reapplied, about 70% of these reapplicants reported second-round preferences that were inconsistent with their first-round reported preferences. Note that PLDAs allow students to report inconsistent preferences in the second round. We believe that some of our insights remain valid if a small fraction of students have an idiosyncratic change in preferences, or if a small number of new students enter in the second round. However, new effects may emerge if students have arbitrarily different preferences in the two rounds. In such settings,

priority drops to the last place of the queue.

strategy-proofness is no longer well defined, and it can be shown that the order condition is no longer sufficient to guarantee type-equivalence and optimality of RLDA. Moreover, in such settings the relative efficiency of the PLDA mechanisms will depend on the details of school supply and student demand.

More than two rounds. Finally, what insights do our results provide for situations in which assignment is done in three or more rounds? For instance, one could consider mechanisms under which the lottery is reversed (or permuted) after a certain number of rounds and thereafter remains fixed. At what stage should the lottery be reversed? Clearly, there are many other mechanisms that are reasonable for this problem, and we leave a more comprehensive study of this question for future work.

Chapter 4

The Cutoff Structure of Top Trading Cycles in School Choice

In this chapter, we provide a framework for optimizing over quantitative objectives when using the Top Trading Cycles (TTC) school choice mechanism. We first develop a characterization of TTC that explains the role of priorities in determining the TTC assignment and can be used to tractably analyze TTC. The TTC assignment can be concisely described by *admissions cutoffs*, with a cutoff p_j^i for each pair of schools i, j that describes the minimal priority a student needs to have at school j in order to use it to attend school i . These cutoffs parallel prices in competitive equilibrium, with students' priorities serving the role of endowments. We show that there is a labeling of schools $\{1, \dots, n\}$ such that for any i the cutoffs are ordered $p_i^1 \geq p_i^2 \geq \dots p_i^i = \dots = p_i^n$. Additionally, to help convey to students that TTC is strategy-proof, we derive cutoffs that are independent of the reported preferences of a given student.

To facilitate tractable analysis of TTC, we then formulate a continuum model of TTC and show how to directly calculate the TTC assignment from the distribution of preferences and priorities by solving a system of equations. We present closed form solutions for parameterized economies. We also show that the discrete TTC model is a particular case of the continuum framework, as for discrete problems the continuum TTC model calculates cutoffs that give the discrete TTC assignment. We establish that the TTC assignment changes smoothly with changes in the underlying

economy, implying that the continuum economy can also be used to approximate sufficiently similar economies.

The tractability of our framework relies on a novel approach to analyzing TTC. A key idea that allows us to define TTC in the continuum is that the TTC algorithm can be characterized by its aggregate behavior over many cycles. Any collection of cycles must maintain *trade balance*, that is, the number of students assigned to each school is equal to the number of students who claimed or traded a seat at that school. For smooth continuum economies we reformulate the trade balance equations into a system of equations that fully characterizes TTC. These equations provide a recipe for calculating the TTC assignment.

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

As an illustration of the framework, we apply it to study the effects of making a school more desirable. As a shorthand, we refer to such changes as an increase in the quality of the school.¹ To evaluate the effects of increasing the quality of a school it is necessary to account for changes in the assignment due to changes in student preferences. First, we derive comparative statics that show how the assignment and student welfare change with changes in a school's quality. We decompose the marginal change in student welfare into the direct increase in utility of students assigned to the more desirable school and the indirect effect that arises from changes in the assignment. A marginal increase in the quality of a popular school can have a negative indirect effect on welfare: as some students switch into the school and gain a marginal utility increase, other students are

¹Examples 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 and Jackson (2017), reduction in class size Krueger (1999); Chetty et al. (2011) and changes in an individual school's funding Dinerstein et al. (2014), but our theoretical model is not specific to any of these examples.

denied admission and can suffer substantial losses. We quantify these effects in a parametric setting, showing that increasing the quality of a popular school can decrease the welfare gains from sorting on idiosyncratic preferences.

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

As another application, we explore the design of priorities for TTC and find that it is "bossy" in the sense that a change in the priority of a student that does not alter her assignment can nonetheless alter the assignment of other students. This implies that it is not possible to determine the TTC cutoffs directly through a supply-demand equation as in Azevedo and Leshno (2016). We characterize the range of possible assignments generated by TTC after changes to the relative priority of high-priority students, and show that a small change to the priorities will only change the assignment of a few students.

A third application of our model provides comparisons between mechanisms in terms of assignments and welfare. We solve for welfare under TTC and DA in a parametric setting and quantify how much welfare is sacrificed due to stability. A comparison between TTC and DA across different school choice environments corroborates a conjecture by Pathak (2016) that the difference between the mechanisms becomes smaller with increased alignment between student preferences and school priorities. We also compare TTC to the Clinch and Trade mechanism of Morrill (2015b) in large economies and find that it is possible for TTC to produce fewer blocking pairs than the Clinch and Trade mechanism.

A few technical aspects of the analysis may be of interest. First, we note that the trade balance equations circumvent many of the measure theoretic complications in defining TTC in the

continuum. Second, a connection to Markov chain theory allows us to show that a solution to the marginal trade balance equations always exists, and to characterize the possible trades.

4.1 Prior Work on TTC

Since Abdulkadiroğlu and Sönmez (2003) introduced school choice as a mechanism design problem and suggested the DA and TTC mechanisms as desirable solutions, TTC has been considered for use in a number of school choice systems. However, while DA has been adopted by many school choice systems, TTC has essentially not been implemented, leading to a number of papers studying the relative merits of the two mechanisms for practical applications. Abdulkadiroğlu et al. (2005b) discuss how the city of Boston debated between using DA and TTC for their school choice systems and ultimately chose DA. Abdulkadiroğlu et al. (2009) compare the outcomes of DA and TTC for the NYC public school system, and show that TTC gives higher student welfare. Kesten (2006) studies the relationship between DA and TTC, and shows that they are equivalent if and only if the priority structure is acyclic.

There are a multitude of characterizations of TTC in the literature. Abdulkadiroğlu et al. (2017c) 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 (2012b) and Morrill (2013, 2015a). These characterizations explore the qualitative properties of TTC, but do not provide another method for calculating the TTC outcome or evaluating quantitative objectives. Ma (1994), Pápai (2000) and Pycia and 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.

Dur and Morrill (2017) show that the outcome of TTC can be expressed as the outcome of a competitive market where there is a price for each priority position at each school, and agents may buy and sell exactly one priority position. Our characterization also provides a connection between TTC and competitive markets, but requires a lower dimensional set of cutoffs and provides a method

for directly calculating these cutoffs. He et al. (Forthcoming) propose an alternative pseudo-market approach for discrete assignment problems that extends Hylland and Zeckhauser (1979) and also uses admission cutoffs. Miralles and Pycia (2014) show a second welfare theorem for discrete goods, namely that any Pareto efficient assignment of discrete goods without transfers can be decentralized through prices and endowments, but require an arbitrary endowment structure.

Several variants of TTC have been suggested in the literature. Morrill (2015b) introduces the Clinch and Trade mechanism, which differs from TTC in that it identifies students who are guaranteed admission to their first choice and assigns them immediately without implementing a trade. Hakimov and Kesten (Forthcoming) introduce Equitable TTC, a variation on TTC that aims to reduce inequity. In Chapter 4.5.2 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 and Molis (2011), Jaramillo and Manjunath (2012) and Saban and Sethuraman (2013) give specific variants of TTC that are strategy-proof and efficient. The continuum model allows us to characterize the possible outcomes from different tie-breaking rules.

Several papers also study TTC in large markets. Hatfield et al. (2016) study the incentives for schools to improve their quality under TTC and find that even in a large market a school may be assigned less preferred students when it improves its quality. Our results in Chapter 4.5.1 quantify these effects. Che and Tercieux (2015, 2018) study the properties of TTC in a large market where the heterogeneity of items grows as the market gets large, whereas our setting considers a large population of agents and a fixed number of item types. The results in Chapter 4.5 show that TTC has different properties in these different large markets.

4.1.1 Practical Implications

When school districts redesigned their school choice mechanisms to improve student welfare, most chose to implement the Deferred Acceptance (DA) mechanism, and essentially none selected the TTC mechanism. This is despite the fact that while TTC is Pareto efficient for students, DA is inefficient, in that it may produce assignments that are Pareto dominated for students. For

example, students are commonly given priority for their neighborhood schools, and DA may assign two students to their respective neighborhood schools even if both students would prefer to swap their assignments.²

One of the possible reasons for the lack of popularity of TTC is the way it has typically been described. While the sequential clearing of trade cycles is simple to state, it obscures the desirable properties of the mechanism and results in an opaque mapping between a student's priorities and their assignment. For example, Boston Public Schools considered both TTC and DA when redesigning its school choice in 2005, and decided in favor of using DA, stating BPS (2005):

The behind the scenes mechanized trading makes the student assignment process less transparent[...] and could lead families to believe they can strategize by listing a school they don't want in hopes of a trade.

Similarly, in Pathak (2016) Pathak writes: *"I believe that the difficulty of explaining TTC, together with the precedent set by New York and Boston's choice of DA, are more likely explanations for why TTC is not used in more districts."* In other words, the combinatorial description of TTC in terms of trading cycles caused users to doubt the strategy-proofness of the mechanism, and eroded trust in the system by making it difficult for parents verify that their children were correctly assigned.

Another major drawback of the algorithmic description of TTC is that it makes it difficult to discern how a student's priorities determine their assignment under TTC. This is exacerbated by the fact that priority at a school has different implications under DA and TTC; under TTC (in contrast to DA) it is possible for a student to gain admission to one school by having priority at another school. This means that school boards could potentially redesign their priority structures to obtain their goals under TTC, but in general the appropriate priority structures under DA and TTC will be different, and current theory provides almost no guidance as to how to design such

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

priority structures under TTC.³

Our cutoff characterization of TTC provides a way to communicate the TTC outcome that is easily verifiable. If students privately know their priorities, publicly publishing the cutoffs $\{p_i^j\}$ allows each student to determine their assignment.⁴ Additionally, we provide TTC cutoffs that are independent of the reported preferences of a given student, which demonstrates that TTC is strategy-proof. Our cutoff characterization also elucidates the role of priorities under TTC. Students can use priority at school i to gain admission to school j if their priority at school i is above the cutoff p_i^j . Each student is assigned to her most preferred school for which she gained admission. As a result, we are hopeful that our framework for understanding the TTC outcome and designing appropriate input (such as school priorities) can increase the adoption of the Pareto efficient TTC in practice.

Finally, cutoff representations have been instrumental for empirical work on DA and variants of DA. Abdulkadiroğlu et al. (2017b) use admission cutoffs to construct propensity score estimates. Agarwal and Somaini (2018); Kapor et al. (2016) structurally estimate preferences from rank lists submitted to non-strategy-proof variants of DA. Both build on the cutoff representation of Azevedo and Leshno (2016). We hope that our cutoff representation of TTC will be similarly useful for empirical work on TTC.

4.2 TTC in School Choice

4.2.1 The Discrete TTC Model

Let \mathcal{S} be a finite set of students, and let $\mathcal{C} = \{1, \dots, n\}$ be a finite set of schools. Each school $i \in \mathcal{C}$ has a finite capacity $q_i > 0$. Each student $s \in \mathcal{S}$ has a strict preference ordering \succ^s over schools.

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

⁴This cutoff representation allows us to give the following non-combinatorial description of TTC. For each school i , each student receives i -tokens according to their priority at school i , where students with higher i -priority receive more i -tokens. The TTC algorithm publishes cutoffs $\{p_i^j\}$. Students can purchase a single school using a single kind of token, and the required number of i -tokens to purchase school j is p_i^j . Theorem 4.1 shows the cutoffs can be observed after the run of TTC. We thank Chiara Margaria, Laura Doval and Larry Samuelson for suggesting this explanation.

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

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

Given an economy E , the discrete Top Trading Cycles algorithm (TTC) calculates an assignment $\mu_{dTTC}(\cdot \mid E) : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$. We omit the dependence on E when it is clear from context. The algorithm runs in discrete steps, as described in Algorithm 1.

Mechanism 1 Top Trading Cycles (TTC)

```

1: procedure TTC( $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \succ^{\mathcal{S}}, \succ^{\mathcal{C}}, q)$ )
2:    $S \leftarrow \mathcal{S}$  ▷ Unassigned students
3:    $C \leftarrow \mathcal{C}$  ▷ Available schools
4:    $\tilde{q} \leftarrow q$  ▷ Residual capacity
5:   while  $|S| > 0, |C| > 0$  do ▷ while there are unassigned students and available schools
6:     for  $i \in C$  do
7:        $i$  points to and offers a seat to highest priority student in  $S$ 
8:     for  $s \in S$  do
9:        $s$  points to most preferred school in  $C$ 
10:    Select at least one trading cycle, i.e. 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 available school. Assign all students in the cycle to the school they point to.6
11:    Remove the assigned students from  $S$ , reduce the capacity of the schools they were assigned to by 1, and remove schools with no remaining capacity from  $C$ .
12:  return  $\mu$ 

```

TTC satisfies a number of desirable properties. An assignment μ is *Pareto efficient for students*

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

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

4.2.2 Cutoff Characterization

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

Theorem 4.1. *Let E be an economy. The TTC assignment is given by*

$$\mu_{dTTC}(s \mid E) = \max_{\succeq^s} \left\{ i \mid r_j^s \geq p_j^i \text{ for some } j \right\},$$

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

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

Now suppose for the sake of contradiction that $i \in B(s, \mathbf{p})$ and student s strictly prefers i to $\mu_{dTTC}(s)$, i.e. $i \succ^s \mu_{dTTC}(s)$. As $i \in B(s, \mathbf{p})$ there exists a school j' such that $r_{j'}^s \geq p_{j'}^i$. Let s' be the student with rank $r_{j'}^{s'} = p_{j'}^i$ at school j' . (Such a student exists by the definition of $p_{j'}^i$.) Then by definition student s' traded a seat at school j' , so since $r_{j'}^s \geq p_{j'}^i = r_{j'}^{s'}$ student s is assigned weakly

before student s' . Additionally, since $i \succ^s \mu_{dTTC}(s)$ school i must reach capacity before student s is assigned, and so since student s' was assigned to school i student s' was assigned strictly before student s . This provides the required contradiction. \square

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 $\mathbf{p} = \{p_j^i\}_{i,j}$ combine with student s 's priorities r^s to give s a *budget set* $B(s, \mathbf{p}) = \{i \mid r_j^s \geq p_j^i \text{ for some } j\}$ of schools she can attend. TTC assigns each student to her favorite school in her budget set.

The cutoffs p_j^i in Theorem 4.1 can be easily identified after the mechanism has been run. Hence Theorem 4.1 provides an intuitive way for students to verify that they were correctly assigned by the TTC algorithm. Instead of only communicating the assignment of each student, the mechanism can make the cutoffs publicly known. Students can calculate their budget set from their privately known priorities and the publicly given cutoffs, allowing them to verify that they were indeed assigned to their most preferred school in their budget set. In particular, if a student does not receive a seat at a desired school i , it is because she does not have sufficiently high priority at *any* school, and so i is not in her budget set. We illustrate these ideas in Example 4.1.

Example 4.1. 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 4.1a illustrates the preferences and priorities of each of the students. Each colored number represents a student. The location of the student in the square indicates their priority, with the horizontal axis indicating priority at school 1 and the vertical axis indicating priority at school 2. The number indicates the student's preferred school, and all students find both schools acceptable. The color indicates the student's assignment under TTC.

The cutoffs \mathbf{p} and resulting budget sets $B(s, \mathbf{p})$ for each student are illustrated in Figure 4.1b. 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

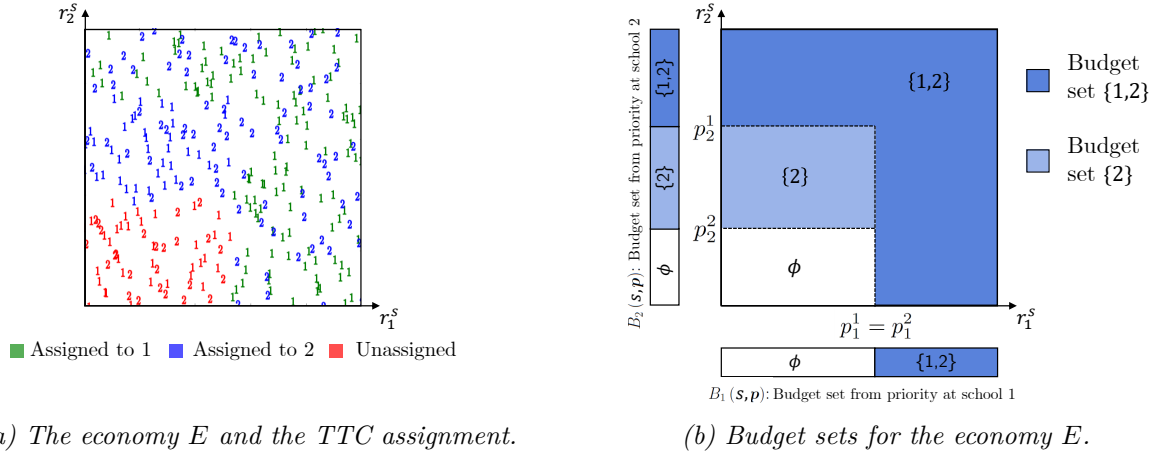


Figure 4.1: The economy and TTC budget sets for Example 4.1.

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 4.1b as for given \mathbf{p} each student's budget set does not depend on her preferences. The assignment of each student is her favorite school in her budget set.

Figure 4.1 shows the role of priorities in determining the TTC assignment in Example 4.1. Students with higher priority have a larger budget set of schools from which they can choose. A student can choose her desired school if her priority for some school is sufficiently high. Priority for each school is considered separately, and priority from multiple schools cannot be combined. For example, a student who has top priority for one school and bottom priority at the other school is assigned to her top choice, but a student who has the median priority at both schools will not be assigned to school 1.

Remark. This example also shows that the TTC assignment cannot be expressed in terms of one cutoff for each school, as the assignment in Example 4.1 cannot be described by fewer than 3 cutoffs.

4.2.3 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 $B_i(s, \mathbf{p}) = \{j \mid r_i^s \geq p_i^j\}$ denote the set of schools that enter student s 's budget set because of her priority at school i . Note that $B_i(s, \mathbf{p})$ depends only on the n cutoffs

$\mathbf{p}_i = \{p_i^j\}_{j \in \mathcal{C}}$. A student's budget set is the union $B(s, \mathbf{p}) = \cup_i B_i(s, \mathbf{p})$. Figure 4.1(b) depicts $B_1(s, \mathbf{p})$ and $B_2(s, \mathbf{p})$ for the economy of Example 4.1 along the x and y axes respectively.

The following proposition shows that budget sets $B_i(s, \mathbf{p})$ can be given by cutoffs \mathbf{p}_i that share the same ordering over schools for every i . We let $\mathcal{C}^{(i)} = \{i, i+1, \dots, n\}$ denote the set of schools that have a higher index than i .

Proposition 4.1. *There exists a relabeling of school indices such that there exist cutoffs $\mathbf{p} = \{p_i^j\}$ that describe the TTC assignment*

$$\mu_{TTC}(s) = \max_{\underset{s}{j}} \{j \mid r_i^s \geq p_i^j \text{ for some } i\},$$

and for any school i the cutoffs are ordered,⁷

$$p_i^1 \geq p_i^2 \geq \dots \geq p_i^i = p_i^{i+1} = \dots = p_i^n. \quad (4.1)$$

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

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

for some $j \leq i$. Moreover, each student's budget set $B(s, \mathbf{p}) = \cup_i B_i(s, \mathbf{p})$ is either $B(s, \mathbf{p}) = \phi$ or $B(s, \mathbf{p}) = \mathcal{C}^{(j)}$ for some j .

Proof. Let the schools be indexed such that they reach capacity in the order $1, 2, \dots, |\mathcal{C}|$. If a student s was assigned (strictly) after school $\ell - 1$ reached capacity and (weakly) before school ℓ reached capacity, we say that the student s was assigned in round ℓ . Given TTC cutoffs p_i^j from Theorem 4.1, we define new cutoffs $\{\tilde{p}_i^j\}$ by setting $\tilde{p}_i^j = \min_{k \leq j} p_i^k$. It evidently holds that $\tilde{p}_i^1 \geq \tilde{p}_i^2 \geq \dots \geq \tilde{p}_i^i = \tilde{p}_i^{i+1} = \dots = \tilde{p}_i^n$ for all i . We show that the cutoffs $\{\tilde{p}_i^j\}$ give the same allocation as the cutoffs $\{p_i^j\}$, i.e. for each student s it holds that

$$\max_{\underset{s}{j}} \{j \mid r_i^s \geq \tilde{p}_i^j \text{ for some } i\} = \mu_{TTC}(s) = \max_{\underset{s}{j}} \{j \mid r_i^s \geq p_i^j \text{ for some } i\}.$$

For each student s let $B(s, \tilde{\mathbf{p}}) = \{j \mid r_i^s \geq \tilde{p}_i^j \text{ for some } i\}$. It suffices to show that for each

⁷The cutoffs \mathbf{p} defined in Theorem 4.1 do not necessarily satisfy this condition. However, the run of TTC produces the following relabeling of schools and cutoffs $\tilde{\mathbf{p}}$ that give the same assignment and satisfy the condition: the schools are relabeled in the order in which they reach capacity under TTC, and the cutoffs $\tilde{\mathbf{p}}$ are given by $\tilde{p}_j^i = \min_{k \leq i} p_j^k$.

student s it holds that $\mu_{dTTC}(s) \in B(s, \tilde{\mathbf{p}})$, and that if $j \in B(s, \tilde{\mathbf{p}})$ then s prefers $\mu_{dTTC}(s)$ to j , i.e. $\mu(s) \succeq^s j$. The former is simple to show, since clearly $\tilde{\mathbf{p}} \leq \mathbf{p}$ and so $B(s, \tilde{\mathbf{p}}) \supseteq B(s, \mathbf{p}) \ni \mu_{dTTC}(s)$ (by Theorem 4.1).

The rest of the proof can be completed in much the same way as the proof of Theorem 4.1. Suppose for the sake of contradiction that $j \in B(s, \tilde{\mathbf{p}})$ and student s strictly prefers j to $\mu_{dTTC}(s)$, i.e. $j \succ^s \mu_{dTTC}(s)$. As $j \in B(s, \tilde{\mathbf{p}})$ there exists a school i' such that $r_{i'}^s \geq \tilde{p}_{i'}^j$. Let s' be the student with rank $r_{i'}^{s'} = \tilde{p}_{i'}^j$ at school i' . (Such a student exists by the definition of the cutoffs $p_{i'}^k$, $k \leq j$.) Then by definition student s' traded a seat at school i' , so since $r_{i'}^s \geq \tilde{p}_{i'}^j = r_{i'}^{s'}$ student s is assigned weakly before student s' . Additionally, since $j \succ^s \mu_{dTTC}(s)$ school j must reach capacity before student s is assigned. Finally, by definition there exists some $k \leq j$ such that $\tilde{p}_{i'}^j = p_{i'}^k$, and so since $r_{i'}^{s'} = p_{i'}^k$ it follows that student s' was assigned to school k . Thus student s' was assigned weakly before school j reached capacity, and so strictly before student s . This provides the required contradiction. The statements about the structure of the set of schools $B_i(s, \mathbf{p})$ student s can afford via her priority at school i and the structure of the budget set $B(s, \mathbf{p}) = \cup_i B_i(s, \mathbf{p})$ follow immediately from the ordered cutoffs. \square

When there exist TTC cutoffs that satisfy inequality (4.1) we say that the schools are *labeled in order*. The cutoff ordering proved in Proposition 4.1 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 4.1 allows us to give a simple illustration for the TTC assignment when there are $n \geq 3$ schools. For each school i , we can illustrate the set of schools $B_i(s, \mathbf{p})$ that enter a student's budget set because of her priority at school i as in Figure 4.2 (under the assumption that schools are labeled in order). This generalizes the illustration along each axis in Figure 4.1(b), and can be used for any number of schools. It is possible that $p_i^j = 1$, meaning that students cannot use their priority at school i to trade into school c .

Dur and Morrill (2017) provide a characterization of TTC as a competitive equilibrium where




Figure 4.2: The schools $B_i(s, \mathbf{p})$ that enter a student's budget set because of her priority at school i . The cutoffs p_i^j are weakly decreasing in j , and are equal for all $j \geq i$ (i.e. $p_i^i = p_i^{i+1} = \dots = p_i^n$). That is, a student's priority at i can add one of the sets $C^{(1)}, C^{(2)}, \dots, C^{(i)}, \emptyset$ to her budget set. If any school enters a student's budget because of her priority at i , then school i must also enter her budget set because of her priority at i .

a priority value function $v(r, i)$ specifies the price of priority r at school i and students are allowed to buy and sell one priority. Given TTC cutoffs $\{p_i^j\}$ where schools are labeled in order, the TTC assignment and priority value function $v(r, i) = n - \min \{j \mid r \geq p_i^j\}$ constitute a competitive equilibrium. We introduce a framework in Chapter 4.3 that allows a direct calculation of this competitive equilibrium as a solution to a set of equations.

4.2.4 Limitations

Although the cutoff structure is helpful in understanding the structure of the TTC assignment, there are several limitations to the cutoffs computed in Theorem 4.1 and Proposition 4.1. First, while the cutoffs can be determined by running the TTC algorithm, Theorem 4.1 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 4.1 does not correspond to the set of possible school assignments that student s can achieve by unilaterally changing her reported preferences.^{8,9} We therefore introduce the continuum model for TTC which allows us

⁸More 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 $\mu_{dTTC}(s \mid E)$ and $\mu_{dTTC}(s \mid E')$ be the assignment of s under the two economies, and let \mathbf{p} be the cutoffs derived by Theorem 4.1 for economy E . Theorem 4.1 shows that $\mu_{dTTC}(s \mid E) = \max_{\succ^s} B(s, \mathbf{p})$ but it may be $\mu_{dTTC}(s \mid E') \neq \max_{\succ'} B(s, \mathbf{p})$.

⁹For example, let E be an economy with three schools $C = \{1, 2, 3\}$, 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 4.1 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

to directly calculate the cutoffs, allowing for comparative statics. Using the continuum model, we present in Chapter 4.3.4 cutoffs that yield refined budget sets which provide for each student the set of schools that she could be assigned to by unilaterally changing her preferences. Thus the appropriate cutoff structure also makes it clear that TTC is strategy-proof.

4.3 Continuum Model and Main Results

4.3.1 Model

We consider the school choice problem with a continuum of students and finitely many schools, as in Azevedo and Leshno (2016). There is a finite set of schools denoted by $\mathcal{C} = \{1, \dots, n\}$, and each school $i \in \mathcal{C}$ has the capacity to admit a mass $q_i > 0$ of students. A student s has a type $\theta \in \Theta$ given by $\theta = (\succ^\theta, r^\theta)$; overloading notation we will sometimes refer to a student s by their type θ .¹⁰ 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_i^θ is the *rank* of student θ at school i , or the i -rank of student θ . Schools prefer students with higher ranks, that is $\theta \succ_i \theta'$ if and only if $r_i^\theta > r_i^{\theta'}$.

Definition 4.1. A *continuum economy* is given by $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ where $q = \{q_i\}_{i \in \mathcal{C}}$ 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, $\sum_{i \in \mathcal{C}} q_i < \eta(\Theta)$. Finally, we make the following technical assumption that ensures that the run of TTC in the continuum economy is sufficiently smooth and allows us to avoid some measurability issues.

Assumption 4.1. *The measure η admits a density ν . That is for any measurable subset of students*

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 .

¹⁰In this continuum model all students of the same type are indistinguishable and we may assume that they are assigned to the same school under TTC.

$A \subseteq \Theta$

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

Furthermore, ν is piecewise Lipschitz continuous everywhere except on a finite grid,¹¹ bounded from above, and bounded from below away from zero on its support.¹²

Assumption 4.1 is general enough to allow embeddings of discrete economies, and is satisfied by all the economies considered throughout the paper. However, it is not without loss of generality, e.g. it is violated when all schools share the same priorities over students.¹³

An immediate consequence of Assumption 4.1 is that a school's indifference curves are of η -measure 0. That is, for any $i \in \mathcal{C}$, $x \in [0, 1]$ we have that $\eta(\{\theta \mid r_i^\theta = x\}) = 0$. This is analogous to schools having strict preferences in the standard discrete model. As r_i^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 $i \in \mathcal{C}$, $x \in [0, 1]$ we have that $\eta(\{\theta \mid r_i^\theta \leq x\}) = x$.

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 $H_i^{j|C}(\mathbf{x})$ be the marginal density of students who are top ranked at school i among all students whose rank at every school k is no better than x_k , and whose top choice among the set of schools C is j .¹⁴ We omit the dependence on C when the relevant set of schools is clear from context, and write $H_i^j(\mathbf{x})$. The marginal densities $H_i^{j|C}(\mathbf{x})$ uniquely determine the distribution η .

¹¹A grid $G \subset \Theta$ is given by $G = \{\theta \mid \exists i \text{ s.t. } r_i^\theta \in D\}$, where $D = \{d_1, \dots, d_L\} \subset [0, 1]$ is a finite set of grid points. Equivalently, ν is Lipschitz continuous on the union of open hypercubes $\Theta \setminus G$.

¹²That is, there exists $M > m > 0$ such that for every $\theta \in \Theta$ either $\nu(\theta) = 0$ or $m \leq \nu(\theta) \leq M$.

¹³We can incorporate an economy where two schools have perfectly aligned priorities by considering them as a combined single school in the trade balance equations, as defined in Definition 4.2. The capacity constraints still consider the capacity of each school separately.

¹⁴Formally

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

where \mathbf{e}^i is the unit vector in the direction of coordinate i . In other words, $H_i^{j|C}(\mathbf{x})$ is the density of students θ with priority $r_i^\theta = x_i$ and $r_k^\theta \leq x_k$ for all $k \in C$ whose most preferred school in C is j .

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 $\mu(i) = \{\theta \mid \mu(\theta) = i\}$ denote the set of students assigned to school i . An assignment μ is *feasible* if it respects capacities, i.e. for each school $i \in \mathcal{C}$ we have $\eta(\mu(i)) \leq q_i$. 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 4.1. 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.

4.3.2 Main Results

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

We remark that directly translating the TTC algorithm to the continuum setting by considering individual trading cycles is challenging, as a direct adaptation of the algorithm would require the clearing of cycles of zero measure. We circumvent the technical issues raised by such an approach by formally defining the continuum TTC assignment in terms of trade balance and capacity equations, which characterize the TTC algorithm in terms of its aggregate behavior over multiple steps. To verify the validity of our definition, we show in Subsection 4.3.3 that continuum TTC can be used to calculate the discrete TTC outcome. We provide further intuition in Section 4.4.

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, $\gamma_i(t)$ is weakly decreasing for all i , 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 $\tilde{\gamma}_i(t)$ is defined only for $t \geq t_0 > 0$ and $i \in \tilde{\mathcal{C}} \subset \mathcal{C}$. For a set $\{t^{(i)}\}_{i \in \mathcal{C}} \in \mathbb{R}_{\geq}^{\mathcal{C}}$ of times we let $t^{(i^*)} \stackrel{def}{=} \min_i [t^{(i)}]$ denote the minimal time.

For a point $\mathbf{x} \in [0, 1]^{\mathcal{C}}$, let

$$D^i(\mathbf{x}) \stackrel{def}{=} \eta \left(\left\{ \theta \mid r^\theta \not\leq \mathbf{x}, Ch^\theta(\mathcal{C}) = i \right\} \right)$$

denote the mass of students whose rank at some school j is better than x_j and their first choice is school i . We will refer to $D^i(\mathbf{x})$ as the demand for i . Recall that $H_j^i(\mathbf{x})$ is the marginal density of students who want i who are top ranked at school j among all students with rank no better than x . Note that $D^i(\mathbf{x})$ and $H_j^i(\mathbf{x})$ 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 $\gamma_i(t)$ is the highest i -priority of any student who remains unassigned by time t . The stopping times $\{t^{(i)}\}_{i \in \mathcal{C}}$ indicate when each school fills its capacity. To verify whether γ and $\{t^{(i)}\}_{i \in \mathcal{C}}$ can correspond to a run of TTC we introduce trade balance conditions and capacity constraints as defined below.

Definition 4.2. Let $E = (\mathcal{C}, \Theta, \eta, q)$ be an economy. We say that the (residual) TTC path $\gamma(t)$ and positive stopping times $\{t^{(i)}\}_{i \in \mathcal{C}} \in \mathbb{R}_{\geq}^{\mathcal{C}}$ satisfy the trade balance and capacity equations for the economy E if the following hold.

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

$$\sum_{k \in \mathcal{C}} \gamma'_k(t) H_k^i(\gamma(t)) = \sum_{k \in \mathcal{C}} \gamma'_i(t) H_i^k(\gamma(t)) \quad (4.2)$$

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

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

$$\begin{aligned} D^{i^*}(\gamma(t^{(i^*)})) &= q_{i^*} \\ D^k(\gamma(t^{(i^*)})) &\leq q_k \quad \forall k \in \mathcal{C} \end{aligned} \quad (4.3)$$

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

3. If $\mathcal{C} \setminus \{i^*\} \neq \emptyset$, define the residual economy $\tilde{E} = (\tilde{\mathcal{C}}, \Theta, \tilde{\eta}, \tilde{q})$ by $\tilde{\mathcal{C}} = \mathcal{C} \setminus \{i^*\}$, $\tilde{q}_i = q_i - D^i(\gamma(t^{i^*}))$ and $\tilde{\eta}(A) = \eta(A \cap \{\theta : r^\theta \leq \gamma(t^{i^*})\})$. Define the residual TTC path $\tilde{\gamma}(\cdot)$ by restricting $\gamma(\cdot) : [t^{i^*}, \infty) \rightarrow [0, 1]^{\tilde{\mathcal{C}}}$ to $t \geq t^{i^*}$ and coordinates within $\tilde{\mathcal{C}}$. Then $\tilde{\gamma}$ and the stopping times $\{t^{(i)}\}_{i \in \tilde{\mathcal{C}}}$ 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. The path $\gamma(t)$ can be thought of as tracking the students who are being offered seats by each school at time t , and $\gamma'_i(t)$ gives the rate at which school i is moving down its priority list at time t . Hence $\gamma'_k(t) H_k^i(\gamma(t))$ gives the rate at which students are assigned to school i at time t due to their priority at school k ,¹⁵ and equation (4.2) states that over every small time increment the mass of students assigned to school i must be equal to the mass of offers made by school i . While all schools have remaining capacity, every assigned student is assigned to his first choice, and thus $D^i(\gamma(t))$ gives the mass of students assigned to school i at time $t \leq t^{(i^*)}$ in the algorithm. The time $t^{(i^*)}$ when school i^* fills its capacity can be calculated as a solution to Equation (4.3). 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 Section 4.4.

Our main result is that the trade balance and capacity equations fully characterize and provide a way to directly calculate the TTC assignment from the problem primitives. We show in Chapter 4.3.3 that this characterization is consistent with the discrete TTC.

Theorem 4.2. *Let $E = (\mathcal{C}, \Theta, \eta, q)$ be an economy. There exist a TTC path $\gamma(\cdot)$ and stopping times $\{t^{(i)}\}_{i \in \mathcal{C}}$ that satisfy the trade balance and capacity equations. Any $\gamma(\cdot), \{t^{(i)}\}_{i \in \mathcal{C}}$ that satisfy the trade balance and capacity equations yield the same assignment $\mu_{\mathcal{C}TTC}$, given by*

$$\mu_{\mathcal{C}TTC}(\theta) = \max_{\succ_{\theta}} \{j : r_i^\theta \geq p_i^j \text{ for some } i\},$$

¹⁵Recall that $H_k^i(\gamma(t))$ gives the marginal density of students who are top ranked at school k when students with priority higher than $\gamma(t)$ have already been assigned.

where the n^2 TTC cutoffs $\{p_i^j\}$ are given by

$$p_i^j = \gamma_i(t^{(j)}) \quad \forall i, j.$$

In other words, Theorem 4.2 provides the following a recipe for calculating the TTC assignment. First, find $\hat{\gamma}(\cdot)$ that solves the marginal trade balance equations (4.2) for all t . Second, calculate $t^{(i^*)}$ from the capacity equations (4.3) for $\hat{\gamma}(\cdot)$. Set $\gamma(t) = \hat{\gamma}(t)$ for $t \leq t^{(i^*)}$. 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 4.2. The TTC path used in this recipe may not be the unique TTC path, but all TTC paths yield the same TTC assignment.

Theorem 4.2 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 4.1), 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 $t^{(i)}$ satisfying the capacity equations (4.3) follows from our assumptions that there are more students than seats and all students find all schools acceptable.

The following immediate corollary of Theorem 4.2 shows that in contrast with the cutoffs given by the discrete model, the cutoffs given by Theorem 4.2 always satisfy the cutoff ordering.

Corollary 4.1. *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,*

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

To illustrate how Theorem 4.2 can be used to calculate the TTC assignment and understand how it depends on the parameters of the economy, we consider the following simple economy. This parameterized economy yields a tractable closed form solution for the TTC assignment. For other economies the equations may not necessarily yield tractable expressions, but the same calculations can be used to numerically solve for cutoffs for any economy satisfying our smoothness

¹⁶Continuity of the TTC path provides an initial condition for $\tilde{\gamma}$, namely that $\tilde{\gamma}_i(t^{(i^*)}) = \gamma_i(t^{(i^*)})$ for all i .




Figure 4.3: The TTC path, cutoffs, and budget sets for a particular instance of the economy \mathcal{E} in Example 4.2. 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 .

requirements.

Example 4.2. We demonstrate how to use Theorem 4.2 to calculate the TTC assignment for a simple parameterized continuum economy. The economy \mathcal{E} has two schools 1, 2 with capacities $q_1 = q_2 = q$ with $q < 1/2$. A fraction $p > 1/2$ of students prefer school 1, and student priorities are uniformly distributed on $[0, 1]$ independently for each school and independently of preferences. This economy is described by

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

where $H_b^c(x)$ 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 4.3. This economy can be viewed as a smoothed continuum version of the economy in Example 4.1.

We start by solving for γ from the trade balance equations (4.2), which simplify to the differential

equation¹⁷

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

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

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

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

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

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

For the residual economy the marginal trade balance equations (4.2) are trivial, and we define the residual TTC path by

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

for $t^{(1)} \leq t \leq t^{(2)}$. Solving the capacity equation (4.3) for $t^{(2)}$ yields that

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

For instance, if we plug in $q = 4/10$ and $p = 2/3$ to match the economy in Example 4.1, the calculation yields the cutoffs $p_1^1 = p_2^1 \approx .54$, $p_1^2 \approx .73$ and $p_2^2 \approx .37$, which are approximately the same cutoffs as those for the discrete economy in Example 4.1.

¹⁷The original trade balance equations are

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

Example 4.2 illustrates how the TTC cutoffs can be directly calculated from the trade balance equations and capacity equations, without running the TTC algorithm. Example 4.2 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 $\{p_i^j\}_{i,j \in \{1,2\}}$, 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 $p_1^1 = p_2^2 = x$, $p_2^1 = (1 - q/p)/x$, $p_1^2 = (1 - 2q)x$ with $x \in [1 - q/p, 1]$ solve these equations, but if $x \neq \left(1 - \frac{q}{p}\right)^p$ then the corresponding assignment is different from the TTC assignment. Section 4.5.2 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.

4.3.3 Consistency with the Discrete TTC Model

In this section we first show that any discrete economy can be translated into a continuum economy, and that the cutoffs obtained using Theorem 4.2 on this continuum economy give the same assignment as discrete TTC. This demonstrates that the continuum TTC model generalizes the standard discrete TTC model. We then show that the TTC assignment changes smoothly with changes in the underlying economy.

To represent a discrete economy $E = (\mathcal{C}, \mathcal{S}, \succ_c, \succ^S, q)$ 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 $I^s = \left\{ \theta \in \Theta \mid \succ^\theta = \succ^s, r^\theta \in [r_i^s, r_i^s + \frac{1}{N}] \forall i \in \mathcal{C} \right\}$ and identify any point $\theta^s \in I^s$ with student s . Formally, for each student $s \in \mathcal{S}$ and school $i \in \mathcal{C}$, we identify each student $s \in \mathcal{S}$ with the N -dimensional cube $I^s = \succ^s \times \prod_{i \in \mathcal{C}} \left[r_i^s, r_i^s + \frac{1}{N} \right)$ of student types with preferences

\succ^s (where $r_i^s = |\{s' \mid s \succ_i s'\}| / |\mathcal{S}|$ is the percentile rank of s at i) and define η to have constant density $\frac{1}{N} \cdot N^N$ on $\cup_s I^s$ and 0 everywhere else.

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. 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.

Proposition 4.2. *Let $E = (\mathcal{C}, \mathcal{S}, \succ_{\mathcal{C}}, \succ^{\mathcal{S}}, q)$ be a discrete economy with $N = |\mathcal{S}|$ students, and let $\Phi(E) = (\mathcal{C}, \Theta, \eta, \frac{q}{N})$ be the corresponding continuum economy. Let \mathbf{p} be the cutoffs produced by Theorem 4.2 for economy $\Phi(E)$. Then the cutoffs \mathbf{p} give the TTC assignment for the discrete economy E , namely,*

$$\mu_{dTTC}(s \mid E) = \max_{\succ^s} \left\{ j \mid r_i^s \geq p_i^j \text{ for some } i \right\},$$

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

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

The idea behind the proof is as follows. 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 $TTC(\gamma \mid \Phi(E))$. Since the assignment of discrete TTC is unique Shapley and Scarf (1974), and the assignment in the continuum model is unique (Proposition 4.2), this proves the theorem.

Proof. Consider a point during the run of discrete TTC when all schools are still available. At this point, denote by x_i the i -rank of the student pointed to by school i for all $i \in \mathcal{C}$, and denote by $S(x)$ the set of assigned students. By construction, $x \in X = \left\{ 0, \frac{1}{N}, \frac{2}{N}, \dots, 1 \right\}^{\mathcal{C}}$. 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 $d_i = \mathbf{1}_{\{i \in K\}}$. Denote by y_i the i -rank of the student pointed to by school i after the cycle is cleared for all $i \in \mathcal{C}$, 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 $H_i^j(\mathbf{x} + \mathbf{1} \cdot \frac{1}{N}) = 1$ if in the discrete TTC school j is the favorite school of the student with i -rank x_i , and $H_i^j(\mathbf{x} + \mathbf{1} \cdot \frac{1}{N}) = 0$ otherwise. On the path $\gamma(t)$ we have that for every $i, j \in K$

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

and if $i \in K$ and $j \notin K$ then $H_i^j(\gamma(t)) = 0$.

Therefore for any $i \in K$

$$\sum_{k \in \mathcal{C}} d_k H_k^i(\gamma(t)) = (1 - (t - t_1)) = \sum_{k \in \mathcal{C}} d_i H_i^k(\gamma(t)),$$

and for any $i \notin K$

$$\sum_{k \in \mathcal{C}} d_k H_k^i(\gamma(t)) = 0 = \sum_{k \in \mathcal{C}} d_i H_i^k(\gamma(t)).$$

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

In other words, Φ embeds a discrete economy into a continuum economy that represents it, and the TTC cutoffs in the continuum embedding give the same assignment as TTC in the discrete model. This shows that the TTC assignment defined in Theorem 4.2 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 Section 4.4.5.

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

In Chapter 4.5.2, 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 proportional to the total variation distance of the two economies, and suggests that the TTC mechanism is fairly robust to small perturbations in preferences.

4.3.4 Proper budget sets

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

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

denote the set of possible school assignments that student s can achieve by unilaterally changing her reported preferences. Note that s cannot misreport her priority.

We observed in Chapter 4.2.4 that in the discrete model the budget set $B(s, \mathbf{p})$ produced by cutoffs $\mathbf{p} = \mathbf{p}(E)$ generated by Theorem 4.1 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 4.3. Let $E = (\mathcal{C}, \mathcal{S}, \succ^S, \succ_C, q)$ be a discrete economy, and let

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

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

Then

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

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

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

Proof. Throughout the proof, we omit the dependence on E and let $B^*(s)$ denote $B^*(s \mid E)$. For brevity, we also let $B(s) = \bigcap_{\mathbf{p} \in \mathcal{P}(E)} B(s, \mathbf{p})$ denote the intersection of all possible budget sets of s in the continuum embedding with some path γ and resulting cutoffs \mathbf{p} . We construct TTC cutoffs $\{(\mathbf{p}^*)_i^j = \gamma_i^*(t^{(j)})\}$ given by a TTC path γ^* and stopping times $\{t^{(j)}\}_{j \in \mathcal{C}}$ that satisfy trade balance and capacity for $\Phi(E)$ such that $B^*(s) \subseteq B(s) \subseteq B(s; \mathbf{p}^*) \subseteq B^*(s)$.

We first show that $B^*(s) \subseteq B(s)$. Suppose $i \notin B(s)$. Then there exists a TTC path γ for E such that $r^s + \frac{1}{|\mathcal{S}|} \mathbf{1} \leq \gamma(t^{(i)})$. Hence for all $\tilde{\succ}$ there exists a TTC path $\tilde{\gamma} \in \mathcal{P}([E_{-s}; \tilde{\succ}])$ such that $r^s + \frac{1}{|\mathcal{S}|} \mathbf{1} \leq \tilde{\gamma}(t^{(i)})$, e.g. the TTC path that follows the same valid directions as γ until it assigns student s . By Proposition 4.2 and Theorem 4.2 for all $\tilde{\succ}$ it holds that $\mu_{dTTC}(s \mid [E_{-s}; \tilde{\succ}]) = \max_{\tilde{\succ}} \left\{ k : r_j^s \geq \tilde{\gamma}(t^{(k)})_j \text{ for some } j \right\}$. Hence for all $\tilde{\succ}$ it holds that $\mu_{dTTC}(s \mid [E_{-s}; \tilde{\succ}]) \neq i$ and so $i \notin B^*(s)$.

We next show that $B(s) \subseteq B(s; \mathbf{p}^*) \subseteq B^*(s)$. Intuitively, we construct the special TTC path γ^* for E by clearing as many cycles as possible that do not involve student s . Formally, let \triangleright be an ordering over subsets of \mathcal{C} where: (1) all subsets containing student s 's top choice available school i^* (under the preferences \succ^s in E) come after all subsets not containing i^* ; and (2) subject to this, subsets are ordered via the shortlex order. Let γ^* be the TTC path for E obtained by selecting valid directions with minimal support under the order \triangleright . (Such a path exists since when using the

shortlex order the resulting valid directions d are piecewise Lipschitz continuous.)

It follows trivially from the definition of $B(s)$ that $B(s) \subseteq B(s; \mathbf{p}^*)$. We now show that $B(s; \mathbf{p}^*) \subseteq B^*(s)$. For suppose $i \in B(s; \mathbf{p}^*)$. Consider the preferences \succ' that put school i first, and then all other schools in the order given by \succ^s . Let E' denote the economy $[E_{-s}; \succ']$. It remains to show that $\mu_{dTTC}(s | E') = i$. Since $i \in B(s; \mathbf{p}^*)$, it holds that $r^s \not\prec \gamma^*(t^{(i)})$. In other words, if we let $\tau^* = \inf \{\tau | \gamma^*(\tau) \not\prec r^s\}$ be the time that the cube I^s corresponding to student s starts clearing, then school i is available at time τ^s . Let γ' be the TTC path for E' obtained by selecting valid directions with minimal support under the order \triangleright , and let $\tau' = \inf \{\tau | \gamma'(\tau) \not\prec r^s\}$. We show that $\tau \leq \tau^*$ and school i is available to student s at time τ' .

Consider the time interval $[0, \min \{\tau^*, \tau'\}]$. During this time the set of valid directions along the TTC path remain the same (i.e. $\frac{d\gamma'}{dt} = \frac{d\gamma^*}{dt}$), as the set of valid directions not involving student s ¹⁸ hasn't changed, and student s ¹⁹ has not yet been assigned under either $TTC(\gamma^*|E)$ or $TTC(\gamma'|E')$ so we do not need to consider the set of valid directions involving student s . Now at worst in going from γ, E to $\gamma', [E_{-s}; \succ']$ we have replaced a valid direction involving s and i^* with a different valid direction involving s and not involving i^* , so student s is assigned sooner in $TTC(\gamma'|E')$ than in $TTC(\gamma^*|E)$, giving $\tau' \leq \tau^*$. Hence $\gamma'(\tau') = \gamma^*(\tau')$ where $\tau' \leq \tau^* \leq t^{(i)}$ and so school i is available to student s when she is assigned. Hence by Proposition 4.2 and Theorem 4.2 it holds that $\mu_{dTTC}(s | E') = i$ and so $i \in B^*(s)$. \square

Proposition 4.3 allows us to construct proper budget sets for each agent that determine not only their assignment given their current preferences, but also their assignment given any other submitted preferences. This particular budget set representation of TTC makes it clear that it is strategy-proof.

¹⁸We say that a valid direction 'involves' a student s if it starts at a point x on the boundary of their cube I^s and points into the interior of the cube.

¹⁹More formally, no points in the cube corresponding to student s are assigned.

4.4 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. This section can be skipped by the reader on a first reading without loss of continuity.

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 (4.2) and (4.3), which determine the characterization given in Theorem 4.2.

4.4.1 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




Figure 4.4: 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.

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 $i \in \mathcal{C}$, let $\gamma_i(\tau) \in [0, 1]$ be the percentile rank of the remaining student with highest i -priority. That is, at time τ in the algorithm each school i is offering a seat to students s for whom $r_i^s = \gamma_i(\tau)$. Let $\gamma(\tau)$ be the vector $(\gamma_i(\tau))_{i \in \mathcal{C}}$. 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 $r_i^s > \gamma_i(\tau)$ for some i must have already been assigned. Likewise, the set of remaining unassigned students is $\{s \mid r^s \leq \gamma(\tau)\}$. See Figure 4.4 for an illustration. Since all assigned students were assigned to their top choice, the remaining capacity at school $i \in \mathcal{C}$ is $q_i - |\{s \mid r^s \not\leq \gamma(\tau) \text{ and } Ch^s(\mathcal{C}) = i\}|$. 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 continuum, the algorithm progresses in continuous time. The state of the algorithm at time $\tau \in \mathbb{R}_{\geq}$ is given by $\gamma(\tau) \in [0, 1]^{\mathcal{C}}$, where $\gamma_i(\tau) \in [0, 1]$ is the percentile rank of the remaining students with highest i -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.

4.4.2 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 i the top student is cleared and $\gamma_i(\cdot)$ 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 $i \in \mathcal{C}$ be some school. For each cycle, the measure of students assigned to school c is equal to the measure of seats offered²⁰ by school i . Therefore, if students are assigned between time τ_1 and τ_2 through clearing a collection of cycles, then the set of students assigned to school i has the same measure as the set of seats offered by school i . 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 $\left\{H_i^j(x)\right\}_{i,j \in \mathcal{C}}$, yielding an equation that determines \mathbf{d} . We provide an illustrative example with two schools in Figure 4.5. For the sake of clarity, we omit technical details in the ensuing discussion. A rigorous derivation can be found in Appendix B.2.

We first identify the measure of students who were offered a seat at a school i or assigned to a school j 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 i

$$|\gamma_i(\tau_2) - \gamma_i(\tau_1)| \approx \varepsilon d_i,$$

that is, during the step from time τ_1 to time τ_2 the algorithm clears students with i -ranks between

²⁰Strictly speaking, the measure of students assigned to each school during the time step is equal to the measure of seats at that school which were claimed by the student offered the seat or traded by the student offered the seat during the time step (not the measure of seats offered). A seat can be offered but not claimed or traded in one of two ways. The first occurs when the seat is offered at time τ but not yet claimed or traded. The second is when a student is offered two or more seats at the same time, and trades only one of them. Both of these sets are of η -measure 0 under our assumptions, and thus the measure of seats claimed or traded is equal to the measure of seats offered.




Figure 4.5: 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 $T_i^j(x, \varepsilon d_i)$ 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.

$\gamma_i(\tau_1) = x$ and $\gamma_i(\tau_2) = x - \varepsilon d_i$. To capture this set of students, let

$$T_i(x, \varepsilon d_i) \stackrel{def}{=} \left\{ \theta \in \Theta \mid r^\theta \leq x, r_i^\theta > x - \varepsilon d_i \right\}$$

denote the set of students with ranks in this range. For all ε , $T_i(x, \varepsilon d_i)$ is the set of top remaining students at i , and when ε is small, $T_i(x, \varepsilon d_i)$ is approximately the set of students who were offered a seat at school i during the step.²¹

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

$$T_i^j(x, \varepsilon d_i) \stackrel{def}{=} \left\{ \theta \in T_i(x, \varepsilon d_i) \mid Ch^\theta(\mathcal{C}) = j \right\},$$

denote the top remaining students on i 's priority list whose top choice is school j . Then the set of students assigned to school j during the step is $\cup_k T_k^j(x, \varepsilon d_k)$, the set of students that got an offer from some school $k \in \mathcal{C}$ and whose top choice is j .

We want to equate the measure of the set $\cup_k T_k^i(x, \varepsilon d_k)$ of students who were assigned to i

²¹The students in the set $T_i(x, \varepsilon d_i) \cap T_k(x, \varepsilon d_k)$ could have been offered a seat at school k and assigned before getting an offer from school i . However, for small ε the intersection is of measure $O(\varepsilon^2)$ and therefore negligible.

with the measure of the set of students who are offered a seat at i , which is approximately the set $T_i(x, \varepsilon d_i)$. By smoothness of the density of η , for sufficiently small δ we have that

$$\eta\left(T_i^j(x, \delta)\right) \approx \delta \cdot H_i^j(x).$$

Therefore, we have that²²

$$\begin{aligned} \eta\left(\cup_k T_k^i(x, \varepsilon d_k)\right) &\approx \sum_{k \in \mathcal{C}} \eta\left(T_k^i(x, \varepsilon d_k)\right) \approx \sum_{k \in \mathcal{C}} \varepsilon d_k \cdot H_k^i(x), \\ \eta\left(T_i(x, \varepsilon d_i)\right) &= \eta\left(\cup_k T_i^k(x, \varepsilon d_i)\right) \approx \sum_{k \in \mathcal{C}} \varepsilon d_i \cdot H_i^k(x). \end{aligned}$$

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,

$$\sum_{k \in \mathcal{C}} \varepsilon d_k \cdot H_k^i(x) \approx \sum_{k \in \mathcal{C}} \varepsilon d_i \cdot H_i^k(x).$$

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

$$\sum_{k \in \mathcal{C}} \gamma'_k(\tau_1) \cdot H_k^i(x) = \sum_{k \in \mathcal{C}} \gamma'_i(\tau_1) \cdot H_i^k(x).$$

4.4.3 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

$$\sum_{k \in \mathcal{C}} d_k \cdot H_k^i(x) = \sum_{k \in \mathcal{C}} d_i \cdot H_i^k(x).$$

²²These approximations make use of the fact that $\eta(T_i(x, \varepsilon d_i) \cap T_k(x, \varepsilon d_k)) = O(\varepsilon^2)$ for small ε .




Figure 4.6: Example of a graph representation for the trade balance equations at x . There is an edge from i to j if and only if $H_i^j(x) > 0$. The two communication classes are framed.

Consider the following equivalent representation. Construct a graph with a node for each school. Let the weight of node i be d_i , and let the flow from node i to node j be $f_{i \rightarrow j} = d_i \cdot H_i^j(x)$. The flow $f_{i \rightarrow j}$ represents the flow of students who are offered a seat at i and wish to trade it for school j when the algorithm progresses down school i 's priority list at rate d_i . Figure 4.6 illustrates such a graph for $\mathcal{C} = \{1, 2, 3, 4\}$. Given a collection of cycles let d_i be the number of cycles containing node i . 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 i 's priority list at rate d_i if and only if \mathbf{d} is a solution to the trade balance equations.

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 i to state j equal to $H_i^j(x)$. The stationary distributions of the Markov chain are characterized by the balance equations, which state that the total probability flow out of state i is equal to the total probability flow into state i . 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 $P_j^i = 0$ for every $j \in K$

and $i \notin K$. See Figure 4.6 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 there is a unique vector \mathbf{d}^K that is a stationary distribution and $\mathbf{d}_i^K = 0$ for any $i \notin K$. 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.²³ 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 4.1. *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 $\tau \leq \min_{i \in \mathcal{C}} \{t^{(i)}\}$ we have that $\gamma(\cdot)$ is given by*

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

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

On the Multiplicity of TTC Paths

In general, there can be multiple solutions to the trade balance equations at x , and therefore multiple TTC paths. The Markov chain and recurrent communication class structure give intuition as to why the TTC assignment is still unique. Each solution \mathbf{d}^K corresponds to the clearing of cycles involving only schools within the set K . The discrete TTC algorithm may encounter multiple disjoint trade cycles, and the outcome of the algorithm is invariant to the order in which these cycles are cleared (when preferences are strict). Similarly here, the algorithm may encounter mutually exclusive combinations of trade cycles $\{\mathbf{d}^{K_1}, \dots, \mathbf{d}^{K_\ell}\}$, which can be cleared sequentially

²³ η has full support if for every open set $A \subset \Theta$ we have $\eta(A) > 0$.




Figure 4.7: Illustration of the gradient field $\mathbf{d}(\cdot)$ and path $\gamma(\cdot)$ (ignoring the capacity equations).

or simultaneously at arbitrary relative rates. Theorem 4.2 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 4.6. 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 4.7 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.

4.4.4 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 i at time τ is

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

Because $\gamma(\cdot)$ is continuous and monotonically decreasing in each coordinate, $D^i(\gamma(\tau))$ 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 $t^{(i^*)}$ that solves the capacity equations

$$\begin{aligned} D^{i^*}(\gamma(t^{(i^*)})) &= q_{i^*} \\ D^k(\gamma(t^{(i^*)})) &\leq q_k \quad \forall k \in \mathcal{C} \end{aligned}$$

where i^* 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 $t^{(i^*)}$ start the path from $\gamma(t^{(i^*)})$ 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.

4.4.5 Comparison between Discrete TTC and Continuum TTC

Table 4.1 summarizes the relationship between the discrete and continuum TTC algorithms, and provides a summary of this section. It presents the objects that define the continuum TTC algorithm with their counterparts in the discrete TTC algorithm. For example, running the continuum TTC algorithm on the embedding $\Phi(E)$ of a discrete economy E performs the same assignments as the discrete TTC algorithm, except that the continuum TTC algorithm performs these assignments continuously and in fractional amounts instead of in discrete steps.

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

Table 4.1: The relationship between the discrete and continuum TTC processes.

Finally, we note that the main technical content of Theorem 4.2 is that there always exists a TTC path γ and stopping times $\{t^{(i)}\}$ that satisfy the trade balance and capacity equations, and that these necessary conditions, together with the capacity equations (4.3), are sufficient to guarantee the uniqueness of the resulting assignment.

Example: Embedding a discrete economy in the continuum model

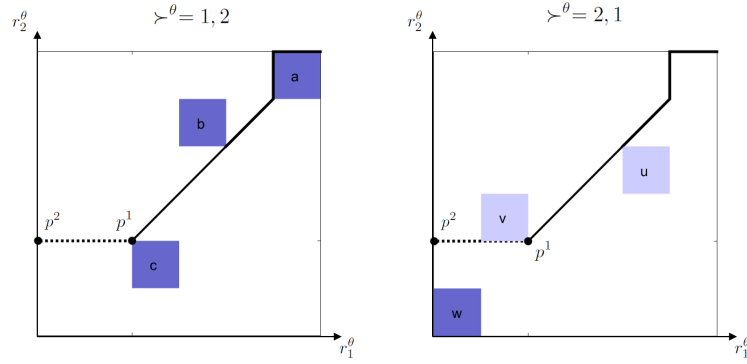
Consider the discrete economy $E = (\mathcal{C}, \mathcal{S}, \succ^{\mathcal{S}}, \succ^{\mathcal{C}}, q)$ with two schools and six students, $\mathcal{C} = \{1, 2\}$, $\mathcal{S} = \{a, b, c, u, v, w\}$. School 1 has capacity $q_1 = 4$ and school 2 has capacity $q_2 = 2$. The school priorities and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

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


In Figure 4.8, 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.

We calculate the TTC paths γ_{all} , γ_1 and γ_2 . We consider only solutions d to the trade balance equations (4.2) 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 4.9.


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



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



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



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

Figure 4.8: Three TTC paths and their cutoffs and allocations for the discrete economy in example 4.4.5. In each set of two squares, students in the left box prefer school 1 and students in the right box prefer school 2. The first round TTC paths are solid, and the second round TTC paths are dotted. The cutoff points p^1 and p^2 are marked by filled circles. Students shaded dark blue are assigned to school 1 and students shaded dark light are assigned to school 2.

At every point (x_1, x_2) with $\frac{2}{3} < x_1, x_2 \leq \frac{5}{6}$ the H matrix is $\begin{bmatrix} 0 & \frac{1}{6} \\ \frac{1}{6} & 0 \end{bmatrix}$, and so $d = \left[-\frac{1}{2}, -\frac{1}{2}\right]$ 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 $\left(\frac{5}{6}, \frac{5}{6}\right)$ and ends at $\left(\frac{2}{3}, \frac{2}{3}\right)$. At every point $x = \left(\frac{1}{3}, x_2\right)$ with




Figure 4.9: The valid directions $d(x)$ for the continuum embedding $\Phi(E)$. Valid directions $d(x)$ are indicated for points x in the grey squares (including the upper and right boundaries but excluding the lower and left boundaries), as well as for points x on the black lines. Any vector $d(x)$ is a valid direction in the lower left grey square. The borders of the squares corresponding to the students are drawn using dashed gray lines.

$\frac{1}{3} < x_2 \leq \frac{2}{3}$ 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.

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 $x = \left(\frac{1}{3}, x_2\right)$ 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 $\left(\frac{1}{3}, \frac{1}{2}\right)$ and ends at $\left(\frac{1}{3}, \frac{1}{3}\right)$ where school 2 fills. All three paths continue until $\left(0, \frac{1}{3}\right)$, where school 1 fills.

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

4.5 Applications

4.5.1 Optimal Investment in School Quality

We apply our model to analyze economies where preferences for schools are endogenously determined by the allocation of resources to schools. Empirical evidence suggests that increased financing affects student achievements Jackson et al. (2016); Lafortune et al. (2018); Johnson and Jackson (2017) as well as demand for housing Hoxby (2001); Cellini et al. (2010), which indicate increased demand for schools. Similarly, Krueger (1999) finds that smaller classes have a positive impact on student performance, and Dinerstein et al. (2014) finds that increased funding for public schools increases enrollment in public schools and reduces demand for private schools.

Under school choice, such resource allocation decisions can change the desirability of schools and therefore change the assignment of students to schools. We explore the implication of such changes in a stylized model. As a shorthand, we refer to an increase in the desirability of a school as an increase in the quality of the school. We explore comparative statics of the allocation and evaluate student welfare. Omitted proofs and derivations can be found in the Appendix B.3.

Model with quality dependent preferences We enrich the model from Chapter 4.3 to allow student preferences to depend on school quality $\delta = \{\delta_i\}_{i \in \mathcal{C}}$, where the desirability of school i is increasing in δ_i . 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 $u^s(i | \delta)$ is the utility of student s for school i given $\delta = \{\delta_i\}_{i \in \mathcal{C}}$ and r_i^s is the student's rank at school i . We assume $u^s(i | \cdot)$ is differentiable, increasing in δ_i and non-increasing in δ_j for any $j \neq i$. The measure ν over Υ specifies the distribution of student types. School capacities are $q = \{q_i\}$, where $\sum q_i < 1$. We will refer to δ_i as the quality of i .

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.²⁴ We assume for all δ that η_δ has a Lipschitz continuous non-negative density ν_δ that is bounded below on its support and depends smoothly on δ . For a given δ , let μ_δ and $\{p_j^i(\delta)\}_{i \in \mathcal{C}}$ denote the TTC assignment and associated cutoffs for the economy \mathcal{E}_δ . We omit the dependence on δ when it is clear from context.

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

Proposition 4.4. *Suppose $\mathcal{E} = (\mathcal{C} = \{1, 2\}, \Upsilon, \nu, q)$ and δ induces an economy \mathcal{E}_δ such that the TTC path γ that, if possible, assigns seats at school 1 before seats at school 2, yields $p_2^1(\delta) > p_2^2(\delta)$.²⁵ Consider $\hat{\delta}$ that increases the quality of school 2, i.e. $\hat{\delta}_2 \geq \delta_2$ and $\delta_1 = \hat{\delta}_1$, and which induces $\mathcal{E}_{\hat{\delta}}$ with TTC path $\hat{\gamma}$ that also assigns seats at 1 before 2 when possible and yields $p_2^1(\hat{\delta}) \geq p_2^2(\hat{\delta})$.*

Then a change from δ to $\hat{\delta}$ induces the cutoffs $p_j^i(\cdot)$ to change as follows:

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

Proposition 4.4 is illustrated in Figure 4.10. 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.

²⁴To make student preferences strict we arbitrarily break ties in favor of schools with lower indices. We assume the




Figure 4.10: The effect of an increase in the quality of school 2 on TTC cutoffs and budget sets. Dashed lines indicate initial TTC cutoffs, and dotted lines indicate TTC cutoffs given increased school 2 quality. The cutoffs $p_1^1 = p_1^2$ and p_2^1 decrease and the cutoff p_2^2 increases. Students in the colored sections receive different budget sets after the increase. Students in dark blue improve to a budget set of $\{1, 2\}$ from \emptyset , students in light blue improve to $\{1, 2\}$ from $\{2\}$, and students in red have an empty budget set \emptyset after the change and $\{2\}$ before.

When there are $n \geq 3$ schools, it is possible to show that an increase in the quality of a school ℓ can either increase or decrease any cutoff. With additional structure we can provide precise comparative statics that mirror the intuition from Proposition (4.4).

Consider the logit economy where students' utilities for each school i are randomly distributed as a logit with mean δ_i , independently of priorities and utilities for other schools. That is, utility for school i is given by $u^s(i | \delta) = \delta_i + \varepsilon_{is}$ with ε_{is} distributed as i.i.d. extreme value shifted to have a mean of 0 McFadden (1973). We assume that the total measure of students is normalized to 1, that there are more students than school seats, i.e. $\sum_i q_i < 1$, and that all students prefer any school to being unassigned²⁶. Schools' priorities are uncorrelated and uniformly distributed. This model combines heterogeneous idiosyncratic taste shocks with a common preferences modifier δ_i . Proposition 4.5 gives the TTC assignment in closed form for the logit economy.

Proposition 4.5. *Under the logit economy schools are labeled in order if $\frac{q_1}{e^{\delta_1}} \leq \frac{q_2}{e^{\delta_2}} \leq \dots \leq \frac{q_n}{e^{\delta_n}}$, and in such cases the TTC cutoffs p_j^i for $i \leq j$ are given by²⁷*

$$p_j^i = \left(R^i\right)^{\frac{e^{\delta_j}}{\pi_i}} \prod_{k < i} \left(R^k\right)^{\frac{e^{\delta_j}}{\pi_k} - \frac{e^{\delta_j}}{\pi_{k+1}}} \quad (4.4)$$

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

Moreover, p_j^i is decreasing in δ_ℓ for $i < \ell$ and increasing in δ_ℓ for $j > i = \ell$.

Figure 4.11 illustrates how the TTC cutoffs change with an increase in the quality of school ℓ . Using equation (4.4), we derive closed form expressions for $\frac{dp_j^i}{d\delta_\ell}$, which can be found in Appendix B.3.

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

²⁵Formally, γ is defined by requiring that for all t it holds that $\gamma'(t)$ is the valid direction at $\gamma(t)$ with support that is minimal under the order $\{1\} < \{1, 2\} < \{2\}$.

²⁶Formally, $u^s(\phi | \delta) = -\infty$. For welfare calculations we only consider assigned students.

²⁷When $i = 1$ we let $\prod_{i' < i} p_{i'}^{i-1} = 1$ and set $\rho_1 = q_1/e^{\delta_1}$.




Figure 4.11: The effects of changing the quality δ_ℓ of school ℓ on the TTC cutoffs p_i^j under the logit economy. If $i < \ell$ then $\frac{dp_i^j}{d\delta_\ell} < 0$ for all $j \geq i$, i.e., it becomes easier to get into the more popular schools. If $i = \ell$ then $\frac{dp_i^j}{d\delta_\ell} = 0$. If $i > \ell$ then $\frac{dp_i^j}{d\delta_\ell} = \frac{dp_i^\ell}{d\delta_\ell} > 0$ for all $j > \ell$, and p_i^ℓ may increase or decrease depending on the specific problem parameters. Note that although p_j^i and p_i^ℓ look aligned in the picture, in general it does not hold that $p_j^i = p_i^\ell$ for all j .

Remark 4.2. Proposition 4.5 can be used to calculate admission probability under multiple tie-breaking as follows. Consider an economy where priorities are determined by a multiple tie-breaking rule where the priority of each student at each school is generated by an independent $U[0, 1]$ lottery draw. As a result, students priorities will be uniformly distributed over $[0, 1]^C$ and uncorrelated with student preferences. If in addition student preferences are given by the MNL model, this is a logit economy. In the logit economy the ex-ante probability that a student will gain admission to school i is given by

$$1 - \prod_{j \in C} p_j^i$$

with p_j^i given by Proposition 4.5.

Comparative statics of student welfare We consider a social planner who can affect quality levels δ of schools in economy \mathcal{E} . We suppose that the social planner wishes to assign students to schools at which they attain high utility, and for the sake of simplicity consider students' social welfare as a proxy for the social planner's objective. Given assignment μ , the social welfare is given by

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

As a benchmark, we first consider neighborhood assignment μ_{NH} which assigns each student to a fixed school regardless of her preferences. We assume this assignment fills the capacity of each school. Social welfare for the logit economy is

$$U_{NH}(\delta) = \sum_i q_i \cdot \delta_i,$$

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

The budget set formulation of TTC allows us to tractably capture student welfare under TTC.²⁸ A student who is offered the budget set $\mathcal{C}^{(i)} = \{i, \dots, n\}$ is assigned to the school $\ell = \arg \max_{j \in \mathcal{C}^{(i)}} \{\delta_j + \varepsilon_{js}\}$, and her expected utility is $U^i = \ln\left(\sum_{j \geq i} e^{\delta_j}\right)$ Small and Rosen (1981). Let N^i be the mass of agents with budget set $\mathcal{C}^{(i)}$. Then social welfare under the TTC assignment given δ simplifies to

$$U_{TTC}(\delta) = \sum_i N^i \cdot U^i.$$

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

Proposition 4.6. *For the logit economy, the change in social welfare $U_{TTC}(\delta)$ under TTC from a marginal increase in δ_ℓ is given by*

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

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

Proposition 4.6 shows that under TTC a marginal increase in the quality of school ℓ will have

²⁸Under 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, $\mathbb{E}[u^s(\mu(s) | \delta)] = \delta_{\mu(s)} + \mathbb{E}[\varepsilon_{\mu(s)s} | \delta_{\mu(s)} + \varepsilon_{\mu(s)s} \geq \delta_i + \varepsilon_{is} \forall i \in B(s, \mathbf{p})]$

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 is²⁹

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

An increase in the quality of school 1 gives higher utility for students assigned to 1, which is captured by the first term. Additionally, it causes some students to switch their preferences to $1 \succ 2$, making school 1 run out earlier in the TTC algorithm, and removing school 1 from the budget set of some students. Students whose budget set did not change and who switched to $1 \succ 2$ are almost indifferent between the schools and hence almost unaffected. Students who lost school 1 from their budget set may prefer school 1 by a large margin, and hence incur significant loss. Thus, there is a total negative effect from changes in the assignment, which is captured by the second term.

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

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

which is larger than the marginal effect under neighborhood assignment.

If admission cutoffs into both schools are equal (that is, $p_1^1 = p_2^2$ and $N^2 = 0$) we say that both schools are equally over-demanded. In such a case, a marginal increase in the quality of either

²⁹Recall that we assume that schools are labeled in order, and thus school 1 is the more selective school. We use that $N^1 = q_1 + q_1 e^{\delta_1 - \delta_2}$, $N^2 = q_2 - q_1 e^{\delta_2 - \delta_1}$.

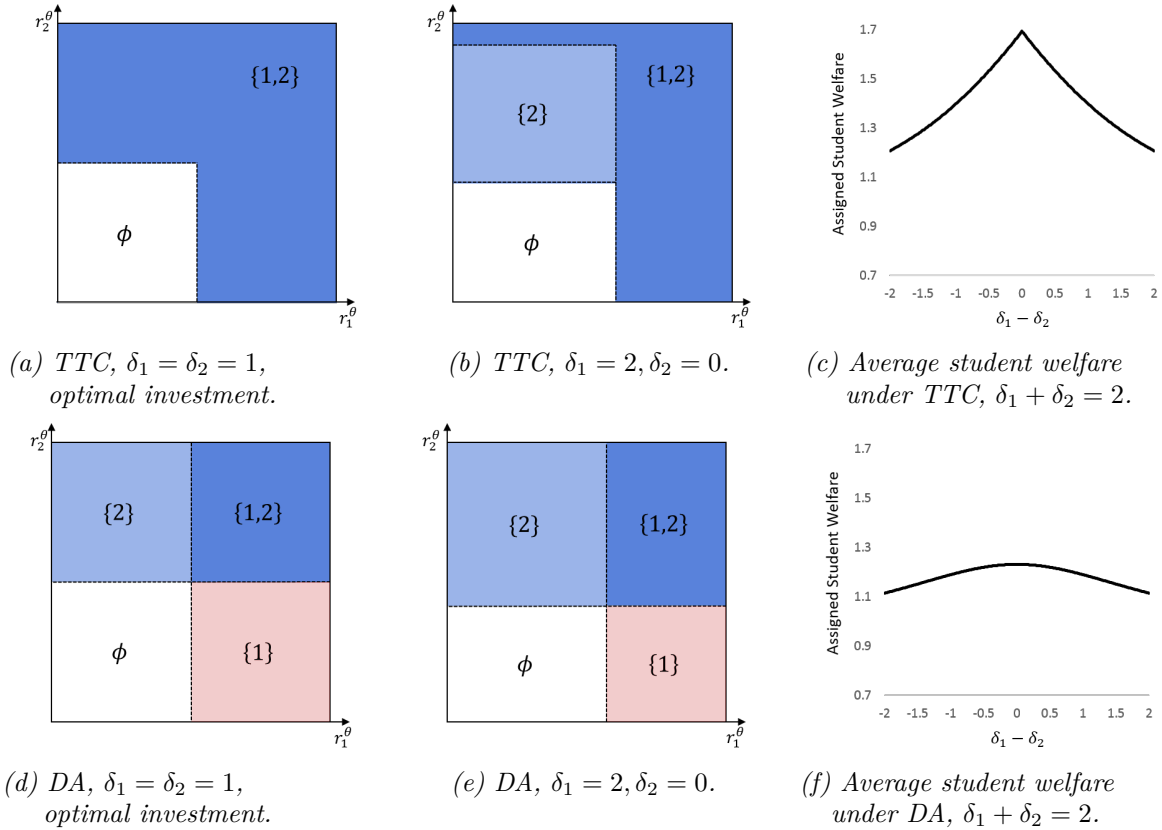


Figure 4.12: Illustration for Example 4.3. Figures (a) and (b) show the budget sets under *TTC* for different quality levels, and Figure (c) shows the average welfare of assigned students under *TTC* for quality levels $\delta_1 + \delta_2 = 2$ for different values of $\delta_1 - \delta_2$. Figures (d) and (e) show the budget sets under *DA*, and Figure (f) shows the average welfare of assigned students under *DA*.

school will have a negative indirect effect on welfare.³⁰

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

Example 4.3. Consider a logit economy with two schools and $q_1 = q_2 = \frac{3}{8}$, and let $Q = q_1 + q_2$ denote the total capacity. Quality levels δ are constrained by $\delta_1 + \delta_2 = 2$ and $\delta_1, \delta_2 \geq 0$.

³⁰That is, if $\delta_1 = \delta_2$ then $\frac{dU_{TTC}}{d\delta_1} < q_1$ and $\frac{dU_{TTC}}{d\delta_2} < q_2$. If we fix $\delta_1 + \delta_2$ and consider $U_{TTC}(\Delta)$ as a function of $\Delta = \delta_1 - \delta_2$ the function $U_{TTC}(\Delta)$ will have a kink at $\Delta = 0$ (see Figure 4.12c).

Under neighborhood assignment $U_{NH}/Q = 1$ for any choice of δ_1, δ_2 . Under TTC the unique optimal quality is $\delta_1 = \delta_2 = 1$, yielding $U_{TTC}/Q = 1 + \mathbb{E}[\max(\varepsilon_{1s}, \varepsilon_{2s})] = 1 + \ln(2) \approx 1.69$. This is because any assigned student has the budget set $B = \{1, 2\}$ and is assigned to the school for which he has higher idiosyncratic taste. Welfare is lower when $\delta_1 \neq \delta_2$, because fewer students choose the school for which they have higher idiosyncratic taste. For instance, given $\delta_1 = 2, \delta_2 = 0$ welfare is $U_{TTC}/Q = \frac{1}{2}(1 + e^{-2}) \log(1 + e^2) \approx 1.20$. Under Deferred Acceptance (DA) the unique optimal quality is also $\delta_1 = \delta_2 = 1$, yielding $U_{DA}/Q = 1 + \frac{1}{3} \ln(2) \approx 1.23$. This is strictly lower than the welfare under TTC because under DA only students that have sufficiently high priority for both schools have the budget set $B = \{1, 2\}$. Two thirds of assigned students have a budget set $B = \{1\}$ or $B = \{2\}$, corresponding to the single school for which they have sufficient priority. If $\delta_1 = 2, \delta_2 = 0$ welfare under DA is $U_{DA}/Q \approx 1.11$.

In Example 4.3, TTC yields higher student welfare by providing all assigned students with a full budget set, thus maximizing each assigned student's contribution to welfare from horizontal taste shocks. However, the assignment it produces is not stable. In fact, both schools admit students whom they rank at the bottom, and thus virtually all unassigned students can potentially block with either school.³¹ 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 efficient, we assume that (despite having more students) school 1 requires the same amount of resources to increase its quality for all its students. Thus, we keep the constraint that $\delta_1 + \delta_2 = 2$. It is straightforward to see that under neighborhood assignment the welfare optimal distribution of quality is $\delta_1 = 2, \delta_2 = 0$. In contrast, we find the welfare optimal distribution under TTC can be closer to egalitarian.

Example 4.4. Consider a logit economy with two schools and $q_1 = 1/2, q_2 = 1/4$, and let $Q = q_1 + q_2$ denote the total capacity. Quality levels δ are constrained by $\delta_1 + \delta_2 = 2$ and $\delta_1, \delta_2 \geq 0$.

Under neighborhood assignment the welfare optimal quality is $\delta_1 = 2, \delta_2 = 0$, yielding $U_{NH}/Q =$

³¹Note that this is not a concern in school choice settings where blocking pairs cannot be assigned outside of the mechanism.

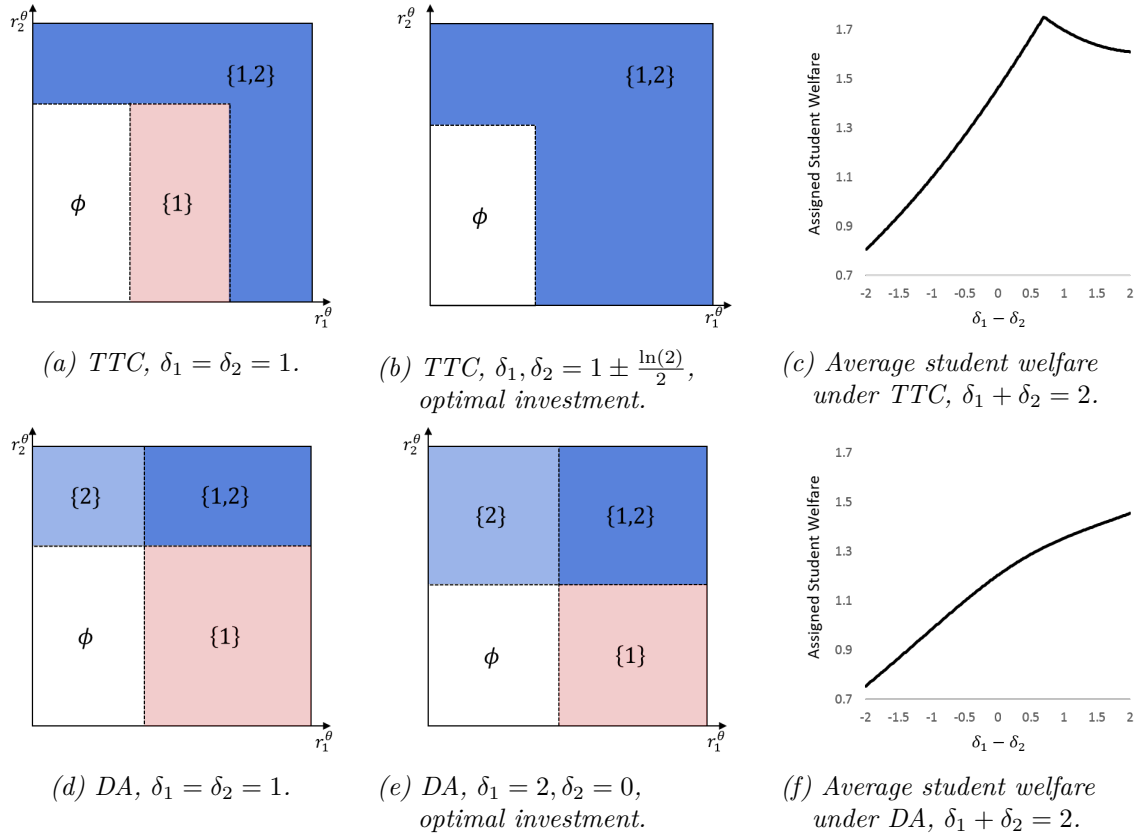


Figure 4.13: Illustration for Example 4.4. Figures (a) and (b) show the budget sets under TTC for different quality levels, and Figure (c) shows the average welfare of assigned students under TTC for quality levels $\delta_1 + \delta_2 = 2$ for different values of $\delta_1 - \delta_2$. Note that $\delta_1 = \delta_2 = 1$ 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.

$4/3 \approx 1.33$. Under TTC assignment the unique optimal quality is $\delta_1 = 1 + \frac{1}{2} \ln(2)$, $\delta_2 = 1 - \frac{1}{2} \ln(2)$, yielding $U_{TTC/Q} = \ln\left(\frac{3e}{\sqrt{2}}\right) \approx 1.75$. Under these quality levels any assigned student has the budget set $B = \{1, 2\}$. Given $\delta_1 = 2, \delta_2 = 0$ welfare is $U_{TTC/Q} \approx 1.61$. The quality levels that are optimal in Example 4.3, namely $\delta_1 = 1, \delta_2 = 1$, yield $U_{TTC/Q} \approx 1.46$. Under DA assignment the unique optimal quality is $\delta_1 = 2, \delta_2 = 0$, yielding $U_{DA/Q} \approx 1.45$. Given $\delta_1 = 1, \delta_2 = 1$ welfare under DA is $U_{DA/Q} \approx 1.20$.

Again in Example 4.4 we find that the optimal quality distribution under TTC provides all assigned students with a full budget set, making all schools equally over-demanded. The optimal quality distribution under neighborhood assignment and DA allocates all resources to the more efficient school. While quality directed to the larger school affects more students and yields more

direct benefit, under TTC an egalitarian distribution leads to more welfare gains from sorting on horizontal tastes. For general parameters the welfare gain from sorting can be lower or higher than the welfare gains from directing all resources to the more efficient school.

Finally, consider a central school board with a fixed amount of resources K to be allocated to the n schools. We assume that the cost of quality δ_i is the convex function $\kappa_i(\delta_i) = e^{\delta_i}$. This specification makes bigger schools more efficient.³² Using Proposition 4.6 we solve for the optimal distribution of school quality. Despite the heterogeneity among schools, social welfare is maximized when all assigned students have a full budget set, which occurs when the amount allocated to each school is proportional to the number of seats at the school.

Proposition 4.7. *Consider a logit economy with cost function $\kappa_i(\delta_i) = e^{\delta_i} \forall i$ and resource constraint $\sum_i \kappa_i(\delta_i) \leq K$. Social welfare is uniquely maximized when the resources κ_i allocated to school i are proportional to the capacity q_i , that is,*

$$\kappa_i(\delta_i) = \frac{q_i}{\sum_j q_j} K$$

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

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.

4.5.2 Design of TTC Priorities

To better understand the role of priorities in the TTC mechanism, we examine how the TTC assignment changes with changes in the priority structure. Notice that any student s whose favorite school is i and who is within the q_i highest ranked students at i is guaranteed admission to i . In the

³²Note that κ_i is the total school funding. This is equivalent to setting the student utility of school i to be to $u^s(i | \kappa_i) = \log(\kappa_i) + \varepsilon_{is} = \log(\kappa_i/q_i) + \log(q_i) + \varepsilon_{is}$, which is the log of the per-student funding plus a fixed school utility that is larger for bigger schools.

following example, we consider changes to the relative priority of such highly ranked students and find that these changes can have an impact on the assignment of other students, without changing the assignment of any student whose priority changed.

Example 4.5. The economy \mathcal{E} has two schools 1, 2 with capacities $q_1 = q_2 = q$, students are equally likely to prefer each school, and student priorities are uniformly distributed on $[0, 1]$ independently for each school and independently of preferences. The TTC algorithm ends after a single round, and the resulting assignment is given by $p_1^1 = p_1^2 = p_2^1 = p_2^2 = \sqrt{1 - 2q}$. The derivation can be found in Appendix B.3.

Consider the set of students $\{s \mid r_i^s \geq m \forall i\}$ 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 $[m, 1]^2$.³³ 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 4.14 illustrates the range of possible TTC cutoffs for \mathcal{E}' and the economy $\bar{\mathcal{E}}$ for which TTC obtains one set of extreme cutoffs.

Example 4.5 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 4.14) 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 $\{s \mid r_i^s \geq m \forall i\}$ 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 4.3 implies that the changes in TTC outcomes are small if $1 - m$ is small.

A second implication is that the TTC priorities can be ‘bossy’ in the sense that changes in the relative priority of high priority students can affect the assignment of other students, even when all high priority students receive the same assignment. Notice that in all the economies considered

³³ The remaining students still have ranks distributed uniformly on the complement of $[m, 1]^2$.




Figure 4.14: The range of possible TTC cutoffs in example 4.5 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 $p_2^2 = \sqrt{(1-2q)\frac{m^2}{1-2m+2m^2}} \approx 0.25$.

in Example 4.5, we only changed the relative priority within the set $\{s \mid \exists i \text{ s.t. } r_i^s \geq m\}$, 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 s with priority $r_1^s = 0.35, r_2^s = 0.1$ could possibly receive his first choice or be unassigned. Such changes to priorities may naturally arise when there are many indifferences in student priorities, and tie-breaking is used. Since priorities are bossy, the choice of tie-breaking between high-priority students can have indirect effects on the assignment of low priority students.

4.5.3 Comparing Mechanisms

In Chapter 4.5.1 we compared the welfare effects of changes in school resource allocation under various school choice mechanisms. Our formulation of TTC also allows us to compare TTC with other school choice mechanisms. In this section, we provide a theoretical explanation for observed similarities between assignments under TTC and Deferred Acceptance (DA), as well as a comparison of the number of blocking pairs induced by TTC and the closely related Clinch and Trade mechanism.

Both TTC and Deferred Acceptance (DA) Gale and Shapley (1962) are strategy-proof, but differ in that TTC is efficient whereas DA is stable. Kesten (2006); Ehlers and Erdil (2010) show the two mechanisms are equivalent only under strong conditions that are unlikely to hold in practice. However, Pathak (2016) evaluates the two mechanisms on application data from school choice in New Orleans and Boston, and reports that the two mechanisms produce similar outcomes. In Chapter 4.5.1 we compared DA and TTC in terms of welfare and assignment and found that large differences were possible.³⁴ Pathak (2016) conjectures that the neighborhood priority used in New Orleans and Boston led to correlation between student preferences and school priorities that may explain the similarity between the TTC and DA allocations in these cities.

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

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

Proposition 4.8. *The proportion of students who have the same assignments under TTC and DA is given by*

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

Proof. We use the methodologies developed in Chapter 4.3.2 and in Azevedo and Leshno (2016) to find the TTC and DA allocations respectively. For each school, students with priority are given a lottery number uniformly at random in $\left[\frac{n-1}{n}, 1\right]$, and students without priority are given a lottery number uniformly at random in $\left[0, \frac{n-1}{n}\right]$, where lottery numbers at different schools are

³⁴Che and Tercieux (2015) show that when there are a large number of schools with a single seat per school and preferences are random both DA and TTC are asymptotically efficient and stable and give asymptotically equivalent allocations. As Example 4.3 shows, these results do not hold when there are many students and a few large schools.

independent. For all values of α , the TTC cutoffs are given by $p_j^i = p_i^j = 1 - \frac{q_i}{nq}$ for all $i \leq j$, and the DA cutoffs are given by $p_i = 1 - \frac{q_i}{nq}$. We now derive these cutoffs.

Consider the TTC cutoffs for the neighborhood priority setting. We prove by induction on ℓ that $p_j^\ell = 1 - \frac{q_\ell}{nq}$ for all ℓ, j such that $j \geq \ell$.

Base case: $\ell = 1$.

For each school i , there are measure q of students whose first choice school is i , αq of whom have priority at i and $\frac{(1-\alpha)q}{n-1}$ of whom have priority at school j , for all $j \neq i$.

The TTC path is given by the diagonal, $\gamma(t) = \left(1 - \frac{t}{\sqrt{n}}, 1 - \frac{t}{\sqrt{n}}, \dots, 1 - \frac{t}{\sqrt{n}}\right)$. At the point $\gamma(t) = (x, x, \dots, x)$ (where $x \geq \frac{n-1}{n}$) a fraction $n(1-x)$ of students from each neighborhood have been assigned. Since the same proportion of students have each school as their top choice, this means that the quantity of students assigned to each school i is $n(1-x)q$. Hence the cutoffs are given by considering school 1, which has the smallest capacity, and setting the quantity assigned to school 1 equal to its capacity q_1 . It follows that $p_j^1 = x^*$ for all j , where $n(1-x^*)q = q_1$, which yields

$$p_j^1 = 1 - \frac{q_1}{nq} \text{ for all } j.$$

Inductive step.

Suppose we know that the cutoffs $\{p_j^i\}_{i,j:i \leq \ell}$ satisfy $p_j^i = 1 - \frac{q_i}{nq}$. We show by induction that the $(\ell+1)$ th set of cutoffs $\{p_j^{\ell+1}\}_{j>\ell}$ are given by $p_j^{\ell+1} = 1 - \frac{q_{\ell+1}}{nq}$.

The TTC path is given by the diagonal when restricted to the last $n-\ell$ coordinates, $\gamma(t^{(\ell)} + t) = \left(p_1^1, p_2^2, \dots, p_\ell^\ell, p_\ell^\ell - \frac{t}{\sqrt{n-\ell}}, p_\ell^\ell - \frac{t}{\sqrt{n-\ell}}, \dots, p_\ell^\ell - \frac{t}{\sqrt{n-\ell}}\right)$.

Consider a neighborhood i . If $i > \ell$, at the point $\gamma(t) = (p_1^1, p_2^2, \dots, p_\ell^\ell, x, x, \dots, x)$ (where $x \geq \frac{n-1}{n}$) a fraction $n(p_\ell^\ell - x)$ of (all previously assigned and unassigned) students from neighborhood i have been assigned in round $\ell+1$. If $i \leq \ell$, no students from neighborhood i have been assigned in round $\ell+1$.

Consider the set of students S who live in one of the neighborhoods $\ell+1, \ell+2, \dots, n$. These are the only students who have priority at one of the remaining schools. Moreover, the same proportion of these students have each remaining school as their top choice out of the remaining schools. This

means that for any $i > \ell$, the quantity of students assigned to school i in round $\ell + 1$ by time t is a $\frac{1}{n-\ell}$ fraction of the total number of students assigned in round $\ell + 1$ by time t , and is given by $\frac{1}{n-\ell} (n - \ell) (p_\ell^\ell - x) nq = n (p_\ell^\ell - x) q$. Hence the cutoffs are given by considering school $\ell + 1$, which has the smallest residual, and setting the quantity assigned to school $\ell + 1$ equal to its residual capacity $q_{\ell+1} - q_\ell$. It follows that $p_j^{\ell+1} = x^*$ for all $j > \ell$ where $n (p_\ell^\ell - x^*) q = q_{\ell+1} - q_\ell$, which yields

$$p_j^{\ell+1} = p_\ell^\ell - \frac{q_{\ell+1} - q_\ell}{nq} = 1 - \frac{q_\ell}{nq} - \frac{q_{\ell+1} - q_\ell}{nq} = 1 - \frac{q_{\ell+1}}{nq} \text{ for all } j > \ell.$$

This completes the proof that the TTC cutoffs are given by $p_j^i = p_i^j = 1 - \frac{q_i}{nq}$ for all $i \leq j$.

Now consider the DA cutoffs. We show that the cutoffs $p_i = 1 - \frac{q_i}{nq}$ satisfy the supply-demand equations. We first remark that the cutoff at school i is higher than all the ranks of students without priority at school i , $p_i \geq \frac{n-1}{n}$. Since every student has priority at exactly one school, this means that every student is either above the cutoff for exactly one school and is assigned to that school, or is below all the cutoffs and remains unassigned. Hence there are $nq(1 - p_i) = q_i$ students assigned to school i for all i , and the supply-demand equations are satisfied.

The students who have the same assignments under TTC and DA are precisely the students at neighborhood i whose ranks at school i are above $1 - \frac{q_i}{nq}$, and whose first choice school is their neighborhood school. This set of students comprises an $\alpha \frac{\sum_i q_i}{nq}$ fraction of the entire student population, which scales proportionally with the correlation between student preferences and school priorities. □

We can also compare TTC with the Clinch and Trade (C&T) mechanism introduced by Morrill (2015b). The C&T mechanism identifies students who are guaranteed admission to their favorite school i by having priority $r_i^s \geq 1 - q$ and assigns them to i by ‘clinching’ without trade. Morrill (2015b) suggests that this design choice is desirable because it can reduce the number of blocking pairs induced by the assignment, and gives an example where the C&T assignment has fewer blocking pairs than the TTC assignment. The fact that allowing students to clinch can change the assignment can be interpreted as another example of the bossiness of priorities under TTC: we can equivalently implement C&T by running TTC on a changed priority structure where students

who clinched at school i have higher rank at i than any other student.³⁵ The following proposition builds on Example 4.5 and shows that C&T may produce more blocking pairs than TTC.

Proposition 4.9. *The Clinch and Trade mechanism can produce more, fewer or an equal number of blocking pairs compared to TTC.*

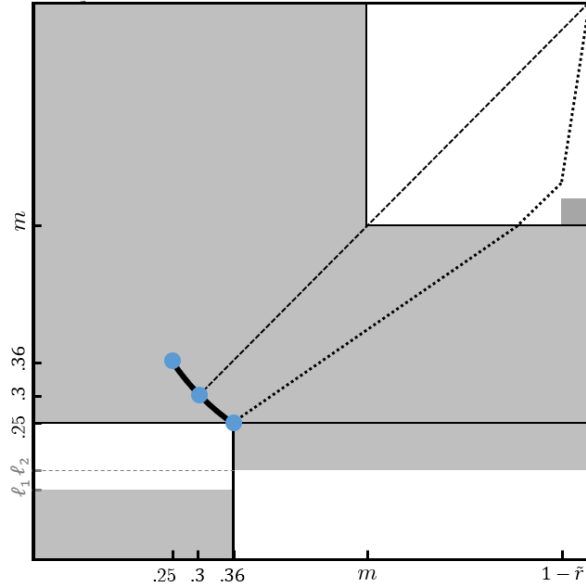


Figure 4.15: Economy \mathcal{E}_1 used in the proof of Proposition 4.9. The black borders partition the space of students into four regions. The density of students is zero on white areas, and constant on each of the shaded areas within a bordered region. In each of the four regions, the total measure of students within is equal to the total area (white and shaded) within the borders of the region.

Proof. Morrill (2015b) provides an example where C&T produces fewer blocking pairs than TTC. Both mechanisms give the same assignment for the symmetric economy in the beginning of Example 4.5. It remains to construct an economy \mathcal{E}_1 for which C&T produces more blocking pairs than TTC. Let economy $\bar{\mathcal{E}}$ be defined as in Section B.3, that is, by taking an economy \mathcal{E} with capacities $q_1 = q_2 = q = 0.455$ where students are equally likely to prefer each school and student priorities are uniformly distributed on $[0, 1]$ independently for each school and independently of preferences, and changing the ranks of top priority students (those with rank $r_1^{\theta}, r_2^{\theta} \geq m = 0.6$) so that they have ranks uniformly distributed in the $\tilde{r} \times \tilde{r}$ square $(1 - \tilde{r}, 1] \times (m, m + \tilde{r}]$ for some $\tilde{r} \leq \frac{(2m-1)(1-m)}{2m}$.

³⁵For brevity, we abstract away from certain details of C&T mechanism that are important when not all schools run out at the same round.

Recall that when running TTC on the economy $\bar{\mathcal{E}}$ the cutoffs are given by $\mathbf{p}^1 = (\bar{p}, \underline{p})$, where $\underline{p} = \sqrt{(1-2q)\frac{m^2}{1-2m+2m^2}}$ and $\bar{p} = \sqrt{(1-2q)\frac{1-2m+2m^2}{m^2}}$. The economy \mathcal{E}_1 is constructed by taking the economy $\bar{\mathcal{E}}$ and redistributing school 2 rank among students with $r_2^\theta \leq \underline{p} \approx 0.25$ so that those with $r_1^\theta \geq \bar{p} \approx 0.36$ have higher school 2 rank.³⁶ The C&T assignment for \mathcal{E}_1 is given by $p_1^1 = p_2^2 = 0.3$, while TTC gives $p_1^1 = \bar{p} \approx 0.36$ and $p_2^2 = \underline{p} \approx 0.25$ (and under both $p_1^1 = p_1^2, p_2^1 = p_2^2$). Under TTC unmatched students will form blocking pairs only with school 2, while under C&T all unmatched students will form a blocking pair with either school. See Figure 4.15 for an illustration. \square

4.6 Discussion

Summary of findings. We have provided a cutoff characterization for the TTC outcome in terms of n^2 cutoffs, one for every pair of schools. The cutoff p_j^i represents the lowest j -priority a student can have in order to use that priority to obtain a seat in school i . In a continuum setting we demonstrated how to compute these cutoffs as the solutions to a system of differential equations. In parametrized economies we provide closed-form expressions for the TTC outcome. We consider the problem of optimal investment in school quality, and show that under TTC the optimal investment levels are equitable, as the greatest utility gains in TTC come from allowing students to choose schools based on their idiosyncratic preferences rather than based on the quality of the school. We are hopeful that our characterization of the TTC assignment can be used to increase the use of TTC in practice and inform optimal decision-making in other aspects of the school assignment problem.

Communicating the TTC outcome. We can simplify how the TTC outcome is communicated to students and their families by using the cutoff characterization. The cutoffs $\{p_i^j\}$ are calculated in the course of running the TTC algorithm. The cutoffs can be published to allow parents to verify their assignment, or the budget set structure can be communicated using the language of tokens (see footnote 4). We hope that these methods of communicating TTC will make the mechanism

³⁶Specifically, select $\ell_1 < \ell_2$. Among students with $r_2^\theta \leq \underline{p}$ and $r_1^\theta \geq \bar{p}$ the school 2 rank is distributed uniformly in the range $[\ell_2, \underline{p}]$. Among students with $r_2^\theta \leq \underline{p}$ and $r_1^\theta < \bar{p}$ the school 2 rank is distributed uniformly in the range $[0, \ell_1]$. Within each range r_1^θ and r_2^θ are still independent. See Figure 4.15 for an illustration.

more palatable to students and their parents, and facilitate a more informed comparison with the Deferred Acceptance mechanism.

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

Optimization and counterfactuals under TTC. Our framework for understanding the TTC outcome can also be used to inform policy decisions and perform counterfactuals using the TTC mechanisms. While examples provided in the paper utilized functional form assumptions to gain tractability, the methodology can be used more generally with numerical solvers. This provides a useful alternative to simulation methods that can be more efficient for large economies, or for calculating an average outcome for large random economies. For example, most school districts uses tie-breaking rules, and current simulation methods perform many draws of the random tie-breaking lottery to calculate the expected outcomes. Our methodology directly calculates the expected outcome from the distribution.

Designing TTC priorities. Arnosti (2015); Ashlagi and Shi (2014); Ashlagi and Nikzad (2016); Feigenbaum et al. show that the optimal choice of tie-breaking lottery can lead to significant welfare gains when using DA. In Chapter 4.5.2 we characterize all the possible TTC outcomes for a class of tie-breaking rules, and find that the choice of tie-breaking rule can have significant effect on the assignment, suggesting that tie-breaking can also play a significant role in determining welfare under TTC. Our cutoff structure for TTC also differs from the one found in Azevedo and Leshno (2016) for DA. This demonstrates that priorities play a different role under these mechanisms, and raises the question of how to best design priorities under TTC. Are there priorities for TTC that, similar to neighborhood priority in DA, allow us to prioritize students for a given school? Can TTC priorities be designed to optimize other global objectives in assigning students to schools, such as maximizing welfare for students from disadvantaged neighborhoods or minimizing busing costs?

Our framework opens the possibility of exploring different tie-breaking schemes when using TTC and also for optimal design of TTC priorities in general, and we leave both questions for future research.

Part II

Mechanisms for School Choice with Incomplete Information

Chapter 5

The Information Acquisition Costs of Matching Markets

Matching markets have been the subject of much academic research as well as substantial interest from practitioners. In these markets agents have preferences over the individuals they are matched to, and the assignment is not determined simply by monetary transfers. Matching theory investigates the role of marketplace rules in determining the allocation, and elucidates how matching markets can and should compute the overall assignment from individual agent preferences. Such models allow us to better understand decentralized markets, such as college admission, or facilitate a better design of centralized assignment mechanisms, such as the medical match and school choice (see, e.g. Roth and Sotomayor, 1992; Roth, 2015).

In this paper, we investigate the effects of mechanisms on how agents form their preferences. The prevalence of incomplete information is well-studied in the context of auction markets (see e.g. Eso and Szentes (2007); Milgrom and Weber (1982)), but is relatively unexplored in matching market settings. This is despite the fact that in matching settings such as medical residency matching and school choice, it is common not only for agents to have incomplete information about their preferences, but also for them to spend a significant amount of effort investigating potential placements before forming their final preferences. For example, in NYC public high school admissions students must submit their preferences over more than 700 programs at more than 400

high schools. Moreover, costly information acquisition is also an important equity problem in school choice, as students from underprivileged backgrounds are often inadequately informed about their options and must exert the most effort to determine their preferences (see, e.g. Hassidim et al., 2015; Kapor et al., 2016).

Thus motivated, we study the effects of market design on costly information acquisition in a many-to-one school choice market. In our model school priorities are common knowledge, and students can acquire costly information about their preferences over school. We model each agent's information acquisition problem using the "Pandora's box" framework of Weitzman (1979) in the tractable continuum matching market of Azevedo and Leshno (2016). Each student knows a prior distribution for each school's utility to them, and must pay a cost to learn the actual utility realization. The utility realizations are independent for each student, and students individually decide on their information acquisition process. The student information acquisition problem admits an optimal solution via a simple index policy, and allows for students to only partially collect information.

We define *stability under incomplete information* for this setting: an outcome for the market is stable with respect to acquired information and information acquisition costs. A blocking pair is a (student, school) pair such that the student (i) has higher priority at the school than another student assigned to that school or the school is undercapacitated, and the student either (ii) prefers the school to their assigned school, given their acquired information, or (ii') does not have enough information to make a decision and is willing to pay the cost to collect further information, and a matching and set of acquired information constitute a stable outcome if there is no blocking pair. This definition extends the standard definition of stability, and is equivalent to the standard definition when students do not incur information acquisition costs and collect all preference information. However, in the presence of information acquisition costs it is possible for different sets of acquired information to lead to different stable outcomes. Hence the design of the market mechanism can induce beliefs that lead students to acquire information differently and implement different outcomes, even when there is a unique stable matching under full information.

In settings with costly information acquisition students need information about their possible

matches in order to optimally acquire information, and students may benefit from waiting for the market to resolve before acquiring information. We refine the set of stable outcomes to the set of *regret-free stable outcomes*, under which the information acquired by each student is the same as if they performed their optimal information acquisition process knowing the preferences and information acquisition processes of all other students. In other words, each student acquires information as if she were the last to enter the market, and no student regrets not waiting for further information about other students' preferences before learning her own preferences. This means that regret-free stable matchings do not depend on student beliefs. We furnish the surprising result that *the set of regret-free stable matchings has a lattice structure*, which it inherits from of the set of stable matchings under complete information (attributed to John Conway in Knuth, 1976), and hence is non-empty and has an outcome that is unambiguously the best for all students.

We then turn to the problem of providing matching mechanisms for implementing regret-free stable outcomes. We show that as regret-free stable matchings are characterized by cutoffs, the student-optimal regret-free stable outcome can be implemented by learning and posting the appropriate admissions cutoffs. For example, given sufficient market structure, school-proposing Deferred Acceptance can be implemented in a sequential manner to learn the regret-free stable cutoffs with regret-free information acquisition. However, we also demonstrate that there exist economies where regret-free stable matchings cannot be computed without incurring additional information acquisition costs, and also where the student-optimality of a regret-free stable matching cannot be verified without incurring additional costly information acquisition. In general settings, standard mechanisms can result in *information deadlocks*, where no information is gathered because every student finds it strictly optimal to wait for others to acquire information first. Hence the presence of costly information acquisition does not affect the structure of the set of stable outcomes but rather the algorithmic questions of *computing* a regret-free stable outcome and *verifying* its optimality.

We show how to approximately compute the market-clearing admissions cutoffs when we have historical information about demand or can estimate it by subsampling, and in such settings provide mechanisms that implement outcomes that are student-optimal regret-free stable with respect to perturbed school capacities. Our results illustrate that, *given sufficient information about aggregate*

student demand for schools, it is possible to approximately implement a regret-free stable matching.

5.1 Prior Work

This paper contributes to the literature of matching markets with incomplete information. The stream of work that is closest to ours is that of Aziz et al. (2016); Rastegari et al. (2013, 2014), which analyze a matching model where there is partial ordinal information on both sides of the market that can be refined through costly interviews. They ask computational questions regarding the minimal number of interviews required to find a stable matching, and find that under a tiered structure an iterative version of DA minimizes the number of required interviews. Our finding that a sequential version of DA implements a regret-free stable matching when agents are willing to inspect all schools they can attend is a particular case of this result where the preferences of one side are known. Drummond and Boutilier (2013, 2014) consider more general algorithms that acquire information through both interviews and comparisons and provide algorithms that achieve approximately stable matchings with low information costs. Lee and Schwarz (2009); Kadam (2015) also study information sharing through interviews.

Several papers consider other aspects of imperfect information in matching markets, without allowing agents to search for information. Liu et al. (2014) suggest a notion of stability under asymmetric information between agents. Chakraborty et al. (2010) consider agents with incomplete information who update their preferences after seeing the matching. Ehlers and Massó (2015) demonstrate that there is a strong connection between ordinal Bayesian Nash equilibria of stable mechanisms under incomplete information and Nash equilibria of the mechanism under corresponding complete information settings. We similarly define a notion of stability under incomplete information and find a strong parallel with the structure of stable matchings under complete information.

Empirical work demonstrates that incomplete information is important in the school choice setting. Kapor et al. (2016) provides empirical evidence that many students participating in a school choice mechanism are not well informed, and make mistakes when reporting their preferences, and

Dur et al. (2015) provides evidence that different parents exert different levels of efforts in learning about school choice.

There is also a growing literature about information acquisition in market design. In an auction setting, Kleinberg et al. (2016) shows that descending price auction create optimal incentives for value discovery. Chen and He (2015) study how the DA and Boston mechanisms give participating agents incentives to learn their preferences and preferences of others, but limit attention to the decision of whether to learn the full ordinal or cardinal valuation for all schools. Bade (2015); Harless and Manjunath (2015) consider information acquisition in assignment problems without stability constraints.

The rational inattention literature that stemmed from the macroeconomic literature also looks at information acquisition by agents. This literature uses a framework introduced by Sims (2003) where the costs of signals are given by information theoretic measures of the informativeness of the signals. Matějka and McKay (2015) shows that in that framework agent's choices can be formulated as a generalized multinomial logit, and Steiner et al. (2017) give a tractable formulation for the choices of agents with endogenous information acquisition in a dynamic setting. Our approach differs in that our model uses a different cost structure, and focuses on the interaction between information acquisition and market mechanisms.

A related question is the communication complexity of transmitting known preference to a mechanism. Gonczarowski et al. (2015) consider the communication complexity of finding a stable matching and show that it requires $\Omega(n^2)$ boolean queries. Ashlagi et al. (2018) find that the communication complexity of finding a stable matching is low under assumptions on the structure of the economy and a Bayesian prior. Their Communication-Efficient Deferred Acceptance protocol utilizes messages about both acceptances and rejections. The analysis in both papers differs from ours in that they assume agent know their full preferences (for example, can report their first choice) and only consider the cost of communicating that information to the mechanism.

Finally, our work contributes to the growing number of papers exploring the use of sequential or multi-round school choice mechanisms. Bo and Hakimov (2017) and Ashlagi et al. (2018) propose the Iterative Deferred Acceptance mechanism (IDAM) and Communication-Efficient Deferred

Acceptance mechanism respectively, which allow for multiple rounds of message-passing where students can learn the set of schools with which they are likely to be matched. Such mechanisms are also currently used in practice; Dur et al. (2015) empirically study a public school system in Wake County that implements an iterative mechanism, Gong and Liang (2016) theoretically and empirically consider a college admissions system in Inner Mongolia that implements an iterative version of Deferred Acceptance, and Bo and Hakimov (2017) propose IDAM in response to a sequential mechanism previously used for college admissions in Brazil. We hope that our findings can be used to better design sequential school choice mechanisms in these cities, and many others across the world.

5.2 Model

We present a model where students learn their preferences through costly information acquisition. The set of schools is denoted by $\mathcal{C} = \{1, \dots, n\}$, and each school $i \in \mathcal{C}$ has capacity to admit $q_i > 0$ students. A student is given by a quadruple $s = (F^s, c^s, r^s, v^s)$. School priorities are publicly known, and captured by the vector $r^s \in [0, 1]^{\mathcal{C}}$. School i prefers student s over student s' if and only if $r_i^s > r_i^{s'}$. We say that r_i^s is the rank of student s at school i . Student s needs to perform costly information acquisition to learn her value for attending each school. Initially student s knows that the value for attending school i is distributed according to prior F_i^s , and may pay a inspection cost of $c_i^s > 0$ to learn the realized value v_i^s . Student s privately knows F^s, c^s (importantly, the designer does not know these parameters). Students must inspect a school in order to attend it. We assume that v_i^s is independently drawn across students and schools.¹

With slight abuse of notation, we use a student type $\theta = \theta(s) = (F^\theta, c^\theta, r^\theta)$ to denote the initially known information of a student $s = (F^\theta, c^\theta, r^\theta, v^s)$. We refer to $\theta \in \Theta$ as a student type, and refer to $s = (\theta, v^s) \in \mathcal{S}$ as the student's realized preference. Formally, $\Theta = \mathcal{F}^{\mathcal{C}} \times \mathbb{R}^{\mathcal{C}} \times [0, 1]^{\mathcal{C}}$, where \mathcal{F} is the set of probability distribution functions, and $\mathcal{S} = \Theta \times \mathbb{R}^{\mathcal{C}}$. We will use s and θ interchangeably to index $F^\theta, c^\theta, r^\theta$. Given a type θ the realized values are randomly distributed

¹This implies that preferences of other students do not provide a student any information about v^s .

$v^s \sim F^\theta$, and with slight abuse of notation we write $s \sim F^\theta$.

Definition 5.1. An *discrete economy* is given by $E = (\mathcal{C}, S, q)$, where $S = \{s_1, \dots, s_N\}$ is the set of students and $q = \{q_i\}_{i \in \mathcal{C}}$ is the vector of quotas at each school.

We make the following assumptions. First, all students and colleges are acceptable. Second, as r_i^s carries only ordinal information, it is normalized to be equal to the percentile rank of student s in college i 's preferences, i.e. $r_i^s = |\{s' \mid s \succ_c s'\}| / |S|$. Third, schools have strict priorities, i.e., $r_i^s \neq r_i^{s'}$ if $s \neq s'$. Fourth, the priors F^θ are such that students have strict preferences, i.e., $\mathbb{P}(v_i^s = v_i^{s'}) = 0$ for all s and $i \neq i'$. Last, we assume there is an excess of students, that is, $\sum_{i \in \mathcal{C}} q_i < |S|$.

It will be useful to consider continuum economies where there is no aggregate uncertainty. The realized preferences v^s of a single student given his type $\theta(s)$ are random. In the continuum economy there is a continuous mass of students of any given type θ , and although the realized preferences of an individual student are random, the aggregate distribution over $s = (\theta, v)$ is known from the initial information F^θ . Formally, a continuum economy is described by a measure η over \mathcal{S} . We require that the measure η is consistent with initial information, that is, for any $A \subset \Theta$ and sets $V_i \subset \mathbb{R}^C$ we have that

$$\eta(\{(\theta, v) \mid \theta \in A, v_i \in V_i\}) = \int_{\theta \in A} \int_{v \in V_1 \times \dots \times V_n} dF^\theta(v) d\eta(\theta).$$

Definition 5.2. A *continuum economy* is given by $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$, where $q = \{q_i\}_{i \in \mathcal{C}}$ is the vector of quotas at each school, and η is a probability measure over \mathcal{S} that is consistent with initial information.

We make the same assumptions about continuum economies as for finite economies: namely that all students and colleges are acceptable; r_i^s is normalized so that for any $i \in \mathcal{C}$ and $x \in [0, 1]$, we have that $\eta(\{(\theta, v) \in \mathcal{S} \mid r_i^\theta \leq x\}) = x$; school priorities are strict, i.e. for any $x \in [0, 1]$ we have $\eta(\{(\theta, v) \in \mathcal{S} \mid r_i^\theta = x\}) = 0$; student preferences are strict, i.e. for any $x \in [0, 1]$ we have $\eta(\{s = (\theta, v^s) \in \mathcal{S} \mid v_i^s = x\}) = 0$; and there is an excess of students, $\sum_{i \in \mathcal{C}} q_i < \eta(\mathcal{S}) = 1$. In what follows, we will define concepts for both the discrete and continuum economy, and let $\eta(\cdot)$ denote the cardinality of a set in the discrete economy, and the measure in the continuum economy.

As in the standard matching model, a *matching* is a mapping $\mu : \mathcal{S} \rightarrow \mathcal{C} \cup \{\emptyset\}$ specifying the assignment of each student. Overloading notation, for school $i \in \mathcal{C}$ let $\mu(i)$ denote the set $\mu^{-1}(i) \subseteq \mathcal{S}$ of students assigned to school i . For each student $s \in \mathcal{S}$ and school $i \in \mathcal{C}$, let the *inspection indicator* χ_i^s be an indicator function that is 1 if student s has inspected school i and 0 otherwise.² We denote the preference information revealed from inspections χ by $v|_\chi = \{v_i^s \mid \chi_i^s = 1\}$.

A matching μ is *feasible* with respect to inspections χ if for each school $i \in \mathcal{C}$ we have that $\mu(i)$ is η -measurable and $\eta(\mu(i)) \leq q_i$, and, for each student s , if $\mu(s) \neq \emptyset$ then then $\chi_{\mu(s)}^s = 1$. This last condition is tantamount to assuming that a student must inspect a school in order to attend it. A *feasible outcome* is a matching and inspection pair (μ, χ) such that μ is feasible with respect to χ . Given (μ, χ) the utility of student s is $u^s(\mu, \chi) = v_{\mu(s)}^s - \sum_{i \in \mathcal{C}} \chi_i^s c_i^s$.

5.2.1 Stability with Costly Information Acquisition

Consider a feasible outcome (μ, χ) . As in the complete information setting, a student-school pair (s, i) forms a blocking pair if: (i) student s has higher priority than some student who is assigned to s or school s did not fill its capacity, namely $r_i^s > \inf \{r_i^{s'} \mid s' \in \mu(i)\}$ or $\eta(\mu(i)) < q_i$; and (ii) student s inspected school i and knows she prefers school i over her assigned school $\mu(s)$, namely $\chi_i^s = 1$ and $v_i^s > v_{\mu(s)}^s$.³ When information acquisition is costly for students there may be a student-school pair (s, i) where (i) holds and student s did not inspect school i . We extend the standard definition and say that (s, i) forms a blocking pair if (i) holds; and (ii') s has not yet inspected school i and prefers to pay the inspection cost c_i^s and be assigned to the better school of i and $\mu(s)$, namely $\chi_i^s = 0$ and $\mathbb{E}_{\tilde{v}_i^s \sim F_i^s} [\max \{v_{\mu(s)}^s, \tilde{v}_i^s\} - c_i^s] \geq v_{\mu(s)}^s$. An outcome (μ, χ) is stable if there are no pairs (s, i) that block by satisfying either (i),(ii) (i.e. the classical stability condition) or (i),(ii').

We remark that stability of (μ, χ) depends only on student's initial information $\theta(s)$ and preferences revealed by inspections $v|_\chi$. Simple examples show that if $\chi \neq \chi'$ are different inspections

²We are implicitly assuming that two students with the same type and values inspect the same schools.

³Recall that feasibility requires that $\chi_{\mu(s)}^s = 1$.

and μ is a matching, it is possible that the outcome (μ, χ) is stable but the outcome (μ, χ') is not.⁴

Given a matching μ define the *budget set* of student s by

$$B^s(\mu) = \left\{ i \in \mathcal{C} \mid r_i^s \geq r_i^{s'} \text{ for some } s' \in \mu(i) \right\} \cup \left\{ i \in \mathcal{C} \mid \eta(\mu(i)) < q_i \right\}.$$

The budget set $B^s(\mu)$ is the set of schools such that (s, i) satisfy condition (i). A stable outcome (μ, χ) must assign student s to a school $i \in B^s(\mu)$ if s so desires, and the student s cannot be assigned to any school in the complement set $\mathcal{C} \setminus B^s(\mu)$. We say that a school i is *available to s* if $i \in B^s(\mu)$, otherwise school i is unavailable to s . The following straightforward lemma characterizes stable outcomes in terms of budget sets.

Lemma 5.1. *A feasible outcome (μ, χ) is stable if and only if for every student we have that*

$$\mu(s) = \arg \max_{i \in \mathcal{C}} \{v_i^s \mid i \in B^s(\mu), \chi_i^s = 1\},$$

and for any $i \in B^s(\mu)$ such that $\chi_i^s = 0$ we have that

$$\mathbb{E}_{\tilde{v}_i^s \sim F_i^s} \left[\max \{v_{\mu(s)}^s, \tilde{v}_i^s\} - c_i^s \right] \leq v_{\mu(s)}^s.$$

5.2.2 Regret-Free Stable Outcomes

To reach an outcome, students must perform inspections to acquire information about their values. These inspections might induce regret. Sometimes this regret is unavoidable: e.g., a student will regret having inspected a school with low value. Other times, regret is avoidable: e.g. a student will regret inspecting a school that is not available to her, or inspecting schools in the wrong order. In other words, a student should carefully select her inspections based on her available information. Below we characterize the information acquisition process that maximizes the student's expected payoff given all potential information, including her initial information and information that could

⁴For example, if there are only two schools, both of which are very costly to inspect compared to the possible values they may yield, then a student who has inspected and is matched to the first but has not inspected the second (χ) might not wish to pay the inspection cost for the second school, causing the current matching to be stable. However, if she had inspected the second school (χ'), she may realize a high value for it and thus form a blocking pair with it.

be provided by the market. We determine how market information can affect the student's information acquisition decision. This allows us to define regret-free stable matchings where agents acquired information optimally.

Consider a student s who possesses initial information $\theta(s) = (F^s, c^s, r^s)$ and needs to select which schools to inspect. Since inspections are costly, student s will want to inspect a school only if inspecting the school can lead to being assigned to that school and receiving a higher value. In particular, student s will not want to inspect a school i if she knows that school i filled its capacity with higher priority students, and therefore she will not be assigned to the school i regardless of her value v_i^s . Thus, the set of schools that student s would like to inspect depends on her potential matches and the preferences of other students.

To fix ideas, first consider the isolated information acquisition problem where student $s = (\theta, v^s)$ is given a subset of schools $C \subseteq \mathcal{C}$ to choose from, each of which guarantees her admission. Student s needs to acquire information to form her preferences and then select her assigned school from C . If χ^s is s 's inspection indicator and $i^* \in C$ is her selected school her utility is $v_{i^*}^s - \sum_{i \in C} \chi_i^s c_i^s$. The adaptive inspection strategy that maximizes the student's expected utility given the initial information F^θ is derived by Weitzman (1979) and is stated in the following lemma.

Lemma 5.2. (Weitzman 1979) Consider a student $s = (\theta, v^s)$ with initial information F^θ and inspection costs c^θ that can adaptively inspect schools and choose a school from $C \subseteq \mathcal{C}$. For each school i , define a index \underline{v}_i^θ to be the unique solution to the equation $\mathbb{E}_{\tilde{v}_i \sim F_i^\theta} [\max\{0, (\tilde{v}_i - \underline{v}_i^\theta)\}] = c_i^\theta$. Sequentially inspect schools one by one in decreasing order of their index \underline{v}_i^θ .⁵ Continue inspecting the following school until the score of the next school to be inspected is below the maximal realized value among inspected schools.⁶

We denote the inspections resulting from this optimal strategy by $\chi^{opt}(F^\theta, c^\theta, v^s; C)$.

The optimal inspection policy is an index policy, where students use the prior information F^θ

⁵In case of multiple schools with equal index $\underline{v}_i^\theta = \underline{v}_{i'}^\theta$, break the tie by first inspecting the school $\min\{i, i'\}$.

⁶That is, if the set of inspected schools is $I = \{i \mid \chi_i^s = 1\}$ then inspect $j^* = \operatorname{argmax}_{j \in C \setminus I} \{\underline{v}_j^\theta\}$ if $\underline{v}_{j^*}^\theta > \max_{i \in I} v_i^s$ and stop otherwise.

to compute indices v_i^θ for each school i and inspect schools in decreasing order of their index.⁷ The set of inspected schools depends on the set of available schools C , the indices $\{v_i^\theta\}_{i \in C}$, and the realized values $\{v_i^s\}_{i \in C}$. The following example illustrates this.

Example 5.1. Suppose that $C = \mathcal{C} = \{1, 2\}$. Let $[x; p]$ denote the probability distribution which assigns probability p to the value x and $1 - p$ to 0. Consider a student with $v_1 \sim F_1 = [10; 1/2]$, and $v_2 \sim F_2 = [4; 3/4]$ and let the inspection costs be $c_1 = 3, c_2 = 1$. Then the optimal inspection strategy is to first inspect school 1, and continue to inspect school 2 only if $v_1 = 0$. If instead $C = \{2\}$ the optimal inspection strategy is to only inspect school 2.

Knowing the set of available schools C helps the student in Example 5.1 to conduct the adaptive information acquisition that maximizes her expected utility. If the student does not know C she her inspection strategy may be sub-optimal in two ways. First, the student may inspect school 1 when it is not available, wasting the cost c_1 . Second, the student should not inspect school 2 before she inspects school 1 or learns that school 1 is not available, because it is likely that she will not choose to inspect school 2 after inspecting school 1.

When student s is part of a matching market, the set of schools that are available to her depends on the resulting matching outcome, and therefore on the preferences of other students. Suppose that student s were to delay her information acquisition until the rest of the market resolved and the matching μ is realized. Arriving last to the market, student s can learn the set of schools available to her, which is $B^s(\mu)$ by Lemma 5.1. The student can optimize her information acquisition by using her initial information F^θ, c^θ as well as the market information $B^s(\mu)$, and applying Lemma 5.2. We say that the outcome (μ, χ) is regret-free stable if every students follows the optimal inspection policy informed by all available market information, that is, every student inspected schools as if she was the last to the market.

Definition 5.3. An outcome (μ, χ) is *regret-free stable* if (μ, χ) is stable and every student s inspected the optimal set of schools given her available set of schools $B^s(\mu)$, that is $\chi^s = \chi^{opt}(F^s, c^s, v^s; B^s(\mu))$ for all $s \in S$. We let $M^{RF}(E)$ denote the set of regret-free stable outcomes for the economy E .

⁷Such a policy can also be constructed by mapping the problem to a multi-armed bandit (MAB) problem; see e.g Olszewski and Weber (2015) for details.

When an outcome is not regret-free stable some students can benefit from delaying their information acquisition until the remainder of the market resolved. Note that while the definition of regret-free stability is ex post in flavor, as it is stated in terms of each student's realized preferences v^s , it only imposes the restriction that the student follows the ex ante optimal inspection strategy given θ and $B^s(\mu) = B^\theta(\mu)$ (before observing v^s). A regret-free stable outcome could be ex post suboptimal, e.g. a student s may inspect a school $i \neq \mu(s)$ with low realized value v_i^s and ex post observe that the inspection cost c_i was wasted, but student s could do no better given all available information from $\theta(s)$ and the market information.

Remark. To verify whether (μ, χ) is regret-free stable it is sufficient to know $v^s|_{\chi}, \chi^s$ and F^s, c^s for each s , and the students' values for uninspected schools is not necessary.

5.3 The Structure of Regret-Free Stable Outcomes

In this section we provide several results about the structure of regret-free stable outcomes. We show that the set of regret-free stable outcomes is a non-empty lattice and give a concise characterization of regret-free stable outcomes in terms of cutoffs.

We begin by exploring how the demand of student s depends on the set of available schools. Consider a student s with available schools $C \subset \mathcal{C}$. If s optimally acquires information, she inspects $\chi^s = \chi^{opt}(F^\theta, c^\theta, v^s; C)$. Denote the resulting demand of s by

$$D^s(C) = \arg \max \{v_i \mid i \in C, \chi^{opt}(F^\theta, c^\theta, v^s; C) = 1\} \in C,$$

which is the most preferred inspected school. Note that $D^s(C)$ depends only on information that is revealed to s . The following lemma shows that $D^s(\cdot)$ satisfies WARP, and we can construct a full preference ordering $\succeq^{\Psi(s)}$ that yields the same demand.

Proposition 5.1 (Reduction to demand from complete information). *Let $s = (F^s, c^s, r^s, v^s)$ be a*

realized student. There exist an ordering $\succeq^{\Psi(s)}$ such that for all $C \subset \mathcal{C}$ we have that

$$D^s(C) = \max_{\succeq^{\Psi(s)}}(C).$$

Proof. Using the indices from Lemma 5.2, define $i \succeq^{\Psi(s)} j$ if and only if $\min\{v_i^\theta, v_i^s\} \geq \min\{v_j^\theta, v_j^s\}$. It is straightforward to verify that $D^s(C) = \max_{\succeq^{\Psi(s)}}(C)$. \square

That is, if we only observe the eventual selection from a set of available schools C , the student s is indistinguishable from a student with complete preference information and preferences $\succeq^{\Psi(s)}$. Given only initial information θ , the demand of θ from a set C is uncertain, as the realized values v^s are unknown. An immediate corollary is the distribution of demand of θ from a set C is identical to the distribution of $\arg\max_{\succeq^{\Psi(s)}}(C)$, where $\succeq^{\Psi(s)}$ is the random preference ordering induced by drawing a random student $s = (F^\theta, c^\theta, r^\theta, v^s)$ from the distribution $v^s \sim F^\theta$.

Proposition 5.1 also allows a characterization of regret-free stable outcomes in terms of cutoffs, as in the complete information model of (Azevedo and Leshno, 2016). Cutoffs $P = \{P_i\}_{i \in \mathcal{C}} \in \mathbb{R}^{\mathcal{C}}$ are admission thresholds for each school. Cutoffs \mathbf{P} determine the budget set of a student s to be

$$B^s(\mathbf{P}) = \{i \in \mathcal{C} \mid r_i^s \geq P_i\},$$

which is the set of schools where s has better rank than the cutoff at that school. Note that $B^s(\mathbf{P})$ depends only on r^s and can be calculated from \mathbf{P} and the initial information $\theta(s)$.

The demand of student s given cutoffs \mathbf{P} is defined to be equal to $D^s(C)$ for a set of available schools equal to his budget set $C = B^s(\mathbf{P})$; for succinctness we will write this as $D^s(\mathbf{P})$. Note that within the definition of $D^s(\mathbf{P})$ we require that student acquire information optimally. Aggregate demand for school i given cutoffs \mathbf{P} is defined to be the mass of students that demand school i ,

$$D_i(\mathbf{P}) = D_i(\mathbf{P}|\eta) = \eta(\{s \in \mathcal{S} \mid D^s(\mathbf{P}) = i\}).$$

We define market-clearing cutoffs as in Azevedo and Leshno (2016) and show there is a one-to-one correspondence between market-clearing cutoffs and regret-free stable outcomes. Note that the

effect of information acquisition is captured within the definition of $D_i(\cdot)$.

Definition 5.4. A vector of cutoffs \mathbf{P} is *market-clearing* if it matches supply and demand for all schools with non-zero cutoffs:

$$D_i(\mathbf{P}) \leq q_i \text{ for all } i \text{ and } D_i(\mathbf{P}) = q_i \text{ if } P_i > 0.$$

We can now state our characterization of regret-free stable outcomes.

Theorem 5.1. *An outcome (μ, χ) is regret-free stable if and only if there exist market-clearing cutoffs \mathbf{P} such that for all s*

$$\mu(s) = D^s(\mathbf{P})$$

and

$$\chi^s = \chi^{opt}(F^\theta, c^\theta, v^s; B^s(\mathbf{P}))$$

Theorem 5.1 shows an equivalence between market clearing cutoffs and regret-free stable outcomes. Because demand $D(\cdot)$ provides us with sufficient information to determine whether \mathbf{P} are market clearing cutoffs, demand $D(\cdot)$ is also sufficient to determine whether a matching μ yields a regret-free stable outcome with some χ . Using Proposition 5.1, for any market with information acquisition $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ we can construct a full information economy E that has the same demand, and therefore the economy \mathcal{E} has the same market clearing cutoffs as E .

Theorem 5.2. *For every continuum economy \mathcal{E} there exists a regret-free stable outcome. Moreover, the set of regret-free stable outcomes (μ, χ) is a non empty lattice under the order \succeq defined by $(\mu, \chi) \succeq (\mu', \chi')$ iff $v_{\mu(s)}^s(\omega) \geq v_{\mu'(s)}^s(\omega) \forall s \in S$.*

Proof. The theorem follows from the reduction shown in Proposition 5.1 to the complete information setting, and analogous results by Blair (1988) on the lattice structure of many-to-one stable matchings in the complete information setting. \square

Uniqueness of the regret-free stable outcome will require that the distribution of student types

be regular. As student types have probabilistic demand, we will need to expand the definition of regularity beyond that found in Azevedo and Leshno (2016).

Definition 5.5. We say that $\theta = (F^\theta, c^\theta, r^\theta)$ is regular if for all $i \neq j$ we have that $\underline{v}_i^\theta \neq \underline{v}_j^\theta$ and $\mathbb{P}_{\tilde{v}_i^s \sim F_i^\theta}(\tilde{v}_i^s = \underline{v}_j^\theta) = 0$.

An measure η is *regular* if $\eta(\{s \mid \theta(s) \text{ is not regular}\}) = 0$ and the image under $D(\cdot \mid \eta)$ of the closure of the set $\{P \in (0, 1)^C \mid D(\cdot \mid \eta) \text{ is not continuously differentiable at } P\}$ has Lebesgue measure 0.

Intuitively, a type θ is regular if there are no ties, and so there is always a unique decision for whether to continue to inspect, and if so which school to inspect. A measure η is regular if there is no positive measure of irregular students and the implied demand is sufficiently smooth.

Theorem 5.3. *Suppose η is a regular measure. Then for almost every q with $\sum_i q_i < 1$ the economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ has a unique regret-free stable outcome.*

Proof. If η satisfies $\eta(\{s \mid \theta(s) \text{ is not regular}\}) = 0$ then for every cutoff P demand $D(P \mid \eta)$ is uniquely specified, and so there is a unique reduction to the complete information setting. The theorem follows from the reduction shown in Proposition 5.1 to the complete information setting, and analogous results by Azevedo and Leshno (2016) in this setting. \square

5.4 Mechanisms

To this point we have discussed properties of regret-free stable outcomes. We now turn to the process by which a market-maker might implement such outcomes. In general, the market arrives at an outcome (μ, χ) following a sequential process in which students provide information to the market, the market provides information to students, students inspect schools to obtain more information, and the process repeats. We can describe any such market procedure as a dynamic mechanism.

The mechanism relies on the information it receives from students. We will be interested in two kinds of mechanisms. First, we consider *direct mechanisms* in which students report all of

their private information and thereby delegate all decision-making. Second, we consider *choice mechanisms*, which are restricted in the nature of information that can be passed between the mechanism and the students. Choice mechanisms can only inform students about the availability of schools, and can only collect ordinal preference information from students.

5.4.1 Direct Mechanisms

In a direct mechanism the students fully delegate their decisions to the mechanism. We can think of a direct mechanism as the following iterative process. At any given state of the mechanism, we can write $\hat{\chi}$ to denote the indicator for the set of inspections the mechanism has conducted so far. Then, for each student s , the mechanism knows $v^s|_{\hat{\chi}}$ and knows F^s, c^s, r^s by assumption. Based on this information the mechanism can either decide to stop acquiring information and output the outcome $(\mu, \hat{\chi})$ for some matching μ , or to decide on the behalf of some students to inspect additional schools. We denote the information available in economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ after inspections χ by $I_{direct}(\mathcal{E}, \chi) = (\nu, q, v|_{\chi}, \chi)$, where ν is defined by $\nu(A) = \eta(\{s \mid \theta(s) \in A\})$ for $A \subset \Theta$. We denote that set of all possible inspection indicators by \mathcal{X} and use $0 \in \mathcal{X}$ to denote the initial state where no student has inspected any school. Let \mathcal{I}_{direct} denote the collection of all possible information sets $I_{direct}(\mathcal{E}, \chi)$.

Definition 5.6. A *direct mechanism* M is a mapping

$$M : \mathcal{I}_{direct} \rightarrow (2^{\mathcal{S}} \times \mathcal{C}) \cup (\mathcal{S}^{\mathcal{C}}, \chi)$$

that takes as input all the information available to the mechanism given previous inspections and returns either a next step of the inspection process, as described by tuple (S, i) of a set $S \subseteq \mathcal{S}$ of students to inspect the school $i \in \mathcal{C}$, or a final outcome (μ, χ) where χ is the current inspection indicator. To ensure termination of the mechanism, we require that iterated applications of the mechanism starting with $I_{direct}(\mathcal{E}, 0)$ will ultimately produce an outcome (μ, χ) , which is the outcome of the mechanism.⁸

⁸More formally, the mapping M induces a mapping $M' : \mathcal{I}_{direct} \rightarrow \mathcal{I}_{direct}$ defined by $M'(I_{direct}(\mathcal{E}, \chi)) =$

Imposing that the mechanism produces a regret-free stable matching ensures that the mechanism makes inspection decisions that are aligned with the optimal solution to each student's single-agent inspection problem.

Definition 5.7. A mechanism M is (*student-optimal*) *regret-free stable* if for any economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ the mechanism outputs a regret-free stable outcome (μ, χ) .

5.4.2 Choice-Based Mechanisms

Direct mechanisms require that students directly report their initial information and all inspected values. However, students may not to communicate detailed cardinal information about their priors and costs. This may preclude the use of direct mechanisms in practice. We therefore consider mechanisms with lower communication requirements, where students provide only information about their preferred choice(s) from given sets of schools.

A choice-based mechanism is an iterative process where the mechanism provides information to students, students choose which schools to inspect, provide information back to the mechanism, and so on. Choice-based mechanisms do not have access to students' private information, and therefore cannot directly inform students which schools they should inspect. The mechanism can only provide information to students about which schools are available to them. Because we are interested in producing regret-free stable outcomes, which do not depend on students' beliefs about other students' preferences, we restrict our attention to mechanisms that only inform a student whether (i) it is certain that school i is available to her (A), (ii) it is certain that school i is unavailable to her (R), or (c) it is uncertain whether school i is available or not (W). We use Accept (A), Reject (R) and Wait-list (W) to denote these three possible messages.

Students receiving an AWR message can choose which schools to inspect, and inform the mechanism of their choices. To simplify notation, we write the response of the student as a refinement of a preference ordering. Given s, χ^s define $\succ^{s|\chi^s}$ by $i \succ^{s|\chi^s} i'$ if $\chi_i^s = \chi_{i'}^s = 1$ and $v_i^s > v_{i'}^s$,

$I_{direct}(\mathcal{E}, \chi')$, where: if $M(I_{direct}(\mathcal{E}, \chi)) = (S, i)$ then we let χ' be the inspections after the students in S have inspected school i , i.e. $(\chi')_j^s = 1 \Leftrightarrow (\chi_j^s = 1 \text{ or } s \in S, j = i)$; and if $M(I_{direct}(\mathcal{E}, \chi)) = (\mu, \chi)$ then we let $\chi' = \chi$. It is sufficient to require that if $I = (\nu, q, v|_{\chi}, \chi)$ is a fixed point of the mapping M' then $M(I) = (\mu, \chi)$ for some matching μ and the same inspections χ .

and $i \sim^{s|\chi^s} \phi$ if $\chi_i^s = 0$. Here we are using symbol ϕ to denote non-inspected schools. Let $\mathcal{L}(\mathcal{C} \cup \{\phi\})$ denote all transitive relations over \mathcal{C} and the non-inspection symbol ϕ . For an economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ and χ an inspection indicator, the information available to a choice-based mechanism is $I_{choice}(\mathcal{E}, \chi) = \left(\left\{ \succ^{s|\chi^s}, r^s \right\}_{s \in \mathcal{S}}, q, \chi \right)$. Let \mathcal{I}_{choice} denote the collection of all possible information sets $\mathcal{I}_{choice}(\mathcal{E}, \chi)$.

Definition 5.8. An *Accept-Waitlist-Reject (AWR) mechanism* M is defined via a mapping

$$M : \mathcal{I}_{choice} \rightarrow \left(\{A, W, R\}^{\mathcal{S} \times \mathcal{C}} \right) \cup \left(\mathcal{S}^{\mathcal{C}}, \chi \right)$$

that takes all the information available to the mechanism given previous inspections and returns either a AWR message for each student about each school, or a outcome (μ, χ) where χ is the current inspection indicator. We require that iterated applications of the mechanism starting with $\mathcal{I}_{choice}(\mathcal{E}, 0)$ ultimately produces an outcome (μ, χ) , which is the outcome of the mechanism.

We formally define general mechanisms, choice-based mechanisms and AWR mechanisms as dynamic games of incomplete information in the appendix.

5.5 Implementing Regret-Free Stable Matchings

We have shown that regret-free stable matchings inherit the lattice structure of of stable matchings in the complete information setting, and can also be characterized using market-clearing cutoffs. In this section, we explore the mechanism design problem of implementing regret-free stable outcomes. We first show that regret-free stable matchings can be implemented by posting market-clearing cutoffs, and that information about these cutoffs is sufficient for regret-free information acquisition.

We then show that in the incomplete information setting, the difficulties lie not in the existence of regret-free stable matchings, but in computing and verifying the stability and optimality of these matchings in a regret-free manner. While standard mechanisms popularized in the complete information setting can discover the market-clearing cutoffs, in many markets they will necessarily incur regret. This is because such mechanisms rely on students gathering and reporting information about their preferences and can result in *information deadlocks*, where no information is gathered

because every student waits for others to acquire and report information first. Moreover, even when these mechanisms discover a regret-free stable matching, they will not be able to check if the matching is student-optimal without incurring regret.

Our conclusion is that information acquisition problems can be mitigated by posting market-clearing cutoffs. Cutoffs provide each agent with sufficient information to perform their inspections in a regret-free stable manner. The natural question, then, is how the market designer should determine which cutoffs to post. Market-clearing cutoffs can be learned and posted by the market designer without incurring regret if there is sufficient information about aggregate demand, either from historical data or from structured demand. We also show that even if market-clearing cutoffs can only be approximated, this is sufficient to implement a matching that is regret-free stable with respect to capacities that are close to the true capacities. Hence learning and posting cutoffs allows us to break the information deadlock and reach a regret-free stable outcome.

5.5.1 All You Need are Cutoffs

Recall from Theorem 5.1 that an outcome (μ, χ) is regret-free stable if and only if there exist market-clearing cutoffs P such that (a) each student follows the Weitzman optimal inspection strategy over her budget set as described by P , and (b) each student is matched to the school in her budget set that is most preferred, given the information revealed by the aforementioned optimal inspection strategy. Note, then, that if a student knows her budget set in advance, then she can optimally solve her information acquisition problem by proceeding as in a single-agent Pandora's Box problem to resolve her own incomplete information. An implication is that any matching mechanism that proceeds by committing to a collection of market-clearing acceptance cutoffs for the schools, then allowing each student to unilaterally optimize her inspection strategy and select her most-demanded school, will necessarily result in a regret-free stable match.

Theorem 5.4. *Let $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ be a continuum economy, and let \mathbf{P} be the student-optimal market-clearing cutoffs in \mathcal{E} . Then Mechanism 2 is regret-free stable.*

Proof. We show that Mechanism 2 produces the student-optimal regret-free stable matching when all students report truthfully, and hence truthful reporting is a Nash equilibrium that produces

Mechanism 2 Acceptance with Market-Clearing Cutoffs (AwMC)

1: **procedure** AwMC($\mathcal{C}, \mathcal{S}, q, P$)

Message Passing from Platform to Students

2: **for** $s \in \mathcal{S}$ **do**

3: **for** $i \in \mathcal{C}$ **do**

4: **if** $r_i^s \geq P_i$ **then**

5: Send message ‘ i accepts’ to s

Message Passing from Students to Platform

6: **for** Student s in \mathcal{S} **do**

7: Student s reports top choice school i^s that accepted them

8: $\mu(s) \leftarrow i^s$

9: **return** μ

the student-optimal regret-free stable matching. Indeed, the mechanism presents each student s with their budget set $B^\mu(s) = \{i \in \mathcal{C} \mid r_i^s \geq P_i(\eta)\}$, and student s is guaranteed to be matched to their reported favorite school $i^s \in B^\mu(s)$. Thus each student is presented precisely the single agent problem on $B^\mu(s)$. Solving this problem yields inspection strategy $\chi^s = \chi^{OPT}$ (by definition of χ^{OPT}), followed by truthfully reporting the students true favorite school: $i^s = D^s(P^*(\eta))$. By construction, demand exactly matches supply under this truthful reporting, so the output μ is the student-optimal regret free stable matching for $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$. \square

This result states that advance knowledge of market-clearing cutoffs are sufficient for regret-free stability. Indeed, posting cutoffs in advance of any information acquisition removes all uncertainty on the part of the agents about which schools they could match with. This, in turn, removes the possibility of regretting one’s choice to explore the value of a match on the grounds that this school was ultimately unattainable. We note that this lack of regret does not depend on the posted cutoffs being market-clearing, but only that the mechanism commits to honoring the implied budget set for each student. Thus, for any economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ and any (not necessarily market-clearing) cutoffs P , there exists a choice of capacities q' such that P are the student-optimal market-clearing cutoffs in $\mathcal{E}' = (\mathcal{C}, \mathcal{S}, \eta, q')$, and hence Mechanism 2 is regret-free stable with respect to \mathcal{E}' . We will make use of this fact when discussing notions of approximation in Section 5.5.3.

5.5.2 Regret-Free Choice-Based Mechanisms and Information Deadlock

Theorem 5.4 shows that knowing market-clearing cutoffs in advance of the market mechanism is sufficient for implementing a regret-free stable matching. Knowing market-clearing cutoffs in advance can be strong requirement. Indeed, aggregate uncertainty about agents' demands might make it difficult for the mechanism designer to know this information before interacting with the students. One might hope to avoid this impasse by way of a mechanism that reaches a stable matching without necessarily determining each student's full budget set. After all, each student's demand is ultimately described by a single ordering over all schools that is consistent across budget sets, which seems to suggest that it might not be necessary to fully learn every student's budget set in order to find a stable match. However, even though the realized demand is described by a single consistent ordering, the student cannot know this ordering a priori precisely because it depends on the values, which are only revealed after costly exploration. In other words, while the realized demand is ordered consistently, the order in which a student would wish to explore depends crucially on her budget set, and hence revealing the budget set can be crucial for avoiding wasteful exploration and regret.

It is perhaps useful to once again consider the school-proposing DA mechanism, which we recall can be interpreted as discovering the market-clearing cutoffs over time. Initially only the highest-ranked students are admitted to the schools of their choice, consistent with implicit cutoffs that are initially high, and these cutoffs then decrease (i.e., lower-ranked students are accepted) until the market clears. This choice-based approach does not post cutoffs in advance, but rather discovers them through repeated interaction with students. This provides hope that a mechanism that proceeds in multiple rounds can elicit enough information to find appropriate market-clearing cutoffs in a regret-free manner.

Indeed, we will show that under certain sufficient conditions on student preferences, the following iterative implementation of the school-proposing Deferred Acceptance is regret-free stable. The key idea is that while students' information acquisition problems are interconnected and can create information deadlocks, school priorities also provide students with partial information, and this may be sufficient to both start and finish the information acquisition process.

Algorithm 5.1 (Iterative Deferred Acceptance.). *At each step:*

- *Each school i proposes to the top q_i students who have not yet rejected them.*
- *Each student (irrevocably) rejects some of the schools that have proposed to them.*

The algorithm terminates when no new proposals are performed, at which point all students are asked for their top choice school among all those that have proposed to them and which they have not rejected, and are assigned to school.

Theorem 5.5. *Suppose that at least one of the following conditions hold:*

1. *All students inspect all schools in their budget set and do not have indifferences, i.e. $\mathbb{E}[v_i^s] = \infty$ for all $s \in S$ and $i \in C$, and $\mathbb{P}(v_i^s = v_j^s) = 0$ for all $s \in S$ and $i \neq j$.*
2. *For all $m \geq \min_i q_i$ the same students occupy the top m ranks at all schools, i.e. for all $i, j \in C$, $\left\{s \mid r_i^s \geq 1 - \frac{m}{|S|}\right\} = \left\{s \mid r_j^s \geq 1 - \frac{m}{|S|}\right\}$. (Recall that r_i^s is normalized to be the percentile rank of student s in school i 's priorities.)*

Suppose all agents only perform regret-free inspections and report truthfully. Then Mechanism 5.1 almost surely implements the school-optimal regret-free stable matching.

Proof. The intuition is that our conditions guarantee that at all stages of proposal, there are students who have sufficient information about their budget set to both inspect some schools and reject some proposals. (This requires that students do not have indifferences in their preferences.) We show this formally for both cases.

Case (1). Suppose that all students inspect all schools in their budget set, i.e. condition (1) holds. First, assuming truthful reporting, it is regret-free for every student to inspect all schools that proposed to them. This is because if school i has proposed to student s , it is not full of students it prefers to s , so for all realized preferences v and for all outcomes $(\mu, \chi) \in M^{RF}(E)$ it holds that $i \in B^s(\mu)$. Since s is willing to inspect any school in her budget set it follows that it is optimal for student s to inspect school i .

Next, suppose students do not have indifferences in their preferences, i.e. $v_i^s \neq v_j^s$ for all s and $i \neq j$. Then it is regret-free for each student to reject all the schools that proposed to them except

the one with the highest observed value. This is because if student s has inspected both i and j and $v_i^s > v_j^s$ then $\mu(s) \neq j$ for all $\mu \in M^{RF}(v)$ and it is optimal for s to reject the school with lower observed value.

When the algorithm terminates either all schools are at capacity or all students are assigned, since during each step all students reject all the schools that proposed to them except one, and so the algorithm terminates with a regret-free stable outcome.

Case (2). Suppose that for all $m \geq \min_i q_i$ the same students occupy the top m ranks at all schools, i.e. condition (2) holds. Let $\underline{q} = \min_i q_i$. It follows that there is a set $S(0)$ of students who are in the top \underline{q} at all schools, and for all $m > \underline{q}$ there is a student s^m who is m -th ranked at all schools. Hence there is a unique regret-free stable matching, and Mechanism 5.1 essentially performs serial dictatorship and tells each student her (unique) budget set in order of her rank. When students don't have indifferences in their preferences, it follows that at each step some students reject all the schools that proposed to them except one, and so when the algorithm terminates it outputs a regret-free stable matching.

We show that there is a unique regret-free stable matching. Fix v and let $\mu, \mu' \in M^{RF}(v)$. Let $s^1, s^2, \dots, s^{\underline{q}}$ be an arbitrary ordering of the students in $S(0)$ and s^m be the m -ranked student for all $m > \underline{q}$. Then for all $m > \underline{q}$ we can define $B^{s^m}(\mu|v)$ by

$$B^{s^m}(\mu|v) = \left\{ i \mid \sum_{s'=s^{m'}, m' < m} \mathbf{1} \{ D^{s'}(B^{s'}(\mu)|v) = i \} < q_i \right\},$$

i.e. the budget set of s^m is the set of schools with residual capacity once all students ranked higher than s^m have chosen their school from their budget set. Hence by induction $B^s(\mu|v) = B^s(\mu'|v)$ for all s and so $\mu(s) = \mu'(s)$ for all s . \square

Theorem 5.5 demonstrates that for certain priorities and preferences iterative Deferred Acceptance, which is a choice-based mechanism, can discover market-clearing cutoffs in a regret-free stable manner. This mechanism is iterative, and one can show that this is necessary: even under the conditions laid out in Theorem 5.5, no one-shot mechanism — choice-based or otherwise — can be regret-free stable. We provide an example in Appendix C.2.1.

Furthermore, the matching found by the school-proposing DA mechanism is not student-optimal. Can one find a student-optimal matching in a regret-free manner? The original proof of Gale and Shapley of the existence of a student-optimal stable matching is constructive: they furnished an algorithm, the Gale-Shapley Deferred Acceptance (DA) algorithm, and demonstrated that it always finds the student-optimal stable matching in polynomial time. An analogous algorithm for identifying the student-optimal regret-free stable matching could also be defined in our incomplete information setting, but would require students to provide information about both their priors and values and may induce students to acquire information in a way that incurs regret. In fact, for many economies, verifying that the student-optimal regret-free stable matching is the student-optimal one necessitates incurring regret with positive probability in the inspection process. This is because in regret-free stable matchings students cannot inspect outside of their budget set, and in many economies with positive probability the student-optimal regret-free stable matching does not provide students with their full budget set. We provide an example that admits no student-optimal regret-free stable mechanism in Appendix C.2.2. The intuition behind the example is that the act of verification requires some student to perform more information acquisition than is allowed under the regret-free stable inspection policy, and if the existing matching is student-optimal it is then costly for them to acquire the necessary information. This mirrors the Grossman-Stiglitz paradox, whereby under costly information acquisition equilibrium market prices cannot be stable, as this would eliminate the benefit of acquiring this information (Grossman and Stiglitz, 1980).

We next show that the conditions of Theorem 5.5 are necessary, in the sense that there is no (even non-choice-based) mechanism that is regret-free stable for general economies. In the more general case where students can suffer regret by inspecting the schools out of order, it may be impossible for any (even multi-round) mechanism to find a regret-free stable matching without incurring regret. Perhaps more fundamentally, this example shows that any mechanism that converges to a stable matching in a regret-free manner (such as school-proposing DA) relies heavily on an assumption that there always exist students willing to inspect some schools in their budget set.

Theorem 5.6. *Let M be a mechanism. Then there exists an economy $E = (\mathcal{C}, \mathcal{S}, q)$ such that,*

when each student reports $I^s = (F^s, c^s)$ truthfully, with positive probability mechanism M does not implement a regret-free stable matching.

Remark. For convenience, we state and prove Theorem 5.6 for finite economies; we note that the result can be extended to continuum economies with minor adjustments.

Proof. Consider an economy E with three schools $\mathcal{C} = \{1, 2, 3\}$ with capacities $q_1 = q_2 = q_3 = 2$ and three students $S = \{x, y, z\}$.⁹

Suppose that school priorities are given by

$$\text{priority at 1 : } r_1^y > r_1^z > r_1^x$$

$$\text{priority at 2 : } r_2^z > r_2^x > r_2^y$$

$$\text{priority at 3 : } r_3^x > r_3^y > r_3^z.$$

and that student values at each school are $U[0, 1]$ variables, i.e. with priors $F_i^s(x) = x \forall x \in [0, 1]$ and student costs for inspection are given by $c_1^x = c_2^y = c_3^z = 0.1$, $c_2^x = c_3^y = c_1^z = 0.2$ and $c_3^x = c_1^y = c_2^z = 0.3$. Note that this means $\underline{v}_1^x = \underline{v}_2^y = \underline{v}_3^z = \sqrt{1 - 0.2} \approx 0.89$, $\underline{v}_2^x = \underline{v}_3^y = \underline{v}_1^z = \sqrt{1 - 0.4} \approx 0.77$ and $\underline{v}_3^x = \underline{v}_1^y = \underline{v}_2^z = \sqrt{1 - 0.6} \approx 0.63$ and so the order in which students $\{x, y, z\}$ wish to inspect schools is exactly the reverse of their priority at each school, e.g. student x wishes to inspect 1 then 2 then 3, and have bottom, middle and top priority out of $\{x, y, z\}$ at those schools respectively.

We will show that for all students s , there exists a school $i = \beta(s)$ such that, with positive probability, in every regret-free stable matching $\mu \in M^{RF}(E)$ student s only inspects i . Also, with positive probability, in every regret-free stable matching $\mu' \in M^{RF}(E)$ school i is not in student s 's budget set $B^{\mu'}(s)$. To see why this implies the theorem, note that under Mechanism M , one of x, y, z must be the first student in $\{x, y, z\}$ to perform an inspection with positive probability. Without loss of generality we may suppose that student is x . If x first inspects $\beta(x)$ then with positive probability, in any regret-free stable matching $\mu' \in M^{RF}(E)$ student x regrets her inspection. If x

⁹Note that strictly speaking, as we assumed that there are more students than seats, the economy should have seven students $S = \{x, y, z, d_1, d_2, d_3, d_4\}$ where the d_i are four dummy students who have lower priority at every school than the students in $\{x, y, z\}$ and who have arbitrary preferences. For simplicity we omit these students in the description of the economy; however note that the proof applies as written to both economies.

first inspects some school other than $\beta(x)$ then with positive probability for any regret-free stable matching $\mu \in M^{RF}(E)$ student x regrets her inspection. Hence with positive probability there exists a student who regrets her inspection process, and so with positive probability M does not implement any regret-free stable matching.

We now turn to proving the claim: for all students s , there exists a school $i = \beta(s)$ such that with positive probability every regret-free stable matching involves student s inspecting only i , and with positive probability no regret-free stable matching has school i in student s 's budget set. In particular, we show that $1 = \beta(x)$ satisfies the required properties. Note that since all priorities and costs are symmetric, the same arguments can be used to show that $2 = \beta(y)$ and $3 = \beta(z)$ satisfy the required properties.

Consider the event that $v_1^x, v_2^y, v_3^z \geq 0.9$, and $v_1^y, v_3^y, v_1^z, v_2^z \leq 0.5$. Note that it then holds that $D^x(\mathcal{C}) = 1$, $D^y(\mathcal{C}) = 2$ and $D^z(\mathcal{C}) = 3$. Now it is easy to check that for all $\omega \in X$ the only regret-free stable matching $\mu \in M^{RF}(E)$ is $(\mu(x), \mu(y), \mu(z)) = (1, 2, 3)$, since if any school i was not assigned to any of x, y, z (i.e. $\mu(i) \cap \{x, y, z\} = \emptyset$) then it would form a blocking pair with the student in $\{x, y, z\}$ whose top choice school is i . Hence $B^\mu(x) = \mathcal{C}$, so x inspects school 1 first, and since $v_1^x \geq 0.9 > \underline{v}_2^x, \underline{v}_3^x$ it follows that x only inspects school 1.

Next consider the event that $v_1^x, v_1^y, v_1^z \geq 0.9$ and $v_2^y, v_3^y, v_2^z, v_3^z \leq 0.5$. Note then that $D^y(\mathcal{C}) = D^z(\mathcal{C}) = 1$. Moreover, since $\underline{v}_1^y \geq \underline{v}_2^y, \underline{v}_3^y \geq 0.5 \geq \underline{v}_2^y, \underline{v}_3^y$ and y has top priority at 1 it follows that in any regret-free stable matching $\mu \in M^{RF}(E)$ student y inspects 1, and since $D^y(\mathcal{C}) = 1$ student y is assigned to 1, i.e. $\mu(y) = 1$. Since $q_1 = 2$ and z has second priority at 1 a similar argument shows that $\mu(z) = 1$. Hence for all regret-free stable matchings $\mu \in M^{RF}(E)$ it follows that 1 is full of students it prefers to x , and $1 \notin B^\mu(x)$.

Note that if no student has performed any inspections then we are unable to discern whether either these events is true, and for any student any inspection they perform will incur regret in either one event or the other, i.e. any inspections incurs regret with positive probability.

This example shows that there does not exist any mechanism that always finds a regret-free stable matching in a regret-free manner. This makes it even more surprising that, for any realization of preferences, the set of regret-free stable matchings $M^{RF}(E)$ not only has a student-optimal

member, but also inherits the lattice structure induced by the deterministic economy with students $\Psi(s)$. The intuition behind the difficulty in finding regret-free stable matchings is that in order to identify the appropriate deterministic preferences that rationalize demand under incomplete information, students need to know both their reservation values and their realized values. \square

To build additional intuition for Theorem 5.6, let us briefly demonstrate why Mechanism 5.1 might not be regret-free in a general economy. In the first round of Mechanism 5.1, every student s is proposed to by all schools except $\beta(s)$. When students were willing to inspect all schools in their budget set this was enough to induce some inspections and rejections. However, in general, there is an inspection *order* that maximizes the resulting expected payoff, and with positive probability students do not inspect some schools in their budget set. Hence even though every student knows some of the schools in their budget set, no student s wants to start inspecting schools until she knows for sure whether $\beta(s)$ is in her budget set. In other words, it is *strictly optimal* for each student to wait until the mechanism forces them to perform inspections.

This intuition illustrates a more general principle: in the presence of costly information acquisition, iterative mechanisms without an activity rule may result in an *information deadlock*, where no actions are taken because every agent can achieve higher utility if another agent acts first.

5.5.3 Regret-Free Learning

While it may not always be possible to discover market-clearing cutoffs through observed choice, the structure of regret-free stable matchings gives us hope that they can still be learned and implemented in an approximate manner. We show that when we have sufficient initial information or market structure, cutoff mechanisms that use estimated cutoffs implement outcomes that are regret-free stable with respect to slightly perturbed school capacities.

Before formalizing these ideas, we first turn to the following question: How do we estimate population demand? In matching markets with one-sided incomplete information, the preferences of the side with full information are a key source of information. For example, in Theorem 5.5, condition (2) guarantees that at any point in iterative Deferred Acceptance there are students who have full information about their budget set. For more general priority structures, there will be

students who have such information at the outset of the mechanism.

Definition 5.9. Let $E = (\mathcal{C}, S, q)$ be an economy. A student s has **free market information** if for all schools i student s is either in the top q_i percentile of students or the bottom $1 - \sum_j q_j$ percentile of students, i.e. $\forall i r_i^s \notin [1 - \sum_j q_j, 1 - q_i)$. We let $S^f(E) = \{s \mid \forall i r_i^s \notin [1 - \sum_j q_j, 1 - q_i)\}$ denote the set of students with free market information in E .

Knowing only their priors and school priorities, students with free market information can determine both their budget sets and their preferences in a regret-free manner.

Hence we may estimate market demand as follows. In some markets historical demand is sufficient for estimating population demand. For example, in college admissions in many countries aggregate student demand for different university courses do not vary much from year to year and historical demand can be used to estimate current demand. Even when such prior information is not available, as long as there are students with free information we can start learning about student preferences. For example, if running iterative Deferred Acceptance assigns some students before reaching a deadlock, the demand of the assigned students could be used to estimate the demand of the remaining students.

Estimated Cutoffs are Robust

We now formalize the claim that outcomes of cutoff mechanisms are robust to errors in estimated demand. The intuition behind these results is that demand with costly information acquisition satisfies WARP (Proposition 5.1), and so all questions about cutoff mechanisms under costly information acquisition reduce to analogous questions about cutoff mechanisms in markets without costly information acquisition.

We first show that the outcome (μ^P, χ^P) from posting cutoffs P when using Mechanism 2 is regret-free stable for capacities q^P that are slightly perturbed from the true capacities, and differs from the regret-free stable assignment under q for only a small number of students.

Theorem 5.7. *Let \mathcal{E} be a continuum economy, let μ be a regret-free stable matching for \mathcal{E} corresponding to market-clearing cutoffs P^* , and let $\bar{q} = D(P^*)$ be the measures of seats assigned under*

μ . Let (μ^P, χ^P) be the outcome of running Mechanism 2 on \mathcal{E} with cutoffs P , and for all i let $q_i^P = |\{s \mid \mu^P(s) = i\}|$. Then (μ^P, χ^P) is regret-free stable with respect to q^P , $\|q^P - \bar{q}\|_2 \leq \|P - P^*\|_2$, and $|\{s \mid \mu(s) \neq \mu^P(s)\}| \leq \|P - P^*\|_2$.

Proof. Note that by definition $q^P \equiv D(P)$. Now it is easy to see that (μ^P, χ^P) is regret-free stable with respect to $D(P)$. Moreover, each student's assignment $\mu(s)$ is equal to their demand $D^s(P)$, which is determined by their budget set $B^s(P)$. Finally, in moving from P to P^* only $\|P - P^*\|_2$ students receive different budget sets. The result follows. \square

When the error in the estimated cutoffs is due to sampling error, the outcome is regret-free stable for capacities that are normally distributed around the market-clearing demand.

Definition 5.10. For a capacity vector $q' = (q'_1, \dots, q'_n)^T$ we let $\Sigma^{q'}$ denote the matrix with entries

$$\Sigma_{ij}^{q'} = \begin{cases} -q_i q_j & \text{if } i \neq j \\ q_i (1 - q_i) & \text{if } i = j. \end{cases}$$

Proposition 5.2 (Distribution of approximately feasible capacities). *Suppose the continuum economy \mathcal{E} admits a unique stable matching μ with cutoffs P^* , and let $\bar{q} = D(P^*|\eta)$.¹⁰ Let $E^k = (\eta^k, q^k)$ be a randomly drawn finite economy, with k students drawn independently according to η , ω drawn independently, and where $q^k = D(P^*|\eta^k)$ is defined so that P^* is a market-clearing cutoff for E^k . Then*

$$\sqrt{k} \cdot (q^k - \bar{q}) \xrightarrow{d} \mathcal{N}(0, \Sigma^{\bar{q}}),$$

where $\mathcal{N}(\cdot|\cdot)$ denotes a C -dimensional normal distribution with given mean and covariance.

Proof. The result follows from the central limit theorem, as $D(P^*|\eta^k) = \frac{1}{k} \sum_{a=1}^k X_a$, where $X_a = D^\theta(P^*)$ is a random variable with $\theta \sim \eta$ capturing the demand of a single student drawn randomly from η and the X_a are independently drawn. \square

¹⁰Note that $\bar{q}_i = q_i$ for all overdemanded schools i , i.e. those such that $P_i^* > 0$.

Hence the mapping from cutoffs to demand is continuous, so approximate cutoffs yield regret-free stable outcomes for approximately feasible capacities. We similarly show in Appendix C.3.1 that the mapping from demand to market-clearing cutoffs is continuous, and estimated cutoffs are robust to errors due to sampling. Thus in order to obtain a desirable estimate of the market-clearing cutoffs P it suffices to furnish an accurate estimate of demand D .

Examples

These results suggest that if we use cutoff mechanisms based on estimated population demand, the resulting outcomes will be robust to small biases or noise due to sampling error. We illustrate this intuition in the following examples.

Example 5.2. In this example, we show how to implement an approximately regret-free stable matching in a setting with historical demand data. Suppose that this year's economy $E = (\mathcal{C}, S^k, \lfloor qk \rfloor)$ is given by drawing k students independently from a distribution η , and last year's economy $E^{hist} = (\mathcal{C}, S^{hist}, \lfloor \alpha qk \rfloor)$ is given by drawing αk students independently also from η for some fixed $\alpha > 0$. Then the student-optimal market-clearing cutoffs \hat{P} for the economy E^{hist} give an unbiased estimator for the student-optimal market-clearing cutoffs both for E and for $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$, and we can show that in this year's economy E posting \hat{P} implements a regret-free stable matching with respect to capacities $\hat{q}^k k$ that are close to qk . Specifically, if we let P^* be the market-clearing cutoffs of $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ and $\bar{q} = D(P^* | \eta)$ then we can use classic results about the convergence of two-step estimators¹¹ to show

$$\sqrt{k} (\hat{q}^k - \bar{q}) \xrightarrow{d} \mathcal{N} \left(0, \frac{(\alpha + 1)}{\alpha} \Sigma_{\bar{q}} \right),$$

for $\Sigma_{\bar{q}}$ defined as in Definition 5.10. The full proof can be found in Appendix C.3.2.

For large economies the capacities that make the outcome regret-free stable converge to the true capacities, and the variance depends only on \bar{q} and α . Moreover, in the absence of historical information ($\alpha = 0$) the cutoff mechanism can perform arbitrarily poorly, whereas more accurate

¹¹See, e.g. Newey and McFadden (1994) for details.

historical information ($\alpha \rightarrow \infty$) leads to smaller perturbations in the capacities.

Example 5.3. Suppose that this year's economy $E = (\mathcal{C}, S^k, qk)$ is given by drawing k students independently from a distribution $\eta(\Gamma^*)$, where student demand $D^s(P|\eta(\Gamma))$ is parametrized by $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$. Suppose also that some positive fraction α of students have free market information. Then by first obtaining preferences from the students with free market information, we can estimate Γ^* and provide an estimate \hat{P} for the student-optimal market-clearing cutoffs for E . We can also show that posting \hat{P} implements a regret-free stable matching with respect to capacities $\hat{q}^k k$ that are close to qk .

Formally, consider the mechanism M^F that runs in two rounds. In the first round it proposes to all students $s \in S^f(E)$, assigns them each to their chosen school and obtains their aggregate demand $\hat{q}_i^f k$ for each school i , and uses this demand to provide an estimate $\hat{\Gamma}$ for Γ^* . In the second round it runs Mechanism 2 with cutoffs \hat{P} on a residual economy E^r computed as follows. The cutoffs \hat{P} are the market-clearing cutoffs for an estimated residual economy $\hat{E}^r = (\mathcal{C}, S^r, \hat{q}^r k)$, where \hat{E}^r is given by drawing $k - |S^f(E)|$ students without free market information in E independently from the distribution $\eta(\hat{\Gamma})$ and the residual capacity for each school i is $\hat{q}_i^r = q_i - \hat{q}_i^f$. The residual economy $E^r = (\mathcal{C}, S^k \setminus S^f(E), \hat{q}^r k)$ is given by removing the students in $S^f(E)$ from E and reducing capacities accordingly at their assigned schools.

Let α denote the measure of students in E who have free market information, let $D_i^f(\Gamma)$ denote the proportion of students in $S^f(E)$ who demand school i as a function of Γ , and let $q_i^f = \alpha D_i^f(\Gamma^*)$ be the target first-round capacities. Define target capacities $\bar{q} = D(P^*|\eta(\Gamma^*))$ in terms of the market-clearing cutoffs P^* of $\mathcal{E} = (\mathcal{C}, S, \eta(\Gamma^*), q)$. We can show that Mechanism M^F implements a regret-free stable matching μ with respect to perturbed capacity $\hat{q}^k k$, where

$$\sqrt{k} (\hat{q}^k - \bar{q}) \xrightarrow{d} \mathcal{N} \left(0, \Sigma^{\bar{q}} + 2 \left(\frac{1}{\alpha} A + I \right) \Sigma^{q^f} A^T \right)$$

as $k \rightarrow \infty$ for $A = \nabla_{\Gamma} D(\Gamma^*) \left(\nabla_{\Gamma} D^f(\Gamma^*) \right)^{-1}$ and $\Sigma^{\bar{q}}, \Sigma^{q^f}$ defined as in Definition 5.10. (Note that $\alpha \leq \min_i q_i + 1 - \sum_i q_i \stackrel{def}{=} \alpha^*$, and $\alpha = \alpha^*$ is achieved when schools have aligned preferences, i.e. condition (2) in Theorem 5.5.) The idea is that since in the first round we assign only students

in $S^f(E)$, the budget set and demand of these students is that same whether we assign them in the first round, or in the second round after posting the cutoffs \hat{P} . Hence the outcome after both rounds is regret-free stable with respect to realized demand. Convergence and variance expressions can be derived using two-step GMM. The full proof can be found in Appendix C.3.2.

For large economies the capacities that make the outcome regret-free stable converge to the true capacities, and the variance depends on \bar{q} , q^f , $A = \nabla_{\Gamma} D(\Gamma^*) \left(\nabla_{\Gamma} D^f(\Gamma^*) \right)^{-1}$ and α . Moreover, in the absence of free market information ($\alpha = 0$) the cutoff mechanism can perform arbitrarily poorly, whereas priorities that yield more students with free market information ($\alpha \rightarrow \alpha^*$) or more accurate estimates of Γ^* ($A \rightarrow 0$) lead to smaller perturbations in the capacities. Finally, the first round in this mechanism corresponds to the first round of iterative school-proposing Deferred Acceptance. If we allow for further rounds of proposals, we can further reduce the noise in the perturbed capacities.

5.6 Discussion

Summary of findings. We have proposed regret-free stability as a suitable solution concept in matching markets with costly information acquisition. We have also shown that, surprisingly, regret-free stable matchings always exist and the set of regret-free stable matchings has a lattice structure. However, we have also shown that the effect of costly information acquisition is that it may be impossible to compute a regret-free stable matching in a regret-free manner, and that standard matching market mechanisms can result in information deadlocks. We have also provided some mechanisms for when we are willing to relax feasibility and provide varying amounts of information in order to achieve a regret-free outcome, and shown that for large economies they can be implemented by perturbing the capacities by $O(\sqrt{k})$ students, where k is the number of students in the market.

Approximation algorithms. Our results demonstrate that in general there is a tradeoff between the regret of a mechanism, the feasibility of the solution, and the amount of information provided to the mechanism. We have provided one class of mechanisms that relax the feasibility constraint

in order to achieve optimal regret. It may also be possible to relax the regret of the mechanism in order to achieve exact feasibility, or to increase the number of rounds of communication in order to better approximate both. We leave these questions open for future work.

Activity rules. We demonstrated that in the presence of costly information acquisition standard matching mechanisms can create situations where it is strictly optimal for every agent to wait for other agents to move first. This illustrates a more general principle, that in the presence of costly information acquisition iterative mechanisms will need an activity rule to converge. Another relevant question is what the appropriate design of activity rules is for such situations.

Stable matchings. We have concentrated our efforts on mechanisms that implement regret-free stable matchings. However, we have also provided a more general notion of stability in incomplete information settings. Is this more general space of outcomes predictive and does it have attractive structural properties? We selected the class of regret-free stable matchings as they compare each agent's utility only with her own utility under other information acquisition strategies. However it may be possible to improve social welfare by moving to a stable matching that transfers utility from one student to another. Are there stable matchings that are more desirable than the student-optimal regret-free stable matching? Our notion of stability under incomplete information can also be naturally extended to settings with two-sided incomplete information, as well as to settings with more general models for costly information acquisition, such as rational inattention models, or other models where agents may refine their priors for a cost. All of these questions become much more interesting in these general settings. We leave them open for future investigation.

Practical market design. Finally, what implications do our results have for practical applications? Colleges in many countries, such as China, India and Australia post historical cutoffs for admission into college programs. Our results on mechanisms with historical cutoffs suggests that if colleges capacities are flexible this can eliminate unnecessary preference formation by applicants. In Israel colleges post a pair of cutoffs for each program; students above the higher cutoff are guaranteed admission, students below the lower cutoff are advised to consider other options, students between the cutoffs are advised to wait for further information on enrolment for that year, and the cutoffs are updated as students register for programs. This very closely mirrors our

Accept-Waitlist-Reject mechanisms and suggests that they can be of practical use. Our result on information deadlock also brings to mind the behavior of participants reacting to activity rules such as deadlines and exploding offers in other markets. In markets such as job markets and Ph.D. admissions, participants often wait until the last minute before expressing their preferences. Clearly costly information acquisition is an important issue in many other markets, and we leave further investigation of the empirical and practical consequences for future work.

Bibliography

Bibliography

- A. Abdulkadirođlu. School choice. *The Handbook of Market Design*, 2013.
- A. Abdulkadirođlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- A. Abdulkadirođlu and T. Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88:233–260, 1999.
- A. Abdulkadirođlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- A. Abdulkadirođlu, P. Pathak, and A. Roth. The New York City high school match. *American Economic Review*, 95(2):364–367, 2005a.
- A. Abdulkadirođlu, P. A. Pathak, A. E. Roth, and T. Sönmez. The Boston public school match. *American Economic Review*, pages 368–371, 2005b.
- A. Abdulkadirođlu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *The American Economic Review*, pages 1954–1978, 2009.
- A. Abdulkadirođlu, Y.-K. Che, and Y. Yasuda. Expanding choice in school choice. *American Economic Journal: Microeconomics*, 7(1):1–42, 2015.
- A. Abdulkadirođlu, N. Agarwal, and P. A. Pathak. The welfare effects of coordinated assignment: Evidence from the New York City high school match. *American Economic Review*, 107(12):3635–89, 2017a.
- A. Abdulkadirođlu, J. D. Angrist, Y. Narita, and P. A. Pathak. Research design meets market design: Using centralized assignment for impact evaluation. *Econometrica*, 85(5):1373–1432, 2017b.
- A. Abdulkadirođlu, Y.-K. Che, P. A. Pathak, A. E. Roth, and O. Tercieux. Minimizing justified envy in school choice: The design of New Orleans’ oneapp. Technical report, National Bureau of Economic Research, 2017c.
- N. Agarwal and P. Somaini. Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica*, 86(2):391–444, 2018.

- N. I. Al-Najjar. Aggregation and the law of large numbers in large economies. *Games and Economic Behavior*, 47(1):1–35, 2004.
- J. Alcalde-Unzu and E. Molis. Exchange of indivisible goods and indifferences: The top trading absorbing sets mechanisms. *Games and Economic Behavior*, 73(1):1–16, 2011.
- R. Anderson, I. Ashlagi, D. Gamarnik, M. Rees, A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange and the alliance for paired donation: Operations research changes the way kidneys are transplanted. *Interfaces*, 45(1):26–42, 2015.
- R. Anderson, I. Ashlagi, D. Gamarnik, and Y. Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017.
- N. Arnosti. Short lists in centralized clearinghouses. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, page 751, 2015.
- N. Arnosti and P. Shi. How (not) to allocate affordable housing. 2017.
- I. Ashlagi and A. Nikzad. What matters in school choice tie-breakings?: How competition guides design. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 767–768. ACM, 2016.
- I. Ashlagi and P. Shi. Improving community cohesion in school choice via correlated-lottery implementation. *Operations Research*, 62(6):1247–1264, 2014.
- I. Ashlagi and P. Shi. Optimal allocation without money: An engineering approach. *Management Science*, 62(4):1078–1097, 2015.
- I. Ashlagi, A. Nikzad, and A. Romm. Assigning more students to their top choices: A tiebreaking rule comparison. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 755–756, 2015.
- I. Ashlagi, M. Burq, P. Jaillet, and V. Manshadi. On matching and thickness in heterogeneous dynamic markets. 2017.
- I. Ashlagi, M. Braverman, Y. Kanoria, and P. Shi. Clearing matching markets efficiently: Informative signals and match recommendations. 2018.
- C. Avery and J. Levin. Early admissions at selective colleges. *American Economic Review*, 100(5):2125–2156, 2010.
- E. M. Azevedo and E. Budish. Strategy-proofness in the large. Technical report, National Bureau of Economic Research, 2017.
- E. M. Azevedo and J. W. Hatfield. Existence of equilibrium in large matching markets with complementarities. 2015.
- E. M. Azevedo and J. D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268, 2016.

- H. Aziz, P. Biró, S. Gaspers, R. de Haan, N. Mattei, and B. Rastegari. Stable matching with uncertain linear preferences. In *International Symposium on Algorithmic Game Theory*, pages 195–206. Springer, 2016.
- S. Bade. Serial dictatorship: The unique optimal allocation rule when information is endogenous. *Theoretical Economics*, 10(2):385–410, 2015.
- M. Balinski and T. Sönmez. A tale of two mechanisms: Student placement. *Journal of Economic Theory*, 84(1):73–94, 1999.
- C. Blair. The lattice structure of the set of stable matchings with multiple partners. *Mathematics of Operations Research*, 13(4):619–628, 1988.
- I. Bo and R. Hakimov. The iterative deferred acceptance mechanism. 2017.
- BPS. Recommendation to implement a new BPS assignment algorithm. Technical report, 2005.
- S. R. Cellini, F. Ferreira, and J. Rothstein. The value of school facility investments: Evidence from a dynamic regression discontinuity design. *The Quarterly Journal of Economics*, 125(1):215–261, 2010.
- A. Chakraborty, A. Citanna, and M. Ostrovsky. Two-sided matching with interdependent values. *Journal of Economic Theory*, 145(1):85–105, 2010.
- Y.-K. Che and O. Tercieux. Efficiency and stability in large matching markets. 2015.
- Y.-K. Che and O. Tercieux. Payoff equivalence of efficient mechanisms in large matching markets. *Theoretical Economics*, 13(1):239–271, 2018.
- Y.-K. Che, J. Kim, and F. Kojima. Stable matching in large economies. *Manuscript, Department of Economics, Columbia University*, 2015.
- Y. Chen and Y. He. Information acquisition and provision in school choice. Technical report, mimeo, 2015.
- R. Chetty, J. N. Friedman, N. Hilger, E. Saez, D. W. Schanzenbach, and D. Yagan. How does your kindergarten classroom affect your earnings? Evidence from project star. *The Quarterly Journal of Economics*, 126(4):1593–1660, 2011.
- J. Combe, O. Tercieux, and C. Terrier. The design of teacher assignment: Theory and evidence. 2016.
- O. Compte and P. Jehiel. Voluntary participation and re-assignment in two-sided matching. Technical report, working paper, 2008.
- J. P. Dickerson and T. Sandholm. Futurematch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI*, pages 622–628, 2015.
- M. Dinerstein, T. Smith, et al. Quantifying the supply response of private schools to public policies. 2014.

- J. Drummond and C. Boutilier. Elicitation and approximately stable matching with partial preferences. In *IJCAI*, 2013.
- J. Drummond and C. Boutilier. Preference elicitation and interview minimization in stable matchings. In *AAAI*, pages 645–653, 2014.
- L. E. Dubins and D. A. Freedman. Machiavelli and the Gale-Shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494, 1981.
- U. Dur. Dynamic school choice problem. Working paper, 2012a.
- U. Dur and T. Morrill. Competitive equilibria in school assignment. *Games and Economic Behavior*, 2017.
- U. Dur, R. G. Hammond, and T. Morrill. Identifying the harm of manipulable school-choice mechanisms. *American Economic Journal: Economic Policy*, 2015.
- U. M. Dur. A characterization of the top trading cycles mechanism for the school choice problem. Available at SSRN 2147449, 2012b.
- L. Ehlers. Top trading with fixed tie-breaking in markets with indivisible goods. *Journal of Economic Theory*, 151:64–87, 2014.
- L. Ehlers and A. Erdil. Efficient assignment respecting priorities. *Journal of Economic Theory*, 145(3):1269–1282, 2010.
- L. Ehlers and J. Massó. Matching markets under (in) complete information. *Journal of Economic Theory*, 157:295–314, 2015.
- A. Erdil and H. Ergin. What’s the matter with tie-breaking? Improving efficiency in school choice. *American Economic Review*, 98(3):669–689, 2008.
- P. Eso and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74:705–731, 2007.
- G. Fack, J. Grenet, and Y. He. Beyond truth-telling: Preference estimation with centralized school choice. 2015.
- I. Feigenbaum, Y. Kanoria, I. Lo, and J. Sethuraman. Dynamic matching in school choice: Efficient seat reallocation after late cancellations. 2018.
- D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.
- Y. A. Gonczarowski, N. Nisan, R. Ostrovsky, and W. Rosenbaum. A stable marriage requires communication. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1003–1017. Society for Industrial and Applied Mathematics, 2015.
- B. Gong and Y. Liang. A dynamic college admission mechanism in Inner Mongolia: Theory and experiment. Technical report, Working paper, 2016. 2, 10, 2016.

- S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408, 1980.
- R. Hakimov and O. Kesten. Equitable top trading cycles mechanisms: Theory and experiments. *International Economic Review*, Forthcoming.
- P. Harless and V. Manjunath. The importance of learning in market design. 2015.
- A. Hassidim, D. Marciano-Romm, A. Romm, and R. I. Shorrer. "Strategic" behavior in a strategy-proof environment. Technical report, working paper, 2015.
- J. W. Hatfield, F. Kojima, and Y. Narita. Improving schools through school choice: A market design approach. *Journal of Economic Theory*, 2016.
- Y. He, A. Miralles, M. Pycia, and J. Yan. A pseudo-market approach to allocation with priorities. *American Economic Journal: Microeconomics*, Forthcoming.
- C. M. Hoxby. All school finance equalizations are not created equal. *Quarterly Journal of Economics*, pages 1189–1231, 2001.
- A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, pages 293–314, 1979.
- C. K. Jackson, R. C. Johnson, and C. Persico. The effects of school spending on educational and economic outcomes: Evidence from school finance reforms. *The Quarterly Journal of Economics*, 131(1):157–218, 2016.
- P. Jaramillo and V. Manjunath. The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory*, 147(5):1913–1946, 2012.
- R. C. Johnson and C. K. Jackson. Reducing inequality through dynamic complementarity: Evidence from head start and public school spending. Technical report, National Bureau of Economic Research, 2017.
- S. Kadam. Interviewing in matching markets. Technical report, Working Paper, 2015.
- S. V. Kadam and M. H. Kotowski. Multi-period matching. Working paper, 2014.
- E. H. Kaplan. Analyzing tenant assignment policies. *Management Science*, 33(3):395–408, 1987.
- A. Kapor, C. Neilson, and S. Zimmerman. Heterogeneous beliefs and school choice. 2016.
- S. Karlin and H. M. Taylor. A first course in stochastic processes. *Academic, San Diego*, 1975.
- O. Kesten. On two competing mechanisms for priority-based allocation problems. *Journal of Economic Theory*, 127(1):155–171, 2006.
- O. Kesten. School choice with consent. *The Quarterly Journal of Economics*, 125(3):1297–1348, 2010.
- B. Kleinberg, B. Waggoner, and E. G. Weyl. Descending price optimally coordinates search. 2016.

- D. Knuth. *Mariages Stables Et Leurs Relations Avec D'Autres Problèmes Combinatoires: Introduction À L'Analyse Mathématique Des Algorithmes*. Presses de l'Université de Montréal, 1976.
- A. B. Krueger. Experimental estimates of education production functions. *The Quarterly Journal of Economics*, 114(2):497–532, 1999.
- J. Lafortune, J. Rothstein, and D. W. Schanzenbach. School finance reform and the distribution of student achievement. *American Economic Journal: Applied Economics*, 10(2):1–26, 2018.
- H.-S. Lee and M. M. Srinivasan. Control policies for the MX/G/1 queueing system. *Management Science*, 35(6):708–721, 1989.
- R. S. Lee and M. Schwarz. Interviewing in two-sided matching markets. Technical report, National Bureau of Economic Research, 2009.
- Q. Liu and M. Pycia. Ordinal efficiency, fairness, and incentives in large markets. 2016.
- Q. Liu, G. J. Mailath, A. Postlewaite, and L. Samuelson. Stable matching with incomplete information. *Econometrica*, 82(2):541–587, 2014.
- J. Ma. Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23(1):75–83, 1994.
- F. Matějka and A. McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.
- D. L. McFadden. Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, pages 105–142, 1973.
- P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, September 1982.
- A. Miralles and M. Pycia. Prices and efficient assignments without transfers. 2014.
- T. Morrill. An alternative characterization of top trading cycles. *Economic Theory*, 54(1):181–197, 2013.
- T. Morrill. Making just school assignments. *Games and Economic Behavior*, 92:18–27, 2015a.
- T. Morrill. Two simple variations of top trading cycles. *Economic Theory*, 60(1):123–140, 2015b.
- Y. Narita. Match or mismatch: Learning and inertia in school choice. 2016.
- W. K. Newey and D. McFadden. Large sample estimation and hypothesis testing. *Handbook of Econometrics*, 4:2111–2245, 1994.
- W. Olszewski and R. Weber. A more general pandora rule? *Journal of Economic Theory*, 160:429–437, 2015.
- S. Pápai. Strategyproof assignment by hierarchical exchange. *Econometrica*, 68(6):1403–1433, 2000.
- P. Pathak. *What really matters in designing school choice mechanisms*. 2016.

- P. A. Pathak. The mechanism design approach to student assignment. *Annual Review of Economics*, 3(1):513–536, 2011.
- J. S. Pereyra. A dynamic school choice model. *Games and Economic Behavior*, 80:100–114, 2013.
- M. Pycia and M. U. Ünver. Incentive compatible allocation and exchange of discrete resources. *Theoretical Economics*, 12(1):287–329, 2017.
- B. Rastegari, A. Condon, N. Immorlica, and K. Leyton-Brown. Two-sided matching with partial information. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pages 733–750. ACM, 2013.
- B. Rastegari, A. Condon, N. Immorlica, R. Irving, and K. Leyton-Brown. Reasoning about optimal stable matchings under partial information. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, pages 431–448. ACM, 2014.
- A. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, pages 617–628, 1982.
- A. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4):748–780, 1999.
- A. E. Roth. *Who Gets What and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt, 2015.
- A. E. Roth and A. Postlewaite. Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4(2):131–137, 1977.
- A. E. Roth and M. Sotomayor. Two-sided matching. *Handbook of Game Theory with Economic Applications*, 1:485–541, 1992.
- A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.
- D. Saban and J. Sethuraman. House allocation with indifference: A generalization and a unified view. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pages 803–820. ACM, 2013.
- L. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.
- P. Shi. Assortment planning in school choice. 2015.
- C. A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.
- K. A. Small and H. S. Rosen. Applied welfare economics with discrete choice models. *Econometrica*, 49(1):105–130, 1981.
- J. Steiner, C. Stewart, and F. Matějka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017.

- X. Su and S. A. Zenios. Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Science*, 52(11):1647–1660, 2006.
- J. Szuffita. Tale of two waitlists: There are too many people at your beloved zoned school. <http://www.nycschoolhelp.com/blog/2016/3/16/tale-of-two-waitlists>. March 16, 2016.
- K. T. Talluri and G. J. Van Ryzin. *The Theory and Practice of Revenue Management*, volume 68. Springer Science & Business Media, 2006.
- M. L. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.
- W. Whitt. *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues*. Springer Science & Business Media, 2002.
- S. A. Zenios. Modeling the transplant waiting list: A queueing model with reneging. *Queueing Systems*, 31(3-4):239–251, 1999.

Appendices

Appendix A

Appendix for Chapter 3

A.1 Definitions and Notation

We begin with some general notation and definitions. Let μ be the initial assignment under DA-STB, and let P be a permutation. We say that a school i *reaches capacity* under a mechanism with output assignment μ if $\eta(\mu(i)) = q_i$.

We re-index the schools in $\mathcal{C} \cup \{n+1\}$ so that $C_i \geq C_{i+1}$. Moreover, we assume that this indexing is done such that if the order condition is satisfied, then $\hat{C}_i^P \geq \hat{C}_{i+1}^P$ (where the cutoffs \hat{C}^P are as defined by PLDA(P)) holds simultaneously for all permutations P .

Recall that in DA each student is given a score $r_i^s = p_i^s + L(s)$, and in PLDA(P) this leads to a second-round score $\hat{r}_i^s = \hat{p}_i^s + P(L(s)) = P(L(s)) + n_i \mathbf{1}_{\{\mu(s)=i\}} + p_i^s \mathbf{1}_{\{\mu(s) \neq i\}}$. Throughout the Appendix, for convenience, we slightly change the second-round score of a student s under PLDA with permutation P to be $\hat{r}_i^s = P(L(s)) + n_i \mathbf{1}_{\{L(s) \geq C_i\}} + p_i^s \mathbf{1}_{\{L(s) < C_i\}}$, meaning that we give each student a guarantee at any school for which she met the cutoff in the first round. By consistency of preferences, it is easily seen that this has no effect on the resulting assignment or cutoffs.

We say that a student can *afford* a school in a round if her score in that round is at least as large as the school's cutoff in that round. We say that the set of schools a student can afford in the second round (with her amended second-round score) is her *affordable set*.

Throughout the Appendix, we let $X_i = \{i, \dots, n+1\}$ be the set of schools at least as affordable as school i , and we let γ_i be the proportion of students whose first-round affordable set is X_i .

A.2 Proof of Theorem 3.1

We first prove Theorem 3.1 in the case where all schools have one priority group. We then show that if the order condition holds, all PLDA mechanisms assign the same number of seats at a given school i to students of a given priority class π . Hence, by restricting to the set of students with priority class π , we can reduce the general problem to the case where all schools have one priority group. This shows that all PLDA mechanisms produce type-equivalent assignments.

Lemma A.1. *Assume that each school has a single priority group, $\mathbf{p} = \mathbf{1}$. If the order condition holds, all PLDA mechanisms produce type-equivalent assignments.*

Proof. Let P be a permutation. Assume that the order condition holds. By Theorem 3.4, we may assume that the global order condition holds. Hence the schools in $\mathcal{C} \cup \{n+1\}$ can be indexed so that $C_i \geq C_{i+1}$ and $\hat{C}_i^P \geq \hat{C}_{i+1}^P$ for all permutations P (simultaneously).

We first present the relevant notation that will be used in this proof. We are interested in sets of schools of the form $X_i = \{i, \dots, n+1\}$. Let

$$\beta_{i,j} = \eta(\{s \in \mathcal{S} : i \text{ is the most desirable school in } X_j \text{ with respect to } \succ^s\})$$

be the measure of the students who, when their set of affordable schools is X_j , will choose i (when following their second-round preferences). Note that $\beta_{i,j} = 0$ for all $j > i$.

Let $E^s(\mathbf{C})$ and $\hat{E}_P^s(\hat{\mathbf{C}}^P)$ be the first-round affordable set and (total) affordable set for student s when running PLDA with permutation P . Note that for each student $s \in \mathcal{S}$, there exists some i such that $E^s(\mathbf{C}) = X_i$, and since the order condition is satisfied, there exists some $j \leq i$ such that $\hat{E}_P^s(\hat{\mathbf{C}}^P) = X_j$. The fact that $\hat{E}_P^s(\hat{\mathbf{C}}^P) = X_j$ for some j is a result of the order condition: students' amended second-round scores guarantee that $E^s(\mathbf{C}) \subseteq \hat{E}_P^s(\hat{\mathbf{C}}^P)$ (every school affordable in the first round is guaranteed in the second) and hence that $j \leq i$. Let $\gamma_i^P = \eta(\{s \in \mathcal{S} : \hat{E}_P^s(\hat{\mathbf{C}}^P) = X_i\})$ be the fraction of students whose (total) affordable set in PLDA(P) is X_i .¹ We note that by definition of PLDA, $\eta(\{s \in \mathcal{S} : \theta^s = \theta, \hat{E}_P^s(\hat{\mathbf{C}}^P) = X_i\}) = \zeta(\{\theta\})\gamma_i^P$; that is, the students whose affordable sets are X_i "break proportionally" into types. For a school i , this means that the measure of students assigned to i is therefore $\sum_{j \leq i} \beta_{i,j} \gamma_j^P$.

Let P' be another permutation, and define $\gamma_i^{P'}$ similarly. We will prove by induction that there exist PLDA(P') cutoffs $\hat{\mathbf{C}}^{P'}$ such that $\gamma_i^{P'} = \gamma_i^P$ for all $i \in \mathcal{C} \cup \{n+1\}$. Note that by the proportional breaking into types of γ_i^P and $\gamma_i^{P'}$, this will imply type-equivalence.

Assume that the PLDA(P') cutoffs $\hat{\mathbf{C}}^{P'}$ are chosen such that $\gamma_j^{P'} = \gamma_j^P$ for all $j < i$, and i is maximal such that this is true. Then we have that $\sum_{j \leq i-1} \beta_{i,j} \gamma_j^P = \sum_{j \leq i-1} \beta_{i,j} \gamma_j^{P'}$. Assume w.l.o.g. that $\gamma_i^P > \gamma_i^{P'}$. It follows that $q_i \geq \sum_{j \leq i} \beta_{i,j} \gamma_j^P \geq \sum_{j \leq i} \beta_{i,j} \gamma_j^{P'}$, where the first inequality follows since i cannot be filled beyond capacity. If the second inequality is strict, then under P' , i is not full,

¹Note that $\eta(\mathcal{S}) = 1$, as η is a probability distribution over \mathcal{S} .

and therefore $\hat{C}_i^{P'} = 0$. However, this means that all students can afford i under P' , and therefore $\gamma_i^{P'} = 1 - \sum_{j < i} \gamma_j^{P'} = 1 - \sum_{j < i} \gamma_j^P \geq \gamma_i^P$, a contradiction. If the second inequality is an equality, then $\beta_{i,i} = 0$ and no students demand school i under the given affordable set structure. It follows that we can define the cutoff $\hat{C}_i^{P'}$ such that $\gamma_i^{P'} = \gamma_i^P$. This provides the required contradiction, completing the proof. \square

Now consider when schools have possibly more than one priority group. We show that if the order condition holds, then all PLDA mechanisms assign the same measure of students of a given priority type to a given school. It is not at all obvious that such a result should hold, since priority types and student preferences may be correlated, and the relative proportions of students of each priority type assigned to each school can vary widely. Nonetheless, the order condition (specifically, the equivalent global order condition) imposes enough structure so that any given priority type is treated symmetrically across different PLDA mechanisms.

Theorem A.1. *If the order condition holds, then for all priority classes π and schools i all PLDA mechanisms assign the same measure of students of priority class π to school i .*

Proof. Fix a permutation P . By Theorem 3.4, we may assume that the global order condition holds.

We show that $\text{PLDA}(P)$ assigns the same measure of students of each priority type to each school i as RLDA. The idea will be to define cutoffs on priority-type-specific economies, and show that these cutoffs are the same as the PLDA cutoffs. However, since cutoffs are not necessarily unique in the two-round setting, care needs to be taken to make sure that the individual choices for priority-type-specific cutoffs are consistent across priority types.

The proof runs as follows. We first define an economy \mathcal{E}_π for each priority class π that gives only as many seats as are assigned to students of priority class π under RLDA. We then invoke the global order condition and Theorems 3.4 and 3.1 to show that all PLDA mechanisms are type-equivalent on each \mathcal{E}_π . We also use the global order condition to argue that it is sufficient to consider affordable sets, and also to select “minimal” cutoffs. Then we construct cutoffs $\mathcal{C}_{\pi,i}^P$ using the economies \mathcal{E}_π and show that they are (almost) independent of priority type. Finally, we show that this means

that $\mathcal{C}_{\pi,i}^P$ also define PLDA cutoffs for the large economy \mathcal{E} and conclude that PLDA(P) assigns the same measure of students of each priority type to each school i as RLDA.

(1) Defining little economies \mathcal{E}_π for each priority type.

Fix a priority class π . Let q_π be a restricted capacity vector, where $q_{\pi,i}$ is the measure of students of priority class π assigned to school i under RLDA. Let \mathcal{S}_π be the set of students s such that $p^s = \pi$, and let η_π be the restriction of the distribution η to \mathcal{S}_π . Let \mathcal{E}_π denote the primitives $(\mathcal{C}, \mathcal{S}_\pi, \eta_\pi, q_\pi)$. Recall that $\hat{\mathbf{C}}^{\mathbf{R}}$ are the second-round cutoffs for RLDA on \mathcal{E} . It follows from the definition of \mathcal{E}_π that $\hat{\mathbf{C}}_\pi^{\mathbf{R}}$ are also the second-round cutoffs for RLDA on \mathcal{E}_π .

Let $\tilde{\mathbf{C}}_\pi^P$ be the second-round cutoffs of PLDA(P) on \mathcal{E}_π . We show that the cutoffs $\tilde{\mathbf{C}}_\pi^P$ defined for the little economy are the same as the consistent second-round cutoffs $\hat{\mathbf{C}}_\pi^P$ for PLDA(P) for the large economy \mathcal{E} , that is, $\tilde{\mathbf{C}}_\pi^P = \hat{\mathbf{C}}_\pi^P$.

(2) Implications of the global order condition.

We have assumed that the global order condition holds. This has a number of implications for PLDA mechanisms run on the little economies \mathcal{E}_π . For all p , the local order condition holds for RLDA on \mathcal{E}_π . Hence, by Theorem 3.4, the little economies \mathcal{E}_π each satisfy the order condition. Moreover, by Theorem 3.1, all PLDA mechanisms produce type-equivalent assignments when run on \mathcal{E}_π . Finally, if we can show that for every permutation P , PLDA(P) assigns the same measure of students of each priority type π to each school i (namely $(q_\pi)_i$) as RLDA, then \mathcal{E} satisfies the global order condition if and only if for all p the little economy \mathcal{E}_π satisfies the global order condition.

The global order condition also allows us to determine aggregate student demand from the proportions of students who have each school in their affordable set. In general, if affordable sets break proportionally across types, and if for each subset of schools $X \subseteq \mathcal{C}$ we know the proportion of students whose affordable set is X , then we can determine aggregate student demand. The global order condition implies that for any pair of permutations P, P' , the affordable sets from both rounds are nested in the same order under both permutations. In other words, for each priority class π there exists a permutation σ_π such that the affordable set of any student in any round of any PLDA mechanism is of the form $\{\sigma_\pi(i), \sigma_\pi(i+1), \dots, \sigma_\pi(n), n+1\}$. Hence when the global order condition holds, to determine the proportion of students whose affordable set is X , it is sufficient

to know the proportion of students who have each school in their affordable set.

Another more subtle implication of the global order condition is the following. In the second round of PLDA, for each permutation P and school i there will generically be an interval that \hat{C}_i^P can lie in and still be market-clearing. The intuition is that there will be large empty intervals corresponding to students who had school i in their first-round affordable set, and whose second-round lottery changed accordingly. When the global order condition holds, we can without loss of generality assume that as many as possible of the cutoffs for a given priority type are 0 or 1, and the global order condition will still hold.

Formally, for cutoffs C we can equivalently define priority-type-specific cutoffs $C_{\pi,i} = (\lfloor C_i - \pi_i \rfloor)^+$. Note the cutoffs C_π are consistent across priority types, namely: (1) cutoffs match for two priority types with the same priority group at a school, $\pi_i = \pi'_i \Rightarrow C_{\pi,i} = C_{\pi',i}$ and $\hat{C}_{\pi,i} = \hat{C}_{\pi',i}$; and (2) there is at most one marginal priority group at each school, $C_{\pi,i}, C_{\pi',i} \in (0, 1) \Rightarrow \pi_i = \pi'_i$. Moreover, if cutoffs C_π are consistent across priority types, then there exist cutoffs C from which they arise.

Suppose that we set as many as possible of the priority-type-specific cutoffs \hat{C}_π^P to be extremal; i.e., we let $\hat{C}_{\pi,i}^P$ be 1 if no students have i in their affordable set, and let $\hat{C}_{\pi,i}^P$ be minimal otherwise. We show that under this new definition, C_π, \hat{C}_π^P satisfies the local order condition consistently with all other PLDAs.

Specifically, let

$$f_{\pi,i}^P(x) = |\{\ell : \ell \geq C_{\pi,i} \text{ or } P(\ell) \geq x\}|$$

be the proportion of students of priority class π who have school i in their affordable set if the first- and second-round cutoffs are $C_{\pi,i}$ and x respectively. Notice that f is decreasing in x . Define cutoffs \tilde{C}_π^P as follows. If $f_{\pi,i}^P(\hat{C}_{\pi,i}^P) = 0$ we set $\tilde{C}_{\pi,i}^P = 1$, and otherwise we let $\tilde{C}_{\pi,i}^P$ be the minimal cutoff satisfying $f_{\pi,i}^P(\tilde{C}_{\pi,i}^P) = f_{\pi,i}^P(\hat{C}_{\pi,i}^P)$.

Since \mathcal{E} satisfies the global order condition, for all π there exists an ordering σ_π such that $C_{\sigma_\pi(1)} \geq C_{\sigma_\pi(2)} \geq \dots \geq C_{\sigma_\pi(n)}$ and $\hat{C}_{\sigma_\pi(1)}^{P'} \geq \hat{C}_{\sigma_\pi(2)}^{P'} \geq \dots \geq \hat{C}_{\sigma_\pi(n)}^{P'}$ for all permutations P' . We show that the global order condition implies that the newly defined cutoffs \hat{C}^P satisfy $\tilde{C}_{\pi,\sigma_\pi(1)}^P \geq$

$\tilde{C}_{\pi, \sigma_\pi(2)}^P \geq \dots \geq \tilde{C}_{\pi, \sigma_\pi(n)}^P$. This is because the global order condition implies that f_π^P is increasing in i ; i.e., for each π , $i < j$, and x it holds that $f_{\pi, \sigma_\pi(i)}^P(x) \leq f_{\pi, \sigma_\pi(j)}^P(x)$. Hence for all $j > i$

$$\begin{aligned} f_{\pi, \sigma_\pi(j)}^P \left(\tilde{C}_{\pi, \sigma_\pi(j)}^P \right) &= f_{\pi, \sigma_\pi(j)}^P \left(\hat{C}_{\pi, \sigma_\pi(j)}^P \right) \\ &\geq f_{\pi, \sigma_\pi(j)}^P \left(\hat{C}_{\pi, \sigma_\pi(i)}^P \right) \quad (\text{since } f \text{ is decreasing}) \\ &\geq f_{\pi, \sigma_\pi(i)}^P \left(\hat{C}_{\pi, \sigma_\pi(i)}^P \right) \quad (\text{since } f \text{ is increasing in } i) \\ &= f_{\pi, \sigma_\pi(i)}^P \left(\tilde{C}_{\pi, \sigma_\pi(i)}^P \right) \end{aligned}$$

and so since we set $\tilde{C}_{\pi, \sigma_\pi(j)}^P$ to be minimal and $f_{\pi, \sigma_\pi(j)}^P(\cdot)$ is decreasing it follows that $\tilde{C}_{\pi, \sigma_\pi(j)}^P \leq \tilde{C}_{\pi, \sigma_\pi(i)}^P$.

(3) Cutoffs $\tilde{C}_{\pi, i}^P$ are (almost) independent of priority type.

We now show that $\tilde{C}_{\pi, i}^P$ depends on π only via π_i , and does not depend on π_j for all $j \neq i$. Since \mathcal{E}_π satisfies the order condition, all PLDA mechanisms on \mathcal{E}_π are type-equivalent, and the proportion of students who have each school in their affordable set is the same across all PLDA mechanisms. Hence for all permutations P , priority classes π , and schools i it holds that $f_{\pi, i}^P(\tilde{C}_{\pi, i}^P) = f_{\pi, i}^P(\hat{C}_{\pi, i}^P) = f_{\pi, i}^R(\hat{C}_{\pi, i}^R)$. This means that $\tilde{C}_{\pi, i}^P$ satisfies the following equation in terms of $\hat{C}_{\pi, i}^R$, $C_{\pi, i}$ and P :

$$f_{\pi, i}^P \left(\tilde{C}_{\pi, i}^P \right) = f_{\pi, i}^R \left(\hat{C}_{\pi, i}^R \right) = 2 - \hat{C}_{\pi, i}^R - C_{\pi, i}. \quad (\text{A.1})$$

(We note that an application of the intermediate value theorem shows that this equation always has a solution in $[0, 1]$, since $f_{\pi, i}^P(0) = 1 - C_{\pi, i}$, $f_{\pi, i}^P(1) = 1$, $f_{\pi, i}^P$ is continuous and decreasing on $[0, 1]$, and we are in the case where $1 - C_{\pi, i} \leq \hat{C}_{\pi, i}^R \leq 1$. Hence $\tilde{C}_{\pi, i}^P$ is defined by $f_{\pi, i}^P$ and $f_{\pi, i}^R$.) In other words, the value of $\tilde{C}_{\pi, i}^P$ is defined by $f_{\pi, i}^P(\cdot)$, $f_{\pi, i}^R(\cdot)$, and $\hat{C}_{\pi, i}^R$, which in turn are defined by $C_{\pi, i}$ and the permutations P or R . Since $C_{\pi, i}$ depends on π only through π_i , it follows that $\tilde{C}_{\pi, i}^P$ depends on π only through π_i . In other words the $\tilde{C}_{\pi, i}^P$ define cutoffs \tilde{C}_i^P that are independent of priority type.

(4) \tilde{C}_i^P are the PLDA cutoffs.

Finally, we remark that \tilde{C}_i^P are market-clearing cutoffs. This is because we have shown that for each priority class π , the number of students assigned to each school i is the same under RLDA and under the demand induced by the cutoffs \tilde{C}_i^P , and we know that the RLDA cutoffs are market-clearing for \mathcal{E} .

Hence \tilde{C}_i^P give the assignments for PLDA on \mathcal{E} , and since \tilde{C}_i^P was defined individually for each priority class π on \mathcal{E}_π , it follows that PLDA(P) assigns the same measure of students of each priority type to each school i as RLDA. \square

We are now ready to prove Theorem 3.1

Proof of Theorem 3.1. Fix a priority class π . By Theorem A.1, for every school i , all PLDA mechanisms assign the same measure $q_{\pi,i}$ of students of priority class π to school i .

Consider the subproblem with primitives $\mathcal{E}_\pi = (\mathcal{C}, q_\pi, \mathcal{S}_\pi, \eta_\pi)$. By Lemma A.1, for all $\theta \in \Theta$ and i ,

$$\eta_\pi(\{s \in \mathcal{S}_\pi : \theta^s = \theta, \hat{\mu}_P(s) = i\}) = \eta_\pi(\{s \in \mathcal{S}_\pi : \theta^s = \theta, \hat{\mu}_{P'}(s) = i\}).$$

Since η_π is the restriction of η to \mathcal{S}_π , it follows that all PLDA mechanisms are type-equivalent. \square

A.3 Proof of Theorem 3.3

We first note that with a single priority class, the first round corresponds to the random serial dictatorship (RSD) mechanism of Abdulkadiroğlu and Sönmez (1998), where the (random) order of students is the single order of tie-breaking. Hence instead of referring to the first-round mechanism as DA-STB, we will sometimes refer to it as RSD.

Recall the cutoff characterization of the set of stable matchings for given student preferences and responsive school preferences (encoded by student scores $r_i^s = p_i^s + L(s)$) (Azevedo and Leshno, 2016). Namely, if $\mathbf{C} \in \mathbb{R}_+^N$ is a vector of cutoffs, let the assignment μ defined by \mathbf{C} be given by assigning each student of type s to her favorite school among those where her score weakly exceeds

the cutoff, $\mu(s) = \max_{\succ s}(\{s_i \in \mathcal{C} : r_i^s \geq C_i\} \cup \{n+1\})$. The cutoffs \mathbf{C} are market-clearing if under the assignment μ defined by \mathbf{C} , every school with a positive cutoff is exactly at capacity, $\eta(\mu(s_i)) \leq q_i$ for all $i \in \mathcal{C}$, with equality if $C_i > 0$. The set of all stable matchings is precisely given by the set of assignments defined by market-clearing vectors (Azevedo and Leshno, 2016).

Under PLDA(P), a student of type s has a second-round score $\hat{r}_i^s = P(L(s)) + \mathbf{1}_{\{L(s) \geq C_i\}}$ at school i for each school $i \in \mathcal{C} \cup \{n+1\}$ (assuming that scores are modified to give guarantees to students who had a school in their first-round affordable set, instead of just students assigned to the school in the first round). In a slight abuse of notation, we will sometimes let $\hat{\mathbf{C}}^P$ refer to the second-round cutoffs from some fixed PLDA(P) (not necessarily corresponding to the student-optimal stable matching given by PLDA).

The proof that any PLDA satisfies the axioms essentially follows from Proposition 3.1. We note that averaging follows from the continuum model, which preserves the relative proportion of students with different reported types under random lotteries and permutations of random lotteries. Hence it suffices to show that any mechanism M satisfying the axioms is a PLDA.

We will show that the reassignment produced by M is type-equivalent to the reassignment produced by some PLDA. If we assume, that conditional on their reports, students' assignments under M are uncorrelated, we are able to explicitly construct a PLDA that provides the same joint distribution over assignments and reassignments as M . We provide a sketch of the proof before fleshing out the details.

Fix a distribution of student types ζ . Since the first round of our mechanism M is DA-STB and M is anonymous, this gives a distribution η of students that is the same (up to relabeling of students) at the end of the first round. For a fixed labeling of students, it also gives a distribution over first-round assignments μ and a distribution over second round assignments $\hat{\mu}$.

We first invoke averaging to assume that all ensuing constructions of aggregate cutoffs and measures of students assigned to pairs of schools in the two rounds are deterministic. Specifically, since the first-round assignment μ is given by STB, and the mechanism satisfies the averaging axiom, we may assume that each pathwise realization of the mechanism gives type-equivalent (two-round) assignments. Hence, for the majority of the proof we perform our constructions of aggregate

cutoffs and measures of students pathwise, and assume that any realization of the lottery numbers produces the same cutoffs and measures of students. (In particular, the quantities $\hat{C}_i, \rho_{i,j}, \gamma_{i,j}$ that we will later define will be the same across all realizations.)

Outline of Proof. We use constrained Pareto efficiency to construct a first-round overdemand ordering $1, 2, \dots, n, n+1$, as in (Ashlagi and Shi, 2014), where school i comes before j in an ordering for the first (second) round if there exists a non-zero measure of students who prefer school i to j in the first (second) round but who are assigned to j in the first (second) round. (In the case of the second-round ordering, we require that these students' second-round assignments j not be the same as their first-round guarantees.) The existence of these orderings follows from the facts that the first-round mechanism, DA-STB, is Pareto efficient, that the two-round mechanism is constrained Pareto efficient. We let $X_i = \{i, i+1, \dots, n, n+1\}$ denote the set of schools after i in the first-round overdemand ordering, and let $\tilde{X}_i = \{\sigma(i), \sigma(i+1), \dots, \sigma(n+1)\}$ denote the set of schools after $\sigma(i)$ in the second-round overdemand ordering.

We next note that instead of assignments μ and $\hat{\mu}$, we can think of giving students first- and second-round affordable sets $E(s), \hat{E}(s)$ so that μ and $\hat{\mu}$ are given by letting each student choose her favorite school in her affordable set for that round. We use weak two-round strategy-proofness and anonymity to show that two students of different types face the same joint distribution over first- and second-round affordable sets. This allows us to construct the permutation P by constructing proportions $\gamma_{i,j}$ of students whose first-round affordable set was X_i and whose (total) affordable set was \tilde{X}_j . This is the most technical step in the proof, and so we separate it into several steps. The crux of the analysis is the fact that for any school i and set $\mathcal{C}' \not\ni i$ of schools, two students with top choices \mathcal{C}' who are assigned to a school they weakly prefer to i the first round have the same conditional probability of being assigned to a school in \mathcal{C}' in the second round.² We term this the “prefix property” and prove it in Lemma A.2.

Finally, we construct the lottery L and verify that if second-round scores are given by first prioritizing all guaranteed students over non-guaranteed students and subsequently breaking ties

²The formal statement also takes into account how demanded the schools they weakly prefer to i are, and is given in terms of student types who were assigned to i , and lottery numbers.

according to the permuted lottery $P \circ L$, then PLDA(P) gives every student the same pair of first- and second- round assignments as M .

Formal Proof. We now present the formal proof. Since we are assuming that the considered mechanism M is weakly strategy-proof, we assume that students report truthfully and so we consider preferences instead of reported preferences. We will explicitly specify when we are considering the possible outcomes from a single student misreporting.

(2a) Definitions

Let the schools be numbered $1, 2, \dots, n$ such that $C_i \geq C_{i+1}$ for all i . The intuition is that this is the order in which they reach capacity in the first round. We observe that all reassignments are index-decreasing. That is, for all i, j , if there exists a non-zero measure of students who are assigned to i in the first round and to j in the second round, and $j \neq n + 1$, then $i \geq j$. This follows since the mechanism respects guarantees, student preferences are consistent, and the schools are indexed in order of increasing first-round affordability. Throughout this section we will denote the outside option $n + 1$ either by 0 or \emptyset , to make it more evident that indices are decreasing.

Next, we define a permutation σ on the schools. We think of this as giving a second-round overdemand (or inverse affordability) ordering, where in the second round the schools fill in the order $s_{\sigma(1)}, s_{\sigma(2)}, \dots, s_{\sigma(n)}$. We will eventually show that M gives the same outcome as a PLDA with cutoffs that are ordered $\hat{C}_{\sigma(1)} \geq \hat{C}_{\sigma(2)} \geq \dots \geq \hat{C}_{\sigma(n)}$. We require that σ satisfies the following property. For all i, j , if there exists a non-zero measure of students with consistent preferences who have second-round preference reports \succ such that $i \succ j$, and who are not assigned to j in the first round, but are assigned to j in the second round, then $i = \sigma(i')$ and $j = \sigma(j')$ for some $i' < j'$. We assume that σ is the unique permutation satisfying this property that is maximally order-preserving. That is, for all pairs of schools i, j for which no non-zero measure of students of the above type exists, $\sigma(i) < \sigma(j)$ iff $i < j$. We also define $\sigma(n + 1) = n + 1$. An ordering σ with the required properties exists since the mechanism is constrained Pareto efficient. In particular, if there is a cycle of schools i_1, i_2, \dots, i_m where for each j there is a set of students \mathcal{S}_j with non-zero measure who prefer i_{j+1} to their second-round assignment i_j and who are not assigned to i_j in the first round, then $\hat{p}_{i_j}^{s_j} = \hat{p}_{i_{j+1}}^{s_j}$ for each $s_j \in \mathcal{S}_j$, and so there is a Pareto-improving cycle that respects

second-round priorities.

Let \mathcal{C}' be a set of schools, and let \succ be a preference ordering over all schools. We say that \mathcal{C}' is a *prefix* of \succ if $i' \succ i$ for all $i' \in \mathcal{C}', i \notin \mathcal{C}'$. For a set of schools \mathcal{C}' , let $i(\mathcal{C}') = \max(\mathcal{C}')$ be the maximum index of a school in \mathcal{C}' . We may think of $i(\mathcal{C}')$ as the index of the most affordable school in \mathcal{C}' in the first round.

For a student type $\theta = (\succ, \hat{\succ})$, an interval $I \subseteq [0, 1]$, and a set of schools \mathcal{C}' , let $\rho^\theta(I, \mathcal{C}')$ be the proportion of students with type θ who, under the mechanism M , have a first-round lottery in the interval I and are assigned to a school in \mathcal{C}' in the second round. When $\mathcal{C}' = \{i'\}$ we will sometimes write $\rho^\theta(I, i')$ instead of³ $\rho^\theta(I, \{i'\})$. In this section, for brevity, when defining preferences \succ we will sometimes write $\succ: [s_{i_1}, s_{i_2}, \dots, s_{i_k}]$ instead of $s_{i_1} \succ s_{i_2} \succ \dots \succ s_{i_k}$.

(2b) Constructing the permutation P .

We now construct the permutation P as follows. For all pairs of indices i, j , we define a scalar $\gamma_{i,j}$, which we will show can be thought of as the proportion of students (of any type) whose first-round affordable set (the set of schools at which their priority meets the first-round cutoffs) is X_i and whose affordable set (the set of schools at which their modified priority, which gives them top priority at all schools in their first-round affordable sets, meets the second-round cutoffs) is \tilde{X}_j .

Now, for all pairs of indices i, j such that $\sigma(j) < i$, we define student preferences $\theta_{i,j} = (\succ_{i,j}, \hat{\succ}_{i,j})$ such that

$$\succ_{i,j}: [s_{\sigma(j)}, s_{i-1}, s_i, n+1] \quad \text{and} \quad \hat{\succ}_{i,j}: [s_{\sigma(j)}, n+1],$$

with all other schools unacceptable. (We remark that in the case where $\sigma(j) = i - 1$, the first two schools in this preference ordering coincide.) We note that the full-support assumption implies that there is a positive measure of such students. Let $\rho_{i,j}$ be the proportion of students of type $\theta_{i,j}$ whose first-round assignment is i and whose second-round assignment is school $\sigma(j)$. Intuitively, $\rho_{i,j}$ is the proportion of students who can deduce that their lottery number is in the interval $[C_i, C_{i-1}]$, and whose second-round affordable set contains \tilde{X}_j .

For a fixed index i , we define $\gamma_{i,j}$ for $j = 1, 2, \dots, n$ to be the unique solutions to the following

³Here we are assuming that this proportion is the same for every realization of the first round of M . This requires non-atomicity and anonymity.

n equations:

$$\begin{aligned}\gamma_{i,j} &= 0 \text{ for all } j \text{ such that } \sigma(j) \geq i \\ \gamma_{i,1} + \dots + \gamma_{i,j} &= \rho_{i,j} \text{ for all } j \text{ such that } \sigma(j) < i.\end{aligned}$$

Note that by this definition it holds that $\gamma_{1,j} = 0$ for all j . We may intuitively think of $\gamma_{i,j}$ as the proportion of students of type $\theta_{i,j}$ whose first-round lottery is in $[C_i, C_{i-1}]$ and whose second-round affordable set contains $\sigma(j)$ but not $\sigma(j-1)$. (This is not quite the case, as we let $\gamma_{i,j} = 0$ for all j such that $\sigma(j) \geq i$. More precisely, if $\sigma(j) < i$ then $\gamma_{i,j}$ is the proportion of students of type $\theta_{i,j}$ whose first-round lottery is in $[C_i, C_{i-1}]$ and whose second-round affordable set contains $\sigma(j)$, but not $\sigma(j')$, where $j' = \max\{j'' : \sigma(j'') < i\}$.) Note that if $\sigma(j) \geq i$ then school $\sigma(j)$ will be in the first-round affordable set for all students whose first-round lottery is in $[C_i, C_{i-1}]$, and we define $\gamma_{i,j} = 0$ and keep track of these students separately.

We also define $\gamma_{i,n+1}$ to be

$$\gamma_{i,n+1} = C_{i-1} - C_i - \sum_{j=1}^n \gamma_{i,j}.$$

Since transfers are index-decreasing, we may intuitively think of $\gamma_{i,n+1}$ as the proportion of students of type $\theta_{i,j}$ assigned to school i in the first round whose only available school in the second round comes from their first-round guarantee.

We define the lottery P from $\gamma_{i,j}$ as follows. We break the interval $[0, 1]$ into $(n+1)^2$ intervals, $\tilde{I}_{i,j}$, where the interval $\tilde{I}_{i,j}$ has length $\gamma_{i,j}$, and the intervals are ordered in decreasing order of the first index i ,⁴

$$\tilde{I}_{n+1,n+1}, \tilde{I}_{n+1,n}, \dots, \tilde{I}_{1,2}, \tilde{I}_{1,1}.$$

The permutation P maps the intervals back into $[0, 1]$ in decreasing order of the second index}

⁴Specifically, let $\tilde{I}_{i,j} = [C_{i-1} - \sum_{j' \leq j} \gamma_{i,j'}, C_{i-1} - \sum_{j' < j} \gamma_{i,j'}]$.




Figure A.1: Constructing the permutation P for $n = 2$ schools, where σ is the identity permutation. The intervals $\tilde{I}_{i,j}$ for $i \leq \sigma(j) = j < n + 1$ are empty by definition, as all transfers are index-decreasing.

j ,⁵

$$P(\tilde{I}_{n+1,n+1}), P(\tilde{I}_{n,n+1}), \dots, P(\tilde{I}_{2,1}), P(\tilde{I}_{1,1}).$$

In Figure A.1, we show an example with two schools, and write \emptyset instead of $n + 1$ for brevity.

We note that $\sum_{j=1}^{n+1} \gamma_{i,j} = C_{i-1} - C_i$, which is the proportion of students whose first-round affordable set is X_i . We may interpret $\gamma_{i,j}$ to be the proportion of students who can deduce that their lottery number is in the interval $[C_i, C_{i-1}]$, and whose second-round affordable set is \tilde{X}_j , and so $\sum_{i=1}^{n+1} \gamma_{i,j}$ is the proportion of students whose second-round affordable set is \tilde{X}_j . We remark that there may be multiple values of i, j for which $\gamma_{i,j} = 0$ (i.e. there are no students whose first-round affordable set is X_i and second-round affordable set is X_j), but that this does not affect our ability to assign students to all possible pairs of schools that are consistent with consistent preferences and the first- and second-round overdemand orderings. For example $\gamma_{1,j} = 0$ for all j , but any student whose first-round affordable set is X_1 is able to attend their top-choice school in round 1 and her second-round affordable set is inconsequential.

We show that there exists a PLDA mechanism with permutation P , where the students with first-round scores in $\tilde{I}_{i,j}$ are precisely the students with a first-round affordable set X_i and a second-round affordable set \tilde{X}_j (where the second-round affordable set is the set of schools for which a student's unmodified second-round score $p_i^s + P(L(s))$ meets the cutoff), and that this PLDA mechanism gives the same joint distribution over first- and second-round assignments as M . To do this, we first show that this distribution of first- and second-round affordable sets gives rise to

⁵Specifically, let $\hat{C}_{\sigma(j)} = 1 - \sum_{i',j':j' \leq j} \gamma_{i',j'}$, and let $P(\tilde{I}_{i,j}) = [\hat{C}_{\sigma(j-1)} - \sum_{i' \leq i} \gamma_{i',j}, \hat{C}_{\sigma(j-1)} - \sum_{i' < i} \gamma_{i',j}]$.

the correct joint first- and second-round assignments over all students. We then use anonymity to construct L in such a way as to have the correct first- and second-round assignment joint distributions for each student. Finally, we verify that these second-round affordable sets give the student-optimal stable matching under the second round school preferences given by P .

(2c) Equivalence of the joint distribution of assignments given by affordable sets and M .

Fix student preferences $\theta = (\succ, \hat{\succ})$. We show that if we let $\gamma_{i,j}$ be the proportion of students with preferences θ who have first- and second-round affordables X_i and \tilde{X}_j respectively, then we obtain the same joint distribution over assignments in the first and second rounds for students with preferences θ as under mechanism M . In doing so, we will use the following “prefix lemma”.

The “prefix lemma” states that for every set of schools \mathcal{C}' , there exist certain intervals of the form $I_i^j = [C_i, C_j]$ such that for any two student types whose top set of acceptable schools under second-round preference reports is \mathcal{C}' , the proportion of students with lotteries in I_i^j who are upgraded to a school in \mathcal{C}' in the second round is the same for each type.

We define a *prefix* of preferences \succ to be a set of schools \mathcal{C}' that is a top set of acceptable schools under \succ ; that is, for all $i' \in \mathcal{C}'$ and $j \notin \mathcal{C}'$, it holds that $i' \succ j$.

Lemma A.2 (Prefix Property). *Let $i < j$ be schools, and let $\mathcal{C}' \not\ni i, j$ be a set of schools such that $i(\mathcal{C}') \leq i$. Let $\theta = (\succ, \hat{\succ})$ and $\theta' = (\succ', \hat{\succ}')$ be consistent preferences such that \mathcal{C}' is a prefix of $\succ, \hat{\succ}$ and some students with preferences θ are assigned to each of schools i and j in the first round, and similarly \mathcal{C}' is a prefix of $\succ', \hat{\succ}'$ and some students with preferences θ' are assigned to each of schools i and j in the first round. Then*

$$\rho^\theta([C_j, C_i], \mathcal{C}') = \rho^{\theta'}([C_j, C_i], \mathcal{C}').$$

That is, the proportion of students of type θ whose first-round lotteries are in the interval $[C_j, C_i]$ and who are assigned to a school in \mathcal{C}' in the second round is the same as the proportion of students of type θ' whose first-round lotteries are in the interval $[C_j, C_i]$ and who are assigned to a school in \mathcal{C}' in the second round.

Sketch of proof of Lemma A.2. The idea of the proof is to use weak strategy-proofness and first-order stochastic dominance to show that the probabilities of being assigned to C' (conditional on certain first-round assignments) are the same for students of type θ or θ' . We then invoke anonymity to argue that proportions of types of students assigned to a certain school are given by the conditional probabilities of individual students being assigned to that school. We present the full proof at the end of Section 3.3.1. \square

We now show that the mechanism M and the affordable set distribution $\gamma_{i,j}$ produce the same joint distribution of assignments.

Students with two acceptable schools.

To give a bit of the flavor of the proof, we first consider student preferences θ of the form $\succ: [k, l, n + 1]$ and $\hat{\succ}: [k, n + 1]$, where all other schools are unacceptable.

There are five ordered pairs of schools that students of this type can be assigned to in the two rounds. Namely, if we let (i, j) denote assignment to i in the first round and to j in the second round, then the ordered pairs are (k, k) , (l, k) , $(l, n + 1)$, $(n + 1, k)$, and $(n + 1, n + 1)$. Since the proportion of students with each first-round assignment is fixed, it suffices to show that the mechanism M and the mechanism that assigns first- and second-round affordable set distributions according to $\gamma_{i,j}$ produce the same proportion of students assigned to (l, k) and the same proportion of students assigned to $(n + 1, k)$.

Let $I_l^k = [C_l, C_k]$, and let $I_{n+1}^{\max\{k,l\}} = [0, C_{\max\{k,l\}}]$. The proportions of students with preferences θ who are assigned to (l, k) and $(n + 1, k)$ under M are given by $\rho^\theta([C_l, C_k], k)$ and $\rho^\theta([0, C_{\max\{k,l\}}], k)$ respectively. We want to show that this is the same as the proportion of students with preferences θ who are assigned to (l, k) and $(n + 1, k)$ respectively when first- and second-round affordable sets are given by the affordable set distribution $\gamma_{i,j}$. We remark that when $k > l$ this holds vacuously, since all the terms are 0. Hence, since for any school k the proportion of students with preferences θ who are assigned to k in the first round does not depend on θ , it suffices to consider the case where $k < l$.

Let $\theta' = (\succ', \hat{\succ}')$ be the preferences given by $\succ': [k, k + 1, \dots, l - 1, l, n + 1]$ and $\hat{\succ}': [k, n + 1]$, where only the schools between k and l are acceptable in the first round, only k is acceptable in

the second round, and all other schools are unacceptable.

For all pairs of indices i, j such that $j < i$, let $\theta'_{i,j} = (\succ_{i,j}, \hat{\succ}_{i,j})$ be the student preferences such that $\succ_{i,j} : [j, i-1, i, n+1]$ and $\hat{\succ}_{i,j} : [j, n+1]$, with all other schools unacceptable. (In the case where $i = j+1$, we let the first two schools under the preference ordering $\succ_{i,j}$ coincide.) We note that $\theta'_{i,j} = \theta_{i,\sigma^{-1}(j)}$, where $\theta_{i,j}$ was defined in (2b), and that for $i > \sigma(j)$ we previously defined $\rho_{i,j} = \sum_{l \leq j} \gamma_{i,l}$ to be the proportion of students of type $\theta_{i,j}$ whose first-round assignment is i and whose second-round assignment is school j .

The proportion of students with preferences θ who are assigned to (l, k) under M is given by

$$\begin{aligned}
\rho^\theta([C_l, C_k], k) &= \rho^{\theta'}([C_l, C_k], k) \text{ (by the prefix property (Lemma A.2), as } k < l) \\
&= \sum_{k < i \leq l} \rho^{\theta'}([C_i, C_{i-1}], k) \\
&= \sum_{k < i \leq l} \rho^{\theta'_{i,k}}([C_i, C_{i-1}], k) \\
&\quad \text{(since the second-round assignment does not depend on the first-round report)} \\
&= \sum_{k < i \leq l} \rho_{i,\sigma^{-1}(k)} \text{ (by the definition of } \rho_{i,\sigma^{-1}(k)}) \\
&= \sum_{k < i \leq l} \sum_{j \leq \sigma^{-1}(k)} \gamma_{i,j} \text{ (by the definition of } \gamma_{i,j}),
\end{aligned}$$

which is precisely the proportion of students with preferences θ who are assigned to (l, k) if the first- and second-round affordable sets are given by $\gamma_{i,j}$. Note that all $\theta'_{i,k}$ and $\rho_{i,\sigma^{-1}(k)}$ in the summation are well-defined, since the sum is over indices satisfying $k < i$.

Similarly, let $\theta'' = (\succ'', \hat{\succ}'')$ be the preferences given by $\succ'' : [k, l, l+1, \dots, n, n+1]$ and $\hat{\succ}'' : [k, n+1]$, where only k and the schools with indices greater than l are acceptable in the first round, only k is acceptable in the second round, and all other schools are unacceptable. Then the

proportion of students with preferences θ who are assigned to $(n + 1, k)$ under M is given by

$$\begin{aligned}
\rho^\theta([0, C_l], k) &= \rho^{\theta''}([0, C_l], k) \text{ (by the prefix property (Lemma A.2))} \\
&= \sum_{l < i \leq n} \rho^{\theta''}([C_i, C_{i-1}], k) \\
&= \sum_{l < i \leq n} \rho^{\theta'_{i,k}}([C_i, C_{i-1}], k) \\
&\quad \text{(since the second-round assignment does not depend on the first-round report)} \\
&= \sum_{l < i \leq n} \rho_{i, \sigma^{-1}(k)} \text{ (by the definition of } \rho_{i, \sigma^{-1}(k)}) \\
&= \sum_{l < i \leq n} \sum_{j \leq \sigma^{-1}(k)} \gamma_{i,j} \text{ (by the definition of } \gamma_{i,j}),
\end{aligned}$$

which is precisely the proportion of students with preferences θ who are assigned to $(n + 1, k)$ if the first- and second-round affordable sets are given by $\gamma_{i,j}$.

(2c.ii.) Students with general preferences.

We now consider general (consistent) student preferences θ of the form $(\succ, \hat{\succ})$, where

$$\succ: [i_1, i_2, \dots, i_k, n + 1] \text{ and } \hat{\succ}: [i_1, i_2, \dots, i_l, n + 1],$$

for some $k > l$ and where all other schools are unacceptable. We wish to show that for every pair of schools $i, i' \in \{i_1, i_2, \dots, i_k, n + 1\}$, the mechanism M and the mechanism that assigns first- and second-round affordable set distributions according to $\gamma_{i,j}$ produce the same proportion of students assigned to (i, i') . It suffices to show that for every prefix \mathcal{C}' of the preferences $\hat{\succ}$ and every school $k' \in \{i_2, \dots, i_k, n + 1\}$, the mechanism M and the mechanism that assigns first- and second-round affordable set distributions according to $\gamma_{i,j}$ produce the same proportion of students assigned to i in the first round and some school in \mathcal{C}' in the second round. We say that the students are assigned to (i, \mathcal{C}') .

Fix a prefix \mathcal{C}' of $\hat{\succ}$ and a school $i = i_{k'}$ satisfying $k' \leq k$. Let $l' \leq l$ be such that $\mathcal{C}' = \{i_1, i_2, \dots, i_{l'}\}$. If $k' \leq l'$ then $k' \in \mathcal{C}'$, and so in any mechanism that respects guarantees, the proportion of students with preferences θ assigned to (i, \mathcal{C}') is the same as the proportion of students

assigned to i in the first round.

Recall that $i(\mathcal{C}')$ is the school in \mathcal{C}' satisfying $i(\mathcal{C}') \geq i' \forall i' \in \mathcal{C}'$, i.e. the school in \mathcal{C}' that was least affordable in the first round. (Note that this is not necessarily i_{θ} , the school in \mathcal{C}' that is least preferred by a student of type θ .) If $k' > l'$ and $i \leq i(\mathcal{C}')$, then in the first round, whenever the school i is available in the first round, so is the preferred school $i(\mathcal{C}')$; thus, for any school i , the proportion of students assigned to i in the first round is 0. It follows that in any mechanism that respects guarantees, the proportion of students assigned to (i, \mathcal{C}') is 0.

From here on, we may assume that $k' > l'$ (i.e., $i \notin \mathcal{C}'$) and $i > i(\mathcal{C}')$. Since $i > i(\mathcal{C}')$, the proportion of students with preferences θ who are assigned to (i, \mathcal{C}') under M is given by $\rho^{\theta}([C_i, C_{i(\mathcal{C}')}], \mathcal{C}')$. Let $i(\sigma(\mathcal{C}'))$ be the school $j \in \mathcal{C}'$ such that $\sigma^{-1}(j)$ is maximal, that the school in \mathcal{C}' that is most affordable in the second round.

Let $\theta' = (\succ', \hat{\succ}')$ be the preferences given by \succ'

$$\hat{\succ}' : [i(\sigma(\mathcal{C}')), \mathcal{C}' \setminus i(\sigma(\mathcal{C}')), i(\mathcal{C}') + 1, i(\mathcal{C}') + 2, \dots, i - 1, i, n + 1],$$

where $i(\sigma(\mathcal{C}'))$ is the most preferred, followed by all schools in \mathcal{C}' and all other schools between $i(\mathcal{C}')$ and i are acceptable in the same order as first round overdemand, and

$$\hat{\succ}' : [i(\sigma(\mathcal{C}')), \mathcal{C}' \setminus i(\sigma(\mathcal{C}')), n + 1],$$

where $i(\sigma(\mathcal{C}'))$ is the most preferred and all other schools in \mathcal{C}' are ordered arbitrarily.

Since $k' > l'$, $i > i(\mathcal{C}')$, and the preferences θ are consistent, the preferences θ' are well defined. Let $\theta'' = (\succ'', \hat{\succ}'')$ be the preferences given by the same first-round preferences $\succ'' = \succ'$ as θ and second-round preference $\hat{\succ}'' : [i(\sigma(\mathcal{C}')), n + 1]$ that only find $i(\sigma(\mathcal{C}'))$ acceptable.

Recall that for all $j > i(\sigma(\mathcal{C}'))$, $\theta'_{j, i(\sigma(\mathcal{C}'))} = (\succ_{j, i(\sigma(\mathcal{C}'))}, \hat{\succ}_{j, i(\sigma(\mathcal{C}'))})$ are the student preferences such that

$$\succ_{j, i(\sigma(\mathcal{C}'))} : [i(\sigma(\mathcal{C}')), j - 1, j, n + 1] \text{ and } \hat{\succ}_{j, i(\sigma(\mathcal{C}'))} : [i(\sigma(\mathcal{C}')), n + 1],$$

with all other schools unacceptable. Additionally, recall that $\rho_{j, \sigma^{-1}(i(\sigma(\mathcal{C}'))})$ is the proportion of

students of type $\theta_{j,i(\sigma(\mathcal{C}'))}$ whose first-round assignment is j and whose second-round assignment is $i(\sigma(\mathcal{C}'))$. Let $\hat{\mathcal{C}} = \{i_1, i_2, \dots, i_{k'-1}\}$ be the schools in \mathcal{C}' that are preferred to school i under \succ and let $i(\hat{\mathcal{C}})$ be the school preferable to i under \succ that is most affordable in the second round.

Then the proportion of students with preferences θ who are assigned to (i, \mathcal{C}') under M is given by $\rho^\theta([C_i, C_{i(\hat{\mathcal{C}})}], \mathcal{C}')$, where

$$\begin{aligned}
\rho^\theta([C_i, C_{i(\hat{\mathcal{C}})}], \mathcal{C}') &= \rho^{\theta'}([C_i, C_{i(\hat{\mathcal{C}})}], \mathcal{C}') \text{ (by the prefix property (Lemma A.2) with prefix } \mathcal{C}') \\
&= \sum_{i(\hat{\mathcal{C}}) < j \leq i} \rho^{\theta'}([C_j, C_{j-1}], \mathcal{C}') \\
&= \sum_{i(\hat{\mathcal{C}}) < j \leq i} \rho^{\theta'}([C_j, C_{j-1}], i(\sigma(\mathcal{C}'))) \\
&\quad \text{(by the definition of the second-round overdemand ordering)} \\
&= \sum_{i(\hat{\mathcal{C}}) < j \leq i} \rho^{\theta''}([C_j, C_{j-1}], i(\sigma(\mathcal{C}'))) \text{ (by the prefix property with prefix } \{i(\sigma(\mathcal{C}'))\}) \\
&= \sum_{i(\hat{\mathcal{C}}) < j \leq i} \rho^{\theta'_{j,i(\sigma(\mathcal{C}'))}}([C_j, C_{j-1}], i(\sigma(\mathcal{C}'))) \\
&\quad \text{(since the second-round assignment does not depend on the first-round report)} \\
&= \sum_{i(\hat{\mathcal{C}}) < j \leq i} \rho_{j,\sigma^{-1}(i(\sigma(\mathcal{C}')))} \text{ (by the definition of } \rho_{j,\sigma^{-1}(i(\sigma(\mathcal{C}')))} \\
&= \sum_{i(\hat{\mathcal{C}}) < j \leq i} \sum_{j' \leq \sigma^{-1}(i(\sigma(\mathcal{C}')))} \gamma_{j,j'} \text{ (by the definition of } \gamma_{j,j'}),
\end{aligned}$$

which is precisely the proportion of students with preferences θ who are assigned to (i, \mathcal{C}') if the first- and second-round affordable sets are given by $\gamma_{j,j'}$. Note that all $\theta'_{j,i(\sigma(\mathcal{C}'))}$ and $\rho_{j,\sigma^{-1}(i(\sigma(\mathcal{C}')))}$ in the summation are well-defined, since the sum is over indices satisfying $j > i(\hat{\mathcal{C}})$, and since $k' > l'$ it follows that $\hat{\mathcal{C}} \supseteq \mathcal{C}'$ and hence $j > i(\hat{\mathcal{C}}) \geq i(\sigma(\mathcal{C}'))$.

(2d) Constructing the lottery L .

Fix a student s who reports first- and second-round preferences $\theta = (\succ, \hat{\succ})$. Suppose that s is assigned to schools (i, j) in the first and second rounds respectively. We first characterize all first- and second-round budget sets consistent with the overdemand orderings that could have led

to this assignment. Let $\underline{i} = \min \{i' \mid \max_{\succ} X_{i'} = i\}$, let $\underline{j} = \min \{j' \mid \max_{\succ} \tilde{X}_{j'} \cup \{i\} = j\}$, and let $\bar{j} = \max \{j' \mid \max_{\succ} \tilde{X}_{j'} \cup \{i\} = j\}$. Then the set of first- and second-round budget sets that student s could have been assigned by the mechanism is given by $\{X_{i'}, X_{j'} \cup \{i\} : \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}\}$. (We remark that the asymmetry in these definitions is due to the existence of the first-round guarantee in the second-round budget sets.)

Conditional on s being assigned to schools (i, j) in the first and second rounds respectively, we assign a lottery number $L(s)$ to s distributed uniformly over the union of intervals $\cup_{i', j': \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}} \tilde{I}_{i', j'}$,

$$(L(s) \mid (\mu(s), \tilde{\mu}(s)) = (i, j)) \sim \text{Unif}\left(\cup_{i', j': \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}} \tilde{I}_{i', j'}\right),$$

independent of all other students' assignments.

We show that this is consistent with the first round of the mechanism being RSD. We have shown in (1) that if for each pair of reported preferences $\theta = (\succ, \hat{\succ}) \in \Theta$, a uniform proportion $\gamma_{i', j'}$ of students with reported preferences θ are given first- and second-round budget sets $X_{i'}, \{i^\theta\} \cup \tilde{X}_{j'}$ (where $i^\theta = \max_{\succ} X_i$ is the first-round assignment of such students), we obtain the same distribution of assignments as M . Since M is anonymous and satisfies the averaging axiom, and since $|\tilde{I}_{i', j'}| = \gamma_{i', j'}$, it follows that each student's first-round lottery number is distributed as $\text{Unif}[0, 1]$.

Given the constructed lottery L , we construct the second-round cutoffs \hat{C}_i for the PLDA and verify that the assignment $\tilde{\mu}$ is feasible and stable with respect to the schools' second-round preferences, as defined by $P \circ L$ and the guarantee structure. Specifically, in PLDA, each student with a first-round score l and a first-round assignment i' has a second-round score $\hat{r}_i = P(l) + \mathbf{1}(i' = i)$ at each school $i \in \mathcal{C}$, and students are assigned to their favorite school i at which their second-round score exceeds the school's second-round cutoff, $\hat{r}_i \geq \hat{C}_i$ (or to the outside option $n + 1$).

Recall that the schools are indexed so that $C_1 \geq C_2 \geq \dots \geq C_{n+1}$, and that the permutation σ is chosen so that the second-round overdemand ordering is given by $\sigma(1), \sigma(2), \dots, \sigma(n + 1) = n + 1$, and so it should follow that the second-round cutoffs \hat{C}_i satisfy $\hat{C}_{\sigma(1)} \geq \hat{C}_{\sigma(2)} \geq \dots \geq \hat{C}_{\sigma(n+1)}$. By the characterization of stable assignments given by Azevedo and Leshno (2016), it suffices

to show that if each student with a first-round assignment i' and second-round lottery number in $[\hat{C}_{\sigma^{-1}(i)}, \hat{C}_{\sigma^{-1}(i-1)}]$ is assigned to her favorite school in $\{i'\} \cup \tilde{X}_i$, where we define $\tilde{X}_i = \{\sigma(i), \sigma(i+1), \dots, \sigma(n+1)\}$, then the resulting assignment $\hat{\mu}$ is equal to the second-round assignment $\tilde{\mu}$ of our mechanism M , and satisfies that $\eta(\hat{\mu}^{-1}(i)) \leq q_i$ for any school i , and $\eta(\hat{\mu}^{-1}(i)) = q_i$ if $\hat{C}_i > 0$.

For fixed i, j , let $\hat{C}_{\sigma(j)} = 1 - \sum_{i', j': j' \leq j} \gamma_{i', j'}$ and let $\hat{C}_{i, \sigma(j)} = \hat{C}_{\sigma(j-1)} - \sum_{i' \leq i} \gamma_{i', j}$. (We remark that since $\gamma_{i, j}$ refers to the i -th school to fill in the first round, i , and the j -th school to fill in the second round, $\sigma(j)$, the \hat{C} are indexed slightly differently than $\gamma_{i, j}$ is.)

We use the averaging assumption and the equivalence of assignment probabilities that we have shown in (I) to conclude that if $\hat{\mu}$ is the assignment given by running DA with round scores \hat{r} and cutoffs \hat{C} , then $\tilde{\mu} = \hat{\mu}$.

This is fairly evident, but we also show it explicitly below. Specifically, consider a student $s \in \mathcal{S}$ with a first-round lottery number $L(s)$ and reported preferences $\theta = (\succ, \hat{\succ})$. Let i, j be such that $L(s) \in \cup_{i', j': \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}} \tilde{I}_{i', j'}$, where $\underline{i} = \min \{i' \mid \max_{\succ} X_{i'} = i\}$, $\underline{j} = \min \{j' \mid \max_{\hat{\succ}} \tilde{X}_{j'} \cup \{i\} = j\}$, and $\bar{j} = \max \{j' \mid \max_{\hat{\succ}} \tilde{X}_{j'} \cup \{i\} = j\}$. Then, because of the way in which we have constructed the lottery L , it holds that $(\mu(s), \tilde{\mu}(s)) = (i, j)$.

Moreover, since

$$\begin{aligned} P(L(s)) &\in P(\cup_{i', j': \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}} \tilde{I}_{i', j'}) \\ &= \cup_{i', j': \underline{i} \leq i' \leq i, \underline{j} \leq j' \leq \bar{j}} P(\tilde{I}_{i', j'}) \end{aligned}$$

where $P(\tilde{I}_{i', j'}) \in [\hat{C}_{\sigma(j')}, \hat{C}_{\sigma(j'-1)}]$, it holds that under $\hat{\mu}$, student s receives her favorite school in $\{i\} \cup \tilde{X}_{j'}$ for some $\underline{j} \leq j' \leq \bar{j}$, which is the school j . Hence $\tilde{\mu}(\lambda) = \hat{\mu}(\lambda) = j$. It follows immediately that the assignment $\hat{\mu}$ is feasible, since it is equal to the feasible assignment $\tilde{\mu}$.

Finally, let us check that the assignment is stable. Suppose that $\hat{C}_j > 0$. We want to show that $\eta(\tilde{\mu}^{-1}(j)) = q_j$. First note that it follows from the definition of \hat{C}_j that

$$1 > \sum_{i', j': j' \leq \sigma^{-1}(j)} \gamma_{i', j'} = \sum_{i'} \rho_{i', \sigma^{-1}(j)}.$$

Consider student preferences $\theta = (\succ, \succ)$ given by $\succ: [j, 1, 2, \dots, j-1, j+1, \dots, n+1]$. Then $\sum_{i'} \rho_{i', \sigma^{-1}(j)}$ is the proportion of students of type θ who are assigned to school j in the second round, which, by assumption, is also the probability that a student with preferences θ is assigned to j in the second round. But since M is non-wasteful, this means that $\eta(\tilde{\mu}^{-1}(j)) = q_j$. It follows from constrained Pareto efficiency that the output of M is the student-optimal stable matching.

Proof of Lemma A.2. Here, we prove the prefix property. We first observe that any schools reported to be acceptable but ranked below j in the first round are inconsequential. Moreover, since M respects guarantees, weak two-round strategy-proofness implies that any schools reported to be acceptable but ranked below j in the second round are inconsequential. Hence it suffices to prove the lemma for first-round preference orderings \succ and \succ' for which j is the last acceptable school.

Suppose that the lemma holds for $i = i(\mathcal{C}')$. Then if $i(\mathcal{C}') = i' < i$ it holds that

$$\begin{aligned} \rho^\theta([C_j, C_i], \mathcal{C}') &= \frac{\rho^\theta([C_j, C_{i(\mathcal{C}')}], \mathcal{C}') (C_{i(\mathcal{C}')} - C_j) - \rho^\theta([C_i, C_{i(\mathcal{C}')}], \mathcal{C}') (C_{i(\mathcal{C}')} - C_i)}{C_i - C_j} \\ &= \frac{\rho^{\theta'}([C_j, C_{i(\mathcal{C}')}], \mathcal{C}') (C_{i(\mathcal{C}')} - C_j) - \rho^{\theta'}([C_i, C_{i(\mathcal{C}')}], \mathcal{C}') (C_{i(\mathcal{C}')} - C_i)}{C_i - C_j} \\ &= \rho^{\theta'}([C_j, C_i], \mathcal{C}'), \end{aligned}$$

where the first and last equalities follow from Bayes' rule, and the second equality holds since the lemma holds for $i = i(\mathcal{C}')$, and the theorem follows. Hence it suffices to prove the lemma for $i = i(\mathcal{C}')$.

Let i_1, \dots, i_k be the indices of the schools in \mathcal{C}' , in increasing order. We observe that $i = i(\mathcal{C}') < j$.

Since we wish to prove that the lemma holds for all pairs θ, θ' satisfying the assumptions, it suffices to show that the lemma holds for a fixed preference θ when we vary only θ' . Therefore, we may, without loss of generality, fix the preferences θ to satisfy that

$$\succ: [i, i_1, \dots, i_{k-1}, j, n+1] \text{ and } \hat{\succ}: [i, i_1, \dots, i_{k-1}, n+1],$$

and all other schools are unacceptable. That is, type θ prefers first school $i = i(\mathcal{C}')$, which is the least overdemanded school in \mathcal{C}' , and then all other schools in \mathcal{C}' in the same order as the overdemand ordering. In the first round school j is also acceptable, and in the second round only schools in \mathcal{C}' are acceptable.

We remark that given the first-round ordering, the worst school in \mathcal{C}' and the school j (namely, i and j) are the only acceptable schools to which students of type θ will be assigned in the first round. Moreover, the proportion of students with preferences θ (or θ') who can deduce that their score is in $[C_j, C_i]$ is precisely $C_i - C_j$, since such students are assigned in the first round to some school not in \mathcal{C}' that they weakly prefer to j , and all such schools are between i and j in the overdemand ordering. Similarly, the proportion of students with preferences θ (or θ') who can deduce that their lottery number is in $[C_i, 1]$ is precisely $1 - C_i$, since such students are assigned in the first round to a school in \mathcal{C}' . (Note that students with preferences θ' may be able to deduce that their lottery number falls in a subinterval of the interval we have specified. However, this does not affect our statements.)

To compare the proportion of students of types θ and θ' whose scores are in $[C_j, C_i]$ and who are assigned to \mathcal{C}' in the second round, we define a third student type θ'' as follows. Let $\theta'' = (\succ', \succ)$ be a set of preferences where the first-round preferences are the same as the first-round preferences of θ' , and the second-round preferences are the same as the second-round preferences of type θ .

Let s be a student with preferences θ , and similarly let s' be a student with preferences θ' . We use the two-round strategy-proofness of the mechanism to show that s has the same probability of being assigned to some school in \mathcal{C}' in the second round as if she had reported type θ'' , and similarly for s' . Since the proportion of students of either type being assigned to a school in \mathcal{C}' in the first round is the same and the mechanism respects guarantees, this is sufficient to prove the prefix property.

Formally, let ρ be the probability that s is assigned to some school in \mathcal{C}' in the second round if she reports truthfully, conditional on being able to deduce that her first-round score is in $[C_j, C_i]$, and let ρ' be the probability that s' is assigned to some school in \mathcal{C}' in the second round if she reports truthfully, conditional on being able to deduce that her first-round score is in $[C_j, C_i]$. (We

note that given her first-round assignment $\mu(\rho')$, the student s' may actually be able to deduce more about her first-round score, and so the interim probability after knowing her assignment that s' is assigned to some school in \mathcal{C}' in the second round if she reports truthfully is not necessarily ρ' .) Let ρ'' be the probability that a student with preferences θ'' and a first-round score in $[C_j, C_i]$ chosen uniformly at random is assigned to some school in \mathcal{C}' in the second round. It follows from the design of the first round and from anonymity that ρ is the probability that a student with preferences θ and a lottery number in $[C_j, C_i]$ chosen uniformly at random is assigned to some school in \mathcal{C}' in the second round, and similarly for ρ' .

Proving the lemma is equivalent to proving $\rho = \rho'$. We show that $\rho = \rho'' = \rho'$. Note that the first equality is between preferences that are identical in the second round, and the second equality is between preferences that are identical in the first round.

We first show that $\rho = \rho''$; that is, changing just the first-round preferences does not affect the probability of assignment to \mathcal{C}' . This is almost immediate from first-order stochastic dominance of truthful reporting, since the second-round preferences under θ and θ'' are identical. (This also illustrates the power of the assumption that the second-round assignment does not depend on first-round preferences. It implies that manipulating first-round reports to obtain a more fine-grained knowledge of the lottery number does not help, since assignment probabilities are conditionally independent of the lottery number.) We present the full argument below.

Let $\hat{\rho}$ be the probability that a student with preferences θ who is unassigned in the first round is assigned to a school in \mathcal{C}' in the second round. We note that since the last acceptable school under preferences θ and θ'' is j , the set of students with preferences θ who are unassigned in the first round is equal to the set of students with preferences θ with lottery number in $[0, C_j]$, and similarly the set of students with preferences θ'' who are unassigned in the first round is equal to the set of students with preferences θ'' with lottery number in $[0, C_j]$. Hence, the fact that θ and θ'' have the same second preferences gives us that $\hat{\rho}$ is also the probability that a student with preferences θ'' who is unassigned in the first round is assigned to a school in \mathcal{C}' in the second round.

The probability of being assigned in the second round to a school in \mathcal{C}' when reporting θ is

given by:

$$(1 - C_i) + (C_i - C_j)\rho + C_j\hat{\rho},$$

The probability of being assigned in the second round to a school in \mathcal{C}' when reporting θ'' is given by:

$$(1 - C_i) + (C_i - C_j)\rho'' + C_j\hat{\rho}.$$

It follows from first-order stochastic dominance of truthful reporting for types θ and θ' that $\rho = \rho''$.

We now show that $\rho' = \rho''$. This is a little more involved, but essentially relies on breaking the set of students with first-round score in $[C_j, C_i]$ into smaller subsets, depending on their first-round assignment, and using first-order stochastic dominance of truthful reporting to show that in each subset, the probability of an arbitrary student being assigned to a school in \mathcal{C}' in the second round is the same for students with either set of preferences θ' or θ'' .

We first introduce some notation for describing the first-round preferences of θ' and θ'' . Let $\{j_1 \leq \dots \leq j_m\}$ be the schools between $i(\mathcal{C}')$ and j in the overdemand ordering, corresponding to schools that a student with preferences θ' and a lottery number in $[C_j, C_i]$ could have been assigned to in the first round. Formally, we define them to be the indices j' for which $j' \notin \mathcal{C}'$, $i(\mathcal{C}') < j' \leq j$, $j' \succeq' j$ and j' is relevant in the first-round overdemand ordering, that is, $k' < j'$ for all k' such that $k' \succ' j'$. We observe that $j_m = j$. For $l = 1, \dots, m$, let ρ'_l be the probability that a student with preferences θ' who was assigned to school j_l is assigned to a school in \mathcal{C}' in the second round.

The set of students with preferences θ' assigned to school j_l in the first round is precisely the set of students with preferences θ' whose first-round lottery number is in $[C_{j_l}, C_{j_{l-1}}]$ and similarly the set of students with preferences θ'' assigned to school j_l in the first round is precisely the set of students with preferences θ'' whose first-round lottery number is in $[C_{j_l}, C_{j_{l-1}}]$. If we define $j_0 = i$, it follows that $(C_i - C_j) = \sum_{l=1}^m (C_{j_{l-1}} - C_{j_l})$, and that

$$(C_i - C_j)\rho' = \sum_{l=1}^m (C_{j_{l-1}} - C_{j_l})\rho'_l.$$

Let ρ''_l be the probability that a student with preferences θ'' who was assigned to school j_l is

assigned to a school in C' in the second round. Then it also holds that

$$(C_i - C_j)\rho'' = \sum_{l=1}^m (C_{j_{l-1}} - C_{j_l})\rho''_l.$$

We show now that $\rho'_l = \rho''_l$ for all l , which implies that $\rho' = \rho''$.

Consider a student s_l who reported $\succ' = \succ''$ in the first round and was assigned to school j_l . Note that such a report is consistent with either reporting θ' or θ'' , and since the first-round reports of these types are the same and the first-round mechanism is DA-STB there exists some set of lottery numbers L_l such that students of type θ' or θ'' are assigned to j_l in the first round if and only if their lottery lies in L_l . The probabilities that this student is assigned in the second round to a school in C' when reporting θ' and θ'' are given by ρ'_l and ρ''_l respectively. Now for any fixed lottery $L(s)$, truthful reporting is a dominant strategy in the second round for types θ and θ' . It follows that $\rho'_l = \rho''_l$.

This completes the proof of the lemma. □

A.4 Proof of Theorem 3.4

Suppose that the order condition holds. In what follows, we will fix a permutation P and show that PLDA(P) satisfies the local order condition and is type-equivalent to the reverse lottery RLDA mechanism. As this holds for every P , it follows that the global order condition holds.

(1) *Every school has a single priority group.*

We first consider the case where $n_i = 1$ for all i ; that is, every school has a single priority group. Recall that the schools are indexed according to the first-round overdemand ordering, so that $C_1 \geq C_2 \geq \dots \geq C_n \geq C_{n+1}$. Since the local order condition holds for RLDA, let us assume that they are also indexed according to the second-round overdemand ordering under RLDA, so that $\hat{C}_1^R \geq \hat{C}_2^R \geq \dots \geq \hat{C}_n^R \geq \hat{C}_{n+1}^R$.

The idea will be to construct a set of cutoffs \tilde{C}^P directly from the permutation P and the cutoffs \hat{C}^R , show that the cutoffs are in the correct order $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \dots \geq \tilde{C}_n^P \geq \tilde{C}_{n+1}^P$, and show that the cutoffs \tilde{C}^P and resulting assignment are market-clearing when school preferences are given by

the amended scoring function with permutation P .

(1a) Definitions.

As in the proof of Theorem 3.1, let $\beta_{i,j} = \eta(\{s \in \mathcal{S} : \operatorname{argmax}_{\zeta} X_j = i\})$ be the measure of students who, when their set of affordable schools is X_j , will choose i . Let $E^s(\mathbf{C})$ be the set of schools affordable for student s in the first round under PLDA with any permutation, let $\hat{E}^s(\hat{\mathbf{C}}^R)$ be the set of schools affordable for type s in the second round under RLDA, and let $\hat{E}^s(\hat{\mathbf{C}}^P)$ be the set of schools affordable for type s in the second round under PLDA(P).

Let $\gamma_i^R = \eta(\{s \in \mathcal{S} : \hat{E}^s(\hat{\mathbf{C}}^R) = X_i\})$ be the fraction of students whose affordable set in the second round of RLDA is X_i , and let $\gamma_i^P = \eta(\{s \in \mathcal{S} : \hat{E}^s(\hat{\mathbf{C}}^P) = X_i\})$ be the fraction of students whose affordable set in PLDA(P) is X_i .

Let \hat{n} be the smallest index such that school \hat{n} does not reach capacity when it is not offered to all the students. In other words, \hat{n} is the smallest index such that every student has school \hat{n} in her affordable set under RLDA, i.e., $\hat{n} \in \hat{E}^s(\hat{\mathbf{C}}^R)$. Since the local order condition holds for RLDA, we may equivalently express \hat{n} in terms of cutoffs as the smallest index such that $(1 - C_{\hat{n}}) + (1 - \hat{C}_{\hat{n}}^R) \geq 1$. Such an \hat{n} always exists, since every student has the outside option $n + 1$ in her affordable set.

(1b) Defining cutoffs for PLDA.

Let us define cutoffs $\tilde{\mathbf{C}}^P$ as follows. For $i \geq \hat{n}$ let $\tilde{C}_i^P = 0$. For each permutation P , define a function

$$f_i^P(x) = |\{\ell : \ell \geq C_i \text{ or } P(\ell) \geq x\}|$$

representing the proportion of students who have i in their affordable set with first- and second-round cutoffs C_i, x under the amended scoring function with permutation P . Since P is measure-preserving, $f_i^P(x)$ is continuous and monotonically decreasing in x .

For $i < \hat{n}$, we inductively define \tilde{C}_i^P to be the largest real smaller than \tilde{C}_{i-1}^P satisfying

$$f_i^P(\tilde{C}_i^P) = f_i^R(\hat{C}_i^R) \tag{A.2}$$

(where we define $\tilde{C}_0^P = 1$). Now $f_i^P(0) = 1 \geq f_i^R(\hat{C}_i^R)$, and

$$\begin{aligned}
f_i^P(\tilde{C}_{i-1}^P) &= f_{i-1}^P(\tilde{C}_{i-1}^P) + |\{l \mid l \in [C_i, C_{i-1}) \text{ and } P(l) \geq \tilde{C}_{i-1}^P\}| \\
&\leq f_{i-1}^R(\hat{C}_{i-1}^R) + (C_{i-1} - C_i) \\
&= (1 - C_i) + (1 - \hat{C}_{i-1}^R) \\
&\leq f_i^R(\hat{C}_i^R) = f_i^P(\tilde{C}_i^P)
\end{aligned}$$

where in the first equality we are using that $C_{i-1} \geq C_i$, the first inequality follows from the definition of \tilde{C}_{i-1}^P , and the last inequality holds since $\hat{C}_{i-1}^R \geq \hat{C}_i^R$.

It follows from the intermediate value theorem that the cutoffs \tilde{C}^P are well defined and satisfy $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \dots \geq \tilde{C}_n^P \geq \tilde{C}_{n+1}^P$.

(1c) The constructed cutoffs clear the market.

We show that the cutoffs \tilde{C}^P and resulting assignment (from letting students choose their favorite school out of those for which they meet the cutoff) are market-clearing when the second-round scores are given by $\hat{r}_i^s = P(L(s)) + n_i \mathbf{1}_{\{L(s) \geq C_i\}} + p_i^s \mathbf{1}_{\{L(s) \geq C_i\}}$. We call the mechanism with this second-round assignment M^P .

The idea is that since the cutoffs \tilde{C}_i^P are decreasing in the same order as C_i and \hat{C}_i^R , the affordable sets are nested in the same order under both sets of second-round cutoffs. It follows that aggregate student demand is uniquely specified by the proportion of students with each school in their affordable set, and we have defined these to be equal, $f_i^P(\tilde{C}_i^P) = f_i^R(\hat{C}_i^R)$. It follows that \tilde{C}^P are market-clearing and give the PLDA(P) cutoffs, and so PLDA(P) satisfies the local order condition (with the indices indexed in the same order as with RLDA). We make the affordable set argument explicit below.

Consider the proportion of lottery numbers giving a (total) affordable set X_i . Since $\hat{C}_1^R \geq \hat{C}_2^R \geq \dots \geq \hat{C}_n^R$, under RLDA this is given by

$$\gamma_i^R = f_{i+1}^R(\hat{C}_{i+1}^R) - f_i^R(\hat{C}_i^R),$$

if $i < \hat{n}$ and by 0 if $i > \hat{n}$, where we define $f_0^P(x) = 1$ for all P and x . Similarly, since $\tilde{C}_1^P \geq \tilde{C}_2^P \geq \dots \geq \tilde{C}_n^P$, under M^P this is given by

$$f_{i+1}^P(\tilde{C}_{i+1}^P) - f_i^P(\tilde{C}_i^P)$$

if $i < \hat{n}$, which is precisely γ_i^R , and by 0 if $i > \hat{n}$.

Hence, for all $i < \hat{n}$, the measure of students assigned to school i under both RLDA and M^P is $\sum_{j \leq i} \beta_{i,j} \gamma_j^R = q_i$, and for all $i \geq \hat{n}$, the measure of students assigned to school i is $\sum_{j \leq \hat{n}} \beta_{i,j} \gamma_j^R < q_i$. It follows that the cutoffs \tilde{C}^P are market-clearing when the second-round scores are given by $\hat{r}_i^s = P(L(s)) + n_i \mathbf{1}_{\{L(s) \geq C_i\}} + p_i^s \mathbf{1}_{\{L(s) \geq C_i\}}$, and so PLDA(P) = M^P satisfies the local order condition.

(2) Some school has more than one priority group.

Now consider when schools have possibly more than one priority group. We show that if RLDA satisfies the local order condition, then PLDA(P) assigns the same number of students of each priority type to each school i as RLDA, and within each priority type assigns the same number of students of each preference type to each school as RLDA. We do this by first assuming that PLDA(P) assigns the same number of students of each priority type to each school i as RLDA, and showing that this gives consistent cutoffs.

We note that this proof uses very similar arguments to the proof of Theorem A.1.

(2a) Defining little economies \mathcal{E}_π for each priority type.

Fix a priority class π . Let q_π be a restricted capacity vector, where $q_{\pi,i}$ is the measure of students of priority class π assigned to school i under RLDA. Let \mathcal{S}_π be the set of students s such that $p^s = \pi$, and let η_π be the restriction of the distribution η to \mathcal{S}_π . Let \mathcal{E}_π denote the primitives $(\mathcal{CS}_\pi, \eta_\pi, q_\pi)$.

Let \tilde{C}_π^P be the second-round cutoffs of PLDA(P) on \mathcal{E}_π . By definition, \hat{C}_π^R are the second-round cutoffs of RLDA on \mathcal{E}_π . We show that the cutoffs \tilde{C}_π^P defined for the little economy are the same as the consistent second-round cutoffs \hat{C}_π^P for PLDA(P) run on the large economy \mathcal{E} , that is, $\tilde{C}_\pi^P = \hat{C}_\pi^P$.

(2b) Implications of RLDA satisfying the local order condition.

Since RLDA satisfies the local order condition for \mathcal{E} , RLDA also satisfies the local order condition for \mathcal{E}_π for all π . It follows from (1) that the global order condition holds on each of the little economies \mathcal{E}_π . Hence by Theorem 3.1 all PLDA mechanisms produce type-equivalent assignments when run on \mathcal{E}_π . Moreover, as in the proof of Theorem A.1, the global order condition on \mathcal{E}_π also allows us to determine aggregate student demand in \mathcal{E}_π from the proportions of students who have each school in their affordable set.

Finally, as in the proof of Theorem A.1, we may assume that for each π and school i the cutoff $\tilde{C}_{\pi,i}^P$ is the minimal real satisfying $f_{\pi,i}^P(\tilde{C}_{\pi,i}^P) = f_{\pi,i}^R(\hat{C}_{\pi,i}^R)$, where for each permutation P , $f_{\pi,i}^P(x) = |\{l : l \geq C_{\pi,i} \text{ or } P(l) \geq x\}|$ is the proportion of students of priority class π who have school i in their affordable set if the first- and second-round cutoffs are $C_{\pi,i}$ and x respectively.

It follows that $\tilde{C}_{\pi,i}^P$ depends on π only via π_i , and does not depend on π_j for all $j \neq i$. This is because $\tilde{C}_{\pi,i}^P$ is defined by $f_{\pi,i}^P(\cdot)$, $f_{\pi,i}^R(\cdot)$, and $\hat{C}_{\pi,i}^R$, which are in turn defined by $C_{\pi,i}$ and the permutations P and R . Moreover, $C_{\pi,i}$ depends on π only through π_i . Hence, if π, π' are two priority vectors such that $\pi_i = \pi'_i$, then $\tilde{C}_{\pi,i}^P = \tilde{C}_{\pi',i}^P$, and so the $\tilde{C}_{\pi,i}^P$ are consistent across priority types and define cutoffs \tilde{C}_i^P that are independent of priority type.

(3) \tilde{C}_i^P are the PLDA cutoffs.

Finally, we show that \tilde{C}_i^P are market-clearing cutoffs. By (1), for each priority class π , the number of students assigned to each school i is the same under RLDA as under the demand induced by the cutoffs \tilde{C}_i^P , and we know that the RLDA cutoffs are market-clearing for \mathcal{E} .

Hence \tilde{C}_i^P give the assignments for PLDA on \mathcal{E} , and since \tilde{C}_i^P was defined individually for each priority class π for \mathcal{E}_π it follows that PLDA(P) assigns the same measure of students of each priority type to each school i as RLDA.

Appendix B

Appendix for Chapter 4

B.1 Omitted Proofs for Section 4.3

Definitions and Notation

We begin with some additional definitions and notation that will be used in the proofs in this section.

In Appendix 4.4.1 we outlined how the TTC path γ can be interpreted as tracking the progression of the algorithm. Throughout the proofs, we make use of this interpretation and will frequently fix an economy \mathcal{E}^1 and a TTC path γ and let $TTC(\gamma|\mathcal{E})$ denote the continuous-time algorithm given by the path γ on the economy \mathcal{E} .² Given a path γ , let $\{t^{(i)}\}_{i \in \mathcal{C}}$ be stopping times such that γ and $\{t^{(i)}\}_{i \in \mathcal{C}}$ satisfy the capacity equations. Let the schools be labeled such that $t^{(i_1)} \leq t^{(i_2)} \leq \dots \leq t^{(i_n)}$, and let $t^{(i_0)} = 0$. We will refer to the progression of the algorithm from time $t^{(i_{\ell-1})}$ to time $t^{(i_\ell)}$ as *Round ℓ of $TTC(\gamma)$* .

Let \underline{x}, \bar{x} be vectors. We let $(\underline{x}, \bar{x}] = \{x : x \not\leq \underline{x} \text{ and } x \leq \bar{x}\}$ denote the set of vectors that are weakly smaller than \bar{x} along every coordinate, and strictly larger than \underline{x} along some coordinate. Let $K \subseteq \mathcal{C}$ be a set of schools. For all vectors x , we let $\pi_K(x)$ denote the projection of x to the coordinates indexed by schools in K .

The following notation is used to incorporate information about the set of available schools. For an economy E and TTC path γ yielding TTC cutoffs \mathbf{p} we let $C(x) = \{j \mid \exists i \text{ s.t. } p_i^j \leq x_i\}$ denote the set of schools available to students with rank x . We denote by

$$\Theta^{i|C} = \{\theta \in \Theta \mid Ch^\theta(C) = i\}$$

the set of students whose top choice in C is i , and denote by $\eta^{i|C}$ the measure of these students. That is, for $S \subseteq \Theta$, let $\eta^{i|C}(S) := \eta(S \cap \Theta^{i|C})$. In an abuse of notation, for a set $A \subseteq [0, 1]^{\mathcal{C}}$, we will often let $\eta(A)$ denote $\eta(\{\theta \in \Theta \mid r^\theta \in A\})$, the measure of students with ranks in A , and let $\eta^{i|C}(A)$ denote $\eta(\{\theta \in \Theta^{i|C} \mid r^\theta \in A\})$, the measure of students with ranks in A whose top choice

¹The economy \mathcal{E} can either be a continuum economy, or a discrete economy E , in which case we let $TTC(\gamma|E)$ denote $TTC(\gamma|\Phi(E))$.

²We will omit the dependence on the economy when it is evident from context.

school in C is i .

We will also find it convenient to define sets of students who were offered or assigned a seat along some TTC path γ . These will be useful in considering the result of aggregating the marginal trade balance equations. For each time τ let

$$\mathcal{T}_i(\gamma; \tau) \stackrel{def}{=} \{\theta \in \Theta \mid \exists \tau' \leq \tau \text{ s.t. } r_i^\theta = \gamma_i(\tau') \text{ and } r^\theta \leq \gamma(\tau')\}$$

denote the set of students who were offered a seat by school i before time τ , let

$$\mathcal{T}^i(\gamma; \tau) \stackrel{def}{=} \{\theta \in \Theta \mid r^\theta \not\leq \gamma(\tau) \text{ and } Ch^\theta(C(r^\theta)) = i\}$$

denote the set of students who were assigned a seat at school i before time τ , and let $\mathcal{T}^{i|C}(\gamma; \tau) \stackrel{def}{=} \{\theta \in \Theta \mid r^\theta \not\leq \gamma(\tau) \text{ and } Ch^\theta(C) = i\}$ denote the set of students who would be assigned a seat at school i before time τ if the set of available schools was C and the path followed was γ .³

For each interval $T = [\underline{t}, \bar{t}]$ let $\mathcal{T}_i(\gamma; T) \stackrel{def}{=} \mathcal{T}_i(\gamma; \bar{t}) \setminus \cup_{t < \underline{t}} \mathcal{T}_i(\gamma; t)$ be the set of students who were offered a seat by school i at some time $\tau \in T$, and let $\mathcal{T}^{i|C}(T; \gamma) \stackrel{def}{=} \mathcal{T}^{i|C}(\gamma; \bar{t}) \setminus \mathcal{T}^{i|C}(\gamma; \underline{t})$ be the set of students who were assigned to a school i at some time $\tau \in T$, given that the set of available schools was $C(\gamma(\tau)) = C$ for each $\tau \in T$. For each union of disjoint intervals $T = \cup_n T_n$ similar define $\mathcal{T}_i(\gamma; T) \stackrel{def}{=} \cup_n \mathcal{T}_i(\gamma; T_n)$ and $\mathcal{T}^{i|C}(T; \gamma) \stackrel{def}{=} \cup_n \mathcal{T}^{i|C}(T_n; \gamma)$. Figure B.1 illustrates examples of \mathcal{T}_i and \mathcal{T}^i for an economy with two schools.

Finally let us set up the definitions for solving the marginal trade balance equations. For a set of schools C and individual schools $i, j \in C$, recall that

$$\begin{aligned} H_i^{j|C}(x) &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \eta \left(\left\{ \theta \in \Theta \mid r^\theta \in [(x_i - \varepsilon) \cdot e^i, x] \text{ and } Ch^\theta(C) = j \right\} \right) \\ &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \eta \left(\left\{ \theta \in \Theta^{j|C} \mid r^\theta \in [(x_i - \varepsilon) \cdot e^i, x] \right\} \right) \end{aligned}$$

is the marginal density of students pointed to by school i at the point x whose top choice school in

³Note that $\mathcal{T}_i(\gamma; \tau)$ and $\mathcal{T}^i(\gamma; \tau)$ include students who were offered or assigned a seat in the school in previous rounds.

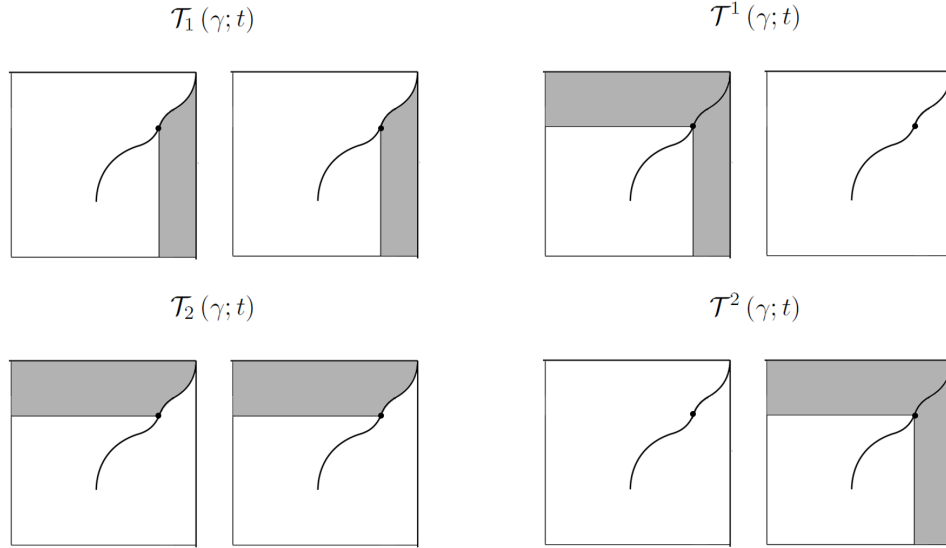


Figure B.1: The sets $\mathcal{T}_i(\gamma; t)$ and $\mathcal{T}^i(\gamma; t)$ for an economy with two schools and a fixed path γ and time t . $\mathcal{T}_i(\gamma; t)$ denotes the set of students who were offered a seat by school i by time t , and $\mathcal{T}^i(\gamma; t)$ denotes the set of students who were assigned to school i by time t . Students in each set are shaded in grey. Note that students are no longer offered seats once they are assigned, and so only students with priorities on the path γ are offered seats by both schools.

\mathcal{C} is j .

Let $H^C(x)$ be the $|\mathcal{C}| \times |\mathcal{C}|$ matrix with (i, j) th entry $H^C(x)_{i,j} = H_i^{j|\mathcal{C}}(x)$. Let $\tilde{H}^C(x)$ be the $|\mathcal{C}| \times |\mathcal{C}|$ matrix with (i, j) th entry

$$\tilde{H}^C(x)_{i,j} = \frac{1}{\bar{v}} H_i^{j|\mathcal{C}}(x) + \mathbf{1}_{i=j} \left(1 - \frac{v_j}{\bar{v}}\right),$$

where $v_j = \sum_{k \in \mathcal{C}} H_j^{k|\mathcal{C}}(x)$ is the row sum of $H(x)$, and the normalization \bar{v} satisfies $\bar{v} \geq \max_j v_j$. $\tilde{H}^C(x)$ is a transformation of $H^C(x)$ that will be convenient for formalizing the connection with continuous time Markov chains presented in Appendix 4.4.3.

Recall that a TTC path γ satisfies the trade balance equations for an economy $E = (\mathcal{C}, \Theta, \eta, q)$ if the following holds:

$$\sum_{k \in \mathcal{C}} \gamma'_k(t) H_k^i(\gamma(t)) = \sum_{k \in \mathcal{C}} \gamma'_i(t) H_i^k(\gamma(t)) \quad \forall i \in \mathcal{C}, \text{ times } t.$$

These may be equivalently stated in terms of the matrix $\tilde{H}(\gamma(t))$ as follows:

$$\gamma'(t) = \gamma'(t) \cdot \tilde{H}(\gamma(t)).$$

Let $\gamma(\tau) = x$. If $\mathbf{d} = -\gamma'(\tau) \geq 0$ solves the trade balance equations for x with available schools C

$$\sum_{k \in C} d_k \cdot H_k^{i|C}(x) = \sum_{k \in C} d_i \cdot H_i^{k|C}(x) \quad \forall i \in C,$$

or equivalently

$$\mathbf{d} = \mathbf{d} \cdot \tilde{H}(x)$$

we say that \mathbf{d} is a *valid gradient at x with available schools C* , and if in addition $\mathbf{d} \cdot \mathbf{1} = -1$ then we say that \mathbf{d} is a *valid direction at x with available schools C* . We omit the references to x and C when they are clear from context.

Let $M^C(x)$ be the Markov chain with state space C , and transition probability from state i to state j equal to $\tilde{H}^C(x)_{i,j}$. We remark that such a Markov chain exists, since $\tilde{H}^C(x)$ is a (right) stochastic matrix for each pair C, x .

We will also need the following definitions. For a matrix H and sets of indices I, J we let $H_{I,J}$ denote the submatrix of H with rows indexed by elements of I and columns indexed by elements of J . Recall that, by Assumption 4.1, the measure η is defined by a probability density ν that is right-continuous and piecewise Lipschitz continuous with points of discontinuity on a finite grid. Let the finite grid be the set of points $\{x \mid x_i \in D_i \forall i\}$, where the D_i are finite subsets of $[0, 1]$. Then there exists a partition \mathcal{R} of $[0, 1]^C$ into hyperrectangles such that for each $R \in \mathcal{R}$ and each face of R , there exists an index i and $y_i \in D_i$ such that the face is contained in $\{x \mid x_i = y_i\}$.

The following notion of continuity will be useful, given this grid-partition. We say that a multivariate function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *right-continuous* if $f(x) = \lim_{y \rightarrow x, y \geq x} f(y)$, where x, y are vectors in \mathbb{R}^n and the inequalities hold coordinate-wise. For an $m \times n$ matrix A , let $\mathbf{1}(A)$ be the

$m \times n$ matrix with entries

$$\mathbf{1}(A)_{ij} = \begin{cases} 1 & \text{if } A_{ij} \neq 0, \\ 0 & \text{if } A_{ij} = 0. \end{cases}$$

We will want some way of comparing two TTC paths γ and $\tilde{\gamma}$ obtained under two continuum economies differing only in their measures η and $\tilde{\eta}$.

Definition B.1. Let γ and $\tilde{\gamma}$ be increasing continuous functions from $[0, 1]$ to $[0, 1]^{\mathcal{C}}$ with $\gamma(0) = \tilde{\gamma}(0)$. We say that $\gamma(\tau)$ is dominated by $\tilde{\gamma}(\tau)$ via school i if

$$\begin{aligned} \gamma_i(\tau) &= \tilde{\gamma}_i(\tau), \text{ and} \\ \gamma_j(\tau) &\leq \tilde{\gamma}_j(\tau) \text{ for all } j \in \mathcal{C}. \end{aligned}$$

We also say that γ is dominated by $\tilde{\gamma}$ via school i at time τ . If γ and γ' are TTC paths, we can interpret this as school j being less demanded under γ , since with the same rank at j , in γ students are competitive with fewer ranks at other schools i . Equivalently, the same rank at j is less valuable under γ than under $\tilde{\gamma}$, as it provides the same opportunities for assignment as lower ranks at other schools (i.e. worse opportunities) under γ compared to $\tilde{\gamma}$. Another interpretation is that more students have been offered seats by the time t at which we reach students with a given j -rank under γ than under $\tilde{\gamma}$. A third interpretation is that fewer students are offered / trade away seats at school i at time t under γ than under $\tilde{\gamma}$.

Basic Lemmas

We will also make use of the following lemmas.

Lemma B.1. Let $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ be a continuum economy such that $\tilde{H}(x)$ is irreducible for all x and C . Then there exists a unique valid TTC path γ . Within each round $\gamma(\cdot)$ is given by

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

where $\mathbf{d}(x)$ is the unique valid direction from $x = \gamma(t)$ that satisfies $\mathbf{d}(x) = \mathbf{d}(x) \tilde{H}(x)$.

Moreover, if we let $A(x)$ be obtained from $\tilde{H}(x) - I$ by replacing the n th column with the all ones vector $\mathbf{1}$, then

$$d(x) = [0, 0, \dots, 0, -1] A(x)^{-1}.$$

Proof. It suffices to show that $d(\cdot)$ is unique. The existence and uniqueness of $\gamma(\cdot)$ satisfying $\frac{d\gamma(t)}{dt} = \mathbf{d}(\gamma(t))$ follows by invoking Picard-Lindelöf as in the proof of Theorem 4.2.

Consider the equations,

$$d(x) \tilde{H}(x) = d(x)$$

$$d(x) \cdot \mathbf{1} = -1.$$

When $\tilde{H}(x)$ is irreducible, every choice of $n - 1$ columns of $\tilde{H}(x) - I$ gives an independent set whose span does not contain $\mathbf{1}$. Therefore if we let $A(x)$ be given by replacing the n th column in $\tilde{H}(x) - I$ with $\mathbf{1}$, then $A(x)$ has full rank, and the above equations are equivalent to

$$d(x) A(x) = [0, 0, \dots, 0, -1],$$

$$\text{i.e. } d(x) = [0, 0, \dots, 0, -1] A(x)^{-1}.$$

Hence $d(x)$ is unique for each x , and hence $\gamma(\cdot)$ is uniquely determined. \square

We now show that any two non-increasing continuous paths $\gamma, \tilde{\gamma}$ starting and ending at the same point can be re-parametrized so that for all t there exists a school $i(\tau)$ such that γ is dominated by $\tilde{\gamma}$ via school $i(\tau)$ at time t . We first show that, if $\gamma(0) \leq \tilde{\gamma}(0)$, then there exists a re-parametrization of γ such that γ is dominated by $\tilde{\gamma}$ on some interval starting at 0.

Lemma B.2. *Suppose $\gamma, \tilde{\gamma}$ are a pair of non-increasing functions $[0, 1] \rightarrow [0, 1]^C$ such that $\gamma(0) \leq \tilde{\gamma}(0)$. Then there exist coordinates i, j , a time \bar{t} and an increasing function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $\gamma_j(g(\bar{t})) = \tilde{\gamma}_j(\bar{t})$, and for all $\tau \in [0, \bar{t}]$ it holds that*

$$\gamma_i(g(\tau)) = \tilde{\gamma}_i(\tau) \text{ and } \gamma(g(\tau)) \leq \tilde{\gamma}(\tau).$$

That is, if we renormalize the time parameter τ of $\gamma(\tau)$ so that γ and $\tilde{\gamma}$ agree along the i th

coordinate, then γ is dominated by $\tilde{\gamma}$ via school i at all times $\tau \in [0, \bar{t}]$, and is also dominated via school j at time \bar{t} .

Proof. The idea is that if we take the smallest function g such that there exists a coordinate i such that for all τ sufficiently small $\gamma_i(g(\tau)) = \tilde{\gamma}_i(\tau)$, then $\gamma(g(\tau)) \leq \tilde{\gamma}(\tau)$ for all τ sufficiently small. The lemma then follows from continuity. We make this precise.

Fix a coordinate i . Let $g^{(i)}$ be the renormalization of γ so that γ and $\tilde{\gamma}$ agree along the i th coordinate, i.e. $\gamma_i(g^{(i)}(\tau)) = \tilde{\gamma}_i(\tau)$ for all τ .

For all τ , we define the set $\kappa_{>}^{(i)}(\tau) = \{j \mid \gamma_j(g^{(i)}(\tau)) > \tilde{\gamma}_j(\tau)\}$ of schools j along which the γ curve renormalized along coordinate i has larger j -value at time τ than $\tilde{\gamma}_j$ has at time τ , and similarly define the set $\kappa_{=}^{(i)}(\tau) = \{j \mid \gamma_j(g^{(i)}(\tau)) = \tilde{\gamma}_j(\tau)\}$ where the renormalized γ curve is equal to $\tilde{\gamma}$. It suffices to show that there exists i, j and a time \bar{t} such that $\kappa_{>}^{(i)}(\tau) = \emptyset$ for all $\tau \in [0, \bar{t}]$ and $j \in \kappa_{=}^{(i)}(\bar{t})$.

Since γ and $\tilde{\gamma}$ are continuous, there exists some maximal $\bar{t}^{(i)} > 0$ such that the functions $\kappa_{>}^{(i)}(\cdot)$ and $\kappa_{=}^{(i)}(\cdot)$ are constant over the interval $(0, \bar{t}^{(i)})$. If there exists i such that $\kappa_{>}^{(i)}(\tau) = \emptyset$ for all $\tau \in (0, \bar{t}^{(i)})$ then by continuity there exists some time $\bar{t} \leq \bar{t}^{(i)}$ and school j such that $j \in \kappa_{=}^{(i)}(\bar{t})$ and we are done. Hence we may assume that for all i it holds that $\kappa_{>}^{(i)}(\tau) = C_{>}^{(i)}$ for all $\tau \in (0, \bar{t}^{(i)})$ for some fixed non-empty set $C_{>}^{(i)}$. We will show that this leads to a contradiction.

We first claim that if $j \in C_{>}^{(i)}$, then $g^{(j)}(\tau) > g^{(i)}(\tau)$ for all $\tau \in (0, \bar{t})$. This is because γ is non-increasing and $\gamma_j(g^{(j)}(\tau)) = \tilde{\gamma}_j(\tau) < \gamma_j(g^{(i)}(\tau))$ for all $\tau \in (0, \bar{t})$, where the equality follows from the definition of $g^{(j)}$ and the inequality since $j \in C_{>}^{(i)}$. But this completes the proof, since it implies that for all i there exists j such that $g^{(j)}(\tau) > g^{(i)}(\tau)$ for all $\tau \in (0, \bar{t})$, which is impossible since there are a finite number of schools $i \in \mathcal{C}$. \square

We are now ready to show that there exists a re-parametrization of γ such that γ always is dominated by $\tilde{\gamma}$ via some school.

Lemma B.3. *Suppose $\bar{t} \geq 0$ and $\gamma, \tilde{\gamma}$ are a pair of non-increasing functions $[0, \bar{t}] \rightarrow [0, 1]^{\mathcal{C}}$ such that $\gamma(0) \leq \tilde{\gamma}(0) = 1$ with equality on at least one coordinate, and $0 = \gamma(1) \leq \tilde{\gamma}(1)$ with equality*

on at least one coordinate. Then there exists an increasing function $g : [0, \bar{t}] \rightarrow \mathbb{R}$ such that for all $\tau \geq 0$, there exists a school i such that $\gamma(g(\tau))$ is dominated by $\tilde{\gamma}(\tau)$ via school i .

Proof. Without loss of generality let us assume that $\bar{t} = 1$. Fix a coordinate i . We define $g^{(i)}$ to be the renormalization of γ so that γ and $\tilde{\gamma}$ agree along the i th coordinate. Formally, let $\underline{t}^{(i)} = \min \{\tau \mid \gamma_i(0) \geq \tilde{\gamma}_i(\tau)\}$ and define $g^{(i)}$ so that $\gamma_i(g^{(i)}(\tau)) = \tilde{\gamma}_i(\tau)$ for all $\tau \in [\underline{t}^{(i)}, 1]$. Let $A^{(i)}$ be the set of times τ such that $\gamma(g^{(i)}(\tau))$ is dominated by $\tilde{\gamma}(\tau)$. The idea is to pick g to be equal to $g^{(i)}$ in $A^{(i)}$. In order to do this formally, we need to show that the sets $A^{(i)}$ cover $[0, 1]$, and then turn (a suitable subset of) $A^{(i)}$ into a union of disjoint closed intervals, on each of which we can define $g(\cdot) \equiv g^{(i)}(\cdot)$.

We first show that $\cup_i A^{(i)} = [0, 1]$. Suppose not, so there exists some time τ such that for all $i \in X \stackrel{\text{def}}{=} \{k : \tau \geq \underline{t}^{(k)}\}$ there exists j such that $\gamma_j(g^{(i)}(\tau)) > \tilde{\gamma}_j(\tau)$. Note that for such i, j , since γ_j is non-increasing this implies that $\gamma_j(0) \geq \tilde{\gamma}_j(\tau)$, and so the function $g^{(j)}(\cdot)$ is defined at τ , i.e. there exists $g^{(j)}(\tau)$ such that $\tilde{\gamma}_j(\tau) = \gamma_j(g^{(j)}(\tau))$. In other words, since γ is non-increasing, for all $i \in X$ there exists j such that $g^{(i)}(\tau) < g^{(j)}(\tau)$, and since $\gamma_j(0) \geq \tilde{\gamma}_j(\tau)$ it also holds that $j \in X$. This is a contradiction since X is finite but non-empty (since $\gamma(0) \leq \tilde{\gamma}(0) = 1$, with equality on at least one coordinate).

We now turn (a suitable subset of $A^{(i)}$) into a union of disjoint closed intervals. By continuity, $A^{(i)}$ is closed. Consider the closure of the interior of $A^{(i)}$, which we denote by $B^{(i)}$. Since the interior of $A^{(i)}$ is open, it is a countable union of open intervals, and hence $B^{(i)}$ is a countable union of disjoint closed intervals. To show that $\cup_{i \in \mathcal{C}} B^{(i)} = [0, 1]$, fix a time $\tau \in [0, 1]$. As $\cup_i A^{(i)} = [0, 1]$, there exists i such that $\gamma(g^{(i)}(\tau)) \leq \tilde{\gamma}(\tau)$. Hence we may invoke Lemma B.2 to show that there exists some school j , time $\bar{\tau} > \tau$ and an increasing function g such that $\gamma_j(g(g^{(i)}(\tau'))) = \tilde{\gamma}_j(\tau')$ and $\gamma(g(g^{(i)}(\tau'))) \leq \tilde{\gamma}(\tau')$ for all $\tau' \in [\tau, \bar{\tau}]$. But by the definition of $g^{(j)}(\cdot)$ this means that $\gamma_j(g(g^{(i)}(\tau'))) = \tilde{\gamma}_j(\tau') = \gamma_j(g^{(j)}(\tau'))$ for all $\tau' \in [\tau, \bar{\tau}]$, and so $g \circ g^{(i)} = g^{(j)}$ and we have shown that $[\tau, \bar{\tau}] \subseteq B^{(j)}$. Hence we may write $[0, 1] = \cup_n T_n$ as a countable union of closed intervals T_n such that any pair of intervals intersects at most at their endpoints, and each interval T_n is a subset of $B^{(i)}$ for some i . For each T_n fix some $i(n) = i$ so that $T_n \subseteq B^{(i)}$. Intuitively, this means that at any time $\tau \in T_n$ it holds that $\gamma(g^{(i(n))}(\tau))$ is dominated by $\tilde{\gamma}(\tau)$ via school i .

We now construct a function g that satisfies the required properties as follows. If $\tau \in T_n \subseteq B^{(i)}$, let $g(\tau) = g^{(i)}(\tau)$. Now g is well-defined despite the possibility that $T_n \cap T_m \neq \emptyset$. This is because if τ is in two different intervals T_n, T_m , then $\gamma_{i(n)}(g^{(i(n))}(\tau)) = \tilde{\gamma}_{i(n)}(\tau) \geq \gamma_{i(n)}(g^{(i(m))}(\tau))$ (by domination via $i(n)$ and $i(m)$ respectively), and $\gamma_{i(m)}(g^{(i(m))}(\tau)) = \tilde{\gamma}_{i(m)}(\tau) \geq \gamma_{i(m)}(g^{(i(n))}(\tau))$ (by domination via $i(m)$ and $i(n)$ respectively), and so $g^{(i(n))}(\tau) \leq g^{(i(m))}(\tau) \leq g^{(i(n))}(\tau)$ and we can pick one value for g that satisfies all required properties. Now by definition $\gamma(g(\tau))$ is dominated by $\tilde{\gamma}(\tau)$ via school i , and moreover g is defined on all of $[0, 1]$ since $\cup_{i \in C} B^{(i)} = [0, 1]$. This completes the proof. \square

Lemma B.4. *Let $C \subseteq \mathcal{C}$ be a set of schools, and let D be a region on which $\tilde{H}^C(x)$ is irreducible for all $x \in D$. For each x let $A(x)$ be given by replacing the n th column of $\tilde{H}^C(x) - I_C$ with the all ones vector $\mathbf{1}$.⁴ Then the function $f(x) = [0, 0, \dots, 0, -1] A(x)^{-1}$ is piecewise Lipschitz continuous in x .*

Proof. It suffices to show that the function which, for each x , outputs the matrix $A(x)^{-1}$ is piecewise Lipschitz continuous in x .

Now

$$H_i^{j|C}(x) = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_{\theta: r^\theta \geq x, r^\theta \not\geq x_i + \varepsilon \cdot e^i, j \succ^\theta C} \nu(\theta) d\theta,$$

where $\nu(\cdot)$ is bounded below on its support and piecewise Lipschitz continuous, and the points of discontinuity lie on the grid. Hence $H_i^{j|C}(x)$ is Lipschitz continuous in x for all i, j , and $\sum_k H_j^{k|C}(x)$ nonzero and hence bounded below, and so $\tilde{H}^C(x)_{i,j}$ is bounded above and piecewise Lipschitz continuous in x , and therefore so is $A(x)$. Finally, since $\tilde{H}^C(x)$ is an irreducible row stochastic matrix for each $x \in D$, it follows that $A(x)$ is full rank and continuous. This is because when $\tilde{H}^C(x)$ is irreducible every choice of $n - 1$ columns of $\tilde{H}^C(x) - I_C$ gives an independent set whose span does not contain the all ones vector $\mathbf{1}_C$. Therefore if we let $A(x)$ be given by replacing the n th column in $\tilde{H}^C(x) - I_C$ with $\mathbf{1}_C$, then $A(x)$ has full rank.

Since $A(x)$ is full rank and continuous, in each piece $\det(A(x))$ is bounded away from 0, and so $A(x)^{-1}$ is piecewise Lipschitz continuous, as required. \square

⁴ I_C is the identity matrix with rows and columns indexed by the elements in C .

Connection to Continuous Time Markov Chains

In this section, we formalize the intuition from Appendix 4.4.3. In Appendix 4.4.3, we appealed to a connection with Markov chain theory to provide a method for solving for all the possible values of $\mathbf{d}(x)$. Specifically, we constructed a continuous time Markov chain with state space \mathcal{C} and transition rates from state i to j equal to $H_i^j(x)$. We argued that if $\mathcal{K}(x)$ is the set of recurrent communication classes of this Markov chain, then the set of valid directions $\mathbf{d}(x)$ is identical to the set of convex combinations of $\{\mathbf{d}^K\}_{K \in \mathcal{K}(x)}$, where \mathbf{d}^K is the unique solution to the trade balance equations (4.2) restricted to K . We present the relevant definitions, results and proofs here in full.

Let us first present some definitions from Markov chain theory.⁵ A square matrix P is a *right-stochastic matrix* if all the entries are non-negative and each row sums to 1. A *probability vector* is a vector with non-negative entries that add up to 1. Given a right-stochastic matrix P , the *Markov chain with transition matrix P* is the Markov chain with state space equal to the column/row indices of P , and a probability P_{ij} of moving to state j in one time step, given that we start in state i . Given two states i, j of a Markov chain with transition matrix P , we say that states i and j *communicate* if there is a positive probability of moving to state i to state j in finite time, and vice versa.

For each Markov chain, there exists a unique decomposition of the state space into a sequence of disjoint subsets C_1, C_2, \dots such that for all i, j , states i and j communicate if and only if they are in the same subset C_k for some k . Each subset C_k is called a *communication class* of the Markov chain. A Markov chain is *irreducible* if it only has one communication class. A state i is *recurrent* if, starting at i and following the transition matrix P , the probability of returning to state i is 1. A communication class is recurrent if it contains a recurrent state.

The following proposition gives a characterization of the stationary distributions of a Markov chain. We refer the reader to any standard stochastic processes textbook (e.g. Karlin and Taylor (1975)) for a proof of this result.

Proposition B.1. *Suppose that P is the transition matrix of a Markov chain. Let \mathcal{K} be the set*

⁵See standard texts such as Karlin and Taylor (1975) for a more complete treatment.

of recurrent communication classes of the Markov chain with transition matrix P . Then for each recurrent communication class $K \in \mathcal{K}$, the equation $\pi = \pi P$ has a unique solution π^K such that $\|\pi^K\| = 1$ and $\text{supp}(\pi^K) \subseteq K$. Moreover, the support of π^K is equal to K . In addition, if $\|\pi\| = 1$ and π is a solution to the equation $\pi = \pi P$, then π is a convex combination of the vectors in $\{\pi^K\}_{K \in \mathcal{K}}$.

To make use of this proposition, define at each point x and for each set of schools C a Markov chain $M^C(x)$ with transition matrix $\tilde{H}^C(x)$. Note that this is equivalent to taking the embedded discrete-time Markov chain of a continuous-time Markov chain with transition rates $H_i^{j|C}(x)$ for $i \neq j$, and transition rates $H_j^{j|C}(x) = \bar{v}$ (where $\bar{v} \geq \max_{j \in C} \left(\sum_{k \in C} H_j^{k|C}(x) \right)$) is the normalization term used to construct $\tilde{H}^C(x)$). We will relate the valid directions $\mathbf{d}(x)$ to the recurrent communication classes of $M^C(x)$, where C is the set of available schools. We will need the following notation and definitions. Given a vector v indexed by \mathcal{C} , a matrix Q with rows and columns indexed by C and subsets $K, K' \subseteq C$ of the indices, we let v_K denote the restriction of v to the coordinates in K , and we let $Q_{K,K'}$ denote the restriction of Q to rows indexed by K and columns indexed by K' .

The following lemma characterizes the recurrent communication classes of the Markov chain $M^C(x)$ using the properties of the matrix $\tilde{H}^C(x)$, and can be found in any standard stochastic processes text.

Lemma B.5. *Let C be the set of available school at a point x . Then a set $K \subseteq C$ is a recurrent communication class of the Markov chain $M^C(x)$ if and only if $\tilde{H}^C(x)_{K,K}$ is irreducible and $\tilde{H}^C(x)_{K,C \setminus K}$ is the zero matrix.*

It is easy to see that the same result holds when we replace \tilde{H}^C by H^C .

The following lemma allows us to characterize the valid directions \mathbf{d} in terms of the matrix $\tilde{H}^C(x)$.

Lemma B.6. *The vector \mathbf{d} is a valid direction at x with available schools C if and only if*

$$\mathbf{d} \cdot \mathbf{1} = -1 \text{ and } \mathbf{d} = \mathbf{d} \cdot \tilde{H}^C(x).$$

Proof. It suffices to show that $\mathbf{d} = \mathbf{d} \cdot \tilde{H}^C(x)$ if and only if

$$\sum_{k \in C} d_k \cdot H_k^{i|C}(x) = \sum_{k \in C} d_i \cdot H_i^{k|C}(x) \quad \forall i \in C.$$

Now

$$\begin{aligned} \mathbf{d} &= \mathbf{d} \cdot \tilde{H}^C(x) \\ \Leftrightarrow d_i &= \sum_{k \in C} d_k \cdot \tilde{H}_k^{i|C}(x) \quad \forall i \in C \\ \Leftrightarrow d_i &= \sum_{k \in C} d_k \cdot \left(\frac{1}{\bar{v}} H_k^{i|C}(x) + \mathbf{1}_{i=k} \left(1 - \frac{v_i}{\bar{v}} \right) \right) \quad \forall i \in C \\ \Leftrightarrow d_i \cdot \frac{v_i}{\bar{v}} &= \sum_{k \in C} d_k \cdot \left(\frac{1}{\bar{v}} H_k^{i|C}(x) \right) \quad \forall i \in C \\ \Leftrightarrow d_i \cdot \sum_{k \in C} H_i^{k|C}(x) &= \sum_{k \in C} d_k \cdot H_k^{i|C}(x) \quad \forall i \in C \end{aligned}$$

which concludes the proof. □

Proposition B.1 and Lemmas B.6 and B.5 allow us to characterize the valid directions $d(x)$.

Theorem B.1. *Let C be the set of available schools, and let $\mathcal{K}(x)$ be the set of subsets $K \subseteq C$ for which $\tilde{H}^C(x)_{K,K}$ is irreducible and $\tilde{H}^C(x)_{K,C \setminus K}$ is the zero matrix. Then for each $K \in \mathcal{K}(x)$ the equation $\mathbf{d} = \mathbf{d} \cdot \tilde{H}^C(x)$ has a unique solution \mathbf{d}^K that satisfies $\mathbf{d}^K \cdot \mathbf{1} = -1$ and $\text{supp}(\mathbf{d}^K) \subseteq K$, and its projection onto its support K has the form*

$$\left(\mathbf{d}^K \right)_K = [0, 0, \dots, 0, -1] A_K^C(x)^{-1},$$

where $A_K^C(x)$ is the matrix obtained by replacing the $(|K| - 1)$ th column of $\tilde{H}^C(x)_{K,K} - I_K$ with the all ones vector $\mathbf{1}_K$.

Moreover, if $\mathbf{d} \cdot \mathbf{1} = -1$ and d is a solution to the equation $\mathbf{d} = \mathbf{d} \cdot \tilde{H}^C(x)$, then \mathbf{d} is a convex combination of the vectors in $\{\mathbf{d}^K\}_{K \in \mathcal{K}(x)}$.

Proof. Proposition B.5 shows that the sets K are precisely the recurrent sets of the Markov chain with transition matrix $\tilde{H}(x)$. Hence uniqueness of the \mathbf{d}^K and the fact that d is a convex combina-

tion of d^K follow directly from Proposition B.1. The form of the solution d^K follows from Lemma B.1. □

This has the following interpretation. Suppose that there is a unique recurrent communication class K , such as when η has full support. Then there is a unique infinitesimal continuum trading cycle of students, specified by the unique valid direction \mathbf{d} satisfying $\mathbf{d} = \mathbf{d} \cdot \tilde{H}(x)$. Moreover, students in the cycle trade seats from every school in K . Any school not in K is blocked from participating, since there is not enough demand to fill the seats they are offering. When there are multiple recurrent communication classes, each of the d^K gives a unique infinitesimal trading cycle of students, corresponding to those who trade seats in K . Moreover, these trading cycles are disjoint. Hence the only multiplicity that remains is to decide the order, or the relative rate, at which to clear these cycles. We will show in Appendix B.1 that, as in the discrete setting, the order in which cycles are cleared does not affect the final allocation.

Proof of Theorem 4.2

We first show that there exist solutions p, γ, t to the marginal trade balance equations and capacity equations. The proof relies on selecting appropriate valid directions $\mathbf{d}(x)$ and then invoking the Picard-Lindelöf theorem to show existence.

Specifically, let C be the set of available schools, fix a point x , and consider the set of vectors \mathbf{d} such that $\mathbf{d} \cdot \tilde{H}^C(x) = \mathbf{d}$. Then it follows from Theorem B.1 that if $\mathbf{d}(x)$ is the valid direction from x with minimal support under the shortlex order, then $\mathbf{d}(x) = \mathbf{d}^{K(x)}$ for the element $K(x) \in \mathcal{K}(x)$ that is the smallest under the shortlex ordering.⁶ As the density $\nu(\cdot)$ defining $\eta(\cdot)$ is Lipschitz continuous, it follows that $\mathcal{K}(\cdot)$ and $K(\cdot)$ are piecewise constant. Hence we may invoke Lemma B.4 and the form of $\mathbf{d}(\cdot)$ as given in Lemma B.1 to conclude that $\mathbf{d}(\cdot)$ is piecewise Lipschitz within each piece, and hence piecewise Lipschitz in $[0, 1]^C$. Since $\mathbf{d}(\cdot)$ is piecewise Lipschitz, it follows from the Picard-Lindelöf theorem that there exists a unique function $\gamma(\cdot)$ satisfying $\frac{d\gamma(t)}{dt} = \mathbf{d}(\gamma(t))$. It follows trivially that γ satisfies the marginal trade balance equations, and since we have assumed

⁶We choose the shortlex ordering to ensure that we choose valid directions corresponding to a single recurrent communication class, rather than unions of recurrent communication classes.

that all students find all schools acceptable and there are more students than seats it follows that there exist stopping times $t^{(i)}$ and cutoffs p_b^i .

Proof of the Uniqueness of the TTC Allocation In this section, we prove the uniqueness claim in Theorem 4.2, that any two valid TTC paths give equivalent allocations. The intuition for the result is the following. The connection to Markov chains shows that having multiple possible valid directions in the continuum corresponds to having multiple possible trade cycles in the discrete model. Hence the only multiplicity in choosing valid TTC directions is whether to implement one set of trades before the others, or to implement them in parallel at various relative rates. We can show that the set of cycles is independent of the order in which cycles are selected, or equivalently that the sets of students who trade with each other is independent of the order in which possible trades are executed. It follows that any pair of valid TTC paths give the same final allocation.

We remark that the crux of the argument is similar to what shows that discrete TTC gives a unique allocation. However, the lack of discrete cycles and the ability to implement sets of trades in parallel both complicate the argument and lead to a rather technical proof.

We first formally define *cycles* in the continuum setting, and a partial order over the cycles corresponding to the order in which cycles can be cleared under TTC. We then define the set of cycles $\Sigma(\gamma)$ associated with a valid TTC path γ . Finally, we show that the sets of cycles associated with two valid TTC paths γ and γ' are the same, $\Sigma(\gamma) = \Sigma(\gamma')$.

Definition B.2. A (*continuum*) *cycle* $\sigma = (K, \underline{x}, \bar{x})$ is a set $K \subseteq \mathcal{C}$ and a pair of vectors $\underline{x} \leq \bar{x}$ in $[0, 1]^{\mathcal{C}}$. The cycle σ is *valid* for available schools $\{C(x)\}_{x \in [0, 1]^{\mathcal{C}}}$ if $K \in \mathcal{K}^{C(x)}(x) \forall x \in (\underline{x}, \bar{x}]$.

Intuitively, a cycle is defined by two time points in a run of TTC, which gives a set of students,⁷ and the set of schools they most desire. A cycle is valid if the set of schools involved is a recurrent communication class of the associated Markov chains.⁸ We say that a cycle $\sigma = (K, \underline{x}, \bar{x})$ *appears*

⁷The set of students is given by taking the difference between two nested hyperrectangles, one with upper coordinate \bar{x} and the other with upper coordinate \underline{x} .

⁸Note that we consider validity only in terms of whether the schools are the appropriate schools for a trading cycle, and not in terms of the feasibility of trade balance for the students in the cycle.

at time t in $TTC(\gamma)$ if $K \in \mathcal{K}^{C(\gamma(t))}(\gamma(t))$ and $\gamma_i(t) = \bar{x}_i$ for all $i \in K$. We say that a *student* θ is in cycle σ if $r^\theta \in (\underline{x}, \bar{x}]^9$, and a *school* i is in cycle σ if $i \in K$.

Definition B.3 (Partial order over cycles). The cycle $\sigma = (K, \underline{x}, \bar{x})$ *blocks* the cycle $\sigma' = (K', \underline{x}', \bar{x}')$, denoted by $\sigma \triangleright \sigma'$, if at least one of the following hold:

(Blocking student) There exists a student θ in σ' who prefers a school in K to all those in K' , i.e. there exist θ and $i \in K \setminus K'$ such that $i \succ^\theta i'$ for all $i' \in K'$.

(Blocking school) There exists a school in σ' that prefers a positive measure of students in σ to all those in σ' , i.e. there exists $i \in K'$ such that $\eta(\theta | \theta \text{ in } \sigma, r_i^\theta > \bar{x}'_i) > 0$.¹⁰

Let us now define the set of cycles associated with a run of TTC. We begin with some observations about $H_i^{b|C}(\cdot)$ and $\tilde{H}^C(\cdot)_{bi}$. For all $b, i \in C$ the function $H_i^{b|C}(\cdot)$ is right-continuous on $[0, 1]^C$, Lipschitz continuous on R for all $R \in \mathcal{R}$ and uniformly bounded away from zero on its support. Hence $\mathbf{1}(H_i^{b|C}(\cdot))$ is constant on R for all $R \in \mathcal{R}$. It follows that $\tilde{H}^C(\cdot)_{bi}$ is also right-continuous, and Lipschitz continuous on R for all $R \in \mathcal{R}$. Moreover, there exists some finite rectangular subpartition \mathcal{R}' of \mathcal{R} such that for all $C \subseteq \mathcal{C}$ the function $\mathbf{1}(\tilde{H}^C(\cdot))$ is constant on R for all $R \in \mathcal{R}'$.

Definition B.4. Let \mathcal{R}' denote the minimal rectangular subpartition of \mathcal{R} such that for all $C \subseteq \mathcal{C}$ the function $\mathbf{1}(\tilde{H}^C(\cdot))$ is constant on R for all $R \in \mathcal{R}'$.

For $x \in [0, 1]^C$ and $C \subseteq \mathcal{C}$, let $\mathcal{K}^C(x)$ be the recurrent communication classes of the Markov chain $M^C(x)$. The following lemma follows immediately from Proposition B.5, since $\mathbf{1}(\tilde{H}^C(\cdot))$ is constant on $R \forall R \in \mathcal{R}'$, and recurrent communication classes depend only on $\mathbf{1}(\tilde{H}^C)$.

Lemma B.7. $\mathcal{K}^C(\cdot)$ is constant on R for every $R \in \mathcal{R}'$.

For each $K \in \mathcal{K}^C(x)$, let $d^K(x)$ be the unique vector satisfying $d = d\tilde{H}^C(x)$, which exists by Theorem B.1.

⁹Recall that since r^θ, \underline{x} and \bar{x} are vectors, this is equivalent to saying that $r^\theta \not\leq \underline{x}$ and $r^\theta \leq \bar{x}$.

¹⁰For i to block the cycle σ it is necessary but not sufficient that $\bar{x}_i > \bar{x}'_i$, since there also need to be students in σ with the intermediate ranks at school i .

Let γ be a TTC path, and assume that the schools are labeled in order. It follows that for all x there exists ℓ such that $C(x) = \mathcal{C}^{(\ell)} \stackrel{def}{=} \{\ell, \ell + 1, \dots, |\mathcal{C}|\}$. For each set of schools $K \subseteq \mathcal{C}$, let $T^{(\ell)}(K, \gamma)$ be the set of times τ such that $C(\gamma(\tau)) = \mathcal{C}^{(\ell)}$ and K is a recurrent communication class for $\tilde{H}^{\mathcal{C}^{(\ell)}}(\gamma(\tau))$. Since γ is continuous and weakly decreasing, it follows from Lemma B.7 that $T^{(\ell)}(K, \gamma)$ is the finite disjoint union of intervals of the form $[\underline{t}, \bar{t}]$. Let $\mathcal{I}(T^{(\ell)}(K, \gamma))$ denote the set of intervals in this disjoint union. We may assume that for each interval T , $\gamma(T)$ is contained in some hyperrectangle $R \in \mathcal{R}'$.¹¹

For a time interval $T = [\underline{t}, \bar{t}] \in \mathcal{I}(T^{(\ell)}(K, \gamma))$, we define the cycle $\sigma(T) = (K, \underline{x}(T), \bar{x}(T))$ as follows. Intuitively, we want to define it simply as $\sigma(T) = (K, \gamma(\underline{t}), \gamma(\bar{t}))$, but in order to minimize the dependence on γ , we define the endpoints $\underline{x}(T)$ and $\bar{x}(T)$ of the interval of ranks to be as close together as possible, while still describing the same set of students (up to a set of η -measure 0). Define

$$\begin{aligned}\underline{x}(T) &= \max \left\{ x : \gamma(\underline{t}) \leq x \leq \gamma(\underline{t}), \eta \left(\theta : Ch_{\theta}(\mathcal{C}^{(\ell)}) \in K, r^{\theta} \in (x, \gamma(\bar{t})) \right) = 0 \right\}, \\ \bar{x}(T) &= \min \left\{ x : \gamma(\underline{t}) \leq x \leq \gamma(\underline{t}) : \eta \left(\theta : Ch_{\theta}(\mathcal{C}^{(\ell)}) \in K, r^{\theta} \in (\gamma(\underline{t}), x] \right) = 0 \right\},\end{aligned}$$

to be the points chosen to be maximal and minimal respectively such that the set of students allocated by γ during the time interval T has the same η -measure as if $\gamma(\underline{t}) = \underline{x}(T)$ and $\gamma(\bar{t}) = \bar{x}(T)$.¹² In other words, $\underline{x}(T)$ and $\bar{x}(T)$ are chosen to be respectively maximal and minimal under the lexicographical order such that

$$\eta \left(\left(\cup_{i \in K} \mathcal{T}^i(\gamma; \bar{t}) \setminus \mathcal{T}^i(\gamma; \underline{t}) \right) \setminus \left\{ \theta : Ch_{\theta}(\mathcal{C}^{(\ell)}) \in K, r^{\theta} \in (\underline{x}(T), \bar{x}(T)] \right\} \right) = 0.$$

In a slight abuse of notation, if $\sigma = \sigma(T)$ we will let $\underline{x}(\sigma)$ denote $\underline{x}(T)$ and $\bar{x}(\sigma)$ denote $\bar{x}(T)$.

Definition B.5. The set of cycles cleared by TTC (γ) in round ℓ , denoted by $\Sigma^{(\ell)}(\gamma)$, is given by

$$\Sigma^{(\ell)}(\gamma) := \bigcup_{K \subseteq \mathcal{C}^{(\ell)}} \bigcup_{T \in \mathcal{I}(T^{(\ell)}(K, \gamma))} \sigma(T).$$

¹¹This is without loss of generality, since if $\gamma(T)$ is not contained we can simply partition T into a finite number of intervals $\cup_{R \in \mathcal{R}'} \gamma^{-1}(\gamma(T) \cap R)$, each contained in a hyperrectangle in \mathcal{R}' .

¹²In order to take the maximum and minimum of the set of possible values for \underline{x} and \bar{x} respectively we order the elements of $[0, 1]^{\mathcal{C}}$ lexicographically.

The set of cycles cleared by $TTC(\gamma)$, denoted by $\Sigma(\gamma)$, is the set of cycles cleared by $TTC(\gamma)$ in some round ℓ ,

$$\Sigma(\gamma) := \bigcup_{\ell} \Sigma^{(\ell)}(\gamma).$$

For any cycle $\sigma \in \Sigma(\gamma)$ and time τ we say that the cycle σ is *clearing at time* τ if $\gamma(\tau) \not\leq \underline{x}(\sigma)$ and $\gamma(\tau) \not\geq \bar{x}(\sigma)$. We say that the cycle σ is *cleared at time* τ or *finishes clearing at time* τ if $\gamma^{(l)}(\tau) \leq \underline{x}(\sigma)$ with at least one equality. We remark that for any TTC path γ there may be multiple cycles clearing at a time τ , each corresponding to a different recurrent set. For any TTC path γ the set $\Sigma(\gamma)$ is finite.

Fix two TTC paths γ and γ' . Our goal is to show that they clear the same sets of cycles, $\Sigma(\gamma) = \Sigma(\gamma')$, or equivalently that $\Sigma(\gamma) \cup \Sigma(\gamma') = \Sigma(\gamma) \cap \Sigma(\gamma')$. We will do this by showing that for every cycle $\sigma \in \Sigma(\gamma) \cup \Sigma(\gamma')$, if all cycles in $\Sigma(\gamma) \cup \Sigma(\gamma')$ that block σ are in $\Sigma(\gamma) \cap \Sigma(\gamma')$, then $\sigma \in \Sigma(\gamma) \cap \Sigma(\gamma')$. We first show that this is true in a special case, which can be understood intuitively as the case when the cycle σ appears during the run of $TTC(\gamma)$ and also appears during the run of $TTC(\gamma')$.

Lemma B.8. *Let $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ be a continuum economy, and let γ and γ' be two TTC paths for this economy. Let $K \subseteq \mathcal{C}$ and \underline{t} be such that at time \underline{t} , γ (γ') has available schools C (C'), the paths γ, γ' are at the same point when projected onto the coordinates K , i.e. $\gamma(\underline{t})_K = \gamma'(\underline{t})_K$, and K is a recurrent communication class of $M^C(\gamma(\underline{t}))$ and of $M^{C'}(\gamma'(\underline{t}))$. Suppose that for all schools $i \in K$ and cycles $\sigma' \triangleright \sigma$ involving school i , if $\sigma' \in \Sigma(\gamma)$, then σ' is cleared in $TTC(\gamma')$, and vice versa. Suppose also that cycle $\sigma = (K, \underline{x}, \bar{x})$ is cleared in $TTC(\gamma)$, $\gamma(\underline{t}) = \underline{x}$, and measure 0 of σ has been cleared by time \underline{t} in $TTC(\gamma')$. Then σ is also cleared in $TTC(\gamma')$.*

Proof. We define the ‘interior’ of the cycle σ by $X = \{x : \underline{x}_i \leq x_i \leq \bar{x}_i \forall i \in K, x_{i'} \geq \underline{x}_{i'} \forall i' \notin K\}$. Fix a time u such that $\gamma'(u) \in X$ and let D' denote the set of available schools at time u in $TTC(\gamma')$. Then we claim that K is a recurrent communication class of $M^{D'}(\gamma'(u))$, and that a similar result is true for γ and a similarly defined D . The claim for γ, D follows from the fact that σ is cleared in $TTC(\gamma)$, $\sigma \in \Sigma(\gamma)$. It remains to show that the claim for γ', D' is true. Formally, by Lemma B.5 it suffices to show that $\tilde{H}^{D'}(x)_{K,K}$ is irreducible and $\tilde{H}^{D'}(x)_{K, D' \setminus K}$ is the zero matrix.

We first examine the differences between the matrices $\tilde{H}^{C'}(\gamma'(t))$ and $\tilde{H}^{D'}(\gamma'(u))$. Since K is a recurrent communication class of $M^{C'}(\gamma'(u))$, it holds that there are no transitions from K to states outside of K , i.e. $\mathbf{1}(\tilde{H}^{C'}(\gamma'(u))_{K,C'\setminus K}) = 0$. Since $K \subseteq D' \subseteq C'$ it follows that $\mathbf{1}(\tilde{H}^{D'}(\gamma'(u))_{K,D'\setminus K}) = 0$. Moreover, since $\mathbf{1}(\tilde{H}^{C'}(\gamma'(u))_{K,C\setminus K}) = 0$, all students' top choice schools out of C' or D' are the same (in K), and so $\tilde{H}^{C'}(\gamma'(u))_{K,K} = \tilde{H}^{D'}(\gamma'(u))_{K,K}$ and both matrices are irreducible. Hence K is a recurrent communication class of $M^{D'}(\gamma'(u))$.

We now invoke Theorem B.1 to show that in each of the two paths, all the students in the cycle σ clear with each other. Specifically, while the path γ is in the 'interior' of the cycle, that is $\gamma(\tau) \in X$, it follows from Theorem B.1 that the projection of the gradient of γ to K is a rescaling of some vector $d^K(\gamma(\tau))$, where $d^K(\cdot)$ depends on $\tilde{H}(\cdot)$ but not on γ . Similarly, while $\gamma'(\tau') \in X$, it holds that the projection of the gradient of γ' to K is a rescaling of the vector $d^K(\gamma'(\tau'))$, for the same function $d^K(\cdot)$. Hence if we let $\pi_K(x)$ denote the projection of a vector x to the coordinates indexed by schools in K , then $\pi_K(\gamma(\gamma^{-1}((\underline{x}, \bar{x})))) = \pi_K(\gamma'(\gamma'^{-1}((\underline{x}, \bar{x}))))$.

Recall that we have assumed that for all schools $i \in K$ and cycles $\sigma' \triangleright \sigma$ involving school i , if $\sigma' \in \Sigma(\gamma)$, then σ' is cleared in $TTC(\gamma')$, and vice versa. This implies that for all $i \in K$, the measure of students assigned to i in time $[0, \underline{t}]$ under $TTC(\gamma)$ is the same as the measure of students assigned to i in time $[0, \underline{t}]$ under $TTC(\gamma')$. Moreover, we have just shown that for any $x \in \gamma(\gamma^{-1}((\underline{x}, \bar{x})))$, $x' \in \gamma'(\gamma'^{-1}((\underline{x}, \bar{x})))$ such that $x_K = x'_K$, if we let $\tau = \gamma^{-1}(x)$ and $\tau' = (\gamma')^{-1}(x')$ then the same measure of students are assigned to i in time $[\underline{t}, \tau]$ under $TTC(\gamma)$ as in time $[\underline{t}, \tau']$ under $TTC(\gamma')$. Since $TTC(\gamma)$ clears σ the moment it exits the interior of σ , this implies that $TTC(\gamma')$ also clears σ the moment it exits the interior. \square

We are now ready to prove that the TTC allocation is unique. As the proof takes several steps, we separate it into several smaller claims for readability.

Proof of uniqueness. Let γ and γ' be two TTC paths, and let the sets of cycles associated with $TTC(\gamma)$ and $TTC(\gamma')$ be $\Sigma = \Sigma(\gamma)$ and $\Sigma' = \Sigma(\gamma')$ respectively. We will show that $\Sigma = \Sigma'$.

Let $\sigma = (K, \underline{x}, \bar{x})$ be a cycle in $\Sigma \cup \Sigma'$ such that the following assumption holds:

Assumption B.1. For all $\tilde{\sigma} \triangleright \sigma$ it holds that either $\tilde{\sigma}$ is in both Σ and Σ' or $\tilde{\sigma}$ is in neither.

We show that if σ is in $\Sigma \cup \Sigma'$ then it is in $\Sigma \cap \Sigma'$. Since Σ and Σ' are finite sets, this will be sufficient to show that $\Sigma = \Sigma'$. Without loss of generality we may assume that $\sigma \in \Sigma$.

We give here an overview of the proof. Let $\Sigma_{\triangleright\sigma} = \{\tilde{\sigma} \in \Sigma : \tilde{\sigma} \triangleright \sigma\}$ denote the set of cycles that are comparable to σ and cleared before σ in $TTC(\gamma)$. Assumption B.1 about σ implies that $\Sigma_{\triangleright\sigma} \subseteq \Sigma'$. We will show that this implies that no students in σ start clearing under $TTC(\gamma')$ until all the students in σ have the same top available school in $TTC(\gamma')$ as when they clear in $TTC(\gamma)$, or in other words, that if some students in σ start clearing under $TTC(\gamma')$ at time t , then the cycle σ appears at time t . We will then show that once some of the students in σ start clearing under $TTC(\gamma')$ then all of them start clearing. It then follows from Lemma B.8 that σ clears under both $TTC(\gamma)$ and $TTC(\gamma')$.

Let ℓ denote the round of $TTC(\gamma)$ in which σ is cleared, $C(x) = \mathcal{C}^{(\ell)} \forall x \in \sigma$. We define the times in $TTC(\gamma)$ and $TTC(\gamma')$ when all the cycles in $\Sigma_{\triangleright\sigma}$ are cleared, by

$$\begin{aligned}\bar{t}_{\triangleright\sigma} &= \min \left\{ t : \gamma(t) \leq (\tilde{x}) \text{ for all } \tilde{\sigma} = (\tilde{K}, \tilde{x}, (\tilde{x})) \in \Sigma_{\triangleright\sigma} \text{ and } \tilde{H}(\gamma(t)) \neq \mathbf{0} \right\}, \\ \bar{t}'_{\triangleright\sigma} &= \min \left\{ t : \gamma'(t) \leq (\tilde{x}) \text{ for all } \tilde{\sigma} = (\tilde{K}, \tilde{x}, (\tilde{x})) \in \Sigma_{\triangleright\sigma} \text{ and } \tilde{H}(\gamma'(t)) \neq \mathbf{0} \right\}.\end{aligned}$$

We define also the times in $TTC(\gamma)$ when σ starts to be cleared and finishes clearing,

$$\underline{t}_{\sigma} = \max \{ t : \gamma(t) \geq \bar{x} \}, \quad \bar{t}_{\sigma} = \min \{ t : \gamma(t) \leq \underline{x} \}$$

and similarly define the times $\underline{t}'_{\sigma} = \max \{ t : \gamma'(t) \geq \bar{x} \}, \bar{t}'_{\sigma} = \min \{ t : \gamma'(t) \leq \underline{x} \}$ for $TTC(\gamma')$.

We remark that part of the issue, carried over from the discrete setting, is that these times $\bar{t}_{\triangleright\sigma}$ and \underline{t}_{σ} might not match up, and similarly for $\bar{t}'_{\triangleright\sigma}$ and \underline{t}'_{σ} . In particular, other incomparable cycles could clear at interwoven times. In the continuum model, there may also be sections on the TTC curve at which no school is pointing to a positive density of students. However, all the issues in the continuum case can be addressed using the intuition from the discrete case.

We first show in Claims (B.1), (B.2) and (B.3) that in both $TTC(\gamma)$ and $TTC(\gamma')$, after all the cycles in $\Sigma_{\triangleright\sigma}$ are cleared and before σ starts to be cleared, the schools pointed to by students in σ and the students pointed to by schools in K remain constant (up to a set of η -measure 0).

Claim B.1. Let $\sigma = (K, \underline{x}, \bar{x}) \in \Sigma$ satisfy Assumption B.1. Suppose there is a school i that some student in σ prefers to all the schools in K . Then school i is unavailable in $TTC(\gamma)$ at any time $t \geq \bar{t}_{\triangleright\sigma}$, and unavailable in $TTC(\gamma')$ at any time $t \geq \bar{t}'_{\triangleright\sigma}$.

Proof. Suppose that school i is available in $TTC(\gamma)$ after all the cycles in $\Sigma_{\triangleright\sigma}$ are cleared. Then there exists a cycle $\tilde{\sigma}$ clearing at time $\tilde{t} \in (\bar{t}_{\triangleright\sigma}, t_\sigma)$ in $TTC(\gamma)$ involving school i . But this means that $\tilde{\sigma} \triangleright \sigma$ so $\tilde{\sigma} \in \Sigma_{\triangleright\sigma}$, which is a contradiction. Hence the measure of students in $\Sigma_{\triangleright\sigma}$ who are assigned to school i is q_i , and the claim follows. \square

Claim B.2. In $TTC(\gamma)$, let $\tilde{\Theta}$ denote the set of students cleared in time $[\bar{t}_{\triangleright\sigma}, t_\sigma)$ who are preferred by some school in $i \in K$ to the students in σ , that is, θ satisfying $r_i^\theta > \bar{x}_i$. Then $\eta(\tilde{\Theta}) = 0$.

Proof. Suppose $\eta(\tilde{\Theta}) > 0$. Then, since there are a finite number of cycles in $\Sigma(\gamma)$, there exists some cycle $\tilde{\sigma} = (\tilde{K}, \tilde{\underline{x}}, (\tilde{\bar{x}})) \in \Sigma(\gamma)$ containing a positive η -measure of students in $\tilde{\Theta}$. We show that $\tilde{\sigma}$ is cleared before σ . Since $\tilde{\sigma}$ contains a positive η -measure of students in $\tilde{\Theta}$, it holds that there exist $t_1, t_2 \in [\bar{t}_{\triangleright\sigma}, t_\sigma)$ and a school $i \in K$ for which $\tilde{x}_i \leq \gamma(t_1)_i < \gamma(t_2)_i \leq (\tilde{\bar{x}})_i$. Hence $\bar{x}_i \leq \gamma(t_\sigma)_i \leq \gamma(t_1)_i < \gamma(t_2)_i \leq \tilde{x}_i$, so $\tilde{\sigma} \triangleright \sigma$ as claimed. But by the definition of t_1, t_2 it holds that $(\tilde{\bar{x}})_i \leq \gamma(t_1)_i < \gamma(t_2)_i \leq \gamma(\bar{t}_{\triangleright\sigma})_i$, so $\tilde{\sigma}$ is not cleared before $\bar{t}_{\triangleright\sigma}$, contradicting the definition of $\bar{t}_{\triangleright\sigma}$. \square

Claim B.3. In $TTC(\gamma')$, let $\tilde{\Theta}$ denote the set of students cleared in time $[\bar{t}'_{\triangleright\sigma}, t'_\sigma)$ who are preferred by some school in $i \in K$ to the students in σ , that is, θ satisfying $r_i^\theta > \bar{x}_i$. Then $\eta(\tilde{\Theta}) = 0$.

Proof. Suppose $\eta(\tilde{\Theta}) > 0$. Then, since there are a finite number of cycles in $\Sigma(\gamma')$, there exists some cycle $\tilde{\sigma} = (\tilde{K}, \tilde{\underline{x}}, (\tilde{\bar{x}})) \in \Sigma(\gamma')$ containing a positive η -measure of students in $\tilde{\Theta}$. We show that $\tilde{\sigma}$ is cleared before σ . Since $\tilde{\sigma}$ contains a positive η -measure of students in $\tilde{\Theta}$, it holds that there exist $t_1, t_2 \in [\bar{t}'_{\triangleright\sigma}, t'_\sigma)$ for which $\tilde{x}_i \leq \gamma'(t_1)_i < \gamma'(t_2)_i \leq (\tilde{\bar{x}})_i$. Hence $\bar{x}_i \leq \gamma'(t'_\sigma)_i \leq \gamma'(t_1)_i < \gamma'(t_2)_i \leq \tilde{x}_i$, so $\tilde{\sigma} \triangleright \sigma$ and must be cleared before σ . Moreover, $(\tilde{\bar{x}})_i \leq \gamma'(t_1)_i < \gamma'(t_2)_i \leq \gamma(\bar{t}'_{\triangleright\sigma})_i$, so it follows from the definition of $\bar{t}'_{\triangleright\sigma}$ that $\tilde{\sigma} \notin \Sigma_{\triangleright\sigma}$, but since we assumed that $\tilde{\sigma} \in \Sigma'$ it follows that $\tilde{\sigma} \in \Sigma' \setminus \Sigma$, contradicting assumption B.1 on σ . \square

We now show in Claims (B.4) and (B.5) that in both $TTC(\gamma)$ and $TTC(\gamma')$ the cycle σ starts clearing when students in the cycle σ start clearing. We formalize this in the continuum model by considering the coordinates of the paths γ, γ' at the time t_σ when the cycle σ starts clearing, and showing that, for all coordinates indexed by schools in K , this is equal to \bar{x} .

Claim B.4. $\gamma_K(t_\sigma) = \bar{x}_K$.

Proof. The definition of t_σ implies that $\gamma(t_\sigma)_i \geq \bar{x}_i$ for all $i \in K$. Suppose there exists $i \in K$ such that $\gamma(t_\sigma)_i > \bar{x}_i$. Since σ starts clearing at time t_σ , for all $\varepsilon > 0$ school i must point to a non-zero measure of students in σ over the time period $[t_\sigma, t_\sigma + \varepsilon]$, whose scores r_i^θ satisfy $\gamma(t_\sigma)_i \geq r_i^\theta \geq \gamma(t_\sigma + \varepsilon)_i$. For sufficiently small ε the continuity of $\gamma(\cdot)$ and the assumption that $\gamma(t_\sigma)_i > \bar{x}_i$ implies that $r_i^\theta \geq \gamma(t_\sigma + \varepsilon)_i > \bar{x}_i$, which contradicts the definition of \bar{x}_i . \square

Claim B.5. $\gamma'_K(t'_\sigma) = \bar{x}_K$.

As in the proof of Claim (B.4), the definition of t'_σ implies that $\gamma'(t'_\sigma)_i \geq \bar{x}_i = \gamma(t_\sigma)_i$ for all $i \in K$. Since we cannot assume that σ is the cycle that is being cleared at time t'_σ in $TTC(\gamma')$, the proof of Claim (B.5) is more complicated than that of the Claim (B.4) and takes several steps.

We rely on the fact that K is a recurrent communication class in $TTC(\gamma)$, and that all cycles comparable to σ are already cleared in $TTC(\gamma')$. The underlying concept is very simple in the discrete model, but is complicated in the continuum by the definition of the TTC path in terms of specific points, as opposed to measures of students, and the need to account for sets of students of η -measure 0.

Let $K_ =$ be the set of coordinates in K at which equality holds, $\gamma'(t'_\sigma)_i = \gamma(t_\sigma)_i$, and let $K_ >$ be the set of coordinates in K where strict inequality holds, $\gamma'(t'_\sigma)_i > \gamma(t_\sigma)_i$. It suffices to show that $K_ >$ is empty. We do this by showing that under $TTC(\gamma')$ at time t'_σ , every school in $K_ >$ points to a zero density of students, and some school in $K_ =$ points to a non-zero density of students, and so if both sets are non-empty this contradicts the marginal trade balance equations. In what follows, let C denote the set of available schools in $TTC(\gamma)$ at time t_σ , and let C' denote the set of available schools in $TTC(\gamma')$ at time t'_σ .

Claim B.6. Suppose that $i \in K_{>}$. Then there exists $\varepsilon > 0$ such that in $TTC(\gamma')$, the set of students pointed to by school i in time $[\underline{t}'_{\sigma}, \underline{t}'_{\sigma} + \varepsilon]$ has η -measure 0, i.e. $\tilde{H}^{C'}(\gamma'(\underline{t}'_{\sigma}))_{ib} = 0$.

Proof. Since $i \in K_{>}$ it holds that $\gamma'(\underline{t}'_{\sigma})_i > \bar{x}_i$, and since γ' is continuous, for sufficiently small ε it holds that $\gamma'(\underline{t}'_{\sigma} + \varepsilon)_i > \bar{x}_i$. Hence the set of students that school i points to in time $[\underline{t}'_{\sigma}, \underline{t}'_{\sigma} + \varepsilon]$ is a subset of those with score r_i^{θ} satisfying $\gamma'(\underline{t}'_{\sigma})_i \geq r_i^{\theta} \geq \gamma'(\underline{t}'_{\sigma} + \varepsilon)_i > \bar{x}_i$. By assumption B.1 and Claim (B.3) any cycle $\tilde{\sigma}$ clearing some of these students contains at most measure 0 of them, since $\tilde{\sigma}$ is cleared after $\Sigma_{\triangleright\sigma}$ and before σ . Since there is a finite number of such cycles the set of students has η -measure 0. \square

Claim B.7. If $i \in K_{=}$, $b \in K$ and $\tilde{H}^C(\gamma(\underline{t}_{\sigma}))_{ib} > 0$, then $\tilde{H}^{C'}(\gamma'(\underline{t}'_{\sigma}))_{ib} > 0$.

Proof. Since every $\tilde{H}^C(\gamma(\underline{t}_{\sigma}))_{ib}$ is a positive multiple of $H_i^{b|C}(\gamma(\underline{t}_{\sigma}))$, it suffices to show that $H_i^{b|C'}(\gamma'(\underline{t}'_{\sigma})) > 0$. Let $\Sigma'_-(\varepsilon) \stackrel{def}{=} (\gamma'(\underline{t}'_{\sigma}) - \varepsilon \cdot e_i, \gamma'(\underline{t}'_{\sigma}))$. We first show that for sufficiently small ε it holds that $\eta^{b|C}(\Sigma'_-(\varepsilon)) = \Omega(\varepsilon)$. Let $\Sigma_-(\varepsilon) \stackrel{def}{=} (\gamma(\underline{t}_{\sigma}) - \varepsilon \cdot e^i, \gamma(\underline{t}_{\sigma}))$. Since $\tilde{H}^C(\gamma(\underline{t}_{\sigma}))_{ib} > 0$, it follows from the definition of $H_i^{b|C}(\cdot)$ that $H_i^{b|C}(x) \doteq \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \eta^{b|C}(\Sigma_-(\varepsilon)) > 0$ and hence $\eta^{b|C}(\Sigma_-(\varepsilon)) = \Omega(\varepsilon)$ for sufficiently small ε . Moreover, at most η -measure 0 of the students in $\Sigma_-(\varepsilon)$ are not in the cycle σ . Finally, $\Sigma'_-(\varepsilon) \supseteq \Sigma_-(\varepsilon) \setminus \Sigma_+(\varepsilon)$, where $\Sigma_+(\varepsilon) \stackrel{def}{=} (\gamma(\underline{t}_{\sigma}) + \varepsilon \cdot e_i, \gamma(\underline{t}_{\sigma}))$. If $\varepsilon < \bar{x}_i - \underline{x}_i$ then η -measure 0 of the students in $\Sigma_+(\varepsilon)$ are not cleared by cycle σ . Hence $\eta^{b|C}(\Sigma'_-(\varepsilon)) \geq \eta^{b|C}(\Sigma_-(\varepsilon)) - \eta^{b|C}(\Sigma_+(\varepsilon)) = \Omega(\varepsilon)$.

Suppose for the sake of contradiction that $H_i^{b|C'}(\gamma'(\underline{t}'_{\sigma})) = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \eta^{b|C'}(\Sigma'_-(\varepsilon)) = 0$, so that $\eta^{b|C'}(\Sigma'_-(\varepsilon)) = o(\varepsilon)$ for sufficiently small ε . Then there is a school $b' \neq b$ and type $\theta \in \Theta^{b|C} \cap \Theta^{b'|C'}$ such that there is an η -measure $\Omega(\varepsilon)$ of students in σ with type θ . Since $b' \in C'$ it is available in $TTC(\gamma')$ at time \underline{t}'_{σ} , and by Claim (B.1) it holds that $b' \in K$. Moreover, $\theta \in \Theta^{b|C}$ implies that θ prefers school b to all other schools in K , so $b = b'$, contradiction. \square

Proof of Claim (B.5). Suppose for the sake of contradiction that $K_{>}$ is nonempty. Since some students in σ are being cleared in $TTC(\gamma')$ at time \underline{t}'_{σ} , by Claim (B.3) there exists $i \in K = K_{=} \cup K_{>}$ and $b \in K$ such that $\tilde{H}^{C'}(\gamma'(\underline{t}'_{\sigma}))_{ib} > 0$. If $i \in K_{>}$ this contradicts Claim (B.6). If $i \in K_{=}$, then $\tilde{H}^C(\gamma(\underline{t}_{\sigma}))_{ib} > 0$ and so by Claim (B.1) $\tilde{H}^C(\gamma(\underline{t}_{\sigma}))_{ib} > 0$. Moreover, $K = K_{=} \cup K_{>}$ is a recurrent

communication class of $\tilde{H}^C(\gamma(\underline{t}_\sigma))$, so there exists a chain $i = i_0 - i_1 - i_2 - \dots - i_n$ such that $\tilde{H}^C(\gamma(\underline{t}_\sigma))_{i_i i_{i+1}} > 0$ for all $i < n$, $i_i \in K_ =$ for all $i < n - 1$, and $i_{n-1} \in K_ >$. By Claim (B.7), $\tilde{H}^{C'}(\gamma'(\underline{t}'_\sigma))_{i_i i_{i+1}} > 0$ for all $i < n$. But since $i_{n-1} \in K_ >$, by Claim (B.6) $\tilde{H}^{C'}(\gamma'(\underline{t}'_\sigma))_{i_{n-1} i_n} = 0$, which gives the required contradiction. \square

Proof that $\Sigma = \Sigma'$. We have shown in Claims (B.4) and (B.5) that for our chosen $\sigma = (K, \underline{x}, \bar{x})$, it holds that $\gamma(\underline{t}_\sigma)_K = \gamma'(\underline{t}'_\sigma)_K = \bar{x}_K$. Invoking Claims (B.2) and (B.3) and Lemma B.8 shows that σ is cleared under both $TTC(\gamma)$ and $TTC(\gamma')$. Hence $\Sigma = \Sigma'$, as required. \square

Proof of Theorem 4.3

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 $\tilde{\eta}$ satisfy the assumptions given in Section 4.3. Suppose also that the measure η and $\tilde{\eta}$ have total variation distance ε and have full support. Let γ be a TTC path for economy \mathcal{E} , and let $\tilde{\gamma}$ be a TTC path for economy $\tilde{\mathcal{E}}$. Consider any school i and any points $x = \gamma(t) \in Im(\gamma)$, $\tilde{x} = \tilde{\gamma}(\tilde{t}) \in Im(\tilde{\gamma})$ such that $x_i = \tilde{x}_i$, and both are cleared in the first round of their respective TTC runs, $t \leq t^{(1)}$ and $\tilde{t} \leq \tilde{t}^{(1)}$. We show that the set of students allocated to school i under $TTC(\gamma)$ from time 0 to t differs from the set of students allocated to school i under $TTC(\tilde{\gamma})$ from time 0 to \tilde{t} by a set of measure $O(\varepsilon|C|)$.

Proposition B.2. *Suppose that $\gamma, \tilde{\gamma}$ are TTC paths in one round of the continuum economies \mathcal{E} and $\tilde{\mathcal{E}}$ respectively, where the set of available schools C is the same in these rounds of $TTC(\gamma)$ and $TTC(\tilde{\gamma})$. Suppose also that γ starts and ends at x, y , and $\tilde{\gamma}$ starts and ends at \tilde{x}, \tilde{y} , where there exist $j, k \in C$ such that $x_j = \tilde{x}_j$, $y_k = \tilde{y}_k$, and $x_a \leq \tilde{x}_a$, $y_a \leq \tilde{y}_a$ for all $a \in C$. Then for all $i \in C$, the set of students with ranks in $(y, x] \cap (\tilde{y}, \tilde{x}]$ who are assigned to i under $TTC(\gamma)$ and not under $TTC(\tilde{\gamma})$ has measure $O(\varepsilon|C|)$.¹³*

Proof. By Lemma B.3, we may assume without loss of generality that γ and $\tilde{\gamma}$ are parametrized such that $x = \gamma(0), y = \gamma(1)$ and $\tilde{x} = \tilde{\gamma}(0), \tilde{y} = \tilde{\gamma}(1)$, and for all times $\tau \in [0, 1]$ there exists a school $i(\tau)$ such that $\gamma(\tau)$ is dominated by $\tilde{\gamma}(\tau)$ via school $i(\tau)$.

¹³This is according to both measures η and $\tilde{\eta}$.

Let $T_i = \{\tau \leq 1 : i(\tau) = i\}$ be the times when γ is dominated by $\tilde{\gamma}$ via school i . We remark that, by our construction in Lemma B.3, we may assume that T_i is the countable union of disjoint closed intervals, and that if $i \neq i'$ then T_i and $T_{i'}$ have disjoint interiors.

Since γ is a TTC path for \mathcal{E} and $\tilde{\gamma}$ is a TTC path for $\tilde{\mathcal{E}}$, by integrating over the marginal trade balance equations we can show that the following trade balance equations hold,

$$\eta(\mathcal{T}_i(\gamma; T_i)) = \eta(\mathcal{T}^{i|C}(\gamma; T_i)) \text{ for all } i \in C. \quad (\text{B.1})$$

$$\tilde{\eta}(\mathcal{T}_i(\tilde{\gamma}; T_i)) = \tilde{\eta}(\mathcal{T}^{i|C}(\tilde{\gamma}; T_i)) \text{ for all } i \in C. \quad (\text{B.2})$$

Since γ is dominated by $\tilde{\gamma}$ via school j at all times $\tau \in T_j$, we have that

$$\mathcal{T}_j(\gamma; T_j) \subseteq \mathcal{T}_j(\tilde{\gamma}; T_j). \quad (\text{B.3})$$

Moreover, by the choice of parametrization, $\cup_j T_j = [0, 1]$ and so, since $x \leq \tilde{x}$,

$$\cup_{i,j} \mathcal{T}^{i|C}(\gamma; T_j) \supseteq \cup_{i,j} \mathcal{T}^{i|C}(\tilde{\gamma}; T_j). \quad (\text{B.4})$$

Now since $\eta, \tilde{\eta}$ have total variation ε , for every school i it holds that

$$\begin{aligned} \eta(\mathcal{T}^{i|C}(\gamma; T_i) \setminus \mathcal{T}^{i|C}(\tilde{\gamma}; T_i)) &\leq \eta(\mathcal{T}^{i|C}(\gamma; T_i)) - \eta(\mathcal{T}^{i|C}(\tilde{\gamma}; T_i)) + \varepsilon \text{ (by (B.4))} \\ &= \eta(\mathcal{T}_i(\gamma; T_i)) - \tilde{\eta}(\mathcal{T}_i(\tilde{\gamma}; T_i)) + \varepsilon \text{ (by (B.1) and (B.2))} \\ &\leq 2\varepsilon \text{ (by (B.3)),} \end{aligned} \quad (\text{B.5})$$

Also, for all schools $i \neq j$, since η has full support and bounded density $\nu \in [m, M]$, it holds that

$$\eta(\mathcal{T}^{i|C}(\gamma; T_j) \setminus \mathcal{T}^{i|C}(\tilde{\gamma}; T_j)) \leq \frac{M}{m} \eta(\mathcal{T}^{j|C}(\gamma; T_j) \setminus \mathcal{T}^{j|C}(\tilde{\gamma}; T_j)). \quad (\text{B.6})$$

Hence, as T_j have disjoint interiors,

$$\begin{aligned}
\eta\left(\mathcal{T}^{i|C}(\gamma; 1) \setminus \mathcal{T}^{i|C}(\tilde{\gamma}; 1)\right) &= \sum_{j \in C} \left(\eta(\mathcal{T}^{i|C}(\gamma; T_j)) - \eta(\mathcal{T}^{i|C}(\tilde{\gamma}; T_j))\right) \quad (\text{by (B.4)}) \\
&\leq \sum_{j \in C} \eta\left(\mathcal{T}^{i|C}(\gamma; T_j) \setminus \mathcal{T}^{i|C}(\tilde{\gamma}; T_j)\right) \\
&\leq \sum_{j \in C} \frac{M}{m} \eta\left(\mathcal{T}^{j|C}(\gamma; T_j) \setminus \mathcal{T}^{j|C}(\tilde{\gamma}; T_j)\right) \quad (\text{by (B.6)}) \\
&\leq 2|C|\varepsilon \frac{M}{m} \quad (\text{by (B.5)}).
\end{aligned}$$

That is, given a school i , the set of students assigned to school i with score $r^\theta \not\leq x$ under γ and not assigned to school i with score $r^\theta \not\leq \tilde{x}$ under $\tilde{\gamma}$ has η -measure $O(\varepsilon|C|)$. The result for $\tilde{\eta}$ follows from the fact that the total variation distance of η and $\tilde{\eta}$ is ε . \square

We are now ready to prove Theorem 4.3.

Proof of Theorem 4.3. Assume without loss of generality that the schools are labeled in order. Let σ be a permutation such that if we reindex school $\sigma(i)$ to be school i then the schools are labeled in order under $TTC(\tilde{\gamma})$. We show by induction on ℓ that $\sigma(\ell) = \ell$ and that for all schools i , the set of students assigned to i under $TTC(\gamma)$ by the end of the ℓ th round and not under $TTC(\tilde{\gamma})$ by the end of the ℓ th round has η -measure $O(\varepsilon\ell|C|)$. This will prove the theorem.

We first consider the base case $\ell = 1$. Let $x = \tilde{x} = \gamma(0)$ and $y = \gamma(t^{(1)})$. Define $\tilde{y} \in Im(\tilde{\gamma})$ to be the minimal point such that $y \leq \tilde{y}$ and there exists i such that $y_i = \tilde{y}_i$. We show that \tilde{y} is near $\tilde{\gamma}(\tilde{t}^{(1)})$, i.e. $|\tilde{y} - \tilde{\gamma}(\tilde{t}^{(1)})|_2 = O(\varepsilon)$. Now by Proposition B.2 the set of students with ranks in $(y, \gamma(0)] \cap (\tilde{y}, \gamma(0)]$ who are assigned to 1 under $TTC(\gamma)$ and not under $TTC(\tilde{\gamma})$ has $\tilde{\eta}$ -measure $O(\varepsilon|C|)$. Hence the residual capacity of school 1 at \tilde{y} under $TTC(\tilde{\gamma})$ is $O(\varepsilon|C|)$, and so since $\tilde{\eta}$ has full support and has density bounded from above and below by M and m , it holds that $|\tilde{y} - \tilde{\gamma}(\tilde{t}^{(1)})|_2 = O\left(\frac{M}{m}\varepsilon|C|\right)$. (If the residual capacity is negative we can exchange the roles of γ and $\tilde{\gamma}$ and argue similarly.)

Let us now show that the inductive assumption holds. Fix a school i . Then by Proposition B.2 the set of students with ranks in $(y, \gamma(0)] \cap (\tilde{y}, \gamma(0)]$ who are assigned to i under $TTC(\gamma)$ and

not under $TTC(\tilde{\gamma})$ has $\tilde{\eta}$ -measure $O(\varepsilon|\mathcal{C}|)$. Moreover, since $\left| \tilde{y} - \tilde{\gamma}(\tilde{t}^{(1)}) \right|_2 = O\left(\frac{M}{m}\varepsilon|\mathcal{C}|\right)$ and $\tilde{\eta}$ has full support and has density bounded from above and below by M and m , the set of students with ranks in $(\tilde{y}, \tilde{\gamma}(\tilde{t}^{(1)})]$ assigned to school i by $TTC(\tilde{\gamma})$ has $\tilde{\eta}$ -measure $O(\varepsilon|\mathcal{C}|)$. Hence the set of students assigned to i under $TTC(\gamma)$ by time $t^{(1)}$ and not under $TTC(\tilde{\gamma})$ by time $\tilde{t}^{(1)}$ has η -measure $O(\varepsilon|\mathcal{C}|)$. Moreover, if $t^{(1)} < t^{(2)}$ then for sufficiently small ε it holds that $\tilde{t}^{(1)} = \min_i \tilde{t}^{(i)}$, and otherwise there exists a relabeling of the schools such that this is true, and so $\sigma(1) = 1$.

We now show the inductive step, proving for $\ell + 1$ assuming true for $1, 2, \dots, \ell$. By inductive assumption, for all i the measure of students assigned to i under $TTC(\gamma)$ and not under $TTC(\tilde{\gamma})$ by the points $\gamma(t^{(\ell)}), \tilde{\gamma}(\tilde{t}^{(\ell)})$ is $O(\varepsilon\ell|\mathcal{C}|)$ for all i .

Let $x = \gamma(t^{(\ell)})$ and $y = \gamma(t^{(\ell+1)})$. Define $\tilde{x} \in Im(\tilde{\gamma})$ to be the minimal point such that $x \leq \tilde{x}$ and there exists b such that $x_b = \tilde{x}_b$. We show that \tilde{x} is near $\tilde{\gamma}(\tilde{t}^{(\ell)})$, i.e. $\left| \tilde{x} - \tilde{\gamma}(\tilde{t}^{(\ell)}) \right|_2 = O(\varepsilon)$. Now by inductive assumption $\eta(\{\theta | r^\theta \in (x = \gamma(t^{(\ell)}), \tilde{\gamma}(\tilde{t}^{(\ell)}))\}) = O(\varepsilon\ell|\mathcal{C}|)$ and so $\left| x - \tilde{\gamma}(\tilde{t}^{(\ell)}) \right|_2 = O(\varepsilon)$. Moreover $\left| \tilde{x}_b - \tilde{\gamma}_b(\tilde{t}^{(\ell)}) \right|_2 = \left| x_b - \tilde{\gamma}_b(\tilde{t}^{(\ell)}) \right|_2$ which we have just shown is $O(\varepsilon)$. Finally, since η has full support and has density bounded from above and below by M and m , it holds that $\max_{b,i,\tau} \frac{\gamma'_b(\tau)}{\tilde{\gamma}'_i(\tau)} = O\left(\frac{M}{m}\right)$ and so for all i it holds that $\left| \tilde{x}_i - \tilde{\gamma}_i(\tilde{t}^{(\ell)}) \right| \leq O\left(\frac{M}{m}\varepsilon\right)$.

The remainder of the proof runs much the same as in the base case, with slight adjustments to account for the fact that $x \neq \tilde{x}$. Define $\tilde{y} \in Im(\tilde{\gamma})$ to be the minimal point such that $y \leq \tilde{y}$ and there exists i such that $y_i = \tilde{y}_i$. We show that \tilde{y} is near $\tilde{\gamma}(\tilde{t}^{(\ell+1)})$, i.e. $\left| \tilde{y} - \tilde{\gamma}(\tilde{t}^{(\ell+1)}) \right|_2 = O(\varepsilon)$. Now by Proposition B.2 the set of students with ranks in $(y, x] \cap (\tilde{y}, \tilde{x}]$ who are assigned to $\ell + 1$ under $TTC(\gamma)$ and not under $TTC(\tilde{\gamma})$ has $\tilde{\eta}$ -measure $O(\varepsilon|\mathcal{C}|)$. This, together with the inductive assumption that the difference in students assigned to school ℓ is $O(\varepsilon\ell|\mathcal{C}|)$, shows that the residual capacity of school $\ell + 1$ at \tilde{y} under $TTC(\tilde{\gamma})$ is $O(\varepsilon(\ell + 1)|\mathcal{C}|)$, and so since $\tilde{\eta}$ has full support and has density bounded from above and below by M and m , it holds that $\left| \tilde{y} - \tilde{\gamma}(\tilde{t}^{(\ell+1)}) \right|_2 = O\left(\frac{M}{m}\varepsilon(\ell + 1)|\mathcal{C}|\right)$. (If the residual capacity is negative we can exchange the roles of γ and $\tilde{\gamma}$ and argue similarly.)

Let us now show that the inductive assumption holds. Fix a school i . Then by Proposition B.2 the set of students with ranks in $(y, x] \cap (\tilde{y}, \tilde{x}]$ who are assigned to i under $TTC(\gamma)$ and not under $TTC(\tilde{\gamma})$ has $\tilde{\eta}$ -measure $O(\varepsilon|\mathcal{C}|)$. Moreover, since $\left| \tilde{y} - \tilde{\gamma}(\tilde{t}^{(\ell+1)}) \right|_2 = O\left(\frac{M}{m}\varepsilon(\ell + 1)|\mathcal{C}|\right)$ and $\tilde{\eta}$ has

full support and has density bounded from above and below by M and m , the set of students with ranks in $(\tilde{y}, \tilde{\gamma}(\tilde{t}^{(\ell+1)}))$ assigned to school i by $TTC(\tilde{\gamma})$ has $\tilde{\eta}$ -measure $O(\varepsilon(\ell+1)|\mathcal{C}|)$. Hence the set of students assigned to i under $TTC(\gamma)$ by time $t^{(\ell+1)}$ and not under $TTC(\tilde{\gamma})$ by time $\tilde{t}^{(\ell+1)}$ has η -measure $O(\varepsilon(\ell+1)|\mathcal{C}|)$. Moreover if $t^{(\ell+1)} < t^{(\ell+2)}$ then for sufficiently small ε it holds that $\tilde{t}^{(\ell+1)} = \min_{i>\ell} \tilde{t}^{(i)}$, and otherwise there exists a relabeling of the schools such that this is true, and so $\sigma(\ell+1) = \ell+1$. \square

B.2 Omitted Proofs for Section 4.4

B.2.1 Derivation of Marginal Trade Balance Equations

In this section, we show that the marginal trade balance equations (4.2) hold,

$$\sum_{k \in \mathcal{C}} \gamma'_k(\tau) \cdot H_k^i(x) = \sum_{k \in \mathcal{C}} \gamma'_i(\tau) \cdot H_i^k(x).$$

The idea is that the measure of students who trade into a school i must be equal to the measure of students who trade out of i .

In particular, suppose that at some time τ the TTC algorithm has assigned exactly the set of students with rank better than $x = \gamma(\tau)$, and the set of available schools is \mathcal{C} . Consider the incremental step of a TTC path γ from $\gamma(\tau) = x$ over ε units of time. The process of cycle clearing imposes that for any school $i \in \mathcal{C}$, the total amount of seats offered by school i from time τ to $\tau + \varepsilon$ is equal to the amount of students assigned to i plus the amount of seats that were offered but not claimed or traded by the student it was over to over that same time period. In the continuum model the set of seats offered but not claimed or traded is of η -measure 0.¹⁴ Hence the set of students assigned to school i from time τ to $\tau + \varepsilon$ has the same measure as the set of students who

¹⁴A student can have a seat that is offered but not claimed or traded in one of two ways. The first is the seat is offered at time τ and not yet claimed or traded. The second is that the student that got offered two or more seats at the same time $\tau' \leq \tau$ (and was assigned through a trade involving only one seat). Both of these sets of students are of η -measure 0 under our assumptions.

were offered a seat at school i in that time,

$$\begin{aligned} & \eta \left(\left\{ \theta \in \Theta^{i|C} \mid r^\theta \in [\gamma(\tau + \epsilon), \gamma(\tau)] \right\} \right) \\ &= \eta \left(\left\{ \theta \in \Theta \mid \exists \tau' \in [\tau, \tau + \epsilon] \text{ s.t. } r_i^\theta = \gamma_i(\tau') \text{ and } r^\theta \leq \gamma(\tau') \right\} \right), \end{aligned}$$

or more compactly,

$$\eta \left(\mathcal{T}^{i|C}(\gamma; [\tau, \tau + \epsilon]) \right) = \eta \left(\mathcal{T}_i(\gamma; [\tau, \tau + \epsilon]) \right). \quad (\text{B.7})$$

We now prove that the marginal trade balance equations follow from equation (B.7). Following the notation in Appendix 4.4.2, for $i, j \in \mathcal{C}$, $x \in [0, 1]^{\mathcal{C}}$, $\alpha \in \mathbb{R}$ we define the set ¹⁵

$$T_i^{j|C}(x, \alpha) \doteq \left\{ \theta \in \Theta^{j|C} \mid r^\theta \in [x - \alpha e^i, x] \right\}.$$

We may think of $T_i^{j|C}(x, \alpha)$ as the set of the next α students on school i 's priority list who are unassigned when $\gamma(\tau) = x$, and want school j . We remark that the sets used in the definition of the $H_i^{j|C}(x)$ are precisely the sets $T_i^{j|C}(x, \alpha)$.

We can use the sets $T_i^{j|C}(x, \alpha)$ to approximate the expressions in equation (B.7) involving $\mathcal{T}_i(\gamma; \cdot)$ and $\mathcal{T}^{i|C}(\gamma; \cdot)$.

Lemma B.9. *Let $\gamma(\tau) = x$ and for all $\epsilon > 0$ let $\delta(\epsilon) = \gamma(\tau) - \gamma(\tau + \epsilon)$. For sufficiently small ϵ , during the interval $[\tau, \tau + \epsilon]$, the set of students who were assigned to school i is*

$$\mathcal{T}^{i|C}(\gamma; [\tau, \tau + \epsilon]) = \bigcup_j T_j^{i|C}(x, \delta_j(\epsilon))$$

and the set of students who were offered a seat at school i is

$$\mathcal{T}_i(\gamma; [\tau, \tau + \epsilon]) = \bigcup_k T_i^{k|C} \left(x - \sum_{i' \neq i} \delta_{i'}(\epsilon) e^{i'}, \delta_i(\epsilon) \right) \cup \Delta$$

for some small set $\Delta \subset \Theta$. Further, it holds that $\lim_{\tau \rightarrow 0} \frac{1}{\tau} \cdot \eta(\Delta) = 0$, and for any $i \neq i', k \neq k' \in \mathcal{C}$ we have $\lim_{\tau \rightarrow 0} \frac{1}{\tau} \cdot \eta \left(T_i^{k|C}(x, \delta_i(\epsilon)) \cap T_{i'}^{k'|C}(x, \delta_{i'}(\epsilon)) \right) = 0$ and $T_i^{k|C}(x, \delta_i(\tau)) \cap T_i^{k'|C}(x, \delta_i(\epsilon)) = \phi$.

¹⁵We use the notation $[\underline{x}, \bar{x}] = \{z \in \mathbb{R}^n \mid \underline{x}_i \leq z_i < \bar{x}_i \forall i\}$ for $\underline{x}, \bar{x} \in \mathbb{R}^n$, and $e^i \in \mathbb{R}^{\mathcal{C}}$ is a vector whose i -th coordinate is equal to 1 and all other coordinates are 0.

Proof. The first two equations are easily verified, and the fact that the last intersection is empty is also easy to verify. To show the bound on the measure of Δ , we observe that it is contained in the set $\bigcup_{i'} \cup_k \left(T_i^{k|C}(x, \delta_i(\epsilon)) \cap T_{i'}^{k|C}(x, \delta_{i'}(\epsilon)) \right)$, so it suffices to show that $\lim_{\tau \rightarrow 0} \frac{1}{\tau} \cdot \eta \left(T_i^{k|C}(x, \delta_i(\epsilon)) \cap T_{i'}^{k|C}(x, \delta_{i'}(\epsilon)) \right) = 0$. This follows from the fact that the density defining η is upper bounded by M , so

$$\eta \left(T_i^{k|C}(x, \delta_i(\epsilon)) \cap T_{i'}^{k|C}(x, \delta_{i'}(\epsilon)) \right) \leq M |\gamma_i(\tau) - \gamma_i(\tau + \epsilon)| |\gamma_{i'}(\tau) - \gamma_{i'}(\tau + \epsilon)|.$$

Since for all schools i the function γ_i is continuous and has bounded derivative, it is also Lipschitz continuous, so

$$\frac{1}{\tau} \eta(\Delta) \leq \frac{1}{\tau} \eta \left(T_i^{k|C}(x, \delta_i(\epsilon)) \cap T_{i'}^{k|C}(x, \delta_{i'}(\epsilon)) \right) \leq M L_i L_{i'} \epsilon$$

for some Lipschitz constants L_i and $L_{i'}$ and the lemma follows. \square

We now now ready to take limits and verify that equation (B.7) implies that the marginal trade balance equations hold. Let us divide equation (B.7) by $\delta_i(\epsilon) = \gamma_i(\tau) - \gamma_i(\tau + \epsilon)$ and take the limit as $\epsilon \rightarrow 0$. Then on the left hand side we obtain

$$\begin{aligned} & \lim_{\epsilon \rightarrow 0} \frac{1}{\delta_i(\epsilon)} \eta \left(\mathcal{T}^{i|C}(\gamma; [\tau, \tau + \epsilon]) \right) \\ &= \lim_{\epsilon \rightarrow 0} \frac{1}{\delta_i(\epsilon)} \eta \left(\bigcup_j T_j^{i|C}(x, \delta_j(\epsilon)) \right) \quad (\text{Lemma B.9}) \\ &= \lim_{\epsilon \rightarrow 0} \left[\sum_{j \in C} \frac{1}{\delta_i(\epsilon)} \eta \left(T_j^{i|C}(x, \delta_j(\epsilon)) \right) + O \left(\frac{(\|\gamma(\tau) - \gamma(\tau + \epsilon)\|_\infty)^2}{\delta_i(\epsilon)} \right) \right] \\ & \quad (\text{as density is bounded, } \nu < M) \\ &= \lim_{\epsilon \rightarrow 0} \left[\sum_{j \in C} \frac{1}{\delta_i(\epsilon)} \eta \left(T_j^{i|C}(x, \delta_j(\epsilon)) \right) \right] \quad (\gamma \text{ Lipschitz continuous}) \\ &= \lim_{\epsilon \rightarrow 0} \left[\sum_{j \in C} \frac{\delta_j(\epsilon)}{\delta_i(\epsilon)} \cdot \frac{1}{\delta_j(\epsilon)} \eta \left(\{\theta \in \Theta^{i|C} \mid r^\theta \in [x - \delta_j(\epsilon) \cdot e^j, x]\} \right) \right] \\ &= \sum_{j \in C} \frac{\gamma'_j(\tau)}{\gamma'_i(\tau)} \cdot H_j^{i|C}(x) \quad (\text{by definition of } \delta \text{ and } H) \end{aligned}$$

as required. Similarly, on the right hand side we obtain

$$\begin{aligned}
& \lim_{\epsilon \rightarrow 0} \frac{1}{\delta_i(\epsilon)} \eta(\mathcal{T}_i(\gamma; [\tau, \tau + \epsilon])) \\
&= \lim_{\epsilon \rightarrow 0} \left[\sum_{k \in C} \frac{1}{\delta_i(\epsilon)} \eta \left(T_i^{k|C} \left(x - \sum_{i' \neq i} \delta_{i'}(\epsilon) e^{i'}, \delta_i(\epsilon) \right) \right) + O \left(\frac{(\|\gamma(\tau + \epsilon) - \gamma(\tau)\|_\infty)^2}{\delta_i(\epsilon)} \right) \right] \text{ (Lemma B.9)} \\
&= \lim_{\epsilon \rightarrow 0} \left[\sum_{k \in C} \frac{1}{\delta_i(\epsilon)} \eta \left(T_i^{k|C} \left(x - \sum_{i' \neq i} \delta_{i'}(\epsilon) e^{i'}, \delta_i(\epsilon) \right) \right) \right] \text{ (\gamma is Lipschitz continuous)} \\
&= \lim_{\epsilon \rightarrow 0} \left[\sum_{k \in C} \frac{1}{\delta_i(\epsilon)} \eta \left(\left\{ \theta \in \Theta^{k|C} \mid r^\theta \in \left[x - \delta(\epsilon), x - \sum_{i' \neq i} \delta_{i'}(\epsilon) e^{i'} \right] \right\} \right) \right] \\
&= \sum_{k \in C} H_i^{k|C}(x) \text{ (by definition of } \delta \text{ and } H)
\end{aligned}$$

as required. This completes the proof.

B.3 Omitted Proofs for Applications (Section 4.5)

Throughout this section, we will say that a vector d is a *valid direction at point x* if d satisfies the marginal trade balance equations at x , and $d \cdot \mathbf{1} = -1$. We will also augment the notation from Section 4.3 to specify the economy. Specifically, for an economy $\mathcal{E} = (\mathcal{C}, \Theta, \eta, q)$ let

$$D^i(x|\mathcal{E}) = \eta \left(\left\{ \theta \mid r^\theta \not\prec x, Ch^\theta(C) = i \right\} \right)$$

denote the mass of students whose rank at some school j is better than x_j and whose first choice is school i .

Effects of Changes in the Distribution of School Quality

In this section, we prove the results stated in Section 4.5.1. We will assume that the total measure of students is 1, and speak of student measures and student proportions interchangeably.

Proof of Proposition 4.4. Given quality δ , let η be the measure over Θ and $\gamma, p, \{t^{(1)}, t^{(2)}\}$ be the TTC path, cutoffs and stopping times. Given quality $\hat{\delta}$, let $\hat{\eta}$ be the measure over Θ and $\hat{\gamma}, \hat{p}, \{\hat{t}^{(1)}, \hat{t}^{(2)}\}$ be the TTC path, cutoffs and stopping times.

For each $x \in [0, 1]^2$ let $\mathbf{d}(x)$ (resp. $\hat{\mathbf{d}}(x)$) denote the valid direction at x under \mathcal{E}_δ (resp. $\mathcal{E}_{\hat{\delta}}$) with support that is minimal under the order $\{1\} < \{1, 2\} < \{2\}$. As there are only two schools, $|d_1(x)| \geq |\hat{d}_1(x)|$ and $|d_2(x)| \leq |\hat{d}_2(x)|$ for all x .¹⁶ It follows that $\hat{\gamma}$ moves faster in the 2 direction than γ does, i.e. if $\gamma_1(t) = \hat{\gamma}_1(\hat{t})$ then $\gamma_2(t) \geq \hat{\gamma}_2(\hat{t})$, and if $\gamma_2(t) = \hat{\gamma}_2(\hat{t})$ then $\gamma_1(t) \leq \hat{\gamma}_1(\hat{t})$. Hence without loss of generality we may assume that the time parameters in the TTC paths are scaled so that at all times t the path $\hat{\gamma}$ is dominated by γ via school 1, i.e. $\gamma_1(t) = \hat{\gamma}_1(t)$ and $\gamma_2(t) \geq \hat{\gamma}_2(t)$ for all t (see Appendix (B.1)).

Suppose for the sake of contradiction that $p_2^1 < \hat{p}_2^1$, i.e. $\gamma_2(t^{(1)}) < \hat{\gamma}_2(\hat{t}^{(1)})$. We may interpret this as it becoming more difficult to use priority at school 2 to trade into 1 after 2 gets more popular. We will show that this will also result in more students being assigned under γ by time $t^{(1)}$ than under $\hat{\gamma}$ by time $\hat{t}^{(1)}$. But since school 1 is also more popular under \mathcal{E} this means that more students are assigned to school 1 under $TTC(\gamma|\mathcal{E})$ than $TTC(\hat{\gamma}|\hat{\mathcal{E}})$, which gives the required contradiction.

More formally, since $\hat{\gamma}$ is dominated by γ via school 1 at time $t^{(1)}$ it follows that $\hat{\gamma}_2(t^{(1)}) \leq \gamma_2(t^{(1)}) < \hat{\gamma}_2(\hat{t}^{(1)})$ and so $\hat{t}^{(1)} < t^{(1)}$, i.e. school 1 now fills earlier. Hence $\hat{\gamma}_1(\hat{t}^{(1)}) \geq \hat{\gamma}_1(t^{(1)}) = \gamma_1(t^{(1)})$, where the equality comes from the assumption that $\hat{\gamma}$ is dominated by γ via school 1 at time $t^{(1)}$. But this gives the necessary contradiction, as $\hat{\gamma}(\hat{t}^{(1)}) \geq \gamma(t^{(1)})$ implies that

$$q_1 = D^1(\hat{\gamma}(\hat{t}^{(1)})|\mathcal{E}_{\hat{\delta}}) < D^1(\gamma(t^{(1)})|\mathcal{E}_\delta) \leq D^1(\gamma(t^{(1)}|\mathcal{E}_\delta)) = q_1,$$

where the first inequality follows from $\hat{\gamma}(\hat{t}^{(1)}) \geq \gamma(t^{(1)})$ and the second inequality holds since $\hat{\delta}_2 \geq \delta_2$ and $\hat{\delta}_1 = \delta_1$.

We now show that $p_1^1 \geq \hat{p}_1^1$, i.e. it becomes easier to use priority at school 1 to be assigned to school 1. Suppose for the sake of contradiction that $p_1^1 < \hat{p}_1^1$, i.e. $\gamma(t^{(1)}) < \hat{\gamma}(\hat{t}^{(1)})$. We will use the marginal trade balance equations to show that this means more students traded into school 1 under γ by time $t^{(1)}$ than under $\hat{\gamma}$ by time $\hat{t}^{(1)}$, which gives the required contradiction.

Since $\hat{\gamma}$ is dominated by γ via school 1 it holds that $\hat{\gamma}(t^{(1)}) = \gamma(t^{(1)}) < \hat{\gamma}(\hat{t}^{(1)})$ and so

¹⁶Note that by definition valid directions have norm 1.

$t^{(1)} > \hat{t}^{(1)}$, i.e. school 1 fills earlier under $TTC(\hat{\gamma}|\hat{\mathcal{E}})$. Hence the sets of students offered seats by school 1 satisfy

$$\mathcal{T}_1(\gamma; t^{(1)}) \supsetneq \mathcal{T}_1(\gamma; \hat{t}^{(1)}) \supseteq \mathcal{T}_1(\hat{\gamma}; \hat{t}^{(1)}),$$

where the first containment follows from the fact that $t^{(1)} > \hat{t}^{(1)}$ and the second containment follows from the fact that $\hat{\gamma}$ is dominated by γ via school 1, and so fewer students are offered/trade away seats at school 1 by time $\hat{t}^{(1)}$ under $\hat{\gamma}$ than under γ .

Moreover, integrating over the marginal trade balance equations gives that under both paths, the set of students who traded a seat at 2 for a seat at 1 has the same measure as the set of students who traded a seat at 1 for a seat at 2,

$$\eta\left(\left\{\theta \in \mathcal{T}_2(\gamma; t^{(1)}) \mid Ch^\theta\{1,2\} = 1\right\}\right) = \eta\left(\left\{\theta \in \mathcal{T}_1(\gamma; t^{(1)}) \mid Ch^\theta\{1,2\} = 2\right\}\right) \quad \text{and} \quad (\text{B.8})$$

$$\hat{\eta}\left(\left\{\theta \in \mathcal{T}_2(\hat{\gamma}; \hat{t}^{(1)}) \mid Ch^\theta\{1,2\} = 1\right\}\right) = \hat{\eta}\left(\left\{\theta \in \mathcal{T}_1(\hat{\gamma}; \hat{t}^{(1)}) \mid Ch^\theta\{1,2\} = 2\right\}\right). \quad (\text{B.9})$$

Hence we can compare the number of students assigned to school 1 using these sets, and find that

$$\begin{aligned} q_1 &= \eta\left(\left\{\theta \in \mathcal{T}_1(\gamma; t^{(1)}) \mid Ch^\theta\{1,2\} = 1\right\}\right) + \eta\left(\left\{\theta \in \mathcal{T}_2(\gamma; t^{(1)}) \mid Ch^\theta\{1,2\} = 1\right\}\right) \\ &= \eta\left(\left\{\theta \in \mathcal{T}_1(\gamma; t^{(1)})\right\}\right) \quad (\text{by (B.8)}) \\ &> \eta\left(\left\{\theta \in \mathcal{T}_1(\hat{\gamma}; \hat{t}^{(1)})\right\}\right) \quad (\text{since the sets are strictly contained}) \\ &= \hat{\eta}\left(\left\{\theta \in \mathcal{T}_1(\hat{\gamma}; \hat{t}^{(1)})\right\}\right) \\ &= \eta\left(\left\{\theta \in \mathcal{T}_1(\gamma; t^{(1)}) \mid Ch^\theta\{1,2\} = 1\right\}\right) + \hat{\eta}\left(\left\{\theta \in \mathcal{T}_2(\gamma; t^{(1)}) \mid Ch^\theta\{1,2\} = 1\right\}\right) \\ &\quad (\text{by (B.8)}) \\ &= q_1 \end{aligned}$$

which gives the required contradiction.

The fact that $\hat{p}_2^2 \geq p_2^2$ follows from the fact that $\hat{p}_1^1 \leq p_1^1$ decreases, since the total number of assigned students is the same. \square

Proof of Proposition 4.5.

In the logit economy we assume that the total measure of students is normalized to 1, and that

$\sum_i q_i < 1$. Recall that we also assume that all students prefer all schools to being unassigned. Note that the logit economy yields that $\mathbb{P}(Ch^\theta(C) = i) = \frac{e^{\delta_i}}{\sum_{k \in C} e^{\delta_k}}$.

We first show that schools are labeled in order if $\frac{q_1}{e^{\delta_1}} \leq \frac{q_2}{e^{\delta_2}} \leq \dots \leq \frac{q_n}{e^{\delta_n}}$. This holds since at any point $\gamma(t) = \mathbf{x}$ in the first round the choice probabilities yield that $\frac{e^{\delta_i}}{\sum_{k \in C} e^{\delta_k}} (1 - \prod_j x_j)$ students are assigned to school i , and so for all i, j the ratio of students assigned to schools j and i respectively is $\frac{e^{\delta_j}}{e^{\delta_i}}$ and if the schools are labeled in order then $\frac{q_1}{e^{\delta_1}} = \min_i \frac{q_i}{e^{\delta_i}}$. The other inequalities hold by induction, since in any round with remaining schools C and $i \in C$ the choice probabilities yield that a fraction $\frac{e^{\delta_i}}{\sum_{k \in C} e^{\delta_k}}$ of the students assigned to schools in C are assigned to school i so again for all $i, j \in C$ the ratio of students assigned to schools j and i respectively in that round (or any preceding round) is $\frac{e^{\delta_j}}{e^{\delta_i}}$.

This also shows that $R^i = 1 - \sum_{i' < i} q_{i'} - \frac{\pi_c}{e^{\delta_i}} q_i$ is the measure of unassigned, or remaining, students after the c th round, since if $i' < i$ then $q_{i'}$ students are assigned to school i , and if $i' \geq c$ then $\frac{e^{\delta_{i'}}}{e^{\delta_i}} q_{i'}$ students are assigned to school i .

TTC Cutoffs We calculate the TTC cutoffs under the logit economy for different student choice probabilities by using the TTC paths and trade balance equations. We show by induction on i that for all i

$$p_j^i = \begin{cases} \left(\prod_{k \leq i} \left(\frac{R^k}{R^{k-1}} \right)^{1/\pi_k} \right) e^{\delta_j} & \text{if } j \geq i, \\ p_j^j & \text{otherwise,} \end{cases} \quad (\text{B.10})$$

where $\pi_i = \sum_{i' \geq i} e^{\delta_{i'}}$, $R^0 = 1$ and for all $i \geq 1$ the quantity $R^i = 1 - \sum_{i' < i} q_{i'} - \frac{\pi_i}{e^{\delta_i}} q_i$ is the measure of unassigned, or remaining, students after the i th round. We note that if we let $\rho_i = \frac{q_i}{e^{\delta_i}} - \frac{q_{i-1}}{e^{\delta_{i-1}}}$, where $q_{i-1} = \delta_{i-1} = 0$, then

$$R^{i-1} - R^i = -\frac{\pi_{i-1}}{e^{\delta_{i-1}}} q_{i-1} + q_{i-1} + \frac{\pi_i}{e^{\delta_i}} q_i = \rho_i \pi_i,$$

and so

$$\sum_{i' \leq i} \rho_{i'} \pi_{i'} = \sum_{i' \leq i} R^{i'-1} - R^{i'} = 1 - R^i.$$

Consider the base case $i = 1$. In round 1, the marginals $H_j^i(x)$ for $i, j \in \mathcal{C}$ at each point $x \in [0, 1]$ are given by $H_j^i(x) = \frac{e^{\delta_i}}{\sum_{k \in \mathcal{C}} e^{\delta_k}} \prod_{i' \neq j} x_{i'}$. As the valid directions $\mathbf{d} = \mathbf{d}(x)$ solve the marginal trade balance equations, they must satisfy $\sum_{k \in \mathcal{C}} d_k H_k^i(x) = \sum_{k \in \mathcal{C}} d_i H_i^k(x)$, or equivalently

$$e^{\delta_i} \sum_{k \in \mathcal{C}} \frac{d_k}{x_k} = \frac{d_i}{x_i} \sum_{k \in \mathcal{C}} e^{\delta_k}.$$

Now the vector $\mathbf{d}(x)$ defined by

$$d_i(x) = -\frac{e^{\delta_i} x_i}{\sum_{j \in \mathcal{C}} e^{\delta_j} x_j}$$

clearly satisfies both the marginal trade balance equations and the normalization $\mathbf{d}(x) \cdot \mathbf{1} = -1$. Moreover since $H(x)$ is irreducible this is the unique valid direction \mathbf{d} .

We now find a valid TTC path γ using the trade balance equations (4.2). Since the ratios of the components of the gradient $\frac{d_j(x)}{d_i(x)}$ only depend on x_j, x_i and the $\delta_{i'}$, for all i we solve for x_i in terms of x_1 , using the marginal trade balance equations and the fact that the path starts at $\mathbf{1}$. This gives the path γ defined by $\gamma_i(\gamma_1^{-1}(x_1)) = x_1^{e^{\delta_i} - \delta_1}$ for all i .

Recall that the schools are indexed so that school i is the most demanded school, that is, $\frac{e^{\delta_1}}{q_1} = \max_i \frac{e^{\delta_i}}{q_i}$. Now school i fills at a time $t^{(1)}$ where the TTC path is given by $\gamma_i(t^{(1)}) = x_1^{e^{\delta_i} - \delta_1}$ and the number of assigned students is given by

$$1 - \prod_i \gamma_i(t^{(1)}) = 1 - R^1$$

where the left hand side is the measure of students with rank at least $\gamma_i(t^{(1)})$ for at least one school i , and the right hand side is the number of assigned students.

This yields that

$$p_j^1 = \gamma_i(t^{(1)}) = \left(\prod_i \gamma_i(t^{(1)}) \right)^{\frac{e^{\delta_j}}{\pi_1}} = (R^1)^{\frac{e^{\delta_j}}{\pi_1}}.$$

where $\pi_1 = \sum_{i' \geq 1} e^{\delta_{i'}}$. This completes the base case.

For the inductive step, suppose that Equation (B.10) holds for the cutoffs in rounds $1, 2, \dots, i-1$.

Consider the residual TTC path during the i th round and let it be denoted by $\tilde{\gamma}$. For all $j \geq i$ let $x_j = \tilde{\gamma}_j(t)$. Recall that by definition $\tilde{\gamma}_j(t) = p_j^{i-1} = p_j^j$ for all $j < i$ and $t \geq t^{(i-1)}$. The residual TTC path is non-constant only for schools j in the set $\mathcal{C}^{(i)} = \{i, i+1, \dots, n\}$, and the marginal trade balance conditions specify that for these schools j and for all $x \leq p^{i-1}$ it holds that $\frac{d_j(x)}{d_i(x)} = \frac{e^{\delta_j x_j}}{e^{\delta_i x_i}}$. Therefore we can solve for x_j in terms of x_i , using the fact that the path starts at p^{i-1} . The marginal trade balance conditions and initial conditions yield that for all $j \geq i$

$$\frac{\tilde{\gamma}_j(t)^{-e^{\delta_j}}}{\tilde{\gamma}_i(t)^{-e^{\delta_i}}} = \frac{(p_j^{i-1})^{e^{-\delta_j}}}{(p_i^{i-1})^{e^{-\delta_i}}} = 1,$$

where the first equality is obtained by integrating over the marginal trade balance equations and providing the initial conditions, and the second equality holds by substituting in the values of p^{i-1} in the inductive assumption. Hence the path $\tilde{\gamma}$ is defined by $\tilde{\gamma}_j(\tilde{\gamma}_i^{-1}(x_i)) = x_i$ for all $j \geq i$, and $\tilde{\gamma}_j(t^{(i)}) = p_j^j$ for all $j < i$.

Now school i fills at a time $t^{(i)}$ where the TTC path is given by $\tilde{\gamma}_j(t^{(i)}) = x_i$ for all $j \geq i$ and $\tilde{\gamma}_j(t^{(i)}) = p_j^j$ for all $j < i$, and the number of students assigned from time $t^{(i-1)}$ to $t^{(i)}$ is given by

$$\prod_{i' \in \mathcal{C}} p_{i'}^{i-1} - \prod_{i' < i} p_{i'}^{i-1} \prod_{j \geq i} \tilde{\gamma}_j(t^{(i)}) = R^{i-1} - R^i, \quad (\text{B.11})$$

where the left hand side is the measure of students with rank at least $\tilde{\gamma}_j(t^{(i)})$ for at least one school j who is not assigned in one of the first $i-1$ rounds. Noting that

$$\begin{aligned} \prod_j p_j^{i-1} &= \left(\prod_{j < i-1} p_j^j \right) \left(\prod_{j \geq i-1} p_j^{i-1} \right) \\ &= \left(\prod_{k < i-1} \left(\frac{R^k}{R^{k-1}} \right)^{1-\pi_{i-1}/\pi_k} \right) \left(\prod_{k \leq i-1} \left(\frac{R^k}{R^{k-1}} \right)^{\pi_{i-1}/\pi_k} \right) \\ &= R^{i-1} \end{aligned}$$

allows us to simplify equation (B.11) to

$$\prod_{j \geq i} x_j = \frac{R^i}{\prod_{i' < i} p_{i'}^{i-1}}.$$

Substituting in $p_{i'}^{i-1} = \left(\prod_{k \leq i'} \left(\frac{R^k}{R^{k-1}} \right)^{1/\pi_k} \right) e^{\delta_{i'}}$ yields

$$x_j = x_i = \left(R^i \prod_{k < i} \left(\frac{R^k}{R^{k-1}} \right)^{-(\pi_k - \pi_i)/\pi_k} \right) e^{\delta_j/\pi_i} = \left(\prod_{k \leq i} \left(\frac{R^k}{R^{k-1}} \right)^{1/\pi_k} \right) e^{\delta_j}$$

as required. \square

TTC Cutoffs - Comparative Statics We perform some comparative statics calculations for the TTC cutoffs under the logit model. For $j \neq \ell$ it holds that the TTC cutoff p_j^1 for using priority at school j to receive a seat at school 1 is decreasing in δ_ℓ . Formally,

$$\begin{aligned} \frac{\partial p_j^1}{\partial \delta_\ell} &= \frac{\partial}{\partial \delta_\ell} \left[(1 - \rho_1 \pi_1) \frac{e^{\delta_j}}{\pi_1} \right] \\ &= -p_j^1 \left(\frac{e^{\delta_\ell + \delta_j}}{(\pi_1)^2} \right) \left[-\ln \left(\frac{1}{1 - \rho_1 \pi_1} \right) + \frac{1}{(1 - \rho_1 \pi_1)} - 1 \right] \end{aligned}$$

is negative, since $0 < \frac{1}{(1 - \rho_1 \pi_1)} < 1$ and $f(x) = x - \ln(x) - 1$ is positive for $x \in [0, 1]$.

For $j = \ell$ the TTC cutoff p_ℓ^1 is again decreasing in δ_ℓ ;

$$\begin{aligned} \frac{\partial p_\ell^1}{\partial \delta_\ell} &= \frac{\partial}{\partial \delta_\ell} \left[(1 - \rho_1 \pi_1) \frac{e^{\delta_\ell}}{\pi_1} \right] \\ &= -p_\ell^1 \left(\frac{e^{\delta_\ell} (\pi_1 - e^{\delta_\ell})}{(\pi_1)^2} \right) \ln \left(\frac{1}{1 - \rho_1 \pi_1} \right) - p_\ell^1 \left(\frac{e^{2\delta_\ell}}{(\pi_1)^2} \right) \left(\frac{1}{(1 - \rho_1 \pi_1)} - 1 \right) \end{aligned}$$

is negative since both terms are negative.

Similarly, for $i < \ell$ and $j \geq i$ the TTC cutoff p_j^i is decreasing in δ_ℓ . We first show that this holds for $i < \ell$ and $j \geq i$, $j \neq \ell$ by showing that $\frac{1}{e^{\delta_j}} \ln p_j^i$ is decreasing in δ_ℓ . Now

$$\begin{aligned} \frac{\partial}{\partial \delta_\ell} \left[\frac{1}{e^{\delta_j}} \ln p_j^i \right] &= \frac{\partial}{\partial \delta_\ell} \left[\sum_{k \leq i} \frac{1}{\pi_k} \ln \left(\frac{R^k}{R^{k-1}} \right) \right] \\ &= \sum_{k \leq i} \left(-\frac{e^{\delta_\ell}}{(\pi_k)^2} \right) \left[\ln \left(\frac{R^k}{R^{k-1}} \right) - \frac{\pi_k}{e^{\delta_\ell}} \cdot \frac{\partial}{\partial \delta_\ell} \left[\ln \left(\frac{R^k}{R^{k-1}} \right) \right] \right] \end{aligned}$$

where

$$\begin{aligned}
\frac{\partial}{\partial \delta_\ell} \left[\ln \left(\frac{R^k}{R^{k-1}} \right) \right] &= \frac{R^{k-1} \frac{\partial R^k}{\partial \delta_\ell} - R^k \frac{\partial R^{k-1}}{\partial \delta_\ell}}{R^{k-1} R^k} \\
&= -\frac{e^{\delta_\ell}}{R^{k-1} R^k} \left[R^{k-1} \left(\frac{q_k}{e^{\delta_k}} \right) - R^k \left(\frac{q_{a-1}}{e^{\delta_{a-1}}} \right) \right] \\
&= -\frac{e^{\delta_\ell}}{R^{k-1} R^k} \left[R^{k-1} \rho_k + \pi_k \rho_k \left(\frac{q_{a-1}}{e^{\delta_{a-1}}} \right) \right] \\
&= -\frac{e^{\delta_\ell} \rho_k}{R^{k-1} R^k} \left(1 - \sum_{i' < k} q_{i'} \right).
\end{aligned}$$

Hence

$$\frac{\partial}{\partial \delta_\ell} \left[\frac{1}{e^{\delta_j}} \ln p_j^i \right] = \sum_{k \leq i} \left(-\frac{e^{\delta_\ell}}{(\pi_k)^2} \right) \left[\ln \left(\frac{R^k}{R^{k-1}} \right) + \left(\frac{1}{R^k} - \frac{1}{R^{k-1}} \right) \left(1 - \sum_{i' < k} q_{i'} \right) \right] \leq 0$$

where the last inequality holds since for all k the first term is negative, and the second term is given by $f_k(R^k) - f_k(R^{k-1})$ where $f_k(x) = (1 - \sum_{i' < k} q_{i'}) \frac{1}{x} + \ln(x)$ has negative derivative $f'_k(x) \leq 0$ for all $x \leq (1 - \sum_{i' < k} q_{i'})$, and $R^k \leq R^{k-1} < (1 - \sum_{i' < k} q_{i'})$ so $f_k(R^k) - f_k(R^{k-1}) \geq 0$.

For $i < \ell$ and $j = \ell$ the TTC cutoff p_ℓ^i is also decreasing in δ_ℓ , since

$$\frac{\partial}{\partial \delta_\ell} \left[\ln p_\ell^i \right] = \frac{\partial}{\partial \delta_\ell} \left[\frac{e^{\delta_\ell}}{e^{\delta_j}} \ln p_i^i \right] = \frac{e^{\delta_\ell}}{e^{\delta_j}} \left(\ln p_i^i + \frac{\partial}{\partial \delta_\ell} \left[\ln p_i^i \right] \right) \leq 0$$

where the last inequality holds since $p_i^i < 1$ and we have shown that $\frac{\partial}{\partial \delta_\ell} \left[\ln p_i^i \right] \leq 0$.

When $i = \ell$ and $j > \ell$, we note that

$$\begin{aligned}
\prod_{k \geq \ell} p_k^\ell \prod_{i' < \ell} p_{i'}^{i'} &= R^\ell, \text{ i.e.} \\
p_j^\ell &= \left(\frac{R^\ell}{\prod_{i' < \ell} p_{i'}^{i'}} \right)^{e^{\delta_j / \pi_\ell}}.
\end{aligned}$$

Hence

$$\begin{aligned}
\frac{\partial}{\partial \delta_\ell} \left[\frac{1}{e^{\delta_j}} \ln p_j^\ell \right] &= \frac{\partial}{\partial \delta_\ell} \left[\frac{1}{\pi_\ell} \left(\ln R^\ell - \sum_{i' < \ell} \ln p_{i'} \right) \right] \\
&= \left(-\frac{e^{\delta_\ell}}{(\pi_\ell)^2} \right) \left(\ln R^\ell - \sum_{i' < \ell} \ln p_{i'} \right) \\
&\quad + \frac{1}{\pi_\ell} \left(\frac{\partial}{\partial \delta_\ell} [\ln R^\ell] - \sum_{i' < \ell} \frac{\partial}{\partial \delta_\ell} [\ln p_{i'}] \right) \\
&\geq 0.
\end{aligned}$$

where the first term is positive since $p_j^\ell < 1$ (from which it follows that $\ln R^\ell - \sum_{i' < \ell} \ln(p) < 0$), and the second term is positive since $\frac{\partial R^\ell}{\partial \delta_\ell} = \frac{\pi_{\ell+1}}{(e^{\delta_\ell})^2} q_\ell > 0$ and we have shown that for all $i' < \ell$ it holds that $\frac{\partial}{\partial \delta_\ell} [\ln p_{i'}] \leq 0$.

Proof of Proposition 4.6.

Welfare Expressions We derive the welfare expressions corresponding to these cutoffs. Let $\mathcal{C}^{(i)} = \{i, i+1, \dots, n\}$. Since the schools are ordered so that $\frac{q_1}{e^{\delta_1}} \leq \frac{q_2}{e^{\delta_2}} \leq \dots \leq \frac{q_n}{e^{\delta_n}}$, it follows that the schools also fill in the order $1, 2, \dots, n$.

Suppose that the total mass of students is 1. Then the mass of students with budget set $\mathcal{C}^{(1)}$ is given by $N^1 = q_1 \left(\frac{\sum_j e^{\delta_j}}{e^{\delta_1}} \right) = \rho_1 \pi_1$, and the mass of students with budget set $\mathcal{C}^{(2)}$ is given by $N^2 = \left(q_2 - \frac{e^{\delta_2}}{\sum_j e^{\delta_j}} N^1 \right) \left(\frac{\sum_{j \geq 2} e^{\delta_j}}{e^{\delta_2}} \right) = \left(\frac{q_2}{e^{\delta_2}} - \frac{q_1}{e^{\delta_1}} \right) \left(\sum_{j \geq 2} e^{\delta_j} \right) = \rho_2 \pi_2$. A straightforward inductive argument shows that the proportion of students with budget set $\mathcal{C}^{(i)}$ is

$$N^i = \left(\frac{q_i}{e^{\delta_i}} - \frac{q_{i-1}}{e^{\delta_{i-1}}} \right) \left(\sum_{j \geq i} e^{\delta_j} \right) = \rho_i \pi_i.$$

which depends only on δ_j for $j \geq i-1$.

Moreover, each such student with budget set $\mathcal{C}^{(i)}$, conditional on their budget set, has expected utility Small and Rosen (1981)

$$U^i = \mathbb{E} \left[\max_{i' \in \mathcal{C}^{(i)}} \{ \delta_j + \varepsilon_{\theta i'} \} \right] = \ln \left[\sum_{j \geq i} e^{\delta_j} \right] = \ln(\pi_i),$$

which depends only on δ_j for $j \geq i$. Hence the expected social welfare from fixed qualities δ_i is given by

$$U_{TTC} = \sum_i N^i \cdot U^i = \sum_i \rho_i \pi_i \ln \pi_i,$$

where $\pi_i = \sum_{j \geq i} e^{\delta_j}$.

Welfare - Comparative Statics Taking derivatives, we obtain that

$$\frac{dU_{TTC}}{d\delta_\ell} = \sum_i \left(\frac{dN^i}{d\delta_\ell} \cdot U^i + N^i \cdot \frac{dU^i}{d\delta_\ell} \right) = \sum_{i \leq \ell+1} \frac{dN^i}{d\delta_\ell} \cdot U^i + \sum_{i \leq \ell} N^i \cdot \frac{dU^i}{d\delta_\ell},$$

where $\sum_{i \leq \ell} N^i \cdot \frac{dU^i}{d\delta_\ell} = \sum_{i \leq \ell} \rho_i \pi_i \cdot \frac{e^{\delta_\ell}}{\pi_i} = e^{\delta_\ell} \sum_{i \leq \ell} \rho_i = q_\ell$. It follows that

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

□

Proof of Proposition 4.7. We solve for the social welfare maximising budget allocation. For a fixed runout ordering (i.e. $\frac{q_1}{e^{\delta_1}} \leq \frac{q_2}{e^{\delta_2}} \leq \dots \leq \frac{q_n}{e^{\delta_n}}$), the central school board's investment problem is given by the program

$$\begin{aligned} \max_{\kappa_1, \kappa_2, \dots, \kappa_n} \quad & \sum_i \left(\frac{q_i}{\kappa_i} - \frac{q_{i-1}}{\kappa_i} \right) \left(\sum_{j \geq i} \kappa_j \right) \ln \left(\sum_{j \geq i} \kappa_j \right) \\ \text{s.t.} \quad & \frac{q_{i-1}}{\kappa_{i-1}} \leq \frac{q_i}{\kappa_{i-1}} \quad \forall i \\ & \sum_i \kappa_i = K \\ & q_0 = 0. \end{aligned} \tag{B.12}$$

We can reformulate this as the following program,

$$\max_{\kappa_2, \dots, \kappa_n} \left(\frac{q_1}{K - \sum_i \kappa_i} \right) K \ln K + \left(\frac{q_2}{\kappa_2} - \frac{q_1}{K - \sum_i \kappa_i} \right) \pi_2 \ln \pi_2 + \sum_{i \geq 3} \left(\frac{q_i}{\kappa_i} - \frac{q_{i-1}}{\kappa_{i-1}} \right) \pi_i \ln \pi_i \tag{B.13}$$

$$\begin{aligned}
s.t. \quad & \frac{q_{i-1}}{\kappa_{i-1}} \leq \frac{q_i}{\kappa_i} \quad \forall i \geq 3 \\
& \frac{q_1}{K - \sum_i \kappa_i} \leq \frac{q_2}{\kappa_2}, \\
& \pi_i = \sum_{j \geq i} \kappa_j.
\end{aligned}$$

Taking the derivatives of the objective U with respect to the budget allocations κ_k gives

$$\frac{\partial U}{\partial \kappa_k} = \left(\frac{q_1}{(K - \sum_i \kappa_i)^2} \right) \ln \left(\frac{K^K}{\pi_2^{\pi_2}} \right) + \sum_{2 \leq i < k} \frac{q_i}{\kappa_i} \ln \frac{\pi_i}{\pi_{i+1}} + \frac{q_k}{(\kappa_k)^2} \ln \left(\frac{\pi_k^{\pi_k}}{\pi_{k+1}^{\pi_{k+1}}} \right),$$

where $\ln \left(\frac{K^K}{\pi_2^{\pi_2}} \right) \geq 0$, $\ln \frac{\pi_i}{\pi_{i+1}} \geq 0$, and $\ln \left(\frac{\pi_k^{\pi_k}}{\pi_{k+1}^{\pi_{k+1}}} \right)$ and so $\frac{\partial U}{\partial \kappa_k} \geq 0 \forall k$.

Moreover, if $\frac{q_{i-1}}{\kappa_{i-1}} = \frac{q_i}{\kappa_i}$, then defining a new problem with $n - 1$ schools, and capacities \tilde{q} and budget $\tilde{\kappa}$

$$\tilde{q}_j = \begin{cases} q_j & \text{if } j < i - 1 \\ q_{i-1} + q_i & \text{if } j = i - 1 \\ q_{j+1} & \text{if } j > i - 1 \end{cases}, \quad \tilde{\kappa}_j = \begin{cases} \kappa_j & \text{if } j < i - 1 \\ \kappa_{i-1} + \kappa_i & \text{if } j = i - 1 \\ \kappa_{j+1} & \text{if } j > i - 1 \end{cases}$$

leads to a problem with the same objective function, since

$$\begin{aligned}
& \left(\frac{q_{i-1}}{\kappa_{i-1}} - \frac{q_{i-2}}{\kappa_{i-2}} \right) \pi_{i-1} \ln \pi_{i-1} + \left(\frac{q_i}{\kappa_i} - \frac{q_{i-1}}{\kappa_{i-1}} \right) \pi_i \ln \pi_i + \left(\frac{q_{i+1}}{\kappa_{i+1}} - \frac{q_i}{\kappa_i} \right) \pi_{i+1} \ln \pi_{i+1} \\
& = \left(\frac{q_{i-1} + q_i}{\kappa_{i-1} + \kappa_i} - \frac{q_{i-2}}{\kappa_{i-2}} \right) \pi_{i-1} \ln \pi_{i-1} + 0 + \left(\frac{q_{i+1}}{\kappa_{i+1}} - \frac{q_{i-1} + q_i}{\kappa_{i-1} + \kappa_i} \right) \pi_{i+1} \ln \pi_{i+1}.
\end{aligned}$$

Hence if there exists i for which $\frac{q_i}{\kappa_i} \neq \frac{q_{i-1}}{\kappa_{i-1}}$, we may take i to be minimal such that this occurs, decrease each of $\kappa_1, \dots, \kappa_{i-1}$ proportionally so that $\kappa_1 + \dots + \kappa_{i-1}$ decreases by ε and increase κ_i by ε and increase the resulting value of the objective. It follows that the objective is maximized when $\frac{q_1}{\kappa_1} = \frac{q_2}{\kappa_2} = \dots = \frac{q_n}{\kappa_n}$, i.e. when the money assigned to each school is proportional to the number of seats at the school. \square

Design of TTC Priorities

We demonstrate how to calculate the TTC cutoffs for the two economies in Figure 4.14 by using the TTC paths and trade balance equations.

Consider the economy \mathcal{E} , where the top priority students have ranks uniformly distributed in $[m, 1]^2$. If $x = (x_1, x_1)$ is on the diagonal, then $\tilde{H}_i^j(x) = \frac{x_1}{2}$ for all $i, j \in \{1, 2\}$, and so there is a unique valid direction $d(\vec{x}) = \begin{bmatrix} -\frac{1}{2} \\ -\frac{1}{2} \end{bmatrix}$. Moreover, $\gamma(t) = (\frac{t}{2}, \frac{t}{2})$ satisfies $\frac{d\gamma(t)}{dt} = \mathbf{d}(\gamma(t))$ for all t and hence Theorem 4.2 implies that $\gamma(t) = (\frac{t}{2}, \frac{t}{2})$ is the unique TTC path. The cutoff points satisfy $p_1^1 = p_2^1 = p_1^2 = p_2^2 = p$ for some constant p , and (by symmetry) the capacity equations $D^1(\mathbf{p}) = D^2(\mathbf{p}) = q$ for $\mathbf{p} = (p, p)$. Since $D^1(\mathbf{p}) + D^2(\mathbf{p}) = 1 - p^2$, it follows that $1 - p^2 = 2q$, or $p = \sqrt{1 - 2q}$. The cutoff points $p_b^c = \sqrt{1 - 2q}$ give the unique TTC allocation.

Consider now the economy $\bar{\mathcal{E}}$, where top priority students have ranks uniformly distributed in the $\tilde{r} \times \tilde{r}$ square $(1 - \tilde{r}, 1] \times (m, m + \tilde{r}]$ for some small \tilde{r} , where $\tilde{r} \leq \frac{(2m-1)(1-m)}{2m}$.

If x is in $(1 - \tilde{r}, 1] \times [m + \tilde{r}, 1]$ then $H_1^j(x) = \frac{1}{2} \left(m + (1 - m) \frac{1-m}{\tilde{r}} \right) \forall j$ and $H_2^j(x) = \frac{m}{2} \forall j$, so there is a unique valid direction $d(x) = \frac{1}{2 + \frac{(1-m)^2}{\tilde{r}m}} \begin{bmatrix} -1 \\ -1 - \frac{(1-m)^2}{\tilde{r}m} \end{bmatrix}$. If x is in $(m, 1 - \tilde{r}] \times (m, 1]$ then $H_i^j(x) = \frac{m}{2}$ for all i, j and there is a unique valid direction $d(x) = \begin{bmatrix} -\frac{1}{2} \\ -\frac{1}{2} \end{bmatrix}$. Finally, if $x = (x_1, x_2)$ is in $[0, 1] \setminus (m, 1]^2$ then $H_1^j(x) = \frac{1}{2}x_2$ and $H_2^j(x) = \frac{1}{2}x_1$ for all j and there is a unique valid direction $d(x) = \frac{1}{x_1 + x_2} \begin{bmatrix} -x_1 \\ -x_2 \end{bmatrix}$.

Hence the TTC path $\gamma(t)$ has gradient proportional to $\begin{bmatrix} -1 \\ -1 - \frac{(1-m)^2}{\tilde{r}m} \end{bmatrix}$ from the point $(1, 1)$ to the point $(1 - \tilde{r}, 1 - \tilde{r} - \frac{(1-m)^2}{m})$, to $\begin{bmatrix} -\frac{1}{2} \\ -\frac{1}{2} \end{bmatrix}$ from the point $(1 - \tilde{r}, 1 - \tilde{r} - \frac{(1-m)^2}{m})$ to the point $(m + \frac{(1-m)^2}{m}, m)$ and to $\begin{bmatrix} -1 - \frac{(1-m)^2}{m^2} \\ -1 \end{bmatrix}$ from the point $(m + \frac{(1-m)^2}{m}, m)$ to the cutoff point (\bar{p}, \underline{p}) .

We find that

$$(\bar{p}, \underline{p}) = \left(\sqrt{(1 - 2q) \frac{(1 - 2m + 2m^2)}{m^2}}, \sqrt{(1 - 2q) \frac{m^2}{1 - 2m + 2m^2}} \right)$$

by observing that $\frac{1}{2}(1 - \bar{p} \cdot \underline{p}) = D^1\left(\left(\bar{p}, \underline{p}\right)\right) = q$ and that (\bar{p}, \underline{p}) lies on the line passing through $\left(m + \frac{(1-m)^2}{m}, m\right)$ with gradient $\frac{1}{1 + \frac{(1-m)^2}{m^2}}$.

We now show that the economy \mathcal{E} is extremal, i.e. if economy \mathcal{E}' is given by perturbing the relative ranks of students in $\{\theta \mid r_c^\theta \geq m \forall c\}$, then the TTC cutoffs for \mathcal{E}' are given by $p_1^1 = p_1^2 = x$, $p_2^1 = p_2^2 = y$ where $x \leq \bar{p} = \sqrt{\frac{1-2q}{1-2m+2m^2}}$ and $y \geq \underline{p} = \sqrt{(1-2q)(1-2m+2m^2)}$. (By symmetry, it follows that $\underline{p} \leq x, y \leq \bar{p}$.)

Let $\bar{\gamma}$ and γ' be the TTC paths for $\bar{\mathcal{E}}$ and \mathcal{E}' respectively. Let (x_{bound}, m) be the point where the TTC path γ' first hits the boundary of the box $[m, 1] \times [m, 1]$ containing all the highly ranked students. We remark that the TTC path γ' for \mathcal{E}' has gradient $\frac{1}{x_{bound}+m} \begin{bmatrix} -x_{bound} \\ -m \end{bmatrix}$ from (x_{bound}, m) to the TTC cutoffs (x, y) .

Consider the aggregate trade balance equations for students assigned before the TTC path reaches (x_{bound}, m) . They stipulate that the measure of students in $[0, m] \times [m, 1]$ who prefer school 1 is at most the measure of students who are either perturbed or in $[x_{bound}, 1] \times [0, m]$, and who prefer school 2. This means that $\frac{1}{2}m(1-m) \leq \frac{1}{2}\left((1-m)^2 + m(1-x_{bound})\right)$, or $x_{bound} \leq m + \frac{(1-m)^2}{m}$. It follows that γ' hits the boundary of the box at a point that is to the left of where $\bar{\gamma}$ hits the boundary box, and hence the path γ' lies above the path $\bar{\gamma}$.¹⁷ It follows that $x \leq \bar{p}$ and $y \geq \frac{1-2q}{\bar{p}} = \underline{p}$.

¹⁷That is, for each x' , if (x', y') lies on γ' and (x', \bar{y}) lies on $\bar{\gamma}$, then $y' \geq \bar{y}$.

Appendix C

Appendix for Chapter 5

C.1 Regret-Free Mechanisms in Extensive Form

In this section, we provide three extensive-form game descriptions of regret-free mechanisms. We first define general mechanisms as extensive-form dynamic games of imperfect information. We note that as regret-free mechanisms are incentive-compatible, we may restrict our attention to direct mechanisms, where students need only report either their type, or given a set of schools to inspect need only inspect that set of schools and truthfully report their inspected values. We then formally define choice-based messages, where we restrict the messages from students to be only choice-based information about their own preferences, such as choice functions and partial orders, and we restrict the actions and messages of the mechanism designer to use only the choice-based information. Finally, we formally define Accept-Waitlist-Reject (AWR) mechanisms, which restrict the mechanism designer to only tell students which schools will definitely be in their budget set, which schools definitely will not, and which schools are uncertain.

C.1.1 General Mechanisms

We formally define general mechanisms as dynamic games of incomplete information. There is a set of players: $s \in \mathcal{S}$, possible actions: $a \in \mathcal{A}$, and possible messages $m \in \mathcal{M}$. There is a set of nodes \mathcal{Z} , with initial node z^0 and terminal nodes T . At each node $z \in \mathcal{Z}$ there is a message history $\mathcal{M}_z = \{\mathcal{M}_z^s\}_{s \in \mathcal{S}}$; at each non-terminal node there is a set of active students $S(z)$ who are sent new messages $\{\iota_z^s = \iota^s(\mathcal{M}_z)\}_{s \in S(z)}$. At each node z every student has private information about their history $h_z^s = (\{a_{z'}^s, \mathcal{M}_{z'}^s, \iota_{z'}^s\})$ of their actions. These actions result in inspections $\chi = \chi_z^s$ and values $v^s|_\chi$, which are also privately known to the student.

There is a partition $\mathcal{H} = \{H_1, H_2, \dots\}$ of the nodes into information sets, which represent the information available to the students. Each student s has a partition $\mathcal{H}^s = \{H_1^s, \dots\}$ such that any two nodes z, z' are in the same information set (i.e. $z, z' \in H_i^s$) for some i if and only if, up to relabeling of other students, they have the same information state $\mathcal{M}_z^s = \mathcal{M}_{z'}^s$ and history $h_z^s = h_{z'}^s$. Note that for a given student s the mechanism designer only knows r^s as well as \mathcal{M}^s , and so if students s, s' satisfy $(r^s, \mathcal{M}_z^s) = (r^{s'}, \mathcal{M}_{z'}^{s'})$ then they are *indistinguishable to the mechanism*

designer at node z . Each information set H_i has a set of active students $S = S(H_i)$ of positive measure such that $s \in S \Rightarrow H_i \in \mathcal{H}^s$ and if two students s, s' are indistinguishable to the mechanism designer at any node $z \in H_i$ then $s \in S \Leftrightarrow s' \in S$.

The available actions $A^s(H_i) \subseteq \mathcal{A}$ for the set of active students $s \in S(H_i)$ at an information set H_i are as follows. Students first inspect some subset of uninspected schools $\{i : \chi_z^s = 0\}$, where the subset can be adaptively chosen based on the observed values of other schools inspected at that node. Students then report a message m to the mechanism.

At terminal nodes $z \in T$ the mechanism outputs a matching μ_z and inspections χ_z , where χ_z^s is consistent with history h_z^s .

We let Σ^s denote the set of strategies for student s , i.e. an action $a \in A^s(H_i)$ for each history H_i such that $s \in S(H_i)$.

Definition C.1. We say that a general mechanism is **regret-free stable** if at all terminal nodes $z = (\mu_z, \chi_z) \in T$ the matching μ_z is regret-free stable with any underlying economy consistent with the mechanism designer's current information state \mathcal{M}_z , and $\chi_z = \chi_z^{RF}(\mu|\cdot)$.

C.1.2 Choice-Based Mechanisms

We formally define choice-based mechanisms as dynamic games of incomplete information as follows.

There is a set of players $s \in \mathcal{S}$, possible actions, $a \in \mathcal{A}$, and possible messages $m \in \mathcal{M}$. The main restriction of a choice-based mechanism is that \mathcal{M} is restricted to be the set of *choice-based information states* $\mathcal{I} = \{\succeq^s\}_{s \in \mathcal{S}}$. There is a set of nodes \mathcal{Z} , with initial node z^0 , and terminal nodes T . At each node $z \in \mathcal{Z}$ there is a *partial message history, given by a single choice-based information state* $\mathcal{I}_z = \{\succeq_z^s\}_{s \in \mathcal{S}}$; at each non-terminal node there is a set of active students $S(z)$ who are sent new messages $\{\iota_z^s = \iota^s(\mathcal{I}_z)\}_{s \in S(z)}$ based on the information state \mathcal{I}_z . At each node each student has private information about their history $h_z^s = (\{a_{z'}^s, \mathcal{I}_{z'}, \iota_{z'}^s\}_i)$ of actions. These actions result in inspections $\chi = \chi_z^s$ and values $v^s|_\chi$, which are also privately known to the student.

There is a $\mathcal{H} = \{H_1, H_2, \dots\}$ of the nodes into information sets. Each student s has a partition $\mathcal{H}^s = \{H_1^s, \dots\}$ such that any two nodes z, z' are in the same information set $z, z' \in H_i^s$ for some i if and only if, up to relabeling of other students, they have the same choice-based information

state $\mathcal{I}_z^s = \mathcal{I}_z^{s'}$ and history $h_z^s = h_z^{s'}$. Note that for a given student s the mechanism designer only knows r^s as well as \mathcal{I}_z^s , and so if students s, s' satisfy $(r^s, \mathcal{I}_z^s) = (r^{s'}, \mathcal{I}_z^{s'})$ then they are indistinguishable to the mechanism designer at node z . Each information set H_i has: (1) a set of active students $S = S(H_i)$ of positive measure such that $s \in S \Rightarrow H_i \in \mathcal{H}^s$ and if two students s, s' are indistinguishable to the mechanism designer at any node $z \in H_i$ then $s \in S \Leftrightarrow s' \in S$; and (2) a set of schools $c(H_i)$.

The available actions $A^s(H_i) \subseteq \mathcal{A}$ for the set of active students $s \in S(H_i)$ at an information set H_i are as follows. Students first inspect some subset of uninspected schools $\{i : \chi_z^s = 0\}$ in $c(H_i)$, where the subset can be adaptively chosen based on the observed values of other schools inspected at that node. Students then report a refinement of \succeq_z^s , which encodes their choice of schools in that set.

At terminal nodes $z \in T$ the mechanism outputs a matching μ_z and inspections χ_z , where χ_z^s is consistent with history h_z^s .

We let Σ^s denote the set of strategies for student s , i.e. an action $a \in A^s(H_i)$ for each history H_i such that $s \in S(H_i)$.

Definition C.2. We say that a choice-based mechanism is **regret-free stable** if at all terminal nodes $z = (\mu_z, \chi_z) \in T$ the matching μ_z is regret-free stable with any underlying economy consistent with the mechanism designer's current information state \mathcal{I}_z , and $\chi_z = \chi_z^{RF}(\mu|\cdot)$

C.1.3 Accept-Waitlist-Reject Mechanisms

We formally define Accept-Waitlist-Reject (AWR) mechanisms as dynamic games of incomplete information as follows.

There is a set of players $s \in \mathcal{S}$, possible actions, $a \in \mathcal{A}$, and possible messages $m \in \mathcal{M}$. The main restriction of an AWR mechanism is that \mathcal{M} is restricted to be the set messages of the form $\mathcal{C}^{\{A,W,R\}} \cup \mathcal{C}$, where each element of $\mathcal{C}^{\{A,W,R\}}$ encodes a possible message from the mechanism to a student about the set of schools that accept them (A) as the school is definitely in their budget set, waitlist them (L) as the school may or may not be in their budget set, or reject them (R) as the school is definitely not in their budget set. There is a set of nodes \mathcal{Z} , with initial node

z^0 , and terminal nodes T . At each node $z \in \mathcal{Z}$ there is a *partial message history* m_z , given by a single element $m_z \in \mathcal{M}^S$, with m_z^s representing the last message sent between the mechanism and each student. At each non-terminal node there is a set of active students $S(z)$ who are sent new messages $\{\iota_z^s \in \mathcal{C}^{\{A,W,R\}}\}_{s \in S(z)}$ based on the message history m_z . At each node each student has private information about their history $h_z^s = (\{a_{z'}^s, m_{z'}^s, \iota_{z'}^s\})$ of actions. These actions result in inspections $\chi = \chi_z^s$ and values $v^s|_\chi$, which are also privately known to the student.

There is a $\mathcal{H} = \{H_1, H_2, \dots\}$ of the nodes into information sets. Each student s has a partition $\mathcal{H}^s = \{H_1^s, \dots\}$ such that any two nodes z, z' are in the same information set $z, z' \in H_i^s$ for some i if and only if they have the same history $h_z^s = h_{z'}^s$. Note that for a given student s the mechanism designer only knows r^s as well as \mathcal{I}^s , and so if students s, s' satisfy $(r^s, m_z^s) = (r^{s'}, m_z^{s'})$ then they are indistinguishable to the mechanism designer at node z . Each information set H_i has: (1) a set of active students $S = S(H_i)$ of positive measure such that $s \in S \Rightarrow H_i \in \mathcal{H}^s$ and $m_z^s \in \mathcal{C}^{\{A,W,R\}}$, and if two students s, s' are indistinguishable to the mechanism designer at any node $z \in H_i$ then $s \in S \Leftrightarrow s' \in S$; and (2) a message in $m_z \in \mathcal{C}^{\{A,W,R\}}$ for those students.

The available actions $A^s(H_i) \subseteq \mathcal{A}$ for the set of active students $s \in S(H_i)$ at an information set H_i are as follows. Students first inspect some subset of uninspected schools $\{i : \chi_z^s = 0\}$ such that i accepts them (i.e. the element of m_z^s corresponding to school i is A), where the subset can be adaptively chosen based on the observed values of other schools inspected at that node. Students then report their favorite school $i \in A$.

At terminal nodes $z \in T$ the mechanism outputs a matching μ_z and inspections χ_z , where χ_z^s is consistent with history h_z^s

We let Σ^s denote the set of strategies for student s , i.e. an action $a \in A^s(H_i)$ for each history H_i such that $s \in S(H_i)$. Note that in an AWR mechanism each student will be an active student at a node exactly once.

Definition C.3. We say that an AWR mechanism is **regret-free stable** if at all terminal nodes $z = (\mu_z, \chi_z) \in T$ the matching μ_z is regret-free stable with any underlying economy consistent with the mechanism designer's current information state \mathcal{I}_z , and $\chi_z = \chi_z^{RF}(\mu|\cdot)$

C.2 Examples demonstrating Impossibility of Regret-Free Stable Mechanisms

We now demonstrate that standard mechanisms can fail spectacularly in learning market-clearing cutoffs and alleviating the costs associated with information acquisition. Intuitively, in choice-based mechanisms students need to know other students' choices in order to determine their optimal inspection strategy, and so in general the student who performs the 'first' inspection will incur additional inspections costs. Standard Deferred Acceptance mechanisms, which are played as one-shot games where students submit their full preference lists, perform especially poorly, as students are given almost no information about their choices before deciding on their inspection strategy. While in some settings regret can be eliminated by allowing for multi-round mechanisms, we prove the stronger result that for general economies even multiple-round mechanisms must either incur regret, or create an information deadlock, where every student waits for others to acquire information first.

C.2.1 Direct One-Shot Mechanisms

To demonstrate the issues in computing regret-free stable matchings, let us first consider the case where all students are willing to inspect any school as long as it is in their budget set. We may view this as a setting where the costs affect *which schools* students are willing to inspect, but not the *order* in which they are willing to inspect them. It is clear that the standard implementation of Deferred Acceptance as a one-shot game will not be regret-free even for such students, as students' budget sets will depend on the preferences of other students, and so students who have low priority at the schools they prefer are likely to incur regret. We illustrate this in the following example.

Example C.1. Consider a discrete economy $E = (\mathcal{C}, \mathcal{S}, q)$ with n students and n schools each with capacity $q_i = 1$. Suppose that school priorities are perfectly aligned, i.e. $r_i^s = r_j^s$ for all $s \in \mathcal{S}$, $i, j \in \mathcal{C}$, and students have random preferences and are willing to incur the cost to attend any school. Such demand can be rationalized e.g. by the priors $F_i^s(x) = 0$ for all $x \in [0, 1)$, $F_i^s(x) = \frac{1}{4}$ for all $x \in [1, 2)$, $F_i^s(x) = 1 - \frac{1}{2^k}$ for all $k \geq 1$ and $x \in [2^k, 2^{k+1})$ and costs $c_i^s = 1$ for all $s \in \mathcal{S}$.

In any one-shot choice-based mechanism, a student s will have no regret only if she chooses to

examine precisely the set of all schools not selected by higher-ranked students. This is because a student is willing to incur the cost to examine any school if and only if it is in her budget set. As student preferences are random, the probability that every student other than the highest-ranked student regrets her inspections is at least $\prod_i \left(1 - \frac{1}{\binom{n}{i}}\right) \geq \left(\frac{n-1}{n}\right)^{n-1} \rightarrow \frac{1}{e}$ all $n \rightarrow 1$. The example can also be modified so that with probability $\rightarrow \frac{1}{e}$ a proportion $\rightarrow 1$ of students incur unbounded regret.¹

This example demonstrates that single-shot choice-based mechanisms cannot hope to find regret-free stable matchings, even in settings where students are willing to incur the costs of searching any number of schools, due to their inability to coordinate the students' search.

C.2.2 Impossibility of Student-Optimal Regret-Free Stable Mechanisms

In this section we provide an example demonstrating that even in settings where it is possible to implement a regret-free stable choice-based mechanism, it may be impossible to verify that a matching is student-optimal without incurring regret.

Example C.2. Consider an economy E with two schools $\mathcal{C} = \{1, 2\}$ with capacities $q_1 = q_2 = 1$ and 2 students $S = \{x, y\}$.²

Suppose that school priorities are given by

$$\text{priority at 1 : } r_1^y > r_1^x$$

$$\text{priority at 2 : } r_2^x > r_2^y.$$

Suppose also that student values at each school have discrete distribution $\mathbb{P}(v_i^s = 1) = \mathbb{P}(v_i^s = 2) = \frac{1}{4}$, $\mathbb{P}(v_i^s = 2^k) = \frac{1}{2^k}$ for all $k > 1$ and $\mathbb{P}(v_i^s = x) = 0$ for all $x \notin \left\{\frac{1}{2^k}\right\}_{k \in \mathbb{N}}$, i.e. with priors $F_i^s(x) = 0$ for all $x \in [0, 1)$, $F_i^s(x) = \frac{1}{4}$ for all $x \in [1, 2)$, $F_i^s(x) = 1 - \frac{1}{2^k}$ for all $k \geq 1$ and $x \in [2^k, 2^{k+1})$, and

¹For each bound K the example can be modified so that with probability $\rightarrow \frac{1}{e}$ all $n \rightarrow 1$ students other than the top priority student incur regret at least K times their utility.

²Strictly speaking, as we assumed that there are more students than seats, the economy should have three students $S = \{x, y, d\}$ where d is a dummy student who has lower priority at every school than the students in $\{x, y\}$ and who has arbitrary preferences. For simplicity we omit these students in the description of the economy; however note that the proof applies as written to both economies.

that student costs for inspection are given by $c_1^x = c_2^y = 1$ and $c_2^x = c_1^y = 2$. As $\mathbb{E}[(v_i^s - \underline{v})] = \infty$ for all s, i and $\underline{v} \in \mathbb{R}$ it follows that both students' optimal strategies are to inspect all the schools that are available to them.³

Note that the matching $\mu = \mu^{school}$ defined by $(\mu(x), \mu(y)) = (2, 1)$ is always regret-free stable, and is the school-optimal regret-free stable matching. Let $\mu' = \mu^{student}$ be defined by $(\mu'(x), \mu'(y)) = (1, 2)$. We will consider two separate events. Let X denote the event that $v_1^x = v_2^y = 2$ and $v_2^x = v_1^y = 4$. Let X' denote the event that $v_1^x, v_2^y > 4$ and $v_2^x = v_1^y(\omega) = 4$. Note that μ^{school} is the student-optimal regret-free stable matching subject to event X , as both x and y obtain their highest valued schools, and that $\mu^{student}$ is the student-optimal regret-free stable matching subject to event X' , as both x and y again obtain their highest valued schools.

Notice that events X and X' are mutually exclusive, and that $\mathbb{P}(X) = \mathbb{P}(X') = \left(\frac{1}{4}\right)^4 > 0$. Furthermore, $v_2^x = v_1^y(\omega) = 4$ in either event. Thus, conditional on one of the events X or X' occurring, the only way to distinguish which event occurred is for student x to inspect school 1 or student y to inspect school 2.

We now first demonstrate why the existence of such X and X' shows that we cannot verify student-optimality in a regret-free manner. Note that if $\mu = \mu^{school}$ is the student-optimal regret-free stable matching then each school is assigned their top choice student, and so based on school preferences alone there are no blocking pairs and the corresponding student budget sets are $B^\mu(x) = \{2\}$, $B^\mu(y) = \{1\}$. Hence under $\chi^{RF}(\mu|\cdot)$ student x only inspects school 2, and student y only inspects school 1. However, if $\mu' = \mu^{student}$ is the student-optimal regret-free stable matching then under $\chi^{RF}(\mu'|\cdot)$ both students inspect both schools. Thus, since it is impossible to distinguish between events X and X' without requiring either student x to inspect school 1 or student y to inspect school 2, one of these inspections must occur in the event $X \vee X'$ in order to determine the student-optimal regret-free stable matching, which incurs regret under event X . In other words, it is impossible to verify that μ^{school} is the student-optimal regret-free stable matching without incurring regret. Since X has positive probability, we conclude that it is impossible to verify that

³It is simple to extend this example so that $v_i^s(\cdot)$ is continuous random variable with continuous density by smoothing the density for 2^k over the interval $[2^{k-1}, 2^k]$.

the student-optimal regret-free stable matching is student-optimal without incurring regret with positive probability.

C.3 Estimating Regret-Free Stable Cutoffs

C.3.1 Continuity and Convergence of Market-Clearing Cutoffs

We first define a metric on the space of economies and on the space of stable matchings. Fix a set of schools \mathcal{C} and a set of students \mathcal{S} . We say that a sequence of continuum economies $\mathcal{E}^k = (\eta^k, q^k)$ converges to the continuum economy $\mathcal{E} = (\eta, q)$ if η^k converges in the weak sense to η , and $q^k \rightarrow q$. We define the distance between stable matchings to be the distance between their associated cutoffs, $d(\mu, \mu') = \max_{P, P': \mathcal{M}(P)=\mu, \mathcal{M}(P')=\mu'} \|P - P'\|$.⁴ Given a finite economy $E = (\mathcal{C}, S, q)$ define the continuum economy $\Phi(E) = (\mathcal{C}, \mathcal{S}, \eta, q)$ by taking the distribution η defined by

$$\eta \left(\left\{ s \in \mathcal{S} \mid \theta^s = \theta^t, v_i^s \in \left\{ v_i^t(\omega) \mid \omega \in X \right\} \right\} \right) = \frac{1}{|\mathcal{S}|} p(X) \quad \forall t \in S, X \subseteq \Omega.$$

We may think of this as first taking the empirical distribution $\sum_{t \in S} \frac{1}{|S|} \delta_t$ and then changing the point distribution δ_t for student t to mirror the possible distribution of values v^t . We say that a sequence of finite economies E^k converges to the continuum economy \mathcal{E} if the embeddings $\Phi(E^k)$ converge to \mathcal{E} .

Theorem C.1. *Suppose the continuum economy \mathcal{E} admits a unique regret-free stable matching μ . Then the regret-free stable matching correspondence mapping economies to regret-free stable matchings is continuous at \mathcal{E} within the set of continuum economies.*

Proof. The theorem follows from the analogous result in Azevedo and Leshno (2016) as well as observing that the set of regular measures is open. \square

Theorem C.2. *Suppose the continuum economy \mathcal{E} admits a unique regret-free stable matching μ , and has a C^1 demand function that is non-singular at the market-clearing cutoffs (i.e. $\partial D(P^*)$)*

⁴Note that if \mathcal{E} is the embedding of a finite economy then there are many cutoffs that give the same matching μ .

non-singular). Let $E^k = (\eta^k, qk)$ be a randomly drawn finite economy, with k students drawn independently according to η , and let P^k be a market-clearing cutoff of E^k . Then

$$\sqrt{k} \cdot (P^k - P^*) \xrightarrow{d} \mathcal{N} \left(0, \partial D(P^*)^{-1} \cdot \Sigma^q \cdot \partial D(P^*)^{-T} \right),$$

where $\mathcal{N}(\cdot|\cdot)$ denotes a C -dimensional normal distribution with given mean and covariance matrix, and

$$\Sigma_{ij}^q = \begin{cases} -q_i q_j & \text{if } i \neq j \\ q_i (1 - q_i) & \text{if } i = j. \end{cases}$$

Theorem C.2 shows that the estimated cutoffs P^k are normally distributed around P^* , and follows directly from the analogous result in Azevedo and Leshno (2016). Another interpretation is that given and underlying population η and cutoffs P^* , if demand is given by sampling k students from η then the resulting market-clearing cutoffs P^k will be normally distributed around P^* .

C.3.2 Omitted Proofs for Section 5.5

C.3.2.1 Example 5.2

We show that

$$\sqrt{k} (\hat{q}^k - \bar{q}) \xrightarrow{d} \mathcal{N} \left(0, \frac{(\alpha + 1)}{\alpha} \Sigma_{\bar{q}} \right).$$

Let X be a random variable that gives a student randomly drawn according to η with corresponding demand D^X , and with probability $\frac{\alpha}{\alpha+1}$ and $\frac{1}{\alpha+1}$ assigns them to be a 'past' student and 'present' student respectively. Let $m(X, P) = \mathbf{1}\{X \text{ is 'past'}\} (D^X(P) - \bar{q})$ and let $g(X, q, P) = \mathbf{1}\{X \text{ is 'present'}\} (D^X(P) - q)$. Note that as $|qk] - qk \leq \frac{1}{k}$ it follows that $\sqrt{k}\hat{P}$ converges in distribution to $\sqrt{k}\hat{P}'$ where \hat{P}' satisfies $\hat{m}(X, \hat{P}') \stackrel{def}{=} \sum_{i=1}^{(\alpha+1)k} \frac{\alpha}{\alpha+1} \left(\frac{D^{(X_i)}(\hat{P}')}{k} - \bar{q} \right) = 0$. Note also that similarly $\sqrt{k}\hat{q}^k$ converges in distribution to $\sqrt{k}\hat{q}^k$ satisfying $\hat{g}(X, \hat{q}^k, P^*) \stackrel{def}{=} \sum_{i=1}^{(\alpha+1)k} \frac{1}{\alpha+1} \left(\frac{D^{(X_i)}(P^*)}{k} - \hat{q}^k \right) = 0$. Hence $\sqrt{k} (\hat{q}^k - \bar{q}) \xrightarrow{d} \mathcal{N}(0, V)$, where

$$V = (1 + \alpha) \text{var} \left(g(X, \bar{q}, P^*) - \frac{1}{\alpha} m(X, P^*) \right).$$

Since $\text{var}(g(\cdot)) = \frac{1}{1+\alpha} \text{var}(D^X(P^*))$, $\text{var}(m(\cdot)) = \frac{\alpha}{1+\alpha} \text{var}(D^X(P^*))$ and $\text{cov}(g(\cdot), m(\cdot)) = 0$, this is equal to $\text{var}(D^X(P^*)) + \frac{1}{\alpha} \text{var}(D^X(P^*)) = \frac{(1+\alpha)}{\alpha} \Sigma^{\bar{q}}$ as required.

C.3.2.2 Example 5.3

We first show that the outcome (μ, χ) after both rounds is regret-free stable with respect to realized demand \hat{q}^k . It suffices to show that for all students $s \in S^f(E)$ with free market information it follows that $\mu(s) = D^s(\hat{P})$. Now since $s \in S^f(E)$ it holds that $\forall i \ r_i^s \notin [1 - \sum_j q_j, 1 - q_i)$, and her first-round budget set is $B^s = \{i \mid r_i^s \geq 1 - q_i\}$. Moreover, if $i \in B^s$ then if $\mu(i) = q_i$ it follows that $\hat{P}_i < r_i^s$, and so $i \in B^s(\hat{P})$. Finally, if $i \in B^s(\hat{P})$ then $1 - r_i^s \leq \sum_j q_j$, as all students find all schools acceptable and so if $r_i^{s'} \geq r_i^s$ then student s' is assigned to some school. In other words, $B^s \subseteq B^s(\hat{P}) \subseteq B^s$, and so $\mu(s) = D^s(B^s) = D^s(B^s(\hat{P})) = D^s(\hat{P})$.

We now show that

$$\sqrt{k}(\hat{q}^k - \bar{q}) \xrightarrow{d} \mathcal{N}\left(0, \Sigma^{\bar{q}} + 2\left(\frac{1}{\alpha}A + I\right)\Sigma^{q^f}A^T\right).$$

Let X be a random variable that gives a student randomly drawn according to η with corresponding demand D^X . Let the first-round cutoffs be $P_i^f = 1 - q_i$, let $m(X, \Gamma) = \mathbf{1}_{\{X \in S^f(E)\}} D^X(P^f | \eta(\Gamma)) - q^f$ and let $g(X, q, \Gamma) = m(X, \Gamma) + \mathbf{1}_{\{X \notin S^f(E)\}} D^X(P^*(\Gamma) | \eta(\Gamma^*)) - (q - q^f)$. Note that the estimated $\hat{\Gamma}$ satisfies $\hat{m}(X, \hat{\Gamma}) \stackrel{def}{=} \sum_{i=1}^k \alpha \frac{(D^{X_i}(P^f | \eta(\hat{\Gamma})) | X \in S^f(E))}{k} - q^f = 0$, and that the estimate demand \hat{q}^k of all students satisfies

$$\hat{g}(X, \hat{q}^k, \hat{\Gamma}) \stackrel{def}{=} \sum_{i=1}^k \alpha \frac{(D^{X_i}(P^f | \eta(\hat{\Gamma})) | X \in S^f(E))}{k} + (1 - \alpha) \frac{(D^{X_i}(P^*(\hat{\Gamma}) | \eta(\Gamma^*)) | X \notin S^f(E))}{k} - \hat{q}^k = 0$$

. Hence $\sqrt{k}(\hat{q}^k - \bar{q}) \xrightarrow{d} \mathcal{N}(0, V)$, where

$$V = \text{var}\left(g(X, \bar{q}, \Gamma^*) - \left(I + \frac{1}{\alpha}A\right)m(X, \Gamma^*)\right)$$

and $A = \mathbb{E}[\nabla_{\Gamma}(g(X, \bar{q}, \Gamma^*) - m(X, \Gamma^*))] \mathbb{E}[\nabla_{\Gamma}m(X, \Gamma^*)]^{-1}$.

Note that $A = \nabla_{\Gamma}D(P^*(\Gamma) | \eta(\Gamma^*))|_{\Gamma=\Gamma^*} \left(\nabla_{\Gamma}D^f(\Gamma^*)\right)^{-1}$. Moreover $\text{cov}(g(\cdot), m(\cdot)) = \text{var}(m(\cdot))$

and so

$$V = \text{var}(g(\cdot)) + \frac{1}{\alpha} \text{var}(m(\cdot)) \left(I + \frac{1}{\alpha} A \right)^T.$$

Since $\text{var}(g(\cdot)) = \text{var}(D^X(\Gamma^*)) = \Sigma^{\bar{q}}$, $\text{var}(m(\cdot)) = 2\alpha \text{var}(D^f(\Gamma^*)) = 2\alpha \Sigma^{q^f}$, this is equal to $\Sigma^{\bar{q}} + 2A\Sigma^{q^f} \left(I + \frac{1}{\alpha} A \right)^T$ as required.