

Multi-Objective Optimization for Public Policy

by

Theodore P. Papalexopoulos

B.S. Physics and Applied Mathematics, Yale University (2015)

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author
Sloan School of Management
April 29, 2022

Certified by.....
Dimitris Bertsimas
Boeing Leaders for Global Operations Professor of Management
Thesis Supervisor

Certified by.....
Nikolaos Trichakis
Associate Professor of Operations Management
Thesis Supervisor

Accepted by
Georgia Perakis
William F. Pounds Professor of Management Science
Co-director, Operations Research Center

Multi-Objective Optimization for Public Policy

by

Theodore P. Papalexopoulos

Submitted to the Sloan School of Management
on April 29, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

Operations research has a storied history of tackling complex problems in public policy, ranging from vaccine distribution to the efficient design of public utility markets. The advent of “big data” analytics, machine learning, and scalable optimization has only expanded the field’s impact, unlocking new research directions and application areas. What makes public policy a challenging domain is a combination of three factors: (i) policymakers must balance multiple objectives that often exist in tension, e.g., tradeoffs in efficiency and fairness; (ii) there are many stakeholders, with often disparate value judgments on how to best balance said objectives; and (iii) those stakeholders may not be technically fluent in analytics.

This thesis develops multi-objective optimization methodologies to support policymakers in designing more efficient, fair, and inclusive policies. We apply our techniques to a range of problems in transplantation policy and public education. A core theme of our work is the need for interpretable decision-support tools, e.g., interactive applications and tradeoff curves, which are crucial in translating abstract policy tradeoffs into actionable insights. Our goal is to provide stakeholders, even those without technical expertise, with an understanding of the range of achievable policy outcomes, so that they can more effectively engage in the policymaking process. We emphasize applications of our work to real-world problems, including an extensive collaboration with the United Network for Organ Sharing (UNOS) to help develop a new national lung allocation policy, which is slated for implementation in 2023.

Chapter 2 addresses a long-standing debate about geographic equity in organ allocation, by using multi-objective optimization to compare efficiency/fairness tradeoffs under different geographic distribution schemes. Chapter 3 introduces a novel optimization-based framework for “ethics-by-design” in scarce resource allocation, aiming to combine data modeling, shareholder input, and ethical theory into a unified approach for policy development in this area. Chapter 4 details our collaboration with UNOS policymakers to apply this framework towards the design of a new national lung allocation policy. Finally, Chapter 5 presents an empirical analysis of school assignment mechanisms for public school districts, investigating tradeoffs between satisfying student preferences and minimizing bus transportation costs.

Thesis Supervisor: Dimitris Bertsimas

Title: Boeing Leaders for Global Operations Professor of Management

Thesis Supervisor: Nikolaos Trichakis

Title: Associate Professor of Operations Management

Acknowledgments

“I may not have gone where I intended to go, but I think I have ended up where I needed to be.” — Douglas Adams

As my time at MIT draws to a close, I would like to pause and reflect on the many people who have contributed, in ways big and small, to this remarkable journey.

First and foremost, I would like to thank my advisors Dimitris Bertsimas and Nikos Trichakis for their unflagging support over the years. My very first conversation with Dimitris is one of those pivotal moments that I will always look back to for inspiration, a crystallization of the core values I aspire to combine in my own working life: academic excellence, real-world impact, and optimism in the face of adversity. Nikos’s mentorship, every bit the epitome of those same values, has been just as formative; his patience and keen judgment have been an invaluable help through the many twists and turns of this Ph.D. Thank you both for believing in me since the day we met and for caring so deeply about my personal and professional development.

I am privileged to count Mohammad Fazel-Zarandi and Jonas Jonasson as committee members on this thesis. Their expertise in public policy analytics, with an emphasis on practical impact, has served as an inspiration for much of the work in this dissertation. I’ve also had the pleasure of teaching alongside Mohammad on three occasions, and in doing so learnt a tremendous amount about engaging and inspiring students. I’m very fortunate to count him as both a mentor and a friend.

Over the years, I have had many positive and impactful interactions with ORC faculty and staff. I would like to thank Prof. Negin Golrezaei for serving on my general exam committee, and Amr Farahat for guiding me through my first teaching experience at Sloan. I am also particularly thankful to Georgia Perakis who has supported and cared for my well-being since I first arrived at MIT. A huge thank you goes to Laura Rose and Andrew Carvalho, without whom the ORC would not function, for all of their support over the years.

There are many collaborators who've contributed to the work presented in this thesis. It was a true privilege to work alongside Darren Stewart and Rebecca Goff at the United Network for Organ Sharing. Their constant encouragement and feedback were instrumental in getting our lung transplantation work in front of policymakers, and their single-minded dedication to improving patient welfare is inspiring beyond words. Yuchen Wang and Drs. Ryutaro Hirose and Parsia Vagefi were supportive and engaging collaborators in my first foray into the world of transplantation policy. I also have tremendous respect for Glenn Cohen and James Alcorn for their open-mindedness and clarity of thought, as well as their willingness to engage with the many complexities of applying analytics to public policy.

My journey as a researcher began at the magical Yale Exoplanet Lab under the supervision of Debra Fischer, whose passion for research and incredible work ethic serve as an inspiration to this day. I was introduced to the wide, wide world of operations research during my time at Analytics Operations Engineering; I would not have pursued this Ph.D. if it weren't for the amazing mentorship I received there from Leon Hsu, Gerry Feigin, Kermit Threatte, and Jim Schor. Finally, my summer collaborating with Christian Tjandraamadja, David Belanger, Ross Anderson, and Juan-Pablo Vielma at Google was a truly wonderful experience. Their open-mindedness, problem-solving prowess, and constant availability have made me a better thinker and a better researcher.

I'm incredibly grateful for all the friends and memories I've made at the ORC. I was lucky to serve as an INFORMS officer alongside three brilliant community builders (and chaotic co-conspirators) in Holly Wiberg, Jess Zhu, and Patricio Foncea Araneda. A tough first few months would have been a lot tougher if it weren't for Julia Yan's kindness and wisdom. Arthur Delarue has been a constant source of optimism, stimulating discussion, and baked goods for as long as I have known him. I would not have made it through any of this without Leann Thayaparan; thank you for listening to me and advising me on anything and everything. And I will always cherish the friends I made as part of the bestest ORC cohort: Peter Cohen, Ryan Cory-Wright, Andreea Georgescu, Lea Kapelevich, Roya Moussapour, and Andy Zheng. Finally,

thank you to the rest of the ORC community, too many friends to count, for making E-40 a place I was excited to arrive at every morning.

I would be remiss if I did not give a special shout out to the MIT Greek Mafia. I am grateful to Yannis Spantidakis, who's always *exo kardia* and the life of the party, and has indulged me in many long conversations and even longer card games; Vassilis Digalakis, who keeps me honest and always pushes me to do bigger and better things; Agni Orfanoudaki whose warmth and vitality is rivaled only by her countless accomplishments; Ilias Zadik, probability guru and mentor extraordinaire; and Alvaro Fernandez Galiana, whose dislike of feta is somehow not enough to disqualify him from being an honorary Greek.

Also the many friends outside of MIT who have supported me tirelessly from afar. John has been there since day one; Nikos and Irini back home; the New York trash crew of Cyril, Maria, David, Aspen, Claire, Julia, and Jordana; and my work family, Megan, Bob, Rainy, Haley, and Matt.

Last but not least, I would like to thank my family. My grandmothers Marina and Marileni are two of the strongest women I know and have been an endless source of comfort and support. My grandfathers Theodoros and Dimitris are no longer with us, but still role models I look up to. My aunts Alexandra and Lena, and cousins Marileni, Leticia, Danae, Anna, Alexandros, and Vassilis have cheered me on from both sides of the Atlantic. And of course Jason Papastavrou, who was the first math professor I ever had back on those childhood walks through Central Park.

I am eternally grateful to my brother Jason. Your apartment is a home away from home and the place I go to destress. Through countless conversations, movie marathons, and fancy meals, your calming presence, relentless optimism, and sage advice have kept me going for 26 years and counting. I am also grateful to my not-so-baby sister Marina, who is wise and empathetic beyond her years and always willing to offer advice.

Finally, my parents Amerimni and Dimitris to whom this thesis is dedicated. You've been supporters, advisors, mentors, and so much more – none of this would have been possible without you. I love you both so much.

To my parents, Amerimni and Dimitris

Contents

1	Introduction	13
1.1	Motivation	13
1.2	Outline and main contributions	15
2	Balancing efficiency and fairness in liver transplant distribution	21
2.1	Introduction	21
2.2	Background	23
2.2.1	Liver allocation in the US	23
2.2.2	Organ distribution concepts	24
2.3	Placement efficiency and geographic equity	26
2.3.1	Data and methods	26
2.3.2	Tradeoff analyses	28
2.3.3	Implications for allocation policy	34
2.4	Redistricting organ allocation	37
2.4.1	Background	37
2.4.2	Overview of the algorithm	38
2.4.3	Mortality oracle model	39
2.4.4	Distance oracle model	41
2.4.5	Local district search	42
2.5	Appendix	44
2.5.1	Distribution concepts in detail	44
2.5.2	Simulation calibration	46
2.5.3	Parameter selection	48

2.5.4	Additional tradeoffs	49
3	Ethics-by-design: efficient, fair and inclusive resource allocation	51
3.1	Introduction	51
3.2	Redesigning organ allocation in the US	53
3.2.1	Policy development and the OPTN	53
3.2.2	The continuous distribution framework	54
3.3	Shifting to ethics-by-design	56
3.3.1	Optimization methodology	56
3.3.2	Implications for policy development	59
3.4	Discussion	63
4	Reshaping US lung allocation through multi-objective optimization	67
4.1	Introduction	67
4.2	Allocation policy in the US	69
4.2.1	Background	69
4.2.2	Continuous distribution	70
4.2.3	Challenges in policy design	72
4.3	Multi-objective optimization framework	74
4.3.1	Mathematical model	75
4.3.2	Methodology	76
4.3.3	Implications for policy design	79
4.4	Reshaping national lung allocation policy	81
4.4.1	Modeling implementation	83
4.4.2	Tradeoffs for placement efficiency	84
4.4.3	Ensuring pediatric priority	86
4.5	Discussion	87
4.6	Glossary	90
4.7	Appendix	94
4.7.1	Robustness to surrogate model approximation error	94
4.7.2	Reformulation of infeasible surrogate optimization	94

4.7.3	Surrogate modeling	96
4.7.4	Simulated outcomes	101
4.7.5	Additional tradeoffs	103
5	Balancing student preferences and transport cost in school choice	107
5.1	Introduction	107
5.2	The school choice problem	110
5.2.1	Mathematical model	110
5.2.2	Literature review	112
5.3	Optimization methodology	114
5.3.1	An MIO formulation for stability	115
5.3.2	Custom pre-solve algorithm	118
5.4	The effect of stability on transport distance	125
5.4.1	Experimental design	125
5.4.2	Results	126
5.5	Conclusions and future work	131
5.6	Appendix	133
5.6.1	Data and methods	133
6	Conclusions	139

Chapter 1

Introduction

1.1 Motivation

Operations research (OR) has a storied history of tackling complex problems in public policy. The so-called diet problem, of historical significance in defining poverty levels in the United States back in the 1960s [29], is still taught in every introductory linear optimization course. So is the school busing problem, which played an integral role in efforts to desegregate public schools in the 1970s [52].

In a thought-provoking plenary delivered to the 2002 International Conference on OR, Jonathan Caulkins suggested that OR’s biggest impact in public policy has come in areas where the “physics” of the system are complex and central (e.g., energy and aviation), or the issues pertain to tactical management or are at the implementation level (e.g., supply chain management and routing). Relatively speaking, he argued, OR’s role has been smaller “in the analysis of strategic issues [...] in general domains” – domains like welfare policy, education policy, or equal employment law [23].

This trend has arguably changed in recent decades, as the appetite for “big data” analytics, machine learning, and scalable optimization has exploded in a wide range of public policy domains. OR techniques have been applied to crowdsource volunteers for NGOs [58], improve inmate assignment mechanisms in penitentiaries [88], design auction markets for farmers in developing countries [50], and evaluate immigration enforcement policy [10]. During the recent COVID-19 pandemic, OR practitioners

helped develop border testing protocols [6], evaluate government intervention policies [51], and propose vaccine allocation systems [8].

What makes the application of OR to such domains challenging, and any impact all the more impressive, is a combination of three key factors. First, policymakers need to balance multiple objectives that often exist in tension, e.g., tradeoffs in efficiency vs. fairness. Second, there are typically many stakeholders, with often disparate value judgments on how to best balance said objectives. And third, these stakeholders are not always technically fluent in analytics.

This thesis develops multi-objective optimization methodologies to support policymakers in designing more efficient, fair, and inclusive policies. We apply our techniques to problems in transplantation policy and public education. A core theme of our work is the need for interpretable decision-support tools, e.g., interactive applications and visualized tradeoff curves, which are crucial in translating abstract policy tradeoffs into actionable insights. The goal is to provide policymakers and stakeholders, even those without technical expertise, with an understanding of the range of achievable policy outcomes, and encourage evidence-driven debate on tradeoffs. We emphasize applications of our work to real-world problems, including an extensive collaboration with the United Network for Organ Sharing (UNOS) to help them develop a new national lung allocation policy, slated for implementation in 2023.

Chapter 2 addresses a long-standing debate about geographic equity in organ allocation, by using multi-objective optimization to compare efficiency/fairness tradeoffs under different geographic distribution schemes. Chapter 3 introduces a novel optimization-based framework for “ethics-by-design” in scarce resource allocation, aiming to combine data modeling, shareholder input, and ethical theory into a unified approach for policy development in this area. Chapter 4 details our collaboration with UNOS policymakers to apply this framework towards the design of a new national lung allocation policy. Finally, Chapter 5 presents an empirical analysis of school assignment mechanisms for public school districts, investigating tradeoffs between satisfying student preferences and minimizing bus transportation costs.

1.2 Outline and main contributions

The summary and main contributions of this thesis follow, listed by chapter.

Chapter 2: Balancing efficiency and fairness in liver transplant distribution

Geographic disparity in access to organs and, by extension, the role of geography in allocation priority, have long been points of contention within the transplantation community [35, 60, 103]. At the core of the debate lies a fundamental tension between efficiency and fairness: on the one hand, allocation policies that prioritize local candidates reduce system-wide transportation burden and can improve post-transplant outcomes by reducing the time organs spend on ice. On the other hand, such “local” policies may result in higher waitlist mortality rates, as organs are not allocated to the highest-risk or most disadvantaged patients if they are far from the donor.

In this chapter, we use optimization-based tradeoff analysis to assess different organ distribution policy concepts. We consider three different distribution schemes that were proposed by the the Organ Procurement & Transplantation Network (OPTN) Ad Hoc Committee on Geography: (1) district-based priority with mathematically optimized boundaries (2) circle-based regions using a fixed distance from the donor hospital, and (3) a novel, boundaryless continuous distribution model. We use counterfactual simulation and surrogate optimization to generate efficiency-fairness trade-off curves for each distribution concept in the context of liver transplantation. We publish our results as an interactive online application (<https://livervis.github.io/>).

Our analysis shows that a Continuous Distribution (CD) concept allows both for the greatest reduction in patient deaths, and the most equitable geographic distribution across comparable organ transportation burden among all others considered. In addition, when applied with an Optimized Prediction of Mortality allocation scheme to evaluate candidates’ medical urgency, continuous distribution allows for a significant reduction in the number of deaths compared to current policy – on the order of 500 lives saved annually in simulation.

In December of 2018, the OPTN Board of Directors issued a directive stating that future policy development for all solid organs would focus on the continuous distribution framework. As a result of this decision, we began a collaboration with researchers at United Network for Organ Sharing (UNOS) that culminated in the work of Chapters 3 and 4.

This chapter also includes an algorithmic contribution to the combinatorial problem of geographic districting. To analyze the district-based distribution concept, we develop a novel multi-objective optimization algorithm for creating of organ-sharing regions, using specialized machine learning models to approximate a simulation-based objective function. The general approach may be of independent interest in other districting applications that require optimization of multiple, black-box objectives, e.g., political districting or land management.

The work in this chapter appeared in *Transplantation* [12].

Chapter 3: Ethics-by-design: efficient, fair and inclusive resource allocation

As has been made clear by current controversies regarding criteria for allocating COVID-19 vaccines in U.S. states and across countries, the distribution of crucial medical goods and services in conditions of scarcity is among the most important, albeit contested, areas for public policy development [42, 43, 66]. Scarce resource allocation, particularly in healthcare settings, is a challenging problem with multiple, often-conflicting objectives and far-reaching ethical and practical implications.

In this chapter, we introduce a novel optimization-based framework for “ethics-by-design” in resource allocation. Our framework seeks to combine data modeling, shareholder input, and ethical theory into a single process for generating sound allocation systems. We use the design of organ allocation policy as a motivating case study, which illustrates the key challenges in the area, namely: (i) the need to accurately predict policy outcomes across a range of utility, efficiency, and fairness dimensions; (ii) the tensions inherent in achieving different policy objectives under scarcity; and

(iii) the need to engage a diverse set of stakeholders, with often conflicting priorities, in evidence-driven debate on tradeoffs.

At a high level, our framework prescribes the use of machine learning and surrogate optimization to approximate the efficient frontier of policy outcomes across all utility, efficiency, and fairness dimensions of interest. Its concrete manifestation is an interactive optimization tool wherein one can specify a set of desired policy outcomes, expressed in the objective and constraints of an optimization problem, and our method will produce a conforming policy in near real-time. This enables policymakers and stakeholders, even those without technical expertise, to quickly iterate on different policy scenarios and refine their value judgments on relevant tradeoffs as they engage in the policy-making process.

We extensively discuss lessons from applying our framework in practice (described in more detail in Chapter 4). First and foremost among them is the need to clearly communicate the limitations of statistical modeling to stakeholders, and allow subject-matter expertise to influence conclusions drawn from the analysis. Second, while the right analytical tools can make tradeoffs more accessible, they do not obviate the need for structured, consensus-seeking processes to engage stakeholders and encourage evidence-driven debate. Finally, given cultural and institutional hurdles, as well as the relative unfamiliarity with machine learning and optimization among policymakers, the adoption of advanced analytical tools like the one we propose requires thoughtful attempts at scientific communication that meets stakeholders “where they live,” rather than a one-size-fits-all strategy.

The work in this chapter has been accepted for publication in the *Journal of Law and the Biosciences* [78].

Chapter 4: Reshaping US lung allocation through multi objective optimization

Following a landmark 2018 decision by the Organ Procurement & Transplantation Network (OPTN) Board of Directors, the OPTN began the process of migrating all of its allocation policies to a continuous distribution (CD) framework [69]. Within this framework, transplant candidates on the national waitlist are ranked not by classification into distinct priority groups, but rather according to a scoring rule. The selection of key policy parameters, namely the score components and their associated weights in the formula, poses a challenging, multi-objective optimization problem with far-reaching ethical and practical implications.

In this chapter, we detail how we applied our ethics-by-design framework (Chapter 3) to the design of a continuous distribution policy for lung allocation, working closely with researchers at the United Network for Organ Sharing (UNOS) and the OPTN Lung Transplantation Committee to help develop their proposal. Using our interactive optimization tool, we explored many different policy options during the design phase, and presented key tradeoff analyses on organ placement efficiency and pediatric access to committee members to help guide their decision-making process. The committee’s official proposal used values for policy parameters that were identified as inflection points in our analysis, highlighting the importance of a global understanding of tradeoffs through optimization.

In December of 2021, the committee’s proposal was approved by the OPTN Board of Directors, and will become national policy starting in 2023, guiding how lungs are allocated in the US for years to come. Our simulation studies suggest that the new policy could reduce waitlist mortality by 21% compared to the status quo, averting 62 waitlist deaths per year, while also improving fairness and equity. Independent simulations performed by the Scientific Registry of Transplant Recipients (SRTR) place the estimated reduction of mortality even higher, at around 40%.

The work in this chapter has been submitted for publication [79].

Chapter 5: Tradeoffs between stability and transportation cost in school choice

School assignment is one of the most challenging problems that public school districts face, in that officials seek to balance student equity and preference considerations with significant logistical and operational costs. As many districts have moved away from neighbourhood districting and towards preference-based assignments, *stable* matching mechanisms have become increasingly popular. First proposed in the seminal work of Gale & Shapley [33], stability reflects a notion of procedural fairness in how matches are formed, by ensuring that there is no student-school pair where both prefer each other over their current assignments (known as a *blocking pair*).

Despite its natural intuition and elegance, stability alone does not suffice to address all of a school district’s objectives. In realistic settings, districts often face hard constraints, e.g., siblings that must be assigned to the same school or legally mandated diversity quotas, that are not easy to incorporate in a stable matching mechanism. The same is true of different policy objectives, e.g., minimizing racial disparities in access to top schools or minimizing bus transportation costs, since common stable matching algorithms do not explicitly model an objective function to select among multiple possible solutions.

In this chapter, we apply global optimization techniques to empirically study tradeoffs in stable matching mechanisms. We formulate stability as a mixed-integer optimization problem, and develop a custom pre-solve algorithm to help scale it. We apply our formulation to synthetic data based on the Boston Public Schools district, optimizing over the space of stable solutions to minimize a proxy of the district’s transportation cost. Our experiments suggest that stability is a particularly constraining property in this setting, and does not allow for significant optimization of the alternative objective. In fact, even in the largest matching instances we consider (with $\approx 16,000$ students and ≈ 100 schools), we find that stable solutions most often number in the single digits, and differ by no more than 0.13% in their objective values.

Motivated by this observation, we propose a natural way to relax stability by allowing a small number of blocking pairs. This enables us to perform tradeoff analyses that balance the level of stability in the matching solution against alternative objectives. In synthetic experiments, we observe a 12.7% average reduction in the distance-based objective when allowing only 1% of student-school pairs to be blocking, and 28.2% reduction for 5% of pairs. These results suggest that districts can, in a sense, get the best of both worlds, by balancing the procedural fairness of stability with alternative objectives and constraints.

Chapter 2

Balancing efficiency and fairness in liver transplant distribution

2.1 Introduction

The US organ allocation system, administered by the United Network for Organ Sharing (UNOS), has historically distributed organs based upon a “local-first” approach, whereby deceased-donor organs are offered first to patients in proximity to the donor location (i.e., within the Donor Service Area, DSA), followed by candidates beyond the DSA but within the donor’s OPTN Region, and finally to candidates residing outside of the donor’s Region. The result has been that a candidate’s access to organ transplant has varied significantly based upon their residence within a specific DSA and OPTN Region.

In an attempt to address this geographic disparity to organ access, on April 2nd, 1998 the “Final Rule” was issued by the Department of Health and Human Services (DHHS) stating that organ allocation should be designed and implemented so as to prioritize waitlisted candidates in order of decreasing medical urgency status, and so that a candidate’s location should not remain a major determinant in access to transplantation [27]. In the case of liver transplantation, the institution of allocation based on the Model for End-Stage Liver Disease (MELD) in 2002 fulfilled the former goal; however, the latter goal has been elusive as nearly two decades later it remains

clear that the geographic disparity in access to liver transplantation has persisted despite the “Final Rule” [3, 34].

On July 13th, 2018, a lawsuit was filed in the US District Court for the Southern District of New York. The plaintiffs – six individuals awaiting liver transplantation – set forth to sue the DHHS, Organ Procurement & Transplantation Network (OPTN), and UNOS, for an “illegal and inequitable liver allocation policy” that had been based on a history of “local-first” distribution [24]. Notably, this lawsuit came on the heels of a court-ordered change in the nationwide lung transplant distribution policy in 2017, which abandoned the use of DSAs as initial areas of distribution in favor of prioritizing transplant centers located within 250 nautical miles of the donor’s hospital [68]. Subsequently, the DHHS directed the UNOS Liver and Intestinal Transplantation Committee in July 2018 to develop a novel liver distribution proposal that eliminated both DSA and OPTN Region as a unit of organ distribution.

In this chapter, we seek to comprehensively analyze the implications of various liver distribution proposals that the OPTN considered in response to the DHHS directive. In particular, we use simulation, machine learning, and optimization to generate efficiency-fairness tradeoff curves for several distribution concepts falling under three frameworks proposed by the OPTN Ad Hoc Geography Committee: (a) fixed geographic areas based on the distance between the organ donor hospital and the transplant candidate’s hospital; (b) mathematical optimization of boundaries; and (c) a continuous distribution model that combines important clinical factors along with proximity to the donor location [70]. In so doing, we provide a visualized demonstration of fairness and efficiency which favors application of a continuous distribution scoring model for candidates awaiting liver transplantation.

2.2 Background

2.2.1 Liver allocation in the US

N.B. The following section serves as motivation and reflects the state of US liver allocation policy at the time of this work (2018). We note that the OPTN has since transitioned to prioritizing transplant centers within a fixed distance from the donor hospital for liver distribution, eliminating the role of both DSAs and OPTN Regions as geographic units in allocation [69].

Liver transplantation candidates are currently prioritized based on a numerical score, ranging between 6 and 40, which reflects their predicted 90-day mortality as determined by the Model for End-Stage Liver Disease (MELD). For distribution of a donor liver, candidates have traditionally been classified as local, regional, or national, depending on their location relative to the donor hospital. Specifically, there are 58 groupings of counties into Donor Service Areas (DSAs), and these DSAs are grouped into 11 broader OPTN Regions. Candidates are classified as local (respectively, regional) if their transplant program is in the same DSA (resp. Region) as the donor hospital. Broadly speaking, for bands of decreasing MELD scores, organs are sequentially offered first to local, then regional, and finally national candidates (with certain exceptions for severely ill patients; see Section 2.5.1). A candidate may opt to accept or decline an offered organ; offers are extended until the organ is accepted or enough time has elapsed for transplantation to be deemed unviable, in which case the organ is discarded.

The extent of geographic sharing of recovered organs – that is, between DSAs and Regions – thus lies at the center of a long-standing debate on efficiency vs. fairness in organ distribution. On the one hand, broader sharing increases access for high urgency candidates who are farther away (urgency-based equity), and balances supply and demand over larger geographic areas (geographic equity). On the other hand, increased organ transport distances shorten the viability window and may have detrimental effects due to increased organ ischemic time (placement efficiency).

Historical disparities in organ supply and demand, both overall and within geographic regions, represent a byproduct of medical and technological advances that have made organ transplantation the standard of care for the treatment of end-stage organ failure. Notably, the geographic areas of distribution, and their boundaries, were developed over three decades ago and were never intended nor designed to be used as areas of distribution for transplanted organs. Despite concerted efforts to address geographic disparity, there has been a persistent inability to achieve consensus from all stakeholders regarding a novel liver distribution scheme. The delicate balance between efficiency and fairness is paramount in the context of this discussion and has previously been used to analyze challenges in resource allocation through the use of tradeoffs with assigned objectives [9].

2.2.2 Organ distribution concepts

In this work, we analyze the efficiency and fairness characteristics of different organ distribution schemes as they relate to the aforementioned sharing tradeoff. We consider several different distribution concepts, broadly classified into three frameworks according to [70]: (a) Fixed Distance from Donor Hospital; (b) Mathematically Optimized Boundaries; and (c) Continuous Distribution. The majority of concepts follow the paradigm of classifying candidates as local/regional, with livers offered first locally, then regionally, and finally nationally, within MELD-score bands (as in current OPTN policy, certain exceptions are applied for severely ill patients). Where the concepts differ is the process by which candidates are classified as local/regional, and/or the MELD score bands used for distribution priority.

At a high level, each concept is associated with a set of parameters that control the level of geographic sharing, and can be varied to yield policies with different efficiency/fairness characteristics. There follows a high-level description of the concepts and their parameters; for more details see Section 2.5.1.

Fixed Distance from the Donor Hospital This framework prescribes using “concentric circles” around the donor hospital to classify candidates as local/regional [72]. In the Acuity Circles (AC) concept, three enlarging circles of fixed radii around the

donor hospital are used to classify candidates using MELD score bands that begin narrow for candidates with increased disease severity, but subsequently become wider at lower MELD thresholds. Similarly, Broader 2-Circle Distribution (B2C) uses three circles of fixed radii, but wider MELD bands throughout. Parameters considered for the AC and B2C concepts were the circles' radii and the MELD score bands' ranges.

Mathematically Optimized Boundaries This framework maintains the usage of distribution boundaries, similar to DSAs and/or OPTN Regions, but calls for their redesign so as to optimize organ distribution. The Optimized Districts (OD) concept is similar to existing policy, wherein existing DSAs are used; however, alternative DSA groupings into a certain number of regions are generated through optimization (see Section 2.4). Parameters considered were the number of regions and the possible groupings of DSAs into regions.

Continuous Distribution This framework seeks to eliminate any type of geographic boundaries, be it circles or regions, and instead accounts for distance between a candidate's transplant program and the donor hospital by incorporating it directly into a candidate's allocation score. In the Continuous Distribution (CD), all candidates nationally are prioritized according to a scoring formula that combines a medical urgency score (MELD) and a proximity score. In particular, the following allocation score is used to rank candidates:

$$(\text{MELD}) - \lambda \cdot (\text{Distance between candidate and donor hospitals})$$

where λ is the proportional factor at which the score decreases per extra unit of distance. Higher values of λ favor proximity over medical urgency. The only parameters considered was the tradeoff weight λ .

Reference Policies In addition to the above, certain baseline policies were selected to serve as references. These were the current (at the time of this work) OPTN 11-Region policy; the AC-style policy with radii of 150/250/500mm (AC 250/500) that was approved in December 2018 by the OPTN for future liver distribution [69]; and two other policies that the OPTN considered for adoption—an AC-style policy

with radii of 150/300/600nm (AC 300/600) and a B2C-style policy with radii of 150/250/500nm (B2C 250/500).

Medical urgency scores While geographic distribution is the focus of this work, it is not the only lever for improving efficiency and fairness in allocation. Previous work has shown that there is significant potential for impact in lives saved through more accurate and objective prioritization of candidate disease severity in liver allocation. A machine-learning based Optimized Prediction of Mortality (OPOM) model was recently developed, which predicts any adult candidate’s three-month waitlist mortality based on 28 risk factors [11]. OPOM has been established to provide significantly more accurate mortality prediction than MELD, reducing mortality by ≈ 400 lives annually in allocation simulation. Given these promising results, we extend our analysis of efficiency and fairness for different distribution concepts to the setting where candidates’ medical urgency is determined by OPOM rather than MELD.

2.3 Placement efficiency and geographic equity

2.3.1 Data and methods

This study used data from the Scientific Registry of Transplant Recipients (SRTR). The SRTR data system includes data on all donor, wait-listed candidates, and transplant recipients in the U.S., submitted by the members of the Organ Procurement & Transplantation Network (OPTN). The Health Resources and Services Administration (HRSA), U.S. Department of Health and Human Services provides oversight to the activities of the OPTN and SRTR contractors.

Evaluation of policies under each distribution concept was based on the Liver Simulated Allocation Model (LSAM, version 2014). LSAM is a discrete-event simulator developed by the SRTR to support studies into alternative allocation policies, and is extensively used by UNOS and the OPTN to evaluate new policy proposals. It uses historical waitlist and transplant data from 2007-2011 to simulate counterfactual allocation under a given prioritization scheme. Of note, LSAM includes a geolocation

model to estimate organ transport distances and times, as well as statistical models to predict a candidate’s acceptance of a given organ and their post-transplant outcomes if accepted (e.g., time to graft failure or relisting).

The benefit of simulation lies in its ability to track patient-level outcomes, and thus provide a holistic view of efficiency and fairness not constrained by any stylized model of allocation. Policies can be evaluated on a range of aggregate allocation outcomes, including:

- mortality rate (on the waitlist, or post-transplant).
- transplant rate, and percent of organs transplanted locally (in the same DSA).
- disparities in the above among patient subpopulations, e.g., by age, sex, race, disease diagnosis, blood type, etc.
- median MELD score at time of transplant (MMaT).
- standard deviation of MMaT across the 58 DSAs (σ MMaT).¹
- transportation and efficiency metrics, e.g., median average organ transport distance, travel time, or % of organs flown.

Because LSAM was unable to simulate all distribution concepts of interest to this study, we implemented an alternative simulator henceforth referred to as OrgSim. OrgSim was built to replicate all of LSAM’s functionalities, utilizing the same historical data and models, but is more computationally efficient and allows for more granular control over the simulated allocation policy. Section 2.5.2 provides calibration results for OrgSim, which is shown to match LSAM within 3.2% on all outcomes considered.

To create tradeoff curves, multiple policies under each distribution concept were generated by varying the associated parameter values, either over a grid of feasible settings or by using mathematical optimization. All generated policies were simulated

¹ σ MMaT is commonly used as a measure geographic equity in allocation. A higher value—that is, higher variance in DSAs’ MMaTs—indicates that certain DSAs are able to provide organs to candidates lower on the waitlist (after high MELD candidates have received offers) while other DSAs are less able to meet high-urgency demand.

over 20 iterations of OrgSim and those selected that lay on the efficient frontier of two primary objectives, namely overall mortality and average organ transport distance. The latter is a commonly used metric for quantifying the level of geographic sharing, as higher values indicate that organs are allocated to more medically urgent candidates who are farther away from the donor’s location. To ensure appropriate generalization, the selection of policies on the efficient frontier was based on simulations of 2007-2010, while 2011 simulation was withheld for out-of-sample validation and analysis. More details on parameter settings and the selection process can be found in Section 2.5.3.

2.3.2 Tradeoff analyses

Given the set of policy instances lying on the mortality vs. distance frontier for each distribution concept, tradeoff curves were generated for a range of different simulation outcomes (including all outcomes listed in Section 2.3.1). The full set of tradeoff curves can be found on our website (livervis.github.io), while this section summarises key analyses to contrast the different distribution concepts.

Figure 2-1 illustrates the tradeoff between annual average deaths (y-axis) and transport distance (x-axis) of selected policies in out-of-sample 2011 simulation. Continuous Distribution (CD) policies dominated all others considered and allowed for the greatest reduction of deaths across all transport distances. Figure 2-2 illustrates the tradeoff between the standard deviation of MMaT across DSAs (y-axis) and average transport distance (x-axis). Again, a continuous distribution model exhibits the lowest level σ MMaT at any given transport distance.

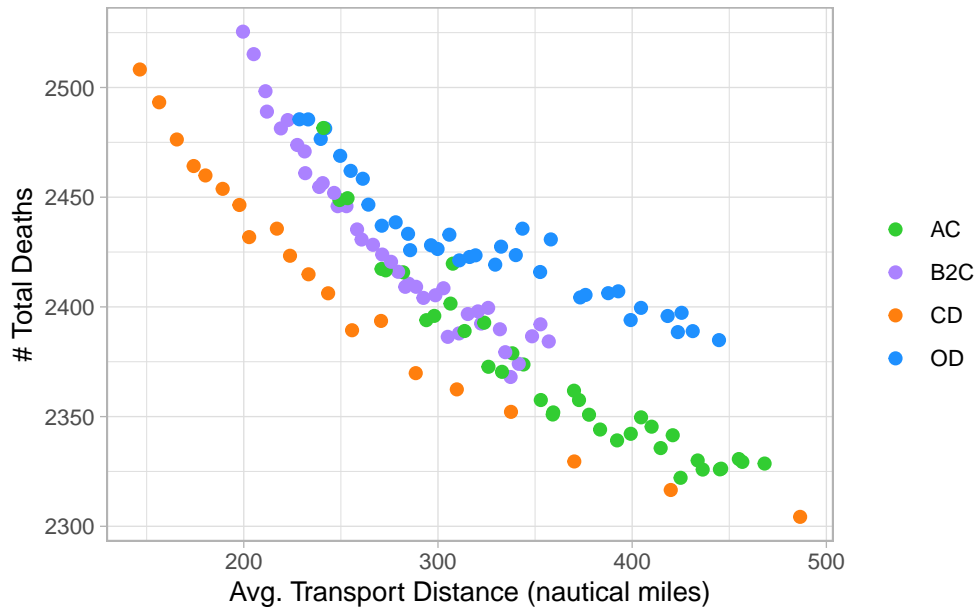


Figure 2-1: Tradeoff between annual average deaths and transport distance for different distribution concepts under MELD-based allocation in 2011 simulation. Individual points correspond to selected policies with different sharing parameters, colored according to the distribution concept they belong to (AC, B2C, CD, OD).




Figure 2-2: Tradeoff between σ MMaT and transport distance for different distribution concepts under MELD-based allocation in 2011 simulation. Individual points correspond to selected policies with different sharing parameters, colored according to the distribution concept they belong to (AC, B2C, CD, OD).

Table 2.1: Transplant allocation outcomes for reference policies versus CD-style policies with comparable transport distance, one using MELD- and one using OPOM-based allocation. Reduction in both annual average deaths and σ MMaT over 20 iterations of 2011 simulation is reported, with the one std. dev. confidence interval in parentheses.

Policy	Avg. Transport Distance (nm)	Deaths (Total)	Reduction in Deaths	σ MMaT	Reduction in σ MMaT
OPTN 11-District	244.7 (2.4)	2510 (24)	-	3.4 (0.1)	-
<i>Comparable CD (MELD)</i>	243.4 (2.3)	2406 (31)	-104 (36)	2.4 (0.1)	-29.8% (3.4%)
<i>Comparable CD (OPOM)</i>	238.5 (1.8)	1952 (32)	-558 (36)	2.0 (0.1)	-39.8% (4.2%)
AC (250,500)	269.1 (2.5)	2420 (27)	-	2.8 (0.1)	-
<i>Comparable CD (MELD)</i>	255.8 (3.3)	2389 (27)	-30 (23)	2.4 (0.1)	-12.0% (4.7%)
<i>Comparable CD (OPOM)</i>	261.5 (2.5)	1933 (22)	-487 (34)	2.0 (0.1)	-27.3% (6.5%)
AC (300,600)	294.0 (2.5)	2394 (24)	-	2.7 (0.1)	-
<i>Comparable CD (MELD)</i>	288.6 (3.1)	2370 (31)	-24 (34)	2.3 (0.1)	-15.5% (3.9%)
<i>Comparable CD (OPOM)</i>	293.6 (2.3)	1920 (27)	-474 (30)	2.0 (0.2)	-26.9% (5.8%)
B2C (250,500)	210.5 (3.9)	2513 (33)	-	3.0 (0.2)	-
<i>Comparable CD (MELD)</i>	202.8 (2.4)	2432 (22)	-81 (31)	2.7 (0.1)	-9.7% (6.4%)
<i>Comparable CD (OPOM)</i>	206.0 (1.8)	1990 (29)	-523 (22)	2.1 (0.1)	-28.0% (5.9%)

To illustrate the benefits more concretely, in Table 2.1 we compared each of the four reference policies (Section 2.2.2) to CD-style policies with comparable transportation burden. In particular, for each reference we selected a CD-style policy—that is, the value of λ —that resulted in the closest average organ transport distance as the reference policy in 2011 simulation. Of note, the current OPTN 11 Region policy resulted in 2510 annual deaths, with an average organ travel distance of 244.7nm. In comparison, utilizing a continuous distribution model while maintaining comparable organ travel distance as current policy (243.4nm), an additional 104 lives were saved every year. In comparison with the OPTN Board-approved AC 250/500 policy, which resulted in organ travel of 269.1nm and 2420 deaths, continuous distribution at comparable travel distance (255.8nm) resulted in a mortality reduction of 31 lives. For the proposed AC 300/600 and B2C 250/500 policies which resulted in 2394 and 2513 deaths respectively, continuous distribution at comparable travel distances resulted in reductions of 24 and 81 deaths respectively. Moreover, all CD-style policies exhibited significant reductions in σ MMaT relative to the corresponding reference policy, on the order of 10-30%.

Figure 2-3 plots the mortality vs. transport distance tradeoff curve of different concepts (paralleling Figure 2-1), but now using OPOM-based instead of MELD-based allocation. Under OPOM, the tradeoffs underlying all distribution concepts persisted and demonstrated the same advantage in a continuous distribution model. Projected mortality reduction due to OPOM was relatively uniform across all distribution concepts and transport distances. Per Table 2.1, a CD-style policy that used OPOM- instead of MELD-based allocation resulted on the order 558 lives saved annually compared to current policy.

Finally, we perform analysis to determine the impact of increased transport distances on utilization rates for marginal grafts from three categories: (a) donors aged ≥ 70 years; (b) donated after cardiac death (DCD); and (c) organs exhibiting macrosteatosis $\geq 30\%$. Table 2.2 does suggests that increased transport requirements were associated with a trend in increased discard rates and decreased average graft survival for marginal grafts.




Figure 2-3: Tradeoff between annual average deaths and transport distance for different distribution concepts under MELD- and OPOM-based allocation for 2011 simulation. Individual points correspond to selected policies with different sharing parameters, colored according to the distribution concept they belong to (AC, B2C, CD, OD), while shape indicates whether MELD or OPOM was used to evaluate candidates' medical urgency in allocation.

Table 2.2: Utilization of grafts for reference policies versus CD-style policies with comparable transport distance, one using MELD- and one using OPOM-based allocation. We use the following donor classifications: (i) overall, (ii) donors aged ≥ 70 years, (iii) donated after cardiac death (DCD), (iv) exhibiting $\geq 30\%$ macrosteatosis. The number of organs of each classification over the simulation horizon is denoted by N . Reported are the percentage of such organs discarded and average graft survival for transplanted candidates over 20 iterations of 2011 simulation, with the one std. dev. confidence interval in parentheses.

Policy	% Organs Discarded			Average Graft Survival (years)				
	Overall ($N = 6760$)	Age ≥ 70 ($N = 241$)	DCD ($N = 393$)	Steatotic ($N = 270$)	Overall ($N = 6760$)	Age ≥ 70 ($N = 241$)	DCD ($N = 393$)	Steatotic ($N = 270$)
OPTN 11-District	8.5 (0.4)	28.7 (2.6)	13.9 (1.6)	17.2 (2.2)	3.83 (0.02)	3.25 (0.16)	3.54 (0.09)	3.73 (0.13)
Comparable CD (MELD)	9.3 (0.3)	31.7 (1.7)	14.3 (1.2)	16.4 (2.1)	3.77 (0.03)	3.22 (0.21)	3.46 (0.11)	3.68 (0.14)
Comparable CD (OPOM)	8.9 (0.2)	27.9 (2.8)	13.2 (2.4)	15.8 (1.8)	3.80 (0.02)	3.20 (0.15)	3.49 (0.12)	3.65 (0.16)
AC (250,500)	9.8 (0.3)	32.2 (3.5)	15.4 (1.5)	18.3 (1.4)	3.77 (0.02)	3.21 (0.18)	3.46 (0.09)	3.68 (0.13)
Comparable CD (MELD)	9.4 (0.3)	31.2 (2.8)	14.4 (1.9)	17.2 (2.2)	3.76 (0.03)	3.20 (0.22)	3.40 (0.09)	3.65 (0.12)
Comparable CD (OPOM)	8.8 (0.2)	29.2 (2.6)	12.1 (1.5)	15.2 (2.0)	3.79 (0.03)	3.22 (0.14)	3.44 (0.10)	3.71 (0.12)
AC (300,600)	9.7 (0.3)	33.1 (2.9)	14.9 (1.7)	18.5 (2.2)	3.76 (0.02)	3.20 (0.20)	3.40 (0.10)	3.67 (0.14)
Comparable CD (MELD)	9.5 (0.3)	33.1 (2.3)	14.3 (1.6)	18.2 (2.2)	3.76 (0.02)	3.24 (0.18)	3.41 (0.07)	3.62 (0.11)
Comparable CD (OPOM)	8.7 (0.3)	27.8 (2.7)	12.8 (1.6)	14.4 (1.4)	3.78 (0.03)	3.20 (0.16)	3.46 (0.11)	3.60 (0.12)
B2C (250,500)	9.9 (0.3)	31.7 (3.2)	15.1 (1.8)	19.5 (2.8)	3.79 (0.03)	3.25 (0.21)	3.47 (0.10)	3.69 (0.16)
Comparable CD (MELD)	9.2 (0.3)	31.9 (2.1)	13.1 (1.8)	16.7 (3.2)	3.78 (0.03)	3.26 (0.19)	3.44 (0.08)	3.71 (0.11)
Comparable CD (OPOM)	9.1 (0.3)	29.9 (3.0)	13.2 (1.9)	15.9 (2.1)	3.80 (0.03)	3.27 (0.15)	3.50 (0.12)	3.67 (0.16)

The data reported here have been supplied by the Minneapolis Medical Research Foundation (MMRF) as the contractor for the Scientific Registry of Transplant Recipients (SRTR). The interpretation and reporting of these data are the responsibility of the authors and in no way should be seen as an official policy of or interpretation by the SRTR or the U.S. Government.

2.3.3 Implications for allocation policy

In order to promote access and efficiency, the Final Rule mandated that organ allocation not be based on the transplant candidate’s place of residence or listing – except as required by sound medical judgment and best use of donated organs to avoid wasting organs and futile transplants [27]. Yet, a geographic disparity in access to organs exists. For liver transplantation, this persistent discrepancy has manifested in differences in median MELD scores at transplant, rates of waitlist mortality, and ultimately has resulted in differential patterns of clinical practice by transplant professionals and candidates awaiting transplantation. The use of living-donor liver transplantation, a high-risk surgical endeavor, has largely been relegated to those areas of the country where the discrepancy between supply and demand is the greatest [98]. In addition, those socioeconomically privileged candidates who can travel, choose to do so by “migrating” to areas with lower MELD scores at transplant in order to achieve liver transplant in an expedited fashion, and subsequently returning with their new liver graft to their home [99].

The liver transplant community has remained divided over the discussions surrounding broader distribution of deceased-donor liver grafts. Indeed, the geographic disparity in access to organs remains a contentious topic that has progressed from a debate among medical professionals to now resulting in a litigious intersection with law, politics, and policy. Although in December of 2018 the UNOS Board of Directors voted to support an acuity circles approach to liver distribution with the hopes of achieving a more consistent and equitable approach, this by no means represents a mathematically optimized approach. Indeed, in December of 2018 the UNOS Board of Directors also approved the recommendation from the UNOS Ad Hoc Geography

Committee to use continuous distribution as a model for developing future organ distribution policies [69]. It is likely that the near future will see all solid organ allocation move to a continuous distribution framework, a direction supported herein by the first demonstration of the potential application of a boundaryless continuous score in liver distribution.

Paramount to achieving consensus on an approach to address geographic disparity is the need for a balance between efficiency and fairness. The dilemma has been previously extensively analyzed as it relates to resource allocation through the use of tradeoffs with assigned objectives [9]; herein we applied this framework to comprehensively analyze outcomes of various liver distribution concepts. This analysis has allowed for an in-depth, mathematically optimized, data-driven analysis of tradeoffs underlying each distribution concept. Tradeoff analyses of this kind have allowed for a methodical comparison of different distribution frameworks in terms of achievable outcomes of interest. The data in Figure 2-1, for example, revealed how many extra lives each framework can save per additional transport mile incurred, while Figure 2-2 reveals how much “fairer” distribution can be by minimizing differences in disease severity scores at transplantation per extra mile.

The generated tradeoff curves have demonstrated that a continuous scoring model achieved the greatest benefits both in terms of efficiency and fairness. Indeed, for any amount of transport distance incurred, there exists a CD-style policy that achieves both the lowest number of deaths, and the lowest standard deviation of MMaT, among all other suggested policies. Tradeoff curves have allowed for a complete assessment of all policies simultaneously to provide the framework for an informed decision – a decision that the transplant community will have to pursue in selecting the single policy that achieves their desired outcome. Although the OPTN’s Liver Intestine Committee accepted MMaT as a metric of geographic disparity in liver allocation in March of 2013, it is important to note that tradeoff curves assessing additional metrics of community interest can be generated in order to aid the decision making process.

It should also be noted that, while the issues of liver distribution have remained in

the forefront of discussions, there is significantly higher potential impact in lives saved through a more accurate and objective prioritization of candidate disease severity in liver allocation. An Optimized Prediction of Mortality (OPOM, www.opom.online) was recently developed utilizing machine learning models trained to predict any adult candidate's three-month waitlist mortality based upon 28 variables [11]. Indeed, OPOM allocation, when compared to MELD, reduced mortality on average by 418 deaths every year in LSAM analysis. An examination of liver distribution policies as applied through a MELD- vs OPOM-based allocation score not only reaffirmed that OPOM results in a significant number of additional lives saved every year, but also that OPOM allocation combined with a continuous scoring distribution policy maximized this potential (558 lives saved annually).

Limitations of this study include that estimating number of deaths averted using LSAM may represent an overestimation, and that LSAM cannot account for changes in practitioner listing or acceptance behavior; however, it is important to note the ability of LSAM to predict the overall directionality of change in assisting organ policy development [36]. In addition, the approach herein was only one possible implementation of continuous distribution, utilizing a simple linear function of distance with disease severity. Snyder et al. eloquently delineated other continuous distribution models that entail non-linear functions of distance [92]. Whereas such more advanced models might not be as simple to communicate to patients and transplant professionals as the linear model considered here, they represent more general models that can only potentially further improve outcomes.

In summary, the transplant community has now accepted the concept of continuous scoring distribution policies to allow for a more equitable, and boundaryless organ distribution. We now demonstrate application of this concept utilizing the model of liver distribution. This first application of a continuous distribution score for liver transplantation demonstrates superiority to all other policies currently employed or considered, and warrants similar consideration for other forms of solid organ transplants.

2.4 Redistricting organ allocation

2.4.1 Background

In this final section, we describe the optimization algorithm we developed to generate Optimized Districts (OD) policies for the tradeoff analysis of Section 2.3.2. Recall that the OD distribution concept seeks to redesign regional boundaries between the 58 existing DSAs so as to optimize organ distribution. The optimization problem at its core is that of geographic districting, and is challenging for two reasons: first, the solution space is highly combinatorial and, second, the objective function is given by a computationally expensive, black-box simulator.

The optimal districting problem falls under the domain of spatial optimization and is typically posed as that of identifying a contiguous k -coloring on a certain planar adjacency graph. The lowest-granularity spatial units (in this case, DSAs), are modeled as nodes in a network that are connected by an edge if they share a geographic border. A valid districting solution is a partition of the nodes into a fixed number of disjoint, node-covering sets (sharing regions), each of which is fully connected and therefore geographically contiguous. The optimizer seeks the optimal such partition with respect to some pre-specified objective, possibly under additional constraints on the structure and features of the partition.

Optimal districting has been long studied in the literature, in applications ranging from political and school districting to sales territory alignment and map generalization [21, 39, 62, 105]. The need to incorporate complex objectives – often multiple at the same time – and hard combinatorial constraints has given rise to a multitude of heuristic approaches in these fields. More recently, exact mixed-integer optimization (MIO) formulations have been proposed that scale to problem sizes comparable to the DSA problem [22, 38, 90, 102]. Notably, such an approach has been previously used to generate DSA sharing districts for organ distribution [34]. As the authors of that study acknowledge, however, the use of MIO poses certain limitations in the objective and constraints that can be modeled, and reduces the complex dynamics of organ allocation to a static problem.

Our goal in this study was to develop a districting algorithm that (i) optimizes over simulated outcomes as directly as possible, and (ii) can efficiently identify a large number of districting solutions that lie on the efficient frontier of mortality and transport distance. Our approach, described below, combines concepts from simulation optimization, graph-based machine learning models, and local optimization methods.

2.4.2 Overview of the algorithm

Given a planar adjacency graph $G = (V, E)$, representing a map of the OPTN’s 58 DSAs, our algorithm seeks a partition Π of the nodes (DSAs) into disjoint, graph-covering sets (regions). We use $|\Pi|$ to denote the number of sets in the partition, and $\Pi(k)$ for $k = 1, \dots, |\Pi|$ for the nodes belonging to set k . Note here the difference between $|\Pi|$, the number of sets in the partition, and $|\Pi(k)|$, the number of nodes in set k .

The goal is to find a partition Π that, when used in conjunction with an OD-style allocation policy (Section 2.2.2), minimizes the weighted sum of simulated candidate deaths and average transport distance:

$$\Pi_{\gamma, K}^* = \arg \min_{\Pi: |\Pi|=K} \underbrace{(\text{Annual Deaths})}_{f_1(\Pi)} + \gamma \cdot \underbrace{(\text{Average Transport Distance})}_{f_2(\Pi)} \quad (2.1)$$

The parameter γ encodes a tradeoff between the two objectives, and can be varied to yield solutions lying across the efficient frontier.

Our algorithm comprises three modules. The first two are “oracle” models that, given any partition Π , predict the two outcomes of interest $f_1(\Pi)$ and $f_2(\Pi)$ under counterfactual OD-style allocation. The third module uses the oracle predictions and a local search procedure to systematically sort through valid partitions and select the most promising candidates for given values of γ and K .

In principle, LSAM/OrgSim simulation could be used directly to provide oracle predictions during the local search. Such an approach, however, is quickly seen to

be computationally intractable, as the space of possible partitions is combinatorially large and the simulator expensive to evaluate. Instead, our approach uses machine learning to create high-fidelity, computationally efficient *surrogates* that replace the simulator at optimization time. At a high level, we start by simulating N randomly-generated DSA partitions to form a data set $\mathcal{D} = \{\Pi^i, f_1(\Pi^i), f_2(\Pi^i)\}_{i=1}^N$, mapping partitions to simulated outcomes. We then train parametric regression models \tilde{f}_j to minimize regression error over \mathcal{D} so that $\tilde{f}_j(\Pi) \approx f_j(\Pi)$ for $j = 1, 2$. Once validated, the models’ predictions obviate the need for repeatedly simulating similar partitions: the search procedure makes incremental changes to an incumbent partition so as to improve the surrogate, as opposed to fully simulated, objective. Of course, the final proposals are full simulated to evaluate the true objective value.

The core contribution of this section lies in the specific formulation and training procedure for the two oracle models, which need to capture the underlying combinatorial structure in the models’ partition input. Simulated outcomes depend on both node-level features – e.g., the number of candidates awaiting transplantation in a given DSA – and the partition structure itself, which determines sharing dynamics between DSAs. In the following two sections we describe our proposed oracle models and describe how they each use the partition input to predict outcomes.

2.4.3 Mortality oracle model

The goal of the mortality oracle model is to predict the total number of patient deaths in simulation for a given partition of the 58 DSAs into sharing regions. As an intermediate step, our model first predicts the number of deaths in each of the underlying regions, and aggregates for a unified estimate. The prediction for each region is akin to classical linear regression, with the exception that regression coefficients are calculated at the node (DSA) level, and averaged to give region-wide coefficients during prediction. In our implementation, independent variables includes the number of donors, number of candidates, and ratio of donors/candidates in the region, as well as number of candidate counts split by MELD status.

More concretely, each region k in partition Π is accompanied by a set of p features,

denoted by $\mathbf{x}_k \in \mathbb{R}^p$ (for a total of $p \cdot |\Pi|$ features). These may include any number of aggregate region attributes, e.g., the total number of candidates or donors in the region, their average urgency score, etc. The dependent variable for region k is denoted by y_k and represents the number of deaths of patients listed in the region over the simulation. Thus the ultimate dependent variable, i.e., the total number of deaths, is given by $\sum_{k=1}^{|\Pi|} y_k$.

Conceptually, our model estimates separate regression coefficients for each node in the graph, globally and independent of any given partition. To capture sharing dynamics within a region, at the prediction step we average the coefficients of all DSAs within the region to get region-level parameters that thus depend on the partition structure. Mathematically, each node $v \in V$ is associated with a coefficient vector $\beta_v \in \mathbb{R}^p$ (for a total of $p \cdot |V|$ parameters), and the number of deaths in region k is predicted as follows:

$$\hat{y}_k(\Pi) = \sum_{j=1}^p \sum_{v \in \Pi(k)} \frac{(\beta_v)_j}{|\Pi(k)|} (\mathbf{x}_k)_j$$

The intuition here is that the β_v are latent parameters belonging to the fixed DSAs, and regions, which change simulation by simulation, “inherit” the average parameter of all DSAs comprising them. For an illustrative example, suppose the j 'th independent variable corresponds to the total number of candidates in a region. Then, ceteris paribus, $(\beta_d)_j$ can be interpreted as a sort of overall mortality rate for DSA v , and the region coefficient $\sum_{v \in \Pi(k)} \frac{(\beta_v)_j}{|\Pi(k)|}$ averages out the mortality rates of the DSAs comprising the region.

For reasons having to do with model training (see below), we define “adjusted” node-level features for each DSA v as $\bar{\mathbf{x}}_v := \frac{\mathbf{x}_k}{|\Pi(k)|} \in \mathbb{R}^p$, where k is the region v belongs to. In a sense, these represent the contribution that v is making to its region’s aggregate features, smaller if there are more DSAs in the region. The overall mortality prediction, summing over regions, can then be written as:

$$\tilde{f}_1(\Pi) = \sum_{k=1}^{|\Pi|} \hat{y}_k(\Pi) = \sum_{j=1}^p \sum_{k=1}^{|\Pi|} \sum_{v \in \Pi(k)} \frac{(\beta_v)_j}{|\Pi(k)|} (\mathbf{x}_k)_j = \sum_{v \in V} \beta_v^\top \bar{\mathbf{x}}_v$$

Note that once the node-level features are computed for a given partition, what remains is a linear function in β . The parameters are estimated by ordinary least squares regression over \mathcal{D} :

$$\beta^* = \arg \min_{\beta} \frac{1}{N} \sum_{i=1}^N \left(f_1(\Pi^i) - \tilde{f}_1(\Pi^i) \right)^2$$

The final model used in our analysis, trained on randomly generated partitions for the 2007-2009 simulation cohorts, yielded a mean absolute error (MAE) of 13 deaths on a range of 2,340 to 2,440 deaths for out-of-sample 2010 simulations.

2.4.4 Distance oracle model

The goal of the distance oracle model is to predict the average organ transport distance in simulation for a given partition of the 58 DSAs into sharing regions. As an intermediate step, our model first predicts flows of organs between each pair of DSAs under a given partition – that is the number of organs recovered in one DSA that were transplanted in another. The pairwise flows are aggregated, weighted by the geographic distance between the DSAs, to yield the overall distance prediction.

Concretely, organ flows are estimated via a separate fractional response regression for each pair of DSAs u and v .² Suppose that s_u organs were recovered in DSA u during simulation, and of those $n_{u,v}$ were transplanted to a candidate residing in DSA v . Then the dependent variable of the logistic regression is defined as $p_{u,v} := \frac{n_{u,v}}{s_u}$. Independent variables capture relevant partition information such as whether u and v are assigned to the same region, the number of other “competing” DSAs in their common region, and their relative candidate and donor ratios. Once validated, the resulting models ($|V|^2$ in total) are then used to predict the percentage of organs recovered in each DSA going to every other DSA under a new partition Π , which we denote by $\hat{p}_{a,b}(\Pi)$.

²Fractional response regression is very similar to logistic regression, with the primary difference being that the dependent variable can take fractional values in $[0, 1]$ rather than being binary.

Finally, the average organ transport distance is estimated by weighting the pairwise estimates by (i) the geographic distance between u and v 's centers, $\delta_{u,v}$, and (ii) the number of recovered organ in each DSA, s_u :

$$\tilde{f}_2(\Pi) = \frac{\sum_{u,v} \delta_{u,v} \cdot s_u \cdot \hat{p}_{u,v}(\Pi)}{\sum_{u,v} s_u \cdot \hat{p}_{u,v}(\Pi)}$$

Since $s_u \cdot \hat{p}_{u,v}(\Pi)$ is the expected number of organ recovered in DSA u that need to be transported to v , the numerator denotes the expected total transport distance, while the denominator is the expected total number of transplants. Their ratio is then the average organ transport distance.

The final model used in our analysis, trained on randomly generated partitions for the 2007-2009 simulation cohorts, yielded a mean absolute error (MAE) of 9.7 miles/transplant on a range of 270 to 435 miles/transplant for out-of-sample 2010 simulations.

2.4.5 Local district search

Armed with the two oracle models for forecasting outcomes, the search procedure seeks to (approximately) solve problem 2.1 for fixed number of regions K and tradeoff parameter γ . Each run starts with an initial, randomly generated partition Π_0 and makes incremental changes that improve the surrogate objective given by the weighted some of the oracle predictions.

To define notation, let \mathcal{P} denote the set of “valid” partitions for the problem at hand, namely those having exactly K geographically contiguous regions. Further, let $N(\Pi)$ denote the local neighbourhood of Π in the sense of a one-unit swap; that is, all $\Pi' \in N(\Pi)$ differs from Π only in the region assignment of a single DSA, which now belongs to the same region as one of its graph neighbours. Note that the set $N(\Pi)$ can be easily enumerated by looking at every edge in the graph that connects two differently assigned nodes and changing the color of either one. In principle, swaps of higher degree are also possible, though we do not use them here.

Initial partitions Π_0 are generated by selecting K graph nodes uniformly at ran-

dom to act as regional centers. Subsequently, the remaining nodes are assigned by randomly selecting a region and one of its edges that connects to an unassigned node randomly. Then, at each optimization step t , all $\Pi' \in N(\Pi_{t-1})$ are enumerated and screened to confirm they are valid in \mathcal{P} (checking that there are K contiguous regions). All valid proposed partitions are evaluated by the two oracle models and the one selected that achieves the greatest reduction in the hybrid objective:

$$\Pi_t = \arg \min_{\Pi' \in \mathcal{P} \cap N(\Pi_{t-1})} \tilde{f}_1(\Pi_{t-1}) + \gamma \cdot \tilde{f}_2(\Pi_{t-1})$$

The process repeats until no improving partition can be found. For the full optimization process, we further implement random restarts, running the local search algorithm with many random initial points and selecting the best ones. Details on the parameter settings and selection process for the efficient frontier can be found in Section 2.5.3.

2.5 Appendix

2.5.1 Distribution concepts in detail

This section describes the detailed implementation of all distribution concepts considered. Applicable to all concepts are the following principles:

1. Top priority is always given to waitlist candidates who are status 1A or 1B³ and are local (as defined by the concept).
2. Offers are then extended to remaining candidates within pre-defined classification groups, as defined by the distribution concept. These are generally based on some combination of medical urgency (MELD or OPOM) and geography.
3. Unless otherwise state, ties within classification groups are broken as in current policy [73], that is in an order dictated by:
 - Blood type compatibility (identical before compatible)
 - Decreasing medical urgency (rounded MELD or OPOM score)
 - Decreasing time on the waitlist at current or higher MELD/OPOM score.
4. Offers are extended only to blood-type compatible patients unless the candidate:
 - has indicated willingness to accept an incompatible donor *and*
 - is designated Status 1A or 1B or has a MELD/OPOM score above 25.

The definition of classification groups for all four distribution concepts follows. In the case of OPOM-based allocation, all references to MELD are replaced by OPOM, which has been scaled to have the same exact distribution as MELD.

³These statuses are special designations granted by the OPTN to indicate utmost medical urgency; e.g., patients with fulminant liver failure and a life expectancy of less than 7 days or similarly dire circumstances.

Optimized Districts (OD): The 58 DSAs are grouped into some number of broader regions, and candidates are classified as local (same DSA), regional (same Region) or national (otherwise) relative to the donor's location. Allocation is modeled after the current OPTN 11-Region distribution policy: offers are extended first to regional candidates designated as Status 1A or 1B, followed by regional candidates with MELD [35,40], with local candidates preceding regional within each integral point of MELD (known as the Share35 rule). There follow local candidates with MELD [29,35), local candidates with MELD [15,29) and then regional candidates with MELD [15,35). Allocation then proceeds to national 1A, then national 1B, and then national MELD [15,40] candidates. Finally offers are made to candidates with MELD [6,15) locally, regionally and nationally. Parameters considered are the number of regions K , the possible groupings of DSAs into regions, and the weighting parameter γ which is varied in the optimization to yield policies along the efficient frontier.

Acuity Circles (AC): Based on the 2018 proposal submitted by the OPTN Liver and Intestine Transplantation Committee [72]. Three concentric circles are defined around each donor hospital (small, medium, large) and candidates assigned to them based on distance to their place of listing. Status 1A candidates within the large circle are offered first, followed by Status 1B within the large circle. Allocation then proceeds in expanding circles (small, medium, large) for each decreasing MELD group [37, 40], [33-37), [29-33) and [X, 29) where X is a threshold parameter typically set to 15 or 20. Candidates outside the large circle are offered next for Status 1A, then 1B and then MELD [X, 40]. Finally, candidates with MELD [6, X) are offered in expanding circles and outside the large circle. Parameters in the model consist of the radii of the three circles in nautical miles, as well as the lower MELD threshold X.

Broader 2-Circle Distribution (B2C): Based on the 2018 proposal submitted by the OPTN Liver and Intestine Transplantation Committee [72]. Three concentric circles are defined around each donor hospital (small, medium, large) and candidates assigned to them based on distance to their place of listing. Status 1A candidates within the large circle are offered first, followed by Status 1B also within the large circle. Candidates within the medium circle with MELD in the [Y, 40] range are

offered next, where Y is a threshold parameter typically set to 29, 32 or 35. Allocation proceeds to candidates in the [15, 40] MELD range in expanding circles (small, medium, large), excluding those high urgency patients already considered. Offers are next extended to candidates outside the large circle for Status 1A, then 1B and then MELD [15, 40]. Finally, candidates with MELD [6, 15) are offered in expanding circles (small, medium, large) and outside the large circle. Parameters in the model consist of the radii of the three circles in nautical miles, as well as the higher MELD threshold Y .

Continuous Distribution (CD): Status 1A candidates within 600 miles of the donor hospital are offered first, followed by Status 1B candidates within the same distance. Within these groups, candidates are ranked by a unified score that weights their blood type compatibility (10 points awarded for identical blood types, 5 for compatible and 0 otherwise) and distance to the donor hospital according to a tradeoff parameter λ – i.e., candidates are sorted in decreasing order of $(\text{compatibility points}) - \lambda \cdot (\text{Distance})$. Ties are broken by waiting time at Status 1A/1B. The third classification group consists of all non-Status 1A/1B candidates nationally, prioritized by a unified score that weights their MELD score and distance to the donor hospital by the same parameter λ , that is in decreasing order of $(\text{MELD}) - \lambda \cdot (\text{Distance})$. Ties are broken by waiting time at current or higher MELD score. Finally, candidates designated as Status 1A/1B beyond 600nm are offered. The primary parameter in the model is the distance tradeoff parameter λ – with larger values favoring proximity over acuity. Notably, when λ is set to zero the result is a “national sharing” policy whereby livers are offered based solely on medical urgency.

2.5.2 Simulation calibration

To model the distribution concepts not available in LSAM 2014, an alternative simulator we refer to as OrgSim was developed in the `python` programming language. OrgSim was built to replicate LSAM’s exact simulation functionalities based on publicly available information [94]. It uses all of the same input and model definition files, and uses the same discrete-event simulation process. It also exactly replicates

LSAM’s geolocation, offer acceptance and post-transplant statistical models.

Table 2.3 contains calibration results for OrgSim, comparing means and standard deviations of several key metrics over 50 iterations of allocation using the OPTN’s current 11-Region policy. Notably, the simulators match within 3.2% on all metrics, and within 0.3% and 0.7% for total deaths and average transport distance – the primary selection criteria for the optimization – respectively.

Table 2.3: Calibration results for OrgSim compared to LSAM on across all aggregate allocation metrics. Both simulators allocate livers according to the current OPTN 11-Region Policy. Values represent the mean of 50 iterations of simulation using either simulator, with standard deviations in parentheses.

Outcome	LSAM	OrgSim	Difference
Deaths (Total)	2513 (25)	2505 (26)	0.3%
Deaths (Waitlist)	1277 (10)	1285 (12)	0.6%
Deaths (Postgraft)	592 (24)	573 (20)	3.2%
Deaths (Removed)	644 (9)	648 (7)	0.5%
Avg. Transport Distance (nm)	244.7 (3.8)	243.1 (3.5)	0.7%
Median Transport Distance (nm)	99.9 (1.8)	100.2 (1.7)	0.3%
Avg. Transport Time (hrs)	1.46 (0.01)	1.43 (0.01)	2.2%
Median Transport Time (hrs)	1.60 (0.01)	1.60 (0.01)	0.1%
Transplant Count (Total)	6265 (18)	6183 (18)	1.3%
Percentage Local Transplants	62.5 (0.6)	62.8 (0.5)	0.4%
Percentage Organs Flown	49.8 (0.5)	49.9 (0.5)	0.1%
MMaT	25.0 (0.0)	25.1 (0.3)	0.3%
σ MMaT	3.4 (0.1)	3.5 (0.1)	1.9%

2.5.3 Parameter selection

This section describes the full set of parameter values that were used to generate candidate policies for the tradeoff analysis (see Sections 2.3.1). For each distribution concept, the goal was to select a set of policies – that is instantiations of the parameters – that lay on the efficient frontier of the two primary metrics of mortality and transport distance. As noted earlier, we assure generalization by selecting the efficient frontier based on simulations over 2007-2010, and base our analysis (including the figures and tables of Section 2.3.2) on out-of-sample simulations of 2011.

Given the small number of parameters for the CD-, AC- and B2C-type policies, here we used a simple grid search. For each concept, a grid of parameters (see below) was simulated using 2010 data. To select the efficient frontier of each concept, the x-axis metric (average transport distance) was split into N equal-sized bins across its simulated range. Then, within each bin, the policy was selected that minimized the y-axis metric (total number of deaths).

For the AC and B2C concepts we used a grid for the radii of the three concentric circles, each ranging between 50 and 1000 nautical miles in increments of 50, and excluding combinations where a larger circle had a radius more than five times that of its immediate predecessor. For the AC concept, we simulated policies with a MELD threshold of $X = 15$ and $X = 20$, while for B2C we considered all three thresholds proposed by the OPTN Liver and Intestine Transplantation Committee, namely $Y = 29, 32$ and 35 . For CD policies we varied the tradeoff parameter λ between 0 and 0.1 in increments of 0.002, capturing national sharing on the one hand, and a policy that equally weighs one point of MELD and 10 nautical miles of distance to the donor hospital on the other.

In the OD concept we used the optimization algorithm described in Section 2.4, seeking solutions with a number of regions that ranged from $K = 4, \dots, 11$, and varying the tradeoff parameter γ between 0 and 10 in increments of 0.25. For each combination of γ and K we conducted 1000 random restarts of the local search, using the trained oracle models and 2010 input data. For each value of γ , the top 10

solutions (across all K) were selected for validation through full simulation. The top 3 of those – now with respect to the fully simulated hybrid objective – were selected to constitute the efficient frontier of OD-style policies. The broad range of γ values ensured exploration of the full frontier.

2.5.4 Additional tradeoffs

To supplement the analysis of Section 2.3.2, we present here a robustness check on the measure of transport burden used as a proxy for the level of geographic sharing. In particular, Figures 2-4 and 2-5 plot tradeoffs between mortality and σ MMaT vs. average organ transport time as calculated by LSAM’s geolocation model. The plots are in direct correspondence with Figures 2-1 and 2-2, and confirm our findings that continuous distribution policies offer the greatest potential for mortality and geographic inequity reduction at any fixed level of transport burden.

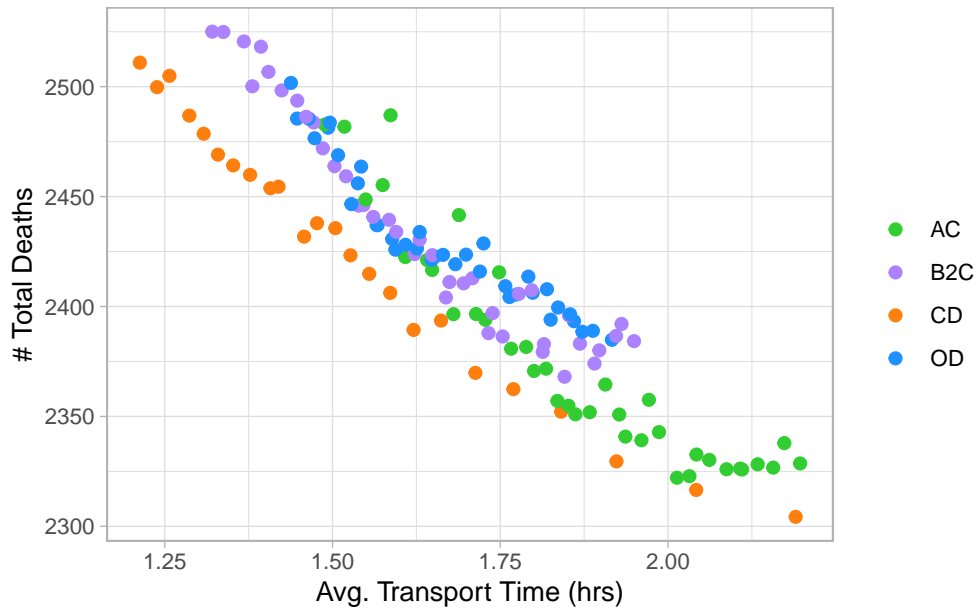


Figure 2-4: Tradeoff between annual average deaths and transport time for different distribution concepts under MELD-based allocation in 2011 simulation. Individual points correspond to selected policies with different sharing parameters, colored according to the distribution concept they belong to (AC, B2C, CD, OD).




Figure 2-5: Tradeoff between σ MMaT and transport time for different distribution concepts under MELD-based allocation in 2011 simulation. Individual points correspond to selected policies with different sharing parameters, colored according to the distribution concept they belong to (AC, B2C, CD, OD).

Chapter 3

Ethics-by-design: efficient, fair and inclusive resource allocation

3.1 Introduction

As has been made clear by current controversies regarding criteria for allocating COVID-19 vaccines in US states and across countries, the distribution of crucial medical goods and services in conditions of scarcity is among the most important, albeit contested, areas for public policy development [42, 43, 66]. COVID-19 vaccine allocation policy questions – for example, how much to prioritize age and measures of vulnerability and whether to consider “indirect benefits” – although in some sense specific to the pandemic, are also representative of a more general challenge in devising allocation policies: how to combine ethical theory, data modeling, and stakeholder input into a single process for generating a sound allocation system.

The importance of deriving and implementing good models for doing this can be seen in what is widely considered as a major failure in this area: the Oregon Medicaid prioritization process of the late 1980s and early 1990s. Oregon’s goal was to expand Medicaid to all persons below the poverty level while keeping costs manageable by restricting which services were reimbursable. What services to cover was determined using a combination of methods including: (i) actuarial estimates of the cost of providing Medicaid coverage; (ii) a state-wide telephone survey formulated

to measure “Quality of Well Being” and calculate the effectiveness of approximately 1600 treatments for particular conditions; and (iii) input from physician specialist panels regarding benefit duration and average age of onset of relevant conditions [81]. The initial process’ focus on cost-benefit analysis generated a list of covered and non-covered procedures that was discomfiting to some, for example by covering tooth-capping but not emergent appendicitis. Ensuing controversy prompted a move to “correct” the process by, among other things, creating larger overarching categories and “moving by hand” into covered categories procedures or treatments that seemed to the commissioner to be common-sense priorities [80].

The Oregon experiment mixed evidence, ethics, and stakeholder input in a witch’s brew that was perceived as ad hoc and political. The process nevertheless had (at least) some of the right ingredients. A purely top-down ethics approach, absent testing of its ramifications through data modeling, is problematic, but so is prioritizing what can be measured rather than measuring what should be prioritized; and both ethics- and data-driven approaches will fail without stakeholder pressure testing and buy in. Can all of this be combined into a better model?

We argue yes, and describe a general framework that we view as a major step forward for ethically informed policymaking. The key idea is to incorporate ethics into a data-driven policy design process from the outset. At the core of our approach is a novel analytical tool, based on machine learning and optimization, that enables stakeholders to assess tradeoffs between different policy objectives. To achieve that aim, users specify their desired system-level outcomes, encompassing diverse ethical and utility considerations, and the tool identifies a conforming policy in near-real time. By exploring the impact of changing user inputs, stakeholders, from ethicists to community leaders, can develop evidence-based value judgments on relevant tradeoffs regardless of their technical expertise, and more effectively engage in the policymaking process. To present our framework, we use as a case study the Organ Procurement & Transplantation Network’s (OPTN) policymaking process for migrating from the current classification-based policy to a continuous distribution model for organ allocation [71, 92].

3.2 Redesigning organ allocation in the US

3.2.1 Policy development and the OPTN

Since 1986, the OPTN has been operated under a federal contract by the United Network for Organ Sharing (UNOS). A core obligation of UNOS is to facilitate development of policies that determine how organs from deceased donors in the United States are allocated to medically suitable candidates on the national waiting list. The unfortunate reality is that demand for transplants far outstrips available supply, and transplant candidates generally wait a significant amount of time before receiving an organ. But not everyone can afford to wait. Some end-stage organ failure patients are so critically ill that without a transplant they will die in a matter of days.

The guiding principles for developing organ allocation policy are set by the OPTN Final Rule [27]. Among other requirements, the Final Rule states that organ allocation policies must be based on “sound medical judgement,” seek to achieve the “best use” of donated organs, and promote “efficient management” of organ placement. Allocation policies should aim to equitably distribute organs to those most in need over as broad a geographic area as is feasible.

In translating the abstract ethical principles established by the Final Rule into concrete policy, the transplant community is confronted with many ethical dilemmas. For example, to which of two equally sick candidates needing an organ should it be offered first? What if one candidate would be transplanted at a hospital closer to the donor’s location, reducing the potentially detrimental effects of organ ischemic time and increasing organ placement efficiency, but the more distant candidate has been waiting longer? Adjudicating such dilemmas in a systematic, objective, and repeatable manner every time a donated organ becomes available (24×7 , 365 days a year) is the role of the organ allocation policies operationalized by UNOS.

Crucially, UNOS does not unilaterally develop policy, but rather serves as a convener of the transplant community to help develop new and refine existing policies. This community consists of a host of stakeholders including transplant surgeons, physicians, Organ Procurement Organization (OPO) professionals, transplant candi-

dates and recipients, living donors and donor families, as well as the general public. The OPTN policy development process is designed to foster careful, wide-ranging discussion and deliberation based on input from these various constituencies [74, 100]. Transplant professionals and patients provide input by serving on OPTN committees that meet regularly to discuss clinical and practical details involved in developing policy and monitoring policy impacts post implementation.

In these discussions, the OPTN employs an evidence-based approach to policy development, guided by subject-matter expertise from transplant professionals and patients and backstopped by public comment and Board approval steps incorporated into the process. Evidence is typically derived from both retrospective analyses of the rich OPTN transplant database (e.g., identifying inequities in access to organs from historical data) and allocation simulation modeling performed by the Scientific Registry of Transplant Recipients (SRTR) contractor to predict outcomes of proposed policy changes [94]. In recent years, more sophisticated approaches incorporating mathematical optimization and simulation modeling have at times played an integral role in the policy development process [34, 63].

3.2.2 The continuous distribution framework

The tension inherent in translating the guidance of the Final Rule into concrete policy is nowhere more evident than in the longstanding policy debates surrounding the role of geography in allocating organs [26, 35, 47, 48, 60, 65, 103]. To what degree should proximity to the donor hospital be considered vis-à-vis a candidate’s medical urgency, especially given the historic precedent of prioritizing “local” candidates (those registered at a transplant hospital within the administrative boundary assigned to OPOs)? With the central aim of increasing transparency and removing hard geographic boundaries that sometimes preclude organs from going to candidates most in need, the OPTN recently embarked on a large-scale initiative to migrate all allocation policies, starting with lung, to the boundaryless Continuous Distribution (CD) framework [69].

Whereas under current policy candidates are prioritized by way of an ordered list

of groups of “similar” candidates, under the CD framework candidates are prioritized by a mathematical formula. In its simplest form, all candidates on the lung waitlist, for example, could be ranked by a weighted sum of their medical priority (quantified by their Lung Allocation Score) and placement efficiency (using distance to the donor hospital as a proxy). The relative weight of these two components would govern how widely organs are offered to those in need, thereby codifying the tradeoff between placement efficiency and need-based equity. Converging on the “right” balance between efficiency and fairness is exactly the sort of ethical dilemma the OPTN seeks to reconcile through the deliberative policy development process.

Converting to the new CD framework also provides an opportunity to revisit value judgements embedded in existing policies [96]. For example, what influence should a candidate attribute like waiting time, which codifies the “first come, first serve” ethic, play vis-à-vis other policy objectives like ensuring equitable access to patients with biological disadvantages (e.g., a harder-to-match blood type)? Or what role should reducing waitlist mortality play compared to maximizing survival time among those who receive a transplant? Additional attributes can be incorporated into the composite allocation score to adjust candidate priority accordingly, and relative weights of each component chosen to strike a defensible balance between the OPTN’s many efficiency, utility, and fairness objectives. Value judgements (weights), as well as attributes used to build the score, will almost certainly differ depending on the organ (kidney, liver, heart, lung, etc.), as each organ-specific policy is expected to have its own formula.

Simulation allocation modeling provides crucial input to the OPTN’s evaluation process by predicting system-wide outcomes of proposed policies. The SRTR’s simulation models use historical waitlist and transplant data to simulate counterfactual allocation under a proposed prioritization scheme and predict aggregated outcomes (e.g., overall mortality rates, transplant rates for different candidate subpopulations, transport metrics, etc.) [94]. The ability to compare policies holistically across multiple objective dimensions facilitates “outcome-driven” policy discussions and helps to identify pain points to be accounted for (e.g., increased inefficiencies, such as more

organ shipments expected to require a flight, or unintended inequities in access to organs under a proposed policy).

Consequently, simulation modeling has typically been conducted via trial-and-error, with simulated outcomes from initial policy ideas used to iteratively refine those ideas until acceptable predicted outcomes are achieved. With continuous distribution, the process might work as follows. Given attributes to be included in the priority formula, an initial set of policies (i.e., corresponding weights) would be proposed, simulated, and their predicted utility/efficiency/fairness outcomes reviewed by stakeholders. The results and committee discussions would inform the selection of a subsequent round of policy options to be simulated, for example by increasing or decreasing certain attribute weights to address a predicted inequity or inefficiency, and the process would be repeated.

Unfortunately, the computationally expensive and time-consuming nature of iteratively running simulations, reviewing and discussing results, and determining the next set of options to simulate means that this crucial evidence-generating phase of policy development can consume many months or even years. Even when a desired improvement is identified during an iteration, it is not a priori clear what, if any, proposed remedy will achieve better outcomes until new simulations are run. Time considerations may constrain the range of options tried, or many iterations may be required until a suitable option is found. The decade-long development of the Kidney Allocation System (KAS), in which well over 30 different policy options were simulated, is a prime example [41, 82, 95].

3.3 Shifting to ethics-by-design

3.3.1 Optimization methodology

In this work, inspired by ethics-by-design principles, we introduce a general analytical framework for outcome-driven policy design when development is constrained by a burdensome evaluation process such as simulation. Our approach leverages the

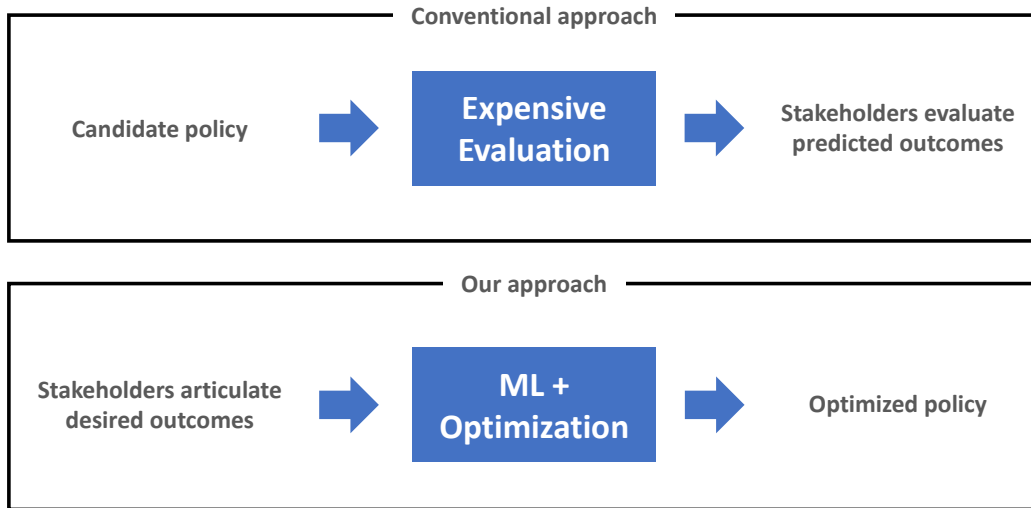


Figure 3-1: One step in the iterative design process under a traditional trial-and-error approach (top) vs. our framework (bottom). In the former, a burdensome evaluation process (e.g., computationally expensive simulation) forms a bottleneck that limits the number of candidate policies stakeholders can use to assess tradeoffs. Moreover, it is not a priori clear whether a proposed remedy aimed at improving outcomes will work until the modified policy is evaluated itself. In our framework, machine learning models remove the evaluation bottleneck and allow optimization to efficiently search for a policy that best achieves a set of pre-specified outcomes.

predictive power of machine learning to replace the bottleneck evaluation step with near-instantaneous model-based evaluation. The ability to more efficiently predict each outcome for a given policy in turn allows us to invert the problem: instead of starting with policy options and iteratively evaluating predicted outcomes, the approach starts with the desired outcomes and returns a specific policy engineered to achieve them (Figure 3-1). The result is an analytical tool that is ideally suited to characterizing tradeoffs between multiple objectives, allowing stakeholders to iterate on what the desired objective balance should be rather what design decisions achieve those outcomes.

We illustrate our approach as it might be applied to the design of a Continuous Distribution (CD) lung allocation policy. The CD case is a particularly potent

example, as the space of possible policies – that is, the space of relative attribute weights for the priority formula – is essentially infinite, and searching over it for the policy that strikes the “right” balance over multiple objectives becomes prohibitively expensive when using conventional simulation. Were it possible to evaluate policies instantaneously, the full efficient frontier of outcomes could be characterized a priori, and the policy instantiations that correspond to different points on it known precisely.

To this end, our methodology seeks first to characterize, using machine learning, the full range of achievable outcomes, and subsequently, through optimization, the efficient frontier. The role of policymakers would then be to select which point on the frontier best achieves the stated policy objectives; that is, decide what relative attribute weights strike their desired balance in the various efficiency, fairness and utility outcomes of interest. To be clear, the end-product of this design process is still a fully transparent CD policy that ranks candidates according to a static allocation formula and can be easily explained to patients, physicians, and other community members.

At a high level, our approach relies on specialized machine learning models to accurately and near instantaneously predict the outcomes of any given policy instantiation. Efficient predictions in turn allow optimization algorithms to sift through the space of possible policies and find the one that best achieves a set of user specified outcomes.

The first step involves using the simulation allocation model to predict the full set of outcomes for a fixed number of randomly generated policies, resulting in a dataset that maps policy instantiations (attribute weights) to each outcome of interest (mortality rate, disparities in transplant rates, transport costs, etc.). We modify the Thoracic Simulated Allocation Model (TSAM, version 2015), publicly available through the SRTR [94], to prioritize patients according to a continuous distribution formula and simulate each policy over the 2009–2011 period. These initial simulations, though computationally expensive, need only be run once and can be parallelized across multiple machines to reduce computational overhead.

The simulated runs are used to train nonlinear regression models that serve as

computationally efficient “approximators” of the simulator. During this offline training process, the models use the initial runs to identify and extrapolate correlations between attribute weights and each simulated outcome; then, given a new set of attribute weights, they can predict outcomes such as mortality rate or transport cost without invoking the computationally expensive simulator. We ensure that the models accurately capture the complexity of the simulator by using piecewise-linear functions that account for nonlinearity in the relationship between attribute weights and outcomes. In validation, we found that these regressions achieved out-of-sample R^2 ranging from 0.90 to 0.99 depending on the outcome, able to predict the simulator’s output with surprising accuracy.

With the ability to predict any simulated outcome accurately and efficiently, we can then use optimization to find new policies that achieve any set of pre-specified outcomes in near real-time. In an optimization problem, one seeks to find the values of decision variables (here, composite score attribute weights) that minimize or maximize an overarching objective function (a simulated outcome, such as mortality or post-transplant survival) subject to additional constraints (efficiency or fairness requirements, such as an upper bound on average organ transport distance or transplant rate disparities for patients with different blood types). Given such a specification, we use an open-source mixed-integer optimization solver to sift through the vast space of possible policies to find one that meets the specified criteria. If no such policy exists (e.g., because the constraints were too stringent), the optimization problem can be modified to instead find a policy that deviates minimally from the stated requirements. Once such a policy has been found, the original simulator is re-invoked to verify the regression models’ predictions and evaluate the optimization-derived policy in full.

3.3.2 Implications for policy development

Because any simulator output modeled in the regression can be designated in either the objective or constraints of an optimization problem, policymakers can quickly iterate through different optimization scenarios – that is, different sets of objectives

Choose Outcome to Optimize




Figure 3-2: Slider-based optimization interface. Any simulation outcome can be selected either as the overarching objective outcome (here, minimizing # Waitlist Deaths), or added as a constraint by providing an upper bound on its value vis-à-vis some reference policy. The optimization methodology produces a conforming allocation policy in seconds, whose predicted outcomes can be used to refine the objective and constraints further.

and constraints involving different policy outcomes – and explore what policies are predicted to achieve them. One might envision a “slider”-based interface, as depicted in Figure 3-2, whereby policymakers iteratively explore scenarios by changing the outcome to optimize and adding or adjusting constraint bounds.

One might begin, for example, by looking for a policy that minimizes waitlist mortality (the overarching objective) without increasing average organ transport distance vis-à-vis current policy (a constraint). The optimization computes a conforming allocation policy within seconds, which is simulated and other outcomes shown. If the resulting policy exhibits some undesirable characteristic, say an unacceptable increase in transplant rate disparities by blood type, an additional constraint could be applied and the optimization resolved to arrive at a new policy.

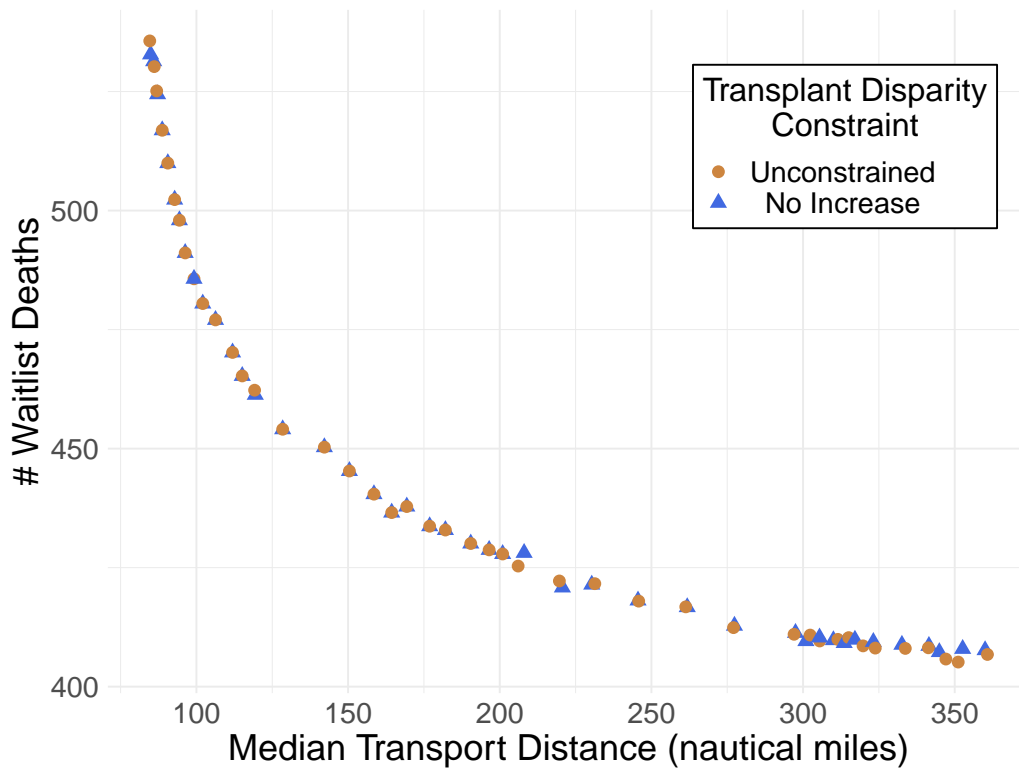


Figure 3-3: Efficient frontier of mortality vs. median organ transport distance in lung CD policy. Each circle represents a different optimization-derived policy, with different relative score attribute weights, selected to minimize waitlist mortality and with a varying upper bound on median organ transport distance. Triangles are similarly selected, but including additional optimization constraints to enforce that the discrepancy in transplant rates across patient blood types and height groups does not increase vis-à-vis current policy.

A natural extension of the above framework is the ability to construct tradeoff curves for simulated outcomes, as depicted in Figure 3-3. Here, the optimization algorithm is invoked a number of times to generate policies on the efficient frontier of two outcomes. Each optimization run attempts to minimize waitlist mortality (y-axis), but the upper bound for median organ transport distance constraint (x-axis) is systematically varied across a range of values, resulting in policies with different placement efficiency characteristics. As the upper bound on distance increases, the optimization selects attribute weights that increasingly favor medical urgency over proximity in the allocation score. Initially, the added flexibility allows policies to realize a significant reduction in mortality as organs are transplanted to sicker patients farther away. At a certain point, however, the reduction exhibits diminishing returns, providing valuable insight into the projected benefits of broader distribution.

Although harder to visualize, the optimization methodology is able to search over the efficient frontier of any number of additional outcomes. We include in the same figure a set of optimized policies, generated as before, but with a “guardrail” constraint that dictates that policies should not increase disparities in transplant rates by candidate blood type and height group vis-à-vis current policy. That the “No Increase” policies fall on the same curve as the “Unconstrained” ones suggests that lung continuous distribution policies are able to achieve essentially the same mortality benefit per mile of additional organ transport distance without exacerbating disparities in access to transplant by candidate blood group.

This study used data from the Scientific Registry of Transplant Recipients (SRTR). The SRTR data system includes data on all donor, wait-listed candidates, and transplant recipients in the US, submitted by the members of the Organ Procurement & Transplantation Network (OPTN). The Health Resources and Services Administration (HRSA), U.S. Department of Health and Human Services provides oversight to the activities of the OPTN and SRTR contractors.

The data reported here have been supplied by the Hennepin Healthcare Research Institute (HHRI) as the contractor for the Scientific Registry of Transplant Recipients (SRTR). The interpretation and reporting of these data are the responsibility of the

author(s) and in no way should be seen as an official policy of or interpretation by the SRTR or the U.S. Government.

3.4 Discussion

As the OPTN case study shows, leveraging the predictive power of machine learning to replace the bottleneck policy evaluation step with near-instantaneous model-based evaluations opens up exciting possibilities for refining ethicists’ and other stakeholders’ input to data-driven policymaking. One of the chief advantages of the proposed approach is that it enables ethicists and communities to offer input at precisely the points at which their views have the most value.

To illustrate, consider the following example of lung allocation policymaking. Asking ethicists or community members whether a lung candidate attribute estimating post-transplant survival and another attribute measuring waitlist survival should both be weighted at 34% of the composite allocation score, or one reduced by 2%, would likely elicit blank stares or weak intuitions at best. Asking instead “would you be willing to tolerate a 10% decrease in average years of post-transplant survival if it meant a 5% reduction in waitlist mortality,” and showing how the tradeoff changes as a “slider” moves, poses a set of considerations for which ethical theory and/or community sentiment is likely to generate more useful feedback. We view this tool as an important way to improve the ability of key stakeholders without medical expertise, including donor and recipient populations, to influence policymaking. This approach also enables ethicists and community groups to clearly pre-specify “guardrails,” for example, that disparities between blood type or racial groups not increase at all, or more than a specified amount, and show in near-real time how including, removing, or modifying those guardrails will change results.

While we view this approach as a major step forward for ethically informed policymaking, it does not solve every problem. First among the difficulties that persist is that modeling is only as good as the data on which it is based and the assumptions upon which it relies. In general, our approach works best when policymakers have

access to high-quality data and analytical tools to accurately predict outcomes, e.g., counterfactual simulation in the case of organ allocation. Even then, an ethicist or a community stakeholder may raise a question about the effect of a policy change – for example, effects on patients with intellectual disabilities – that is not captured in the input data and cannot be accurately modeled in a tradeoff curve. Similarly, predictions in applications areas where practitioners have little historical experience (e.g., COVID-19 vaccine allocation), are likely to rely on assumptions that should be closely examined. Recent work has shown how particular modeling choices in, say, curation of the training data or selection of objective metric, can result in biased model recommendations. Fortunately, the machine learning and optimization communities have made significant headway in developing methods to identify and mitigate such model-based biases [5, 44, 97].

In practice, data and model limitations should also be clearly communicated to stakeholders during the design process, to avoid over-reliance on imperfect predictions and to allow clinical and subject-matter expertise to influence conclusions drawn from the analysis. The ideal solution is to use representative, high-quality data sets, when that is not possible to be transparent about the limitations of the data sets being used, and when those limitations are serious enough to reconsider whether the model can do the work we want out of it. That being said, we note that such limitations also beset older, simpler approaches to policymaking, and one advantage of our framework is that it streamlines determining where current modeling practices fall short in addressing ethically important questions; this in turns sets up the possibility of changing those practices to fill the gaps.

Second, when done right, the approach discussed here seeks input from a multiplicity of stakeholders including ethicists, transplant professionals, donor communities, potential recipients, and sub-stratifications such as racial and ethnic minorities. The introduction of new analytical tools does not obviate the need for wide-ranging, cross-cutting deliberation to reach consensus, particularly as engaging with these stakeholders also means engaging with their biases. We believe our approach makes it easier for stakeholders to more meaningfully assess and weigh relevant tradeoffs, but it does

not guarantee more convergence between the groups. How should a policymaker respond when various communities champion different tradeoffs? Should all views be treated equally? Many ethicists will chafe at the idea of “ethics by headcount” in the sense of aggregating, without exploring the reasons behind, preferences for tradeoffs and potentially dismissing some as inconsistent or problematic. Though the federal regulation governing US organ allocation policy provides a final safeguard to rule out legally inconsistent policy options, this Final Rule does not precisely dictate which tradeoffs are out of bounds. Again, this problem is not new; in making tradeoffs more visible and accessible, our approach may make the problem more common – though some may view this more a feature (by inviting greater engagement with the tradeoffs) than a weakness.

Finally, there often exist significant cultural and institutional hurdles that must be overcome for the adoption of advanced analytical tools like the one we propose. In the case of organ allocation, US policymakers over the last two decades have most often used trial-and-error simulation to explore policy options. The shift to include advanced machine learning and optimization methods for helping develop a lung continuous distribution policy challenged the community’s flexibility to accommodate a new approach. Despite the introduction of new methodology, many of the core components of the age-old policymaking process remained in place including, committee discussion, public comment feedback from the community, and SRTR simulation modeling. The additional analyses, particularly optimized tradeoff curves (Figure 3-3), provided extra scrutiny of the proposal and helped stakeholders home-in on a final set of policy options worth consideration. Leveraging these tradeoff curves also provided a gentle introduction to the use of advanced mathematical methods that could pave the way for broader community acceptance of even greater reliance on goal-driven optimization in future allocation policy development.

Although some in the transplantation community might struggle with more sophisticated analytical tools, others see their adoption as an opportunity to improve the historically time-consuming policy development process. For years, committee discussion and retrospective data analysis informed policy proposals that were then

modeled by the SRTR. However, with these new tools, the community and committee can feel more confident about their chosen allocation policy options before the final, confirmatory simulation modeling is conducted. Over time, these more complex models can gain public confidence as new-and-improved policies are implemented and demonstrated to meaningfully improve outcomes for patients awaiting organ transplantation.

More generally, one might worry that, given relative unfamiliarity with machine learning among ethicists and stakeholder communities of interest, some may find it difficult to understand precisely what the approach we describe here “does” or be concerned about “not seeing the whole picture.” To remedy this gap will require thoughtful attempts at scientific communication that meets stakeholders “where they live,” rather than a one-size-fits all strategy. Policymakers need to be sensitive to algorithmic aversion and key opportunities to manage it including the use of interactive tools like “sliders” that enable stakeholders to see how AI/ML works even if they will never get “under the hood” [28]. For healthcare policymakers focused on allocation controversies, it might also be worth considering, depending on the setting, combining the machine learning methods discussed here with deliberative democracy techniques, such as deliberative polling, consensus conferences, and citizen juries [20, 31, 32].

Chapter 4

Reshaping US lung allocation through multi-objective optimization

4.1 Introduction

Following a landmark 2018 decision by the Organ Procurement & Transplantation Network Board of Directors, the OPTN began the process of migrating all of its allocation policies to a new framework termed “continuous distribution” [69]. Within this framework, transplant candidates on the national waitlist are ranked not by classification into distinct priority groups, but rather according to a scoring rule. The selection of key policy parameters, namely the score components and their associated weights in the formula, poses a challenging, multi-objective optimization problem with far-reaching ethical and practical implications. Given the limited supply of available organs, policies should seek to distribute organs to those most in need, over as broad a geographic area as feasible, while also safeguarding equity in access for different communities.

Crucially, policymakers in this setting must reconcile the views of a large and diverse set of stakeholders who often champion different tradeoffs. For example, some transplant professionals might favour local transplants, aiming to reduce the potentially detrimental effects of longer organ transport times, while others might favour broader distribution to reach more medically urgent candidates who are farther

away. Federal regulation, in the form of a “Final Rule” governing transplant policy [27], provides nonnegotiable, but often imprecise, guardrails, such as limitations on the extent to which geography can be incorporated into the rank-ordering algorithm. Similar debates arise in calibrating pediatric priority, or in striving to reduce access disparities based on race, sex, age, blood group and other patient characteristics.

Adjudicating these debates in a rigorous manner necessitates the development of flexible decision support tools that can shed light on relevant tradeoffs and promote convergence towards more effective and inclusive allocation policies. To this end, we introduce a novel framework for policy design based on multi-objective optimization. At a high level, we use machine learning and mixed-integer representable surrogate optimization models to approximate allocation outcomes, allowing us to characterize the (approximate) efficient frontier of policy design decisions. By varying the objective and constraints of the optimization, one can design policies with different efficiency and fairness characteristics. The result is a decision-support tool that enables stakeholders, even those without technical expertise, to quickly iterate on different policy scenarios and refine their value judgments on relevant tradeoffs as they engage in the policymaking process.

Herein we describe how we applied our methodology to the design of a continuous distribution policy for lung allocation, working closely with researchers at the United Network for Organ Sharing (UNOS) and the OPTN Lung Transplantation Committee to help develop their proposal. Using our optimization tool, we explored many different policy options during the design phase, while the resultant tradeoff analysis was presented to committee members and helped guide their decision-making process [54, 55]. The committee’s official continuous distribution proposal used values for policy parameters that were identified as inflection points in our analysis for organ placement efficiency and pediatric priority. The proposal was distributed for public comment in August of 2021 [57], ratified in October of the same year [56], and approved by the OPTN Board of Directors in December 2021 [75]. Thus, it will become national policy starting in 2023, guiding how lungs are allocated in the US for years to come. We perform simulation studies that suggest the new policy could reduce

waitlist mortality by 21% compared to the status quo, averting 62 waitlist deaths per year, while also improving fairness and equity. Independent simulations performed by the Scientific Registry of Transplant Recipients (SRTR) place the estimated reduction of mortality even higher, at around 40% or 87 waitlist deaths averted per year.

4.2 Allocation policy in the US

4.2.1 Background

In the United States, transplantation policy is overseen by the Organ Procurement & Transplantation Network (OPTN), which since 1986 has been operated under federal contract by the United Network for Organ Sharing (UNOS). Any patient seeking a organ transplant is registered on a national waitlist maintained by the OPTN. When an organ is recovered from a deceased donor, all candidates are ranked based on a fixed set of rules (the allocation policy) and the organ is offered sequentially until a candidate accepts.

The guiding principles for developing organ allocation policy are set forth by the OPTN Final Rule [27]. Among other requirements, the Final Rule states that organ allocation policies must be based on “sound medical judgment,” seek to achieve the “best use” of donated organs, and promote “efficient management” of organ placement. Allocation policies must aim to equitably distribute organs to those most in need, over as broad a geographic area as feasible. Translating these abstract ethical principles into concrete policy presents many ethical dilemmas for the transplant community. For example, what if one candidate would be transplanted at a hospital closer to the donor’s location, reducing potentially detrimental effects of increased organ ischemic time and improving organ placement efficiency, but another, more distant candidate is sicker and likely to die sooner? Which of these two patients should an organ be offered to first?

Crucially, UNOS does not unilaterally determine allocation policy, but rather serves as a convener of the transplant community to help develop new and refine ex-

isting policy [74]. This community consists of transplant surgeons, physicians, Organ Procurement Organization professionals, transplant candidates and recipients, living donors and donor families, as well as the general public. Transplant professionals and patients provide input by serving on OPTN Committees that meet regularly to discuss the clinical and practical details involved in developing policy. Committees work with UNOS to perform quantitative and qualitative analyses, solicit input from key constituencies, and formulate proposals to change how allocation works. These proposals are distributed to the broader community for public comment and possibly refined as a result, before being presented to the OPTN Board of Directors for final approval.

4.2.2 Continuous distribution

Current organ allocation policies operate within a classification-based framework, whereby candidates are separated into “similar” groups and each group assigned a distinct priority level (with additional ranking rules defined within each group [76]). For example, an organ from a pediatric donor might be offered first to all pediatric candidates within some fixed distance of the donor hospital, ordered by their time on the waitlist. If none accepts, the organ is then offered to high-urgency adult candidates within the same boundary, followed by pediatric candidates who are farther away, and so forth. Such classification-based allocation systems have long been criticized as being unfair, particularly as regards to geographic distribution, since candidates’ access to organs can vary highly depending on their inclusion or not in a particular group [26, 35, 60, 65, 103].

With the central aim of increasing transparency and removing hard geographic (and other types of) boundaries that sometimes preclude organs from going to candidates most in need, the OPTN has recently embarked on a large-scale initiative to migrate all allocation policies – starting with lungs – from the current classification-based system to a boundaryless Continuous Distribution (CD) framework [53]. For exposition purposes, we will introduce the CD framework in the context of lung allocation, using the OPTN Lung Transplantation Committee’s recent development

process as a reference, but the general principles apply to all organs.

Under the CD paradigm, all candidates on the waitlist are prioritized according to a scoring rule. Relevant candidate and donor attributes are combined into a single Composite Allocation Score (CAS) which precisely determines a candidate's waitlist rank for the given organ. The two primary design decisions involved in developing such a policy are (1) what patient and donor attributes should be included as score components in the CAS and with what functional form, and (2) their relative weighting in the composite score. Typically there is broad consensus on the former, while the latter is a central point of debate given its importance in balancing different policy objectives (see Section 4.2.3).

As regards to the first design decision, the Lung Transplantation Committee's initial CD development process determined that the following patient and donor attributes would be included in the lung CAS:

- Post-transplant outcomes (PTAUC, “Post-Transplant Area Under (the survival) Curve”), a survival analysis-based measure of how long a patient is expected to live if they receive a transplant of median quality.
- Medical urgency (WLAUC, “Waitlist Area Under (the survival) Curve”), a survival analysis-based measure of how long a patient is expected to live if they do not receive a transplant and remain on the waitlist.
- Placement efficiency, a function of distance between the donor and recipient hospitals, reflecting the resources needed to match, transport and transplant the organ to a given patient.
- Biological disadvantage points for candidates who are medically harder to match based on blood type, height and immune system sensitization level.¹
- Patient access points awarded to pediatric patients (aged < 18 years) and prior living donors.

¹Potential recipients are “sensitized” if their immune system makes antibodies against potential donors. Sensitization usually occurs as a consequence of pregnancy, blood transfusions, or previous transplantation. Highly sensitized patients are more likely to reject an organ transplant than are unsensitized patients.

More details on the selected attributes and how they are measured can be found in the committee’s proposal [57].

The second design decision, and focus of this work, is the selection of a set of relative weights for each attribute. Whenever an organ o becomes available, each compatible candidate p is assigned a vector $a_{o,p} \in \mathbb{R}^m$ reflecting their priority across of each of the m aforementioned attributes. To create a scoring rule, we seek a vector of score weights $\lambda \in \mathbb{R}^m$ so that $CAS_{o,p} = \sum_m \lambda_m a_{o,p,m}$. We next focus on why selecting these weights is an exceedingly difficult problem.

4.2.3 Challenges in policy design

The selection of score weights presents two key challenges. First, it is fundamentally a multi-objective problem, in which policymakers must strike a balance between multiple efficiency, utility and fairness-related objectives. Moreover, different stakeholders might champion different tradeoffs in these objectives that must be reconciled during the development process. Second, given a vector of score weights, predicting policy outcomes under counterfactual allocation is challenging in and of itself, and requires, in this context, complex simulations to capture the many interacting effects of a given prioritization scheme. Simulation is both computationally expensive and non-transparent, making it hard to anticipate what outcomes would result for a given policy, or how they would compare across different proposals.

In the case of continuous distribution, the score weights are the primary levers through which policymakers can influence tradeoffs to meet the Final Rule’s many utility (post-transplant outcomes), efficiency (transport burden and costs), and fairness (mortality and equity in access to organs) mandates. For example, placing a low relative weight on placement efficiency results in policies that favour broad geographic distribution to highly urgent or otherwise disadvantaged patients, but might exhibit increased logistical burdens and worse post-transplant outcomes as organs need to be transported longer distances. Similarly, a higher weight on biological disadvantage might result in more equitable distribution to patients who are medically harder to match, but increase waitlist mortality as organs are not allocated to the most severely

ill patients.

Given a set of scoring attributes then, their weights should ideally be chosen to reflect the community’s consensus value judgments on the importance of different utility, equity, and efficiency considerations. Since different stakeholders often champion different tradeoffs, the OPTN’s policy development process is designed to foster wide-ranging, evidence-based discussion on tradeoffs and solicit input from the community’s many and diverse constituencies [74].

Simulation modeling plays an integral role in this process by allowing stakeholders to assess policy proposals across a wide range of metrics. The Simulated Allocation Model (SAMs), developed and maintained by the Scientific Registry of Transplant Recipients (SRTR), use historical waitlist and transplant data to simulate counterfactual allocation under a proposed policy, e.g., a continuous distribution policy with given attribute weights [94]. Through discrete event simulation, SAMs can be used to predict system-wide outcomes such as:

- Waitlist and post-transplant mortality rate.
- Transplant and discarded organ rate.
- Disparities in the above among patient subpopulations, e.g., by age, blood type, geography, sex, race, etc.
- Transportation and efficiency metrics, e.g., median organ transport distance or % of organs flown.

While this ability to compare policy options across multiple metrics facilitates “outcome-driven” policy debates, the black-box nature of the simulator introduces a significant challenge in selecting score weights. The mapping of weights to simulated outcomes is both computationally expensive and non-transparent, rendering optimization over them particularly challenging. In other words, it is hard to anticipate what the simulated outcomes would be for a given set of weights – which in turn makes it hard to find what weights would achieve a desired balance in multiple objectives.

The OPTN committees’ simulation modeling has typically been conducted via trial-and-error, with simulated outcomes from initial policy ideas used to iteratively refine those ideas until acceptable predicted outcomes are achieved. The time consuming process of iteratively running simulations, reviewing and discussing results, and determining the next set of options to simulate means that this crucial evidence-generating phase of policy development can consume many months or even years. Worse, the limited number of options that can be simulated might not reflect the full range of tradeoffs, while often it is difficult to identify what, if any, proposed remedy would achieve some desired improvement during an iteration. The decade-long development of the Kidney Allocation System (KAS), in which well over 30 different policy options were simulated, is a prime example [82, 95].

4.3 Multi-objective optimization framework

In this work, we introduce a general optimization-based framework for policy design that seeks to address the two challenges we discussed in the previous section, namely (i) that predicting a policy’s outcomes is the result of some complex, non-transparent process that complicates optimization, and (ii) that policymakers seek to balance tradeoffs in multiple objectives, reconciling potentially disparate value judgments from diverse stakeholders. To facilitate exposition, we introduce our methodology in the context of continuous distribution for lung allocation; however, the model we develop can be applied more generally for the design of any type of parameterized policy, be it in organ allocation or beyond (see relevant discussion in Section 4.5).

The basic building block of our framework is an approach to solve what is essentially an inverse control problem: given a set of desired outcomes, our model seeks the policy parameters that best achieve them. Crucially, we use machine learning to approximate the simulator when predicting allocation outcomes, allowing us to bypass explicit simulation during optimization. As a result, policymakers can directly design policies with given efficiency/fairness characteristics, in near real-time, unconstrained by the expensive, non-transparent simulator. Moreover, this more efficient

optimization enables outcome-driven tradeoff analysis by varying the optimization objective and constraints. The result is a flexible decision support tool for policy-makers, including those with non-technical backgrounds, to evaluate policy tradeoffs and understand the impact of design decisions on their outcomes of interest.

4.3.1 Mathematical model

We consider a function $B : \mathcal{X} \mapsto \mathcal{Y}$, with \mathcal{X} a bounded domain and $\mathcal{Y} \subseteq \mathbb{R}^d$. The domain \mathcal{X} represents the set of valid policy parameter settings, e.g., possible score weights λ for a continuous distribution allocation policy. For the purposes of optimization, we require that \mathcal{X} is Mixed Integer Linear Optimization (MILO) representable; that is, it can be expressed in terms of continuous- and integer-valued decision variables and linear inequality constraints on them. The multi-dimensional output \mathcal{Y} encodes a set of d outcomes that may be of interest to decision-makers, e.g., mortality rate, transplant rate, transport burden, equity metrics, etc. We assume without loss of generality that lower values are preferable for all outcomes (negating the sign of the outcome if necessary).

The function B is not assumed to have known functional form, and may only be queried at particular inputs $x \in \mathcal{X}$ to observe $y = B(x)$. Such functions are typically referred to as “black-box” functions, in the sense that they do not provide explicit gradients or structural properties that might aid in optimization over \mathcal{X} . Moreover, they are typically computationally expensive to evaluate, which renders finite-difference gradient methods impractical to run at large scale. In lung allocation for example, B would represent the counterfactual allocation simulator that predicts transplant outcomes for a given parameterized policy.

We now define the concept of a *problem instance* on the function B ; that is, an optimization problem over \mathcal{X} with a fixed set of outcomes appearing in the objective and constraints. For lung allocation, a problem instance might correspond to, say, finding the score weights λ that minimize waitlist mortality (primary outcome), with pre-specified upper bounds on organ transport distance and transplant rate disparities (constrained outcomes). Mathematically, a problem instance $\pi = (j, \mathbf{b})$ consists of (i)

a *primary objective* index $j \in [d]$ and (ii) a *requirement* vector $\mathbf{b} \in \mathbb{R}^d$. These induce the following optimization problem:

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{X}}{\text{minimize}} && B_j(\mathbf{x}) \\ & \text{subject to} && B_i(\mathbf{x}) \leq b_i \quad \forall i \neq j \end{aligned} \quad P(\pi)$$

The primary objective is the index of an outcome to be minimized directly. Other outcomes appear in the constraints, with right-hand side given by the requirement vector \mathbf{b} . Individual outcome requirements b_i are allowed to be $+\infty$, to denote that the outcome should be unconstrained.

Crucially, the set of problem instances of interest to the decision makers is not known a priori, and may change as they explore the space of solutions and achievable outcomes. If B is computationally expensive, sequentially solving multiple instances as users refine their requirements and explore tradeoffs quickly becomes impractical. This serves as a primary motivation for our work; we seek an efficient tool to compute solutions to arbitrary problem instances, or provide reasonable alternative solutions should those prove infeasible.

4.3.2 Methodology

Our approach, which comprises three phases, relies on surrogate modeling techniques [45] to formulate tractable approximations of $P(\pi)$. First, in the *Sample Design* phase, we generate a sample of representative inputs and evaluate each using the black-box function B . The paired inputs and outputs serve as a supervised machine learning dataset, which is used to train high-fidelity, mixed-integer representable surrogate models of B during the *Surrogate Modeling* phase. Third, in the *Optimization* phase, we dynamically formulate and solve mixed-integer linear optimization problems (MILOs) to solve $P(\pi)$ for various instances π , using the surrogate models in lieu of B .

Sample design

In the sample design phase, we query B at a representative sample of N design points and track the vector of outcomes, to be used in training of downstream surrogate models. We denote this dataset by $\{\mathbf{x}^n, \mathbf{y}^n = B(\mathbf{x}^n)\}_{n=1}^N$, with $\mathbf{x}^n \in \mathcal{X}$ and $\mathbf{y}^n \in \mathcal{Y}$ for all n . Concretely, in the lung allocation setting each data point would correspond to a randomly sampled set of scoring weights λ , as well as the vector of simulated outcomes using the corresponding scoring rule.

We sample each \mathbf{x}^n uniformly from \mathcal{X} to ensure broad generalization of the surrogate models across the domain, while the sample size N controls the quality of interpolation and must be traded off against the computational complexity of repeated queries of B . In practice, the queries of B are independent and can be effectively parallelized to create a larger dataset.

In general, the sampling procedure used to generate training samples is highly dependent on application and the form of \mathcal{X} . The problem of sampling uniformly from different domains is well studied in the experimental design literature, e.g., Latin hypercube sampling for hyper-rectangles [61], custom samplers for structured polyhedral sets like the unit simplex [91], or hit-and-run methods for more general polyhedral sets [104].

Surrogate modeling

Using the dataset $\{\mathbf{x}^n, \mathbf{y}^n\}_{n=1}^N$, we train machine learning models to predict each individual outcome; that is, parametric approximations $f_i(\mathbf{x}; \boldsymbol{\theta}^i) \approx B_i(\mathbf{x})$ for each outcome $i \in [d]$. Parameters $\boldsymbol{\theta}^i$ are estimated by minimizing an application-specific loss function, e.g., least-squares error, over the training set, and using cross-validation to select between different models for each outcome. We evaluate goodness-of-fit based on R^2 on a held-out test set, and retrain on the entire dataset to better interpolate the sampled design points.

Crucially, since the surrogate models will be optimized over, we require that their output is mixed-integer representable in the design variables \mathbf{x} . The choice of para-

metric class is once again dependent on the application, and should balance the need for high predictive power (for accurate approximation) and low parametric complexity (for efficient optimization over the trained models). MILO formulations have been studied for a broad range of machine learning methods with varying power and complexity, including regression trees, tree-based ensembles, and neural networks [1, 14, 59, 64]. In our implementation, we found that relatively simple piece-wise linear functions achieved very high accuracy while allowing for near real-time optimization (see Section 4.4.1).

Optimization phase

In the optimization phase, given any problem instance π , we solve a surrogate optimization problem to find an approximate solution to $P(\pi)$. In particular, we replace B with the trained approximations, and consider:

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{X}}{\text{minimize}} && f_j(\mathbf{x}; \boldsymbol{\theta}^j) \\ & \text{subject to} && f_i(\mathbf{x}; \boldsymbol{\theta}^i) \leq b_i \quad \forall i \end{aligned} \tag{S(\pi)}$$

Given our assumptions on \mathcal{X} and f_i , $S(\pi)$ is a Mixed Integer Linear Optimization problem. Its size scales with the number of policy parameters and outcomes, as well as the parametric complexity of the surrogate models. In our implementation (Section 4.4.1), the relatively small number of parameters and outcomes (in the tens) resulted in problems that could be solved in seconds by open-source MILO solvers.

The solution of $S(\pi)$ is ultimately evaluated with a single query of B and its true outcomes computed. Due to errors in the approximation models, it is not guaranteed that the solution is in fact feasible for $P(\pi)$. However the magnitude of constraint violations, if any, is not expected to be large if the surrogates are accurate. Should the violations prove unacceptably large we propose adding a safety term to the approximation (see Appendix 4.7.1).

Finally, we note that, based on the choice of \mathbf{b} , $S(\pi)$ could be infeasible – indicat-

ing that the desired requirements \mathbf{b} were overly stringent. To address this possibility, the model can be readily modified using slack variables, and appropriately reformulated to produce a feasible solution that comes “as close as possible” to meeting the requirements \mathbf{b} . We formalize this in Appendix 4.7.2.

4.3.3 Implications for policy design

There are several important advantages to our optimization approach. First, it enables policymakers to dynamically explore the space of policies and outcomes as they iterate towards consensus. The formulation of problem instances might be implemented as an interactive application (e.g., see Figure 4-1), allowing users to intuitively and flexibly define the objective and constraints for their outcomes of interest. Once the instance of interest has been specified, the application will automatically formulate and solve the corresponding MILO to display optimized policy parameters and their simulated outcomes. As a result, stakeholders, even those without technical expertise, can efficiently explore different policy options and refine their value judgments on relevant tradeoffs.

To exemplify, without our framework, policymakers for lung allocation would need to contemplate and debate how much weight to award for each CAS attribute, thereby “mixing apples & oranges” with limited intuition. Within our framework, and using a tool like the one shown in Figure 4-1, policymakers can instead focus on debating appropriate target outcomes; a setting where ethical theory and community sentiment is more likely to generate useful feedback.

In particular, policymakers might iteratively solve multiple $P(\pi)$ to produce policies with varying efficiency/fairness characteristics. If a generated policy exhibits some undesirable characteristic, e.g., an unintended increase in transplant rate disparities for patients of different blood types, a constraint on that outcome could be added to the optimization to address it. Conversely, policymakers might perform sensitivity analysis by removing or modifying a particular constraints to see how other outcomes are affected; for example, does loosening an equity constraint decrease wait-list mortality, and if so by how much?

Choose Outcome to Optimize

Minimize # Waitlist Deaths ▾




Figure 4-1: Illustrative interface for specifying problem instances. Any simulation outcome can be selected either as the overarching objective outcome (here, minimizing # Waitlist Deaths), or added as a constraint by providing an upper bound on its value vis-à-vis some reference policy. The surrogate optimization produces a conforming allocation policy in seconds, whose predicted outcomes can be used to refine the objective and constraints further.

More generally, our framework allows for global tradeoff analysis. By systematically varying the constraint bound on one outcome (e.g., organ transport distance) while optimizing a second (e.g., minimizing waitlist mortality), policymakers can visualize the (approximate) efficient frontier of outcomes. Inflection points, where marginal gains in one outcome diminish significantly relative to another, can provide guidance for how to balance different objectives. Crucially, the advantage of an optimization-based approach here is that policies on the tradeoff curves might be additionally constrained. For example, policymakers might examine the trade-off in transport distance and waitlist mortality, while ensuring that transplant rate disparities for different subpopulations do not exceed a certain level.

4.4 Reshaping national lung allocation policy

We discuss how we collaborated with UNOS to apply our framework to the design of a continuous distribution policy for lung allocation. Among others, we performed tradeoff analyses to address two questions confronting the OPTN Lung Transplantation Committee during the policy development process: (i) how to reconcile the Final Rule’s competing mandates of placement efficiency and broader geographic distribution to the most urgent patients, and (ii) how to ensure that pediatric patients maintained the same high level of access to organs as they do under the current classification-based system.

Both of these questions have far-reaching implications for fairness, equity, and the welfare of patients, as we discuss next. In particular, the first question has been at the forefront of policy debate for decades (see Section 4.2.2), and in fact served as a primary motivation for the OPTN’s decision to migrate all organ allocation to continuous distribution. Under continuous distribution, the relative weight of placement efficiency in the CAS formula controls the level of geographic distribution. It thus directly encodes the tradeoff between increased transportation burden on the one hand, and waitlist mortality and geographic equity on the other. To this end, we generated tradeoff curves for waitlist mortality vs. various transportation burden metrics, identifying inflection points that could guide the committee’s selection of a weight for placement efficiency.

The second question reflects the committee’s desire to ensure high levels of access for pediatric patients (aged less than 18 years). Beyond ethical motivations for prioritizing this vulnerable population, pediatric patients are also generally considered medically harder to match due to organ size compatibility constraints. Under the previous classification-based system, pediatric patients were given the highest priority for organs recovered from pediatric donors, ensuring that they were offered compatible organs at a high rate. Under continuous distribution, the committee sought to identify the minimum value for pediatric weight that would result in the highest possible transplant rates for pediatric patients, and our framework was applied to this end.

The results of our analyses, which we detail in the following sections, were presented to the OPTN Lung Transplantation Committee in March 2021 [54, 55]. We identified 10% placement efficiency weight as an inflection point in the tradeoff of waitlist mortality vs. transportation burden, and a range of 15%-20% pediatric weight as the minimum necessary to maintain high levels of pediatric access. After independently validating our results via additional simulation studies performed by the SRTR [93], the Committee eventually adopted our recommendations for the placement efficiency and pediatric access weights. In particular, the committee’s official continuous distribution proposal, depicted in Figure 4-2, recommended values of 10% and 20% respectively for the two attributes. The proposal was distributed for public comment in August of 2021 [57], ratified in October of the same year [56], and approved by the OPTN Board of Directors in December 2021 [75]. Thus, it will be implemented as the OPTN’s first continuous distribution policy in 2023, dictating how deceased-donor lungs are allocated in the US for years to come.




Figure 4-2: Continuous distribution attribute weights used in the OPTN Lung Transplantation Committee’s official proposal, as ratified by the OPTN Board of Directors in December 2021 [57, 75].

4.4.1 Modeling implementation

We first describe the implementation of our framework (Section 4.3) to the design of a lung continuous distribution policy in more detail. Here, the primary policy parameters to be optimized are the CAS weights λ for the six attributes selected by the committee (Section 4.2.2). In addition to the attributes themselves, the committee previously worked with UNOS to develop a set of *rating scales*, consisting of linear functions for some attributes and nonlinear functions for others, which translated the inherent value associated with each attribute value to a comparable $[0,1]$ scale [57]. Thus, our optimization domain consists of weight vectors λ drawn from the unit simplex, $\mathcal{X} = \{\lambda \in [0, 1]^6 : \sum_i \lambda_i = 1\}$, so that candidates' CAS scores also lie on a unit scale $[0, 1]$.

The black-box function B represents the SRTR's Thoracic Simulated Allocation Model (TSAM, version 2015), which simulates counterfactual allocation of 3,326 recovered lungs to 6,546 waitlist candidates over a two-year period from 2009-2011 [94]. We modify the simulator to accept a CAS weight vector λ as input, and implement a continuous distribution allocation scheme using the committee's selected attributes. The output domain \mathcal{Y} consists of 47 simulated outcomes encompassing a wide range of efficiency, utility and fairness metrics, including waitlist mortality, transplant rates, transport burden and equity measures for different patient classifications (see Appendix 4.7.4 for the full list). We applied the optimization methodology of Section 4.3.2 as follows:

Sample Design We sampled 10,000 weight vectors λ uniformly at random from the unit simplex to generate a training dataset [91]. Each generated policy was simulated 20 times, and the average value of each outcome was recorded.

Surrogate Modeling We fit piece-wise linear surrogate models separately for each outcome. Our hypothesis class consisted of functions that were the minimum or maximum of K affine functions of λ , with parameters θ^i corresponding to the coefficients and intercept of each function. We used a random 80-10-10% training/validation/testing split and selected the best-performing hyper-parameters based

on R^2 on the held-out validation set. Out-of-sample R^2 of the selected models ranged from 0.85-0.99 for all outcomes, with an average of 0.96. Further implementation details and results are given in Appendices 4.7.3 and 4.7.4.

Optimization We implemented problem instance formulation as an interactive tool, as depicted in Figure 4-1, and used the open-source CBC solver [30] to solve the corresponding MILOs. More information on formulations of our chosen surrogate models can be found in Appendix 4.7.3.

This study used data from the Scientific Registry of Transplant Recipients (SRTR). The SRTR data system includes data on all donor, wait-listed candidates, and transplant recipients in the US, submitted by the members of the Organ Procurement & Transplantation Network (OPTN). The Health Resources and Services Administration (HRSA), U.S. Department of Health and Human Services provides oversight to the activities of the OPTN and SRTR contractors.

4.4.2 Tradeoffs for placement efficiency

To address the committee’s question regarding geographic distribution, we applied our framework to generate tradeoff curves for waitlist mortality and different organ transport metrics. Figure 4-3 plots the total number of waitlist deaths in simulation vs. median organ transport distance for a series of optimized policies. Each point corresponds to a policy obtained by our model where the objective was to minimize the number of waitlist deaths (y-axis), and the upper bound on transport distance (y-axis) was varied across a grid. All policies used a fixed value of 20% for pediatric access weight (see Section 4.4.3), and were required to place equal weight on the WLAUC and PTAUC factors (per committee guidance). Additional constraints enforced that transplant rate disparities for patients of different blood type and height group did not increase vis-à-vis current policy. The secondary x-axis at the top of the figure displays the non-linear (but monotonic) mapping between optimized policies’ placement efficiency weight and resulting median organ transport distance in simulation. Current policy is also depicted through a distinctively labeled point.

On the left part of Figure 4-3, we observe policies with a relatively high weight




Figure 4-3: Tradeoff of waitlist mortality and placement efficiency in continuous distribution for lungs. Each point represents a different policy with optimized weights. The corresponding placement efficiency weight is shown on the top axis.

on placement efficiency. These result in higher waitlist mortality as nearby candidates are prioritized more than medically-urgent ones. Moving to the right of the graph, the placement efficiency weight decreases and organs are allocated to more medically urgent patients who are farther away, decreasing waitlist mortality but increasing transport burden. Of note, we identify a point of diminishing returns at $\approx 10\%$ placement efficiency weight (median organ transport distance of ≈ 175 nautical miles). We observed qualitatively similar results in tradeoff curves for which the x-axis represented two other simulated transport metrics, namely estimated transport cost and percentage of organs expected to be flown, with diminishing returns at around 10% placement efficiency weight (Appendix 4.7.5).

The robustness of the 10% finding across different transport metrics was independently validated through additional simulations performed by the SRTR on an updated 2018-2019 patient cohort [93].² The final round of official simulation modeling varied placement efficiency weight between of 10% , 15% , and 20% , and confirmed

²The updated simulation model and cohort used by the SRTR was not publicly available at the time of this work.

our findings of diminishing returns to waitlist mortality as a function of transport burden, culminating in the committee’s official proposal of a 10% weight on placement efficiency (Figure 4-2). Of note, an optimized CD policy with a 10% placement efficiency weight is seen to reduce waitlist mortality by $\approx 21\%$ compared to current policy in Figure 4-3, averting ≈ 118 deaths over the 2-year simulation horizon. The SRTR’s independent simulation study estimated an even greater reduction of $\approx 40\%$ in waitlist mortality, or 175 deaths averted over a 2-year period.³ At the same time, it is also more equitable as it no longer relies on hard geographic boundaries in allocation priority.

4.4.3 Ensuring pediatric priority

We also applied our framework to examine the impact of increasing pediatric access weight on pediatric transplant rates. In Figure 4-4, each point corresponds to an optimized policy with the objective of minimizing pediatric transplant rates (y-axis), while varying the lower bound on pediatric access weight (x-axis) across a grid. Optimizations were otherwise unconstrained. Thus, the plot illustrates the worst-case guaranteed pediatric transplant rate under a continuous distribution policy with a given pediatric access weight, regardless of any other requirements.

Figure 4-4 shows that a pediatric weight of $\approx 10\%$ suffices to ensure that pediatric access stays at the same level as current policy. Moreover, pediatric transplant rate seems to stabilize at a weight of about 15-20%. Although some preliminary, informal analysis by UNOS had suggested that a 30% weight would ensure that pediatric patients were always prioritized above adult patients (maximizing their transplant rate), our results indicate that a lower weight is sufficient to achieve the same goal. The committee’s subsequent simulation requests to the SRTR, as well as their final proposal, fixed pediatric access weight to 20%.

³We note that the absolute values from our study cannot be directly compared to the SRTR’s, as they simulate allocation in different patient cohorts.




Figure 4-4: Guaranteed pediatric transplant rate for patients aged 0-11 years vs. pediatric access weight. Rates were similarly evaluated for patients aged 12-17 years (Appendix 4.7.5).

The data reported here have been supplied by the Hennepin Healthcare Research Institute (HHRI) as the contractor for the Scientific Registry of Transplant Recipients (SRTR). The interpretation and reporting of these data are the responsibility of the author(s) and in no way should be seen as an official policy of or interpretation by the SRTR or the U.S. Government.

4.5 Discussion

In this work we present a novel analytical framework for multi-objective policy design when prediction of policy outcomes is constrained by a burdensome, non-transparent evaluation process. We leverage machine learning and mixed-integer optimization to enable dynamic design of policies with given outcomes. The result is a flexible decision support tool for policymakers, even those without technical expertise, to understand the range of achievable outcomes and evaluate tradeoffs. One of the chief advantages of this approach is that it enables stakeholders and community groups to offer input

at precisely the points at which their views have the most value; in determining the set of desired outcomes of a policy, rather than struggling to determine the exact design parameters that achieve them.

We note that the multi-objective policy design setting we address is hardly unique to the organ allocation problem we studied. Our machine-learning and surrogate optimization approach can be readily applied to the design of any type of parametric allocation policy, not just a scoring rule. Beyond organ allocation, policymakers in a broad range of application domains face similar challenges in policy design, seeking to model complex system dynamics and balance multiple objectives. For example, educational districts seek to assign students to schools so as to best balance family preferences, transportation and operational logistics, and equity and diversity requirements. In public housing allocation, newly available units must be assigned to waitlist applicants with diverse preferences, incorporating broader equity and diversity considerations. Our framework might be applied to evaluate tradeoffs in either domain, when policy outcomes depend on assignments made by a computationally-expensive, black-box algorithm (e.g., mixed-integer matching formulation) with parameterized objective and constraints. Using a meta-optimization tool such as the one we designed for organ allocation, policymakers in school choice or public housing allocation can dynamically explore the range of possible assignment outcomes and evaluate tradeoffs in their objectives of choice.

While we view our approach as a major step forward in effective policy design, we also note that it is subject to certain limitations. First and foremost, policies designed through surrogate optimization are only as good as the underlying black-box function’s predictions. In organ allocation for example, current simulation tools do not account for behavioural changes to patient or clinician practice when allocation policy changes, which might result in prediction errors [36, 37, 49]. Machine-learning models trained to approximate the simulator, and any optimization over them, would likely inherit any such predictive bias. In practice, these limitations should be pointed out to stakeholders to avoid over-reliance on imperfect predictions, and to allow clinical and subject-matter expertise to influence the conclusions drawn from the tool. Secondly,

surrogate optimization works well only if the trained models are highly accurate approximations of the black-box function. In our implementation, relatively simple piece-wise linear models sufficed to produce extremely high-fidelity approximations; however, other applications might require more complex surrogate models, which may be harder to train and optimize over to achieve similar performance. Finally, our approach is best suited to application domains where the problem dimension (number of policy parameters and/or outcomes) is not exceedingly high, to ensure both accuracy of the surrogate models with a relatively small number of training samples, as well as reasonable computational overhead when solving repeated mixed-integer programs.

4.6 Glossary

We provide here a glossary of relevant terms and definitions for our lung allocation work. This includes common metrics used to determine patients' allocation priority, as well as measure system-wide outcomes during simulation.

Active patient years: Total active waitlist time for a group of patients, measured in years. For a single patient, active time is measured as time elapsed between their listing (addition to the waitlist) and their removal (due to death, receiving a transplant, becoming too ill to transplant, or other reason), excluding any periods when they were listed as inactive. The total wait time, summed over groups of patients, is used as a normalizing factor to compare transplant/mortality rates across populations with different sizes.

Continuous Distribution (CD): An allocation policy whereby candidates for a given organ are ranked according to Composite Allocation Score (CAS) that incorporates both patient and donor attributes (see Section 4.2.2).

Diagnosis Group: A categorization of disease diagnosis for lung waitlist candidates into four broad groups: (A) obstructive lung disease, (B) pulmonary vascular disease, (C) cystic fibrosis and immunodeficiency disorder (D) restrictive lung disease. The classification is standard under OPTN policy [76].

Lung Allocation Score (LAS): a primary measure of lung allocation priority, used by the OPTN to prioritize lung candidates under current (non-CD) policy [76]. Defined as $PTAUC - 2 \cdot WLAUC$, and normalized to lie on a 0-100 scale.

Median LAS at Transplant (MLaT): The median transplanted patient's LAS score at time-of-transplant for given population group. The population-Weighted Mean Absolution Deviation (WMAD) of MLaT between different OPOs is used as a measure of geographic equity, in the sense that larger variances in MLaT occur when high-LAS (i.e., high urgency/benefit) patients are receiving most transplants in some OPOs but not others.

Net Benefit: A measure of a lung patient's expected benefit from receiving a transplant of median quality, calculated as the difference between their PTAUC and

WLAUC at any given point in time. Given the definition of those two metrics, net benefit measures the expected additional days they are expected to live within the next year if they receive a transplant rather than remain on the waitlist. Average net-benefit at time-of-transplant is used as a measure of system-wide utility, in the sense that higher values occur when organs are allocated to patients who would gain the most life expectancy from receiving a transplant.

Organ Procurement Organization (OPO): not-for-profit organizations responsible for recovering deceased-donor organs and collecting patient and donor clinical information to the OPTN. There are 57 OPOs in the United States, each covering different geographic donation service area.

Organ Transport Distance: The geographic distance between the donor hospital where an organ is recovered and the transplant center where a candidate is listed (and would be transplanted if allocated the organ), measured in nautical miles. Average or median organ transport distance across all transplants is used as a system-wide measure of efficiency, in the sense that higher values indicate higher transport burden.

Organ Transport Cost: The estimated cost of transporting an organ from the donor hospital to the transplant center where a candidate is listed (and would be transplanted if allocated the organ), measured in dollars. In reality, organ transport cost can vary greatly based on circumstance. For the purposes of simulation, an expected cost is calculated as a non-linear function of organ transport distance, according to a model developed by the SRTR. Given the approximate nature of this estimate, we express total transport costs over a simulated period only as a percentage relative the total transport cost estimated for current (non-CD) policy over the same simulated period.

Organs Flown: The number of transplants (expressed as a count or percentage of the total) requiring organ transport by plane. In reality, cases when a recovered organ must be flown can vary greatly based on circumstance. According to UNOS guidance, simulations assumed that an organ would be flown if the distance between the donor hospital and transplant center was greater than 75 nautical miles.

Post-TX Mortality: Patient deaths (expressed as a count, or normalized by total

wait time) that occurred due to graft failure after the patient received a transplant. In our simulations, we only consider post-transplant deaths within the first year, as predicted by a survival analysis model included with the version of TSAM we obtained.

PTAUC: “Post-Transplant Area Under (the survival) Curve”, a survival analysis-based measure of a lung patient’s *post-transplant outcomes*. PTAUC measures the expected number of days the patient is expected to survive within the next year if they receive a transplant of median quality. Its calculation is based on a patient’s clinical and demographic data at a given point in time, and is standard according to OPTN Policy [76].

Transplant Rate (TX Rate): The number of transplants received by a group of patients divided by their total active waitlist time, measured in transplants per patient-year. We report transplant rates broken down by six different patient classifications, according to blood type (A, B, O, AB), diagnosis group (A, B, C, D), height sextile (<157cm, 157-163cm, 163-168cm, 168-173cm, 173-178cm, 178cm), age (<12, 12-17, 18-34, 35-50, 51-65, >66 years old), biological sex (Male, Female) and race (Asian, Black, Hispanic, White).

TX Rate Disparity: The Weighted Mean Absolution Deviation (WMAD) of transplant rates of different patient groups, measured in transplants per patient-year (see Section 4.7.3 for a mathematical definition of WMAD). For example, transplant rate disparities by patient blood type would be the population-weighted WMAD of transplant rates for patients of type A, B, O and AB. These metrics are used as a system-wide measure of equity by different patient classifications, in the sense that smaller discrepancies between groups occur when different sub-populations are receiving transplants at similar rates. Weights are included in the calculation, based on the number of patients within each sub-population, so that the measure is not dominated by small groups whose transplant rates tend to be more volatile.

Waitlist Mortality: Patient deaths (expressed as a count, or normalized by total wait time) that occurred while a patient was on the waitlist or after they were removed from the waitlist due to being too ill to receive a transplant.

WLAUC: “Waitlist Area Under (the survival) Curve”, a survival analysis-based measure of a lung patient’s *waitlist urgency*. WLAUC measures the expected number of days the patient is expected to survive within the next year if they remain on the waitlist and do not receive a transplant. Its calculation is based on a patient’s clinical and demographic data at a given point in time, and is standard according to OPTN Policy [76].

4.7 Appendix

4.7.1 Robustness to surrogate model approximation error

As mentioned in Section 4.3.2, the optimal solution to the surrogate optimization $S(\pi)$ might not in fact be feasible for the instance $P(\pi)$, due to errors in the surrogate models' approximations. In particular, if surrogate model's prediction f_i for a constrained outcome was an overestimation of the true function's value B_i at the computed solution, then the requirement b_i may be violated. Of course, the magnitude of constraint violations, if any, is not expected to be large if the surrogate models are accurate.

Should the violations prove unacceptably large, we propose adding a safety term to the approximation inspired by robust optimization principles. Given a measure of the approximation error ϵ_i in f_i , e.g., the out-of-sample root mean squared error (rMSE), we make the requirement on outcome i stricter by replacing the constraints in $S(\pi)$ by:

$$f_i(\mathbf{x}; \boldsymbol{\theta}^i) + \gamma\epsilon_i \leq b_i$$

where γ is a hyper-parameter controlling the level of conservatism. Higher values of γ make the model more robust to approximation errors, but they also decrease the feasible space and might result in sub-optimality in the objective.

4.7.2 Reformulation of infeasible surrogate optimization

As mentioned in Section 4.3.2, it may be the case that a set of desired outcomes, given by the requirement vector \mathbf{b} , might be too stringent, rendering the surrogate optimization problem $S(\pi)$ infeasible. To address this possibility, we formulate a modified problem $\hat{S}(\pi)$, where slack variables allow the optimizer to selectively violate constraints, and seek to minimize the total amount of violation.

We first perform a pre-processing step that serves to standardize the requirement vector \mathbf{b} , so that constraint violations for outcomes with different numerical scales are

comparable. We use the sampled dataset $\{\mathbf{x}^n, \mathbf{y}^n\}_{n=1}^N$ from the *Sample Design* phase (Section 4.3.2), to compute lower and upper bounds for each outcome $i \in [d]$, namely $l_i := \min_n y_i^n$ and $u_i := \max_n y_i^n$. The bounds are informal in the sense that they do not constitute explicit constraints on an outcome, but are rather used to gain a sense of its range across the “representative” samples generated during the Sample Design phase. In particular, they allow us to rewrite requirements b_i in terms of a unitless quantity σ_i :

$$\sigma_i := \frac{b_i - l_i}{u_i - l_i}$$

Intuitively, σ_i expresses the user’s requirement on outcome i as a relative distance to its minimum/maximum. Typically, we expect σ_i to lie in the interval $[0, 1]$ as requirements should be achievable by some point in the domain; however, our approach does not depend on this as an assumption.

Since l_i and u_i can be computed during the sample design phase, and are therefore fixed at optimization time, $S(\pi)$ is exactly equivalent to:

$$\begin{aligned} \min_{\mathbf{x} \in \mathcal{X}} \quad & f_j(\mathbf{x}; \boldsymbol{\theta}^j) \\ \text{s.t.} \quad & f_i(\mathbf{x}; \boldsymbol{\theta}^i) \leq l_i + \sigma_i(u_i - l_i) \quad \forall i \neq j \end{aligned} \tag{S'(\pi)}$$

where all requirements b_i have been replaced by their relative counterparts σ_i . We emphasize that the optimization problem has not changed. In particular, if $S(\pi)$ is infeasible then so is $S'(\pi)$.

If this is the case, we can now introduce a modified formulation of $S'(\pi)$ whose optimal solution deviates from the given requirements in some minimum sense, and has a trivially non-empty feasible set:

$$\begin{aligned}
& \min_{\mathbf{x} \in \mathcal{X}, \mathbf{s} \geq 0} \sum_{i=1}^d s_i \\
& \text{s.t. } f_i(\mathbf{x}; \boldsymbol{\theta}^i) \leq l_i + (\sigma_i + s_i)(u_i - l_i) \quad \forall i \neq j \\
& \quad f_j(\mathbf{x}; \boldsymbol{\theta}^j) \leq l_j + (0 + s_j)(u_j - l_j)
\end{aligned} \tag{5}$$

Here we’ve introduced slack variables $\mathbf{s} \in \mathbb{R}_+^d$. Intuitively, s_i is a decision variable that is added to the relative requirement σ_i , allowing the optimizer to loosen any constraint at its discretion. The objective seeks to minimize the total amount of slack, i.e., the sum total deviation from the desired constraints. Because the relative requirements σ_i are on the comparable scale induced by the ranges (l_i, u_i) , minimizing the sum of slacks does not unduly prioritize outcomes with a larger magnitude.

The second constraint provides necessary incentive to minimize the primary objective j , albeit indirectly. Intuitively, the objective’s relative requirement is set to $\sigma_j = 0$, and also adjusted by some positive slack, denoting that the outcome should deviate minimally from the lower bound l_j that was observed over the entire dataset.

4.7.3 Surrogate modeling

As described in Section 4.3.2, the *Surrogate Modeling* phase of our methodology involves fitting Mixed Integer Linear Optimization (MILO) representable surrogate models to approximate the black-box function B . We use machine learning to predict each output of B individually, fitting a parametric approximation $f_i(\mathbf{x}; \boldsymbol{\theta}^i) \approx B_i(\mathbf{x})$ for each $i \in [d]$ using the sampled dataset from the *Sample Design* phase.

Concretely, in our lung allocation case study, we fit a model to predict each simulated outcome (e.g., number of waitlist deaths, average organ transport distance) as a function of the CAS scoring weights λ . In what follows we mathematically describe the hypothesis class we used in our implementation, explain the fitting procedure, and formulate the MILO representation of these functions.

Hypothesis Class

In our lung allocation implementation, we use a hypothesis class of functions we term *affine extremum functions*. These functions are the point-wise minimum or maximum of K affine transformations of the input. Mathematically, given an input $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^p$, we fit models of the form:

$$f_i(\mathbf{x}; \boldsymbol{\beta}) = \max_{k=1, \dots, K} \{\beta_{0,k} + \boldsymbol{\beta}_k^\top \mathbf{x}\} \quad \text{or} \quad f_i(\mathbf{x}; \boldsymbol{\beta}) = \min_{k=1, \dots, K} \{\beta_{0,k} + \boldsymbol{\beta}_k^\top \mathbf{x}\} \quad (4.1)$$

We refer to whether the left (max) or right (min) version is used as the function’s *sense*. This class of functions is piece-wise linear over \mathbf{x} (and therefore MILO-representable, as we show below), and its members are convex/concave in x depending on their sense. For simplicity in exposition, we refer to each affine transformation $\beta_{0,k} + \boldsymbol{\beta}_k^\top \mathbf{x}$ as a “piece” of the overall function.

The sense and number of pieces K are hyper-parameters, to be selected through cross-validation. For fixed sense and K , there are a total of $K \cdot (p+1)$ parameters that need to be estimated, consisting of a linear coefficient vector $\boldsymbol{\beta}_k \in \mathbb{R}^p$ and intercept $\beta_{0,k}$ for each of the K pieces. We use shorthand $\boldsymbol{\beta}$ to refer to a concatenation of all these parameters.

Fitting procedure

During model fitting, we estimate $\boldsymbol{\beta}$ using standard supervised machine learning techniques. We use the dataset generated during the Sample Design phase (Section 4.3.2) to create a dataset for each individual outcome, denoted by $\mathcal{D}_i = \{\mathbf{x}^n, y_i^n = B_i(\mathbf{x}^n)\}_{n=1}^N$. We randomly split \mathcal{D}_i into training $\mathcal{D}_i^{\text{train}}$, validation $\mathcal{D}_i^{\text{val}}$ and testing $\mathcal{D}_i^{\text{test}}$ sets, using an 80%-10%-10% split.

For a function of the form (4.1) with fixed K and sense, we estimate best-fit parameters on $\mathcal{D}_i^{\text{train}}$ by minimizing standard ℓ_2 loss:

$$\min_{\boldsymbol{\beta}} \sum_{n=1}^{|\mathcal{D}_i^{\text{train}}|} (y_i^n - f_i(\mathbf{x}^n; \boldsymbol{\beta}))^2$$

Both inputs and outputs are rescaled to have zero mean and unit standard deviation over the training set to improve numerical performance. Optimization is performed using (sub-)gradient descent with random restarts, starting the procedure from 10 randomly initialized $\boldsymbol{\beta}$ and picking the final iterate that achieved minimum training error. We implement fitting in the Julia programming language, version 1.1.0, computing gradients using the `ForwardDiff` package and using the gradient descent implementation of the `Optim` package.

We use cross-validation to select the best-performing hyper-parameters (K and sense). For each individual outcome, we search over a grid of K and sense, and select the model trained on $\mathcal{D}_i^{\text{train}}$ that achieved the maximum validation set R^2 on $\mathcal{D}_i^{\text{val}}$. We compute out-of-sample R^2 for the selected model on the test set $\mathcal{D}_i^{\text{test}}$ to evaluate the model’s true performance, and retrain on the full dataset \mathcal{D}_i for use in the optimization.

MILO formulation

In the *Optimization* phase (Section 4.3.2), we formulate problems $S(\pi)$ where (multiple) surrogate models of the form (4.1) appear in the objective and constraints, depending on the problem instance π . We present here the MILO formulation for a model with *maximum* sense (first option in Equation 4.1), and note that the other case is similar up to some signs. The full formulation of $S(\pi)$ is obtained by combining the surrogate formulations of all relevant outcomes.

Consider a trained surrogate $f_i(\mathbf{x}; \boldsymbol{\beta}) = \max_{k=1, \dots, K} \{\beta_{0,k} + \boldsymbol{\beta}_k^\top \mathbf{x}\}$ with fixed K . At optimization time, model parameters $\boldsymbol{\beta}$ are fixed to their trained values, while inputs \mathbf{x} are decision variables. Our goal is to model a decision variable $y \in \mathbb{R}$ so that, in any feasible solution, $y = f_i(\mathbf{x}; \boldsymbol{\beta})$. We use the following big-M formulation to linearize the function, introducing binary auxiliary variables $z_k \in \{0, 1\}$ for all $k \in [K]$:

$$\sum_{k=1}^K z_k = 1 \tag{4.2}$$

$$y \geq \beta_{0,k} + \boldsymbol{\beta}_k^\top \mathbf{x} \quad \forall k \in [K] \tag{4.3}$$

$$y \leq \beta_{0,k} + \boldsymbol{\beta}_k^\top \mathbf{x} + M(1 - z_k) \quad \forall k \in [K] \tag{4.4}$$

where M is some sufficiently large fixed value (see below). Intuitively, the binary variables z_k encode which of the K pieces achieves the maximum value for a given \mathbf{x} ; that is, in any feasible solution we have $z_k = 1$ if $\boldsymbol{\beta}_k^\top \mathbf{x} = \max_{k'} \boldsymbol{\beta}_{k'}^\top \mathbf{x}$ and 0 otherwise. Constraint 4.2 naturally enforces that one of the pieces must achieve the maximum. Constraints 4.3 ensure that y is lower bounded by the value of each piece, and is therefore greater than their maximum. Finally, Constraints 4.4 upper bound y to enforce exact equality to the maximum piece; that is, for $\{k : z_k = 1\}$ the upper bound matches the lower bound, while the remaining upper bounds are assumed to be even greater for large enough M .

For the formulation to be valid then, M must be selected so that, for *any* \mathbf{x} in the domain, the upper bounds for the non-maximum pieces do not conflict with the maximum lower bound. In our implementation, we compute a value for M during pre-processing as follows:

$$M = \max_{k,k'} \max_{\mathbf{x} \in \mathcal{X}} \{(\beta_{0,k} + \boldsymbol{\beta}_k^\top \mathbf{x}) - (\beta_{0,k'} + \boldsymbol{\beta}_{k'}^\top \mathbf{x})\}$$

In short, we compute the maximum difference between any two pieces over the entire domain \mathcal{X} . We note that, when \mathcal{X} is a continuous domain (as in our case), the inner optimization (i.e., for fixed k, k') is a simple linear program (LP). We can then obtain a suitable value for M by solving $(K^2 - K)$ linear optimization problems, one for each pair of pieces, and take the maximum objective values as M .

Derived models: Mean Absolute Deviation (MAD)

In our lung allocation implementation, certain outcomes of interest were defined as mathematical functions of other outcomes. Specifically, transplant rate disparities for a given patient classification (e.g., by patient blood type) were defined as the Weighted Mean Absolution Deviation (WMAD) of the individual transplant rates of each relevant subclass of patients (blood type A, B, AB and O), weighted by number of candidates in that subclass.

We found that modeling this dependence explicitly in the MILO formulations yielded better performance (i.e., more accurate surrogate optimization) than training a new surrogate model for the disparity metric. More concretely, instead of fitting an affine extremum model to directly predict a WMAD-based disparity metric, we create what we refer to as a *derived* surrogate model, whose prediction is obtained by first calculating transplant rate predictions for each subclass (from the corresponding surrogate models) and subsequently computing their WMAD.

Mathematically, suppose y_j , $j \in [J]$ are the predicted transplant rates for J classes of patient. The derived surrogate model that predicts transplant rate disparities among the classes is defined as:

$$y = \text{WMAD}(y_1, \dots, y_J) = \frac{1}{J} \sum_{j=1}^J w_j \left| y_j - \frac{1}{J} \sum_{j=1}^J w_j y_j \right| \quad (4.5)$$

where $|\cdot|$ is the absolute value operator, and w_j is a weight associated with the j 'th class (e.g., percentage of total population corresponding to the class). Weights are assumed to sum to one, i.e., $\sum_j w_j = 1$.

To use the derived model during optimization, our goal is to model a decision variable $y \in \mathbb{R}$ so that in any feasible solution $y = \text{WMAD}(y_1, \dots, y_J)$. At optimization time, the weights w_j are fixed, while $y_j \forall j \in [J]$ are decision variables corresponding to the outputs of other surrogate models (formulated, e.g., as in Appendix 4.7.3). Here, we use common techniques to linearize each absolute value in the sum of equation 4.5, and optimization can proceed via MILO [13].

4.7.4 Simulated outcomes

Table 4.1 lists the full set of simulated outcomes that were modeled, and therefore allowed to appear in the objective/constraints of our optimization tool during our collaboration with UNOS and the OPTN. More detailed definitions of the outcomes and related terminology are provided in Appendix 4.6. The table includes detailed results from the Surrogate Modeling phase; in particular, for each simulated outcome j it shows:

- The minimum and maximum value observed for the outcome over the training set. These numbers are intended to illustrate the range and scale of values for each outcome.
- The hyper-parameters of the surrogate model that was ultimately selected through cross-validation. Models are labeled by their hyper-parameters, namely whether they used Min/Max sense, and the number of pieces in the affine extremum class (see Appendix 4.7.3). For example, a “Min7” model indicates that the function was the point-wise minimum of 7 affine functions. Outcomes that were defined as the WMAD of other outcomes, i.e., all transplant rate disparity outcomes, are labeled as “Derived” (see Appendix 4.7.3).
- The out-of-sample R^2 and root mean squared error (rMSE) of the selected model.

As noted in Section 4.4.1, our surrogate models exhibit remarkably high out-of-sample R^2 's for all outcomes, ranging from 0.85-0.99 with an average of 0.96. Such high fidelity over the entire domain is key in enabling accurate surrogate optimization regardless of the user's desired requirement vector.

Outcome	l_j	u_j	Model	R2	rMSE
Overall Deaths #	904	1305	Max5	0.98	12
Waitlist Deaths #	399	847	Max5	0.98	12
Post-TX Deaths #	445	538	Min9	0.92	5
Overall Transplants #	3414	3496	Min5	0.94	4
Avg. Net Benefit at TX (days)	10.6	58.3	Min9	0.99	1.0
MLaT	36.0	44.1	Min3	0.98	0.3
MLaT Disparity - OPO	1.1	4.2	Min9	0.93	0.2
Med. Organ Transp. Distance (nm)	74.3	726.3	Max9	0.99	12.3
Avg. Organ Transp. Distance (nm)	132.4	849.7	Max9	0.98	15.7
Organ Transp. Cost (% of current policy)	69.9%	216.6%	Max5	0.98	3.0%
Organs Flown (% of total)	49.7%	96.9%	Max5	0.99	0.9%
Overall TX Rate	1.25	1.57	Max5	0.98	0.01
TX Rate Disparity - OPO	0.27	1.35	Max5	0.98	0.03
TX Rate - Blood Type A	0.32	1.36	Max5	0.97	0.04
TX Rate - Blood Type AB	0.27	1.39	Max5	0.97	0.04
TX Rate - Blood Type B	0.88	1.48	Max5	0.89	0.03
TX Rate - Blood Type O	1.24	6.81	Max9	0.98	0.12
TX Rate Disparity - Blood Type	0.03	3.15	Derived	0.98	0.06
TX Rate - Age 0-11	0.83	6.28	Min7	0.91	0.19
TX Rate - Age 12-17	1.64	55.26	Min7	0.94	2.65
TX Rate - Age 18-34	1.25	1.96	Max9	0.97	0.02
TX Rate - Age 35-50	1.24	2.00	Max12	0.98	0.02
TX Rate - Age 51-65	0.92	1.49	Max5	0.97	0.02
TX Rate - Age 66+	1.24	1.95	Max5	0.92	0.04
TX Rate Disparity - Age Group (18+)	0.10	0.37	Derived	0.95	0.01
TX Rate - Diagnosis Group A	0.35	1.54	Max5	0.96	0.05
TX Rate - Diagnosis Group B	0.36	1.71	Max5	0.97	0.04
TX Rate - Diagnosis Group C	1.38	2.89	Max7	0.97	0.04
TX Rate - Diagnosis Group D	1.62	3.46	Max7	0.98	0.04
TX Rate Disparity - Diagnosis Group	0.07	1.43	Derived	0.98	0.03
TX Rate - Female	0.97	1.26	Max7	0.99	0.01
TX Rate - Male	1.50	2.01	Min7	0.94	0.02
TX Rate Disparity - Sex	0.22	0.40	Derived	0.85	0.01
TX Rate - Asian	1.39	2.50	Max5	0.95	0.05
TX Rate - Black	1.15	1.67	Max7	0.97	0.01
TX Rate - Hispanic	1.47	2.69	Max5	0.98	0.03
TX Rate - White	1.17	1.59	Max5	0.96	0.02
TX Rate Disparity - Race	0.02	0.22	Derived	0.99	0.00
TX Rate - Height 0-157cm	0.96	2.14	Max7	0.97	0.04
TX Rate - Height 157-163cm	0.81	1.13	Max7	0.97	0.01
TX Rate - Height 163-168cm	0.79	1.37	Max5	0.95	0.03
TX Rate - Height 168-173cm	0.98	1.53	Max7	0.94	0.02
TX Rate - Height 173-178cm	1.38	1.99	Max7	0.91	0.04
TX Rate - Height 178+cm	1.44	4.17	Min5	0.94	0.15
TX Rate Disparity - Height Group	0.12	0.86	Derived	0.95	0.03

Table 4.1: Surrogate modeling results for simulated lung allocation outcomes. Acronyms are defined in the Glossary.

4.7.5 Additional tradeoffs

In Figure 4-5 we provide two additional plots to evaluate the tradeoff between waitlist mortality and transport burden. The plots exactly parallel Figure 4-3 from the main paper (see Section 4.4.2 for details), except that two different transport burden metrics are used in lieu of median organ transport distance in the x-axis: (top) estimated organ transport cost as a percentage of current policy, and (bottom) percentage of organs expected to fly. The non-linear (but monotonic) relationship between the transport metric and placement efficiency in the optimized policies is plotted on the secondary x-axis at the top. Of note, we observe qualitatively similar results in terms of diminishing returns to reducing waitlist mortality as placement efficiency is decreased beyond 10%.

Finally, in Figure 4-6 we plot the worst-case guaranteed transplant rate for adolescent patients aged 12-17 years as a function of the pediatric priority weight, paralleling Figure 4-4 from the main paper (see Section 4.4.3 for details). We note that transplant rates are very high for this population due to its relatively small size, which results in short active wait times when they are highly prioritized. A pediatric weight of 2% is already enough to guarantee a greater transplant rate than under current policy.

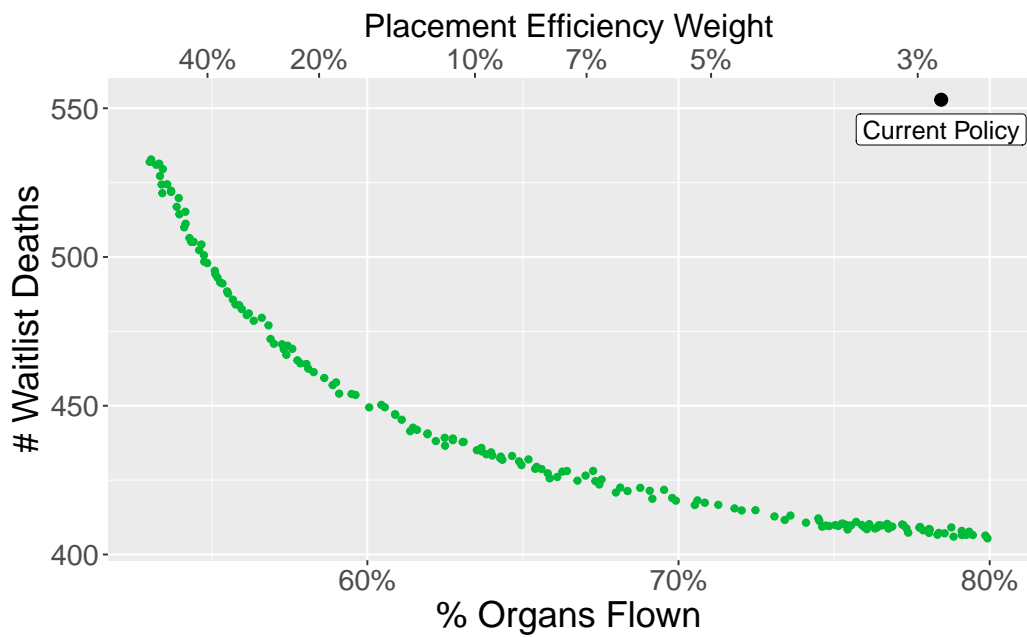



Figure 4-5: Tradeoff of waitlist mortality and placement efficiency in continuous distribution.




Figure 4-6: Guaranteed pediatric transplant rate for patients aged 12-17 years vs. pediatric access weight.

Chapter 5

Balancing student preferences and transport cost in school choice

5.1 Introduction

In the fall of 2020, nearly 50 million students were enrolled in public schools across the United States, representing more than 15% of the country's population [67]. Given the extraordinary effect that education has on student welfare, often well beyond one's school years, the operation of public school districts is widely considered one of the most impactful areas for national public policy in the country. Districts regularly face a range of challenging problems, from how to design stimulating curricula that address the needs of a diverse student body, to ensuring equity in access to education for marginalized and disadvantaged communities.

Crucially, constraints on public education spending often mean that policymakers need to carefully prioritize between different objectives. This is perhaps nowhere more evident than in the problem of school assignment, wherein districts must decide which students will attend which school every year. An assignment mechanism should accomplish multiple objectives: address students' individual learning needs, satisfy family preferences for particular schools, create racially and socio-economically diverse learning environments, and more. At the same time, it must also consider significant operational constraints like staffing requirements, school capacity, and transportation

cost. Assigning students to schools closer to their home, for example, can produce significant savings for districts that have to bus students to and from school every day – savings that can then be reinvested in improving other aspects of the system. Indeed, keeping transportation costs manageable is a primary focus of many districts’ assignment mechanisms, as for example in Boston, where transportation accounts for over 10% of the district’s \$1.3 billion budget [19].

Recent years have seen many districts implement “school-choice”-type assignment mechanisms, whereby student express preferences for (and can be assigned to) schools across the district. The goal is to create a system that is both more responsive to students’ needs and more equitable, particularly as the historical precedent of basing assignments on students’ neighbourhoods has been shown to propagate racial and socio-economic inequities [52]. Advocates of these systems also argue that schools respond constructively to the competition induced by a preference-based system, raising their achievement and productivity to attract better students [40].

School-choice mechanisms typically rely on some variant of the celebrated Deferred Acceptance (DA) algorithm of Gale & Shapley [33]. Given a (possibly incomplete) set of student preferences for schools, and school priorities for students, the DA algorithm guarantees a desirable property termed *stability*; namely, that there exists no student-school pairs where both agents prefer each other over their current assignments (known as a *blocking pair*). This reflects a notion of procedural fairness in how matches are formed, in that there is no possible change to the assignment that would make both the student and school happier. The problem of stable matching has been widely studied in the literature, both theoretically and empirically, with applications ranging from medical residency matching to school choice [86].

Despite its theoretical elegance and popularity in practice however, stability alone does not suffice to address all of a school district’s objectives. In realistic settings, districts often face hard constraints, e.g., siblings that must be assigned to the same school, or legally-mandated diversity quotas, that are not easy to incorporate in a DA algorithm. The same is true of different policy objectives, such as minimizing racial disparities in access to top schools or minimizing transportation cost, since DA does

not explicitly model an objective function to select among multiple possible stable solutions. Some objectives may be indirectly optimized by engineering schools’ priorities in the DA algorithm—for example, ranking nearby students higher in order to decrease transportation requirements—however, such approaches tend to be heuristic and hard to adapt to multiple objectives. As a result, districts have typically resorted to elaborate choice mechanisms, with varying degrees of success in achieving the stated policy objectives in practice [77].

In this work, we apply global mixed-integer optimization (MIO) to empirically study the problem of school choice. MIO formulations provide a natural way to model (one or more) assignment objectives, as well as external constraints, while maintaining the notion of stability. To this end, we formulate stable matching as a MIO and develop a custom pre-solve algorithm to help scale it. We then perform experiments to characterize the space of stable solutions. In particular, we generate synthetic preferences for students in the Boston Public Schools (BPS) district, incorporating both distance to school and school quality in students’ rankings. We then optimize over the space of stable solutions to minimize a proxy of bus transportation cost, namely the sum of distances from students’ homes to their assigned schools (henceforth, transport distance).

Our experiments suggest that stability is a particularly constraining property in this setting, and does not allow for significant optimization of transport distance. Even in the largest-sized problems we consider (16,255 students and 92 schools), we find that stable solutions most often number in the single digits, and differ by no more than 0.13% in their objective values. This observation motivates a second set of experiments, where we relax stability by allowing a fixed number of blocking pairs in the matching formulation. Here, we explore through tradeoff curves how the optimal transport distance varies as we allow more blocking pairs, and observe that our hypothetical district can decrease transport distance by 12.7% (respectively 28.2%) on average if it allows 1% (resp. 5%) of student-school pairs to be blocking. These results suggest that, while full stability is quite constraining, it does not need to be relaxed significantly to realize considerable gains in alternative objectives.

5.2 The school choice problem

In this section, we introduce a mathematical model of the school choice problem as originally formulated by Gale & Shapley [33]. We then summarise key results from the literature that are relevant to this work.

5.2.1 Mathematical model

We first present the mathematical model of the stable school assignment problem. A matching problem is defined by a set of students \mathcal{I} and a set of schools \mathcal{J} , each forming one side of the two-sided matching market. A student can be assigned to no more than one school, and each school j is associated with a maximum capacity $C_j \in \mathbb{N}$ of students it can accommodate. We will often refer to students and schools as agents in the market.

All agents provide ranked preference lists for their possible matches on the other side. We encode these by matrices of real-valued scores, namely values $p_{ij} \in \mathbb{R}$ and $q_{ij} \in \mathbb{R}$ for all $(i, j) \in \mathcal{I} \times \mathcal{J}$, denoting respectively student i 's utility from being assigned to school j and school j 's utility from being assigned student i . Agents' rankings need not be exhaustive, in which case we refer to the problem as having *incomplete preferences* and set $p_{ij} = q_{ij} = -\infty$ for any inadmissible pair. Finally, we assume that all preferences are *strict*, which is to say that no agent is indifferent between two of their ranked choices.

For convenience in exposition, we define the following notation: the set of admissible student-school pairs is denoted by $\mathcal{S} = \{(i, j) \in \mathcal{I} \times \mathcal{J} : p_{ij} \neq -\infty, q_{ij} \neq -\infty\}$. We use $\mathcal{P}(i) = \{j \in \mathcal{J} : p_{ij} \neq -\infty\}$ to refer the set of all schools ranked by student i , also referred to as i 's valid schools, and similarly $\mathcal{Q}(j) = \{i \in \mathcal{I} : q_{ij} \neq -\infty\}$ for school j 's valid students. Finally, $\mathcal{P}_{<}(i, j) = \{k \in \mathcal{P}(i) : p_{ik} < p_{ij}\}$ is the set of schools that student i prefers *less* than school j , and $\mathcal{Q}_{<}(i, j) = \{k \in \mathcal{Q}(j) : q_{kj} < q_{ij}\}$ the set of students that school j prefers less than student i .

A matching M is many-to-one mapping from \mathcal{I} to \mathcal{J} . If a student-school pair $(i, j) \in \mathcal{I} \times \mathcal{J}$ is matched in M we write that $M(i) = j$ and $i \in M^{-1}(j)$. If a student i

is not matched to any school in M , we write that $M(i) = 0$. A matching is considered valid if all of the following hold: (1) every student i is either unmatched or matched to a valid school, i.e., $M(i) \in \mathcal{P}(i) \cup \{0\}$; (2) every school j is assigned only valid students, i.e., $M^{-1}(j) \subseteq \mathcal{Q}(j)$; (3) no school exceeds its capacity, i.e., $|M^{-1}(j)| \leq C_j$. The sets $\mathcal{I}_u(M) = \{i \in \mathcal{I} : M(i) = 0\}$ and $\mathcal{J}_u(M) = \{j \in \mathcal{J} : |M^{-1}(j)| < C_j\}$ denote the set of students who are unmatched and the set of schools that are strictly under capacity in M respectively.

We next build up the notion of a γ -stable matching, which we find to be a conceptually useful generalization of Gale & Shapley's original definition of stability [33].

Definition 1 (Student unhappiness) *Given a matching M , we say that student $i \in \mathcal{I}$ is unhappy with respect to school $j \in \mathcal{J}$ if they are not matched to each other and one of the following is true:*

- a) i is unmatched, i.e., $i \in \mathcal{I}_u(M)$.
- b) i prefers j over their current match, i.e., $M(i) \in \mathcal{P}_<(i, j)$.

Definition 2 (School unhappiness) *Given a matching M , we say that school $j \in \mathcal{J}$ is unhappy with respect to student $i \in \mathcal{I}$ if they are not matched to each other and either of the following is true:*

- a) j is strictly under capacity, i.e., $j \in \mathcal{J}_u(M)$.
- b) j is matched with at least one student they prefer less than i , i.e., $M^{-1}(j) \cap \mathcal{Q}_<(i, j) \neq \emptyset$.

Intuitively, being unhappy with respect to a possible match means that, all other things equal, one would prefer an outcome where the match happens over their current outcome. The motivation for this definition is natural if one takes a game theoretic perspective; a situation where two parties are mutually unhappy is “unstable” in the sense that the two parties have an incentive to disrupt the overall assignment mechanism by unilaterally accepting each other.

Definition 3 (Blocking pair) *Given a matching M , we say that a student-school pair $(i, j) \in \mathcal{S}$ is blocking if both parties are unhappy with respect to each other.*

We use $N_b(M)$ to denote the number of blocking pairs in M , and $P_b(M)$ to denote the same as a percentage of the number of admissible pairs in \mathcal{S} , i.e., $P_b(M) = N_b(M)/|\mathcal{S}|$.

Definition 4 (γ -Stability) *A matching M is said to be γ -stable if $P_b(M) \geq \gamma$.*

Note that a 1-stable matching exactly corresponds to Gale & Shapley’s original formulation of stability [33], which requires that no pairs are blocking. We introduce the more general definition here because it provides a principled, tunable mechanism to relax stability, as we discuss in Section 5.3.1. For convenience, we still use the term *stable* to refer to 1-stable solutions. Finally, we use \mathcal{M}_γ to denote the space of all γ -stable matchings, and note that the size of this set is non-increasing in γ .

5.2.2 Literature review

As mentioned, the notion of stability in matching markets was first proposed in the seminal work of Gale & Shapley [33], who refer to the above as the *college admissions problem*. In the original paper, the authors prove the existence of at least one stable solution for any preference set (as long as there are no extrinsic constraints), and propose an algorithm that is guaranteed to find such a solution. Since then, the college admissions problem and its variants have been studied extensively in the literature, both theoretically and empirically, with perhaps the most notable real-world application being the matching of graduating medical students to residency programs [84]. In this section, we summarise key results from the literature that are relevant to our study, and refer the reader to [86] for a more extensive review.¹

The first key result is, of course, Gale & Shapley’s iterative Deferred Acceptance (DA) algorithm itself, which comes in two variants. The *student-proposing* DA algorithm works as follows: initially, each student applies to their top-ranked school, and each school places its top C_j (i.e., its capacity) applicants on a waitlist, rejecting the rest. Rejected applicants then apply to their second-ranked school, and schools update their waitlists to include the top C_j students among either new applicants or their previous waitlist. The remaining students are rejected (including some who were

¹Unless otherwise noted, references for all of the stated results can be found in [86].

previously on a waitlist) and apply to their next-ranked school that has not already rejected them. The process repeats until there are no more rejections, at which point every student is either waitlisted or has been rejected by all of their ranked schools (in which case they will remain unmatched). The resulting waitlists form a provably stable matching.

The same is true of the second, *school-proposing* variant of the DA algorithm. Here, the procedure is inverted; at each iteration, schools offer admission to any top-ranked students who have not already rejected them. Students tentatively accept an offer if it is better than the one they currently hold, in which case they reject their previous offer and the corresponding school must find a new student to fill the position. Again, the process repeats until there are no more rejections, at which point the standing offers form a provably stable matching.

Interestingly, the two algorithms may (and often do) produce different stable matchings, each optimal for the proposing side. This result is again due to Gale & Shapley, who in their original paper show that: (1) there are matching instances for which the algorithms produce different solutions; (2) the student-proposing algorithm gives a *student-optimal* assignment, meaning that all students are assigned to their top choice among all schools they would possibly get in any stable assignment; and (3) the school-proposing algorithm is similarly *school-optimal*.

The non-uniqueness of stable solutions serves as a key motivation for our work. We seek to characterize the space of stable solutions in realistic application settings, and explore the extent to which they can accommodate different objectives, e.g., maximizing equity in access or reducing transportation cost. Moreover, we ask whether relaxing the stability requirement, by allowing a small number of blocking pairs, can improve other policy objectives. The formulation of stable matching as a mixed-integer linear optimization (MILO) problem is crucial here, as it allows us to flexibly model different objectives and constraints.

Optimization techniques have been previously studied in the context of school choice. Baiou & Balinski provide a description of the stable admissions polytope in terms of linear constraints [4], and Sethuraman et al. analyze its geometric prop-

erties [87]. Bodoh-Creed [15] use optimization to investigate the tensions between student welfare, encouraging neighbourhood schools, and diversity under stability; however, their approach analyzes system-wide outcomes using a large-scale continuum approximation that does not result in an assignment mechanism per se. Delorme et al. develop pre-processing heuristics to scale an MILO formulation of the stable admissions problem [25]. Finally, optimization techniques have been used to design assignment mechanisms that optimize for particular objectives other than stability [2, 7, 52, 89].

The final result we wish to highlight is the so-called Rural Hospitals Theorem (RHT), first proved in [85]:

Theorem 1 (Rural Hospitals Theorem) *When all preferences are strict, the set of matched students and filled positions is the same in every stable matching. Furthermore, any school that is strictly under capacity in some stable matching is matched with exactly the same set of students in every stable matching.*

This remarkable result should already provide some intuition for how constraining a property stability really is. In particular, the second part of the theorem implies that a large number of students (i.e., all those assigned to an under capacity school) only have one possible assignment option under any stable solution. Notably, this property will form the basis of the pre-solving methodology that allows us to scale global optimization to realistic problem sizes (Section 5.3.2).

5.3 Optimization methodology

We next present a novel formulation of the stable matching problem in terms of binary decision variables and linear constraints, which will allow us to optimize over the space of stable solutions using Mixed Integer Linear Optimization (MILO). Our formulation readily generalizes to the case of γ -stability, and can incorporate a wide range of objective functions that may be of interest to school districts. We also introduce here a custom pre-solve algorithm to help scale the formulation to realistic problem sizes.

5.3.1 An MIO formulation for stability

We first introduce the MILO formulation for the case of 1-stable matchings. We use this particular formulation over others commonly found in the literature, e.g., [4, 46], because it readily generalizes to our notion γ -stability, as we discuss next. The basic formulation uses five sets of binary decision variables:

- $x_{ij} = \{\text{student } i \text{ is matched to school } j\}$ for all $(i, j) \in \mathcal{S}$.
- $y_j = \{\text{school } j \text{ is strictly under capacity}\}$ for all $j \in \mathcal{J}$.
- $z_i = \{\text{student } i \text{ is unmatched}\}$ for all $i \in \mathcal{I}$.
- $u_{ij} = \{\text{student } i \text{ is unhappy w.r.t. school } j\}$ for all $(i, j) \in \mathcal{S}$.
- $w_{ij} = \{\text{school } j \text{ is unhappy w.r.t. student } i\}$ for all $(i, j) \in \mathcal{S}$.

where the definition of unhappiness is as in Section 5.2.1. The space of stable matchings is then described by the following set of linear constraints:

$$\sum_{j \in \mathcal{P}(i)} x_{ij} \leq 1, \quad \forall i \in \mathcal{I} \quad (5.1)$$

$$\sum_{i \in \mathcal{Q}(j)} x_{ij} \leq C_j, \quad \forall j \in \mathcal{J} \quad (5.2)$$

$$z_i = 1 - \sum_{k \in \mathcal{P}(i)} x_{ik}, \quad \forall i \in \mathcal{I} \quad (5.3)$$

$$y_j \leq C_j - \sum_{i \in \mathcal{Q}(j)} x_{ij} \leq C_j y_j, \quad \forall j \in \mathcal{J} \quad (5.4)$$

$$u_{ij} = z_i + \sum_{k \in \mathcal{P}_{<}(i,j)} x_{ik}, \quad \forall (i, j) \in \mathcal{S} \quad (5.5)$$

$$w_{ij} \geq x_{kj}, \quad \forall (i, j) \in \mathcal{S}, k \in \mathcal{Q}_{<}(i, j) \quad (5.6)$$

$$w_{ij} \geq y_j, \quad \forall (i, j) \in \mathcal{S} \quad (5.7)$$

$$w_{ij} \leq y_j + \sum_{k \in \mathcal{Q}_{<}(i,j)} x_{kj}, \quad \forall (i, j) \in \mathcal{S} \quad (5.8)$$

$$u_{ij} + w_{ij} \leq 1, \quad \forall (i, j) \in \mathcal{S} \quad (5.9)$$

Constraints 5.1 require that each student i is assigned to at most one school (they may of course be unassigned), while constraints 5.2 require that each school do not exceed its allotted capacity. Constraints 5.3 enforce that $z_i = 1$ if and only if student i is unmatched, and similarly 5.4 require that $y_j = 0$ if and only if school j is exactly at capacity. Constraints 5.5 enforce the definition of u_{ij} , since the right-hand size equals 1 either when student i is unassigned ($z_i = 1$), or they are assigned to a school k they prefer less than j . This disjunction can be modeled compactly with an equality because it is not possible for both student unhappiness conditions to hold at the same time. Constraints 5.6 – 5.8 perform a similar function for w_{ij} , but cannot be modeled with a single equality because a school can be assigned multiple students they prefer less than i . Instead their combination enforces that $w_{ij} = 0$ if and only if school j is exactly at capacity (i.e., $y_j = 0$) and it is not matched with any student k that it prefers less than i . Finally, constraints 5.9, which we call *stability constraints*, enforce that student i and school j do not form a blocking pair.

We comment here that this formulation has a large number of variables and constraints, and it is not a priori clear that even state-of-the-art MILO solvers could scale to problems of realistic size. Concretely, the number of variables scales as $O(|\mathcal{S}|)$, while the number of constraints is driven by 5.6, and is $O(\sum_{j \in \mathcal{J}} |\mathcal{Q}(j)|^2)$.² As we discuss in the Section 5.3.2, however, the size of the problem can be significantly reduced, so that solving problems with tens of thousands of students and hundreds of schools to optimality is possible.

We also note that the formulation can be readily extended to describe the space of γ -stable solutions. Doing so involves adding new binary variables α_{ij} for each $(i, j) \in \mathcal{S}$ that encode whether a given pair is blocking, and replacing constraint 5.9 by the following:

²This follows by a counting argument: consider an admissible pair $(i, j) \in \mathcal{S}$, and suppose that student i is ranked n 'th on school j 's preference list. Then i appears as a less-preferred alternative for $n - 1$ of school j 's admissible students. Hence, the total number of type 5.6 constraints for school j is $\sum_{n=1}^{|\mathcal{Q}(j)|} (n - 1) = O(|\mathcal{Q}(j)|^2)$.

$$u_{ij} + w_{ij} \leq 1 + \alpha_{ij}, \quad \forall (i, j) \in \mathcal{S} \quad (5.10)$$

$$\sum_{(i,j) \in \mathcal{S}} \alpha_{ij} \leq \lfloor (1 - \gamma)|\mathcal{S}| \rfloor \quad (5.11)$$

Intuitively, constraint 5.10 allows (i, j) to be a blocking pair – that is for both the student and school to be unhappy – if $\alpha_{ij} = 1$, while constraint 5.11 enforces that no more than $(1 - \gamma)\%$ of pairs are blocking.

In practice, it is prudent (though not strictly necessary) to include the constraint:

$$\sum_{(i,j) \in \mathcal{S}} x_{ij} \geq N \quad (5.12)$$

where N is a minimum number of students that should be matched to a school, given as input. We note that, in the basic formulation with $\gamma = 1$, this constraint is either trivially satisfied or trivially infeasible by virtue of the Rural Hospitals Theorem, which guarantees that the number of matched students is the same in all stable solutions. The same is not true when we relax stability, and so we might ask that at least a certain number of students be matched in the final solution. This would be desirable in the case where the objective function incentivizes leaving students unmatched, as is the case for our experiments minimizing total distance between students and their assigned schools. For the purposes of our study, we set N to equal the (unique) number of students matched when $\gamma = 1$, readily available by using a DA algorithm to compute a stable solution.

Finally, we note that the above formulation allows us to express a range of useful objectives as linear functions, which can then be optimized using global MILO solvers. In our experiments, we use the general form:

$$\min \sum_{(i,j) \in \mathcal{S}} c_{ij} x_{ij} \tag{5.13}$$

to minimize total driving distance between students and their assigned schools, by setting c_{ij} to equal the distance between i 's home and school j . The same form could be used to maximize some notion of assignment utility; for example, setting $c_{ij} = n$, where n is school j 's rank on i 's preference list, would produce a student-optimal matching in the Gale & Shapley sense (see Section 5.2.2). One could also minimize the number of unmatched students ($-\sum_{ij} x_{ij}$), under capacity schools ($\sum_j y_j$), blocking pairs ($\sum_{ij} \alpha_{ij}$), or weighted sum of any of the above.

5.3.2 Custom pre-solve algorithm

In pre-solving for combinatorial optimization, the goal is to improve computational performance of solvers by inferring the optimal values for certain variables or eliminating unnecessary constraints according to problem structure. Given the size of formulation 5.1 – 5.9, this is necessary for us to scale our experiments to realistic problem sizes, which can have tens of thousands of students and hundreds of schools. In this section, we describe the pre-solve algorithm we developed for the case of fully stable matching ($\gamma = 1$). Possible extensions to general γ are discussed in Section 5.5.

At a high-level, our pre-solver begins by running the DA algorithm to produce a (possibly suboptimal) stable solution to the given matching problem. It relies on the Rural Hospitals Theorem (RHT) to infer a set of assignments that are guaranteed to appear in any stable solution, and then iterates a series of constraint propagation rules that may force further assignments, or eliminate certain student-school pairs from contention. As a result, we are left with an “induced” subproblem consisting of: (i) a subset of students that still need be matched; (ii) a subset of schools that still have positions to fill; (iii) a reduced set of admissible pairs for those students and schools; (iv) a reduced set of potential blocking pairs for which stability constraints are needed. The induced problem is typically much smaller than in the original, and

can be passed to an MILO solver to determine the optimal matching according to some objective.

Step 1: Using the Rural Hospitals Theorem

Given access to a stable solution M_{DA} , readily available by running either variant of Gale & Shapley's DA algorithm, the RHT forces the matching outcomes of a subset of students; namely, those who are either unmatched in M_{DA} , or matched to a school that is strictly under capacity. This first pre-solving step is detailed in Algorithm 1, which we describe next.

The first part of the RHT states that the set of matched students and filled positions is the same in every stable solution. We can therefore immediately fix $z_i = 1$ for any student that is unmatched under M_{DA} , as well as all $x_{ij} = 0$ for those students. Analogously, we can fix $y_j = 1$ for any school that did not reach its full capacity under M_{DA} , as the unfilled positions must remain so in any stable solution. Conversely, all matched students under M_{DA} must have $z_i = 0$ (though we may not be able to infer an exact assignment) and all remaining schools must have $y_j = 0$.

The second part of the theorem, significantly more powerful, states that all schools that did not fill their positions in M_{DA} will receive *exactly* the same set of students in all stable solutions. We can therefore fix $x_{ij} = 1$ for any students assigned to said schools, and also $x_{ij} = 0$ for the rest. In other words, assignment outcomes are fully determined for this subset of schools.

Thus, by the end of Algorithm 1, we need only optimize over the set of schools that filled all positions, and students who were assigned to them. Moreover, we can strengthen our formulation by the following straightforward corollary of the RHT.

Algorithm 1 `fix_rht_variables(x, y, z)`

Input optimization variables (x, y, z) as per Section 5.3.1

Output Fixes the values of some (x, y, z) variables as implied by the RHT.

$M_{DA} \leftarrow \text{run_deferred_acceptance}()$ # produces a stable solution

for i in \mathcal{I} **do** # for each student

if $i \in \mathcal{I}_u(M_{DA})$ **then** # i is unmatched

`fix($z_i = 1$)`

`fix($x_{ij} = 0$)` $\forall j \in \mathcal{P}(i)$

else # i is matched

`fix($z_i = 0$)`

end if

end for

for $j \in \mathcal{J}$ **do** # for each school

if $j \in \mathcal{J}_u(M_{DA})$ **then** # j is under capacity

`fix($y_j = 1$)`

`fix($x_{ij} = 1$)` $\forall i \in M_{DA}^{-1}(j)$

`fix($x_{ij} = 0$)` $\forall j \in \mathcal{Q}(j) \setminus M_{DA}^{-1}(j)$

else # j is at capacity

`fix($y_j = 0$)`

end if

end for

Corollary 1 *For any student that was matched in M_{DA} but not to an under capacity school, constraint 5.1 can be changed to an exact equality. Moreover, for any school that was at capacity in M_{DA} , constraint 5.2 can be changed to an exact equality and constraint 5.4 eliminated.*

$$\sum_{j \in \mathcal{P}(i)} x_{ij} = 1, \quad \forall i \text{ s.t. } M_{DA}(i) \neq 0, M_{DA}(i) \notin \mathcal{J}_u(M_{DA}) \quad (5.14)$$

$$\sum_{i \in \mathcal{Q}(j)} x_{ij} = C_j, \quad \forall j \notin \mathcal{J}_u(M_{DA}) \quad (5.15)$$

Step 2: Propagating constraint implications

Given the assignments made as part of the previous step, we are able to make additional deductions about variable values by considering problem structure. We do so through *domain propagation*, a common technique in constraint programming for combinatorial optimization [83]. Algorithm 2 presents a set of rules, each associated with a given constraint in our formulation, that analyzes all variables values inferred so far to propose new variables values to fix. Each rule is applied successively until no rule produces a new deduction, at which point we say that the algorithm has converged to an irreducible subproblem.

In particular, note that any variable x_{ij} that has been fixed to 1 fully determines student i 's outcome, and reduces school j 's capacity by one. In addition, any x_{ij} that has been fixed to 0 implies that the corresponding student-school pair is, for all intents and purposes, inadmissible and can be removed from \mathcal{S} . The following logical rules can then be used to deduce additional variable values:

1. By Corollary 1, a student that has only one admissible school must be assigned to it. Similarly, a school whose remaining capacity equals the number of its admissible students must be assigned all of those students.
2. A student i that is forcibly unmatched (by the RHT) or assigned to a school k they prefer less than j must be unhappy w.r.t. to j ($u_{ij} = 1$). Conversely, if

they must be matched (by Corollary 1) and none of their admissible schools is preferred less than j , they will never be unhappy w.r.t. j ($u_{ij} = 0$).

3. Analogously, a school j that is forced to be under capacity (by the RHT) or assigned any student k they prefer less than i must be unhappy w.r.t. i ($w_{ij} = 1$). Conversely, if they must fill all positions (by Corollary 1) and none of their admissible students is preferred less than i , they will never be unhappy w.r.t. i ($w_{ij} = 0$).
4. To avoid forming a blocking pair, any student i that is unhappy w.r.t. to school j implies that j cannot be unhappy w.r.t. to i ($w_{ij} = 0$). This in turn implies that j cannot be assigned any student k they prefer less than i ($x_{kj} = 0$).
5. Conversely, any school j that is unhappy w.r.t. to student i implies that i cannot be unhappy w.r.t. school j ($u_{ij} = 0$). This in turn implies that i cannot be assigned to any school k they prefer less than j ($x_{ik} = 0$).

Note that these rules interact in a way that may cause “cascades” of inferences. For example, a pair (i, j) that is eliminated because it would make j unhappy (4) may force i 's assignment to a different school j' (1), which in turn forces j' to be unhappy with respect to a different student i' (3), and so forth. Only once all rules fail to produce a new deduction does the pre-solve algorithm terminate.

Step 3: Extracting the induced subproblem

At this point, we can extract a partial solution to the problem, consisting of those students that must remain unmatched by the RHT (i.e., $x_{ij} = 0$ for all j), and those whose matching has been determined after pre-solving (i.e., $x_{ij} = 1$ for some j). Certain schools' outcomes may also have been fully determined, either because they will never fill their remaining capacity (by the RHT), or because the inferred assignments exhaust their capacity.

As a result, we are left with a reduced set of students $\bar{\mathcal{I}} \subseteq \mathcal{I}$ and schools $\bar{\mathcal{J}} \subseteq \mathcal{J}$ whose outcomes have yet to be fully determined. In addition, by eliminating any

Algorithm 2 `propagate_implications`(x, y, z, u, w)

Input optimization variables (x, y, z, u, w) as per Section 5.3.1

Output Fixes the values of some (x, y, z, u, w) variables by propagation rules.

Note 1: In what follows, preference information (\mathcal{P} s and \mathcal{Q} s) exclude pairs found to be inadmissible in earlier iterations (i.e. when x_{ij} is known to be 0). Similarly, C_j denotes remaining capacity after subtracting the number of students known to be assigned to j .

Note 2: Conversely, the set \mathcal{S} includes all possible pairs, even those deemed inadmissible, since students and schools whose assignments have been fixed can still be unhappy (and therefore form blocking pairs) depending on other assignments.

while new values are still being inferred **do**

for $i \in \mathcal{I}$ **do** # for each student

$$\mathcal{P}(i) = \{k\} \implies \mathbf{fix}(x_{ik} = 1) \quad (\text{by 5.14})$$

$$x_{ik} = 1 \implies \mathbf{fix}(x_{ij} = 0) \quad \forall j \in \mathcal{P}(i), j \neq k \quad (\text{by 5.14})$$

end for

for $j \in \mathcal{J}$ **do** # for each school

$$|\mathcal{Q}(j)| = C_j \implies \mathbf{fix}(x_{kj} = 1) \quad \forall k \in \mathcal{Q}(j) \quad (\text{by 5.15})$$

$$C_j = 0 \implies \mathbf{fix}(x_{ij} = 0) \quad \forall i \in \mathcal{Q}(j) \quad (\text{by 5.15})$$

end for

for (i, j) in \mathcal{S} **do** # for each admissible pair

$$z_i = 1 \text{ or } x_{ik} = 1 \text{ for some } k \in \mathcal{P}_{<}(i, j) \implies \mathbf{fix}(u_{ij} = 1) \quad (\text{by 5.5})$$

$$z_i + \sum_{k \in \mathcal{P}_{<}(i, j)} x_{ik} = 0 \implies \mathbf{fix}(u_{ij} = 0) \quad (\text{by 5.5})$$

$$u_{ij} = 0 \implies \mathbf{fix}(x_{ik} = 0) \quad \forall k \in \mathcal{P}_{<}(i, j) \quad (\text{by 5.5})$$

$$y_j = 1 \text{ or } x_{kj} = 1 \text{ for some } k \in \mathcal{Q}_{<}(i, j) \implies \mathbf{fix}(w_{ij} = 1) \quad (\text{by 5.6-5.7})$$

$$y_j + \sum_{k \in \mathcal{Q}_{<}(i, j)} x_{kj} = 0 \implies \mathbf{fix}(w_{ij} = 0) \quad (\text{by 5.8})$$

$$w_{ij} = 0 \implies \mathbf{fix}(x_{kj} = 0) \quad \forall k \in \mathcal{Q}_{<}(i, j) \quad (\text{by 5.8})$$

$$u_{ij} = 1 \implies \mathbf{fix}(w_{ij} = 0) \quad (\text{by 5.9})$$

$$w_{ij} = 1 \implies \mathbf{fix}(u_{ij} = 0) \quad (\text{by 5.9})$$

end for

end while

combination $(i, j) \in \bar{\mathcal{I}} \times \bar{\mathcal{J}}$ where $x_{ij} = 0$ after pre-solving, we are left with a reduced set of admissible pairs $\bar{\mathcal{S}} \subseteq \mathcal{S}$ for said students/schools. We can apply the MILO formulation to the reduced problem with one slight modification: it is actually possible for the unhappiness of student or school with fully determined outcomes to depend on assignments that have not been fully determined. We must therefore change constraint 5.9 to include all potential blocking pairs rather than just $\bar{\mathcal{S}}$. In practice, we can eliminate many pairs that would never be blocking under the current assignments, i.e., those where either u_{ij} or w_{ij} has been fixed to 0. We refer to the set of remaining, potentially blocking pairs as $\bar{\mathcal{B}}$.

Discussion and extensions

The induced subproblem is typically much smaller than the original (see Section 5.4.2), and can be passed to an MILO solver to produce the remaining optimal assignments for a given objective function. As long as there are no extrinsic constraints, the feasible set is guaranteed to be non-empty since there always exists a stable solution. And in the case where that solution is unique, the pre-solve algorithm may even terminate having found it – though this is not guaranteed.

Finally, we note that when there are extrinsic constraints – e.g., two siblings that must be assigned to the same school – the feasible set may in fact be empty [84]. In such cases, it is possible to extend the pre-solve algorithm to detect infeasibility early by checking that propagation rules do not produce conflicting inferences. As an illustrative example, suppose we have inferred that assigning student i to school j would cause a blocking pair to form, so that $x_{ij} = 0$ is forced. If the pre-solve algorithm then also infers that i 's sibling must attend j , e.g., because it is the only admissible school left, we have detected infeasibility as x_{ij} cannot also equal 1.

5.4 The effect of stability on transport distance

In this section, we present experimental results on a range of synthetic school-choice problems to empirically evaluate the impact of stability on alternative assignment objectives. To this end, we consider the total distance between students and their assigned schools as a motivating objective function, since it serves as a proxy for districts’ transportation costs for bussing students to and from school. Our experimental goals are three-fold: first, to evaluate the effectiveness of the pre-solve algorithm (Section 5.3.2) in scaling MILO. Second, to characterize the stable solution space and the extent to which it can accommodate optimization of the alternate objective function. And, third, to examine tradeoffs between stability and the alternative objective when stability is relaxed to varying degrees.

5.4.1 Experimental design

Our experiments use synthetic data based on the Boston Public Schools (BPS) district 2017 Transportation Challenge dataset [18]. The dataset contains the addresses of 16,255 hypothetical BPS students and the 92 schools they must be transported to. We use this information to calculate pair-wise driving distances d_{ij} between each student i and school j , as well as approximate capacities for each school. We then generate synthetic preferences for both students and schools as follows: we use a parametric utility function for students that is the weighted sum of driving distance to school, school quality, and a random component. Students express preferences for (a random number of) their top-ranked schools according to this utility function, while the remaining schools are considered inadmissible. At the same time, schools rank all students that ranked them, based on a utility function that includes only distance to the student and a random component. Further details on the data generating process can be found in Appendix 5.6.1.

By varying the weight of the different factors in students’ and schools’ utility functions, we get preference profiles that range from being fully determined by distance and quality on the one hand, to being completely random on the other. Thus, problem

instances we generate vary considerably in terms of how much each side’s preferences correlate with the alternative matching objective, i.e., the sum of driving distances between each student and their assigned school. We report all results averaged over several random instances for each preference profile.

Experiments are all parallelized on a cluster of single-CPU machines with ≤ 32 GB of memory. Data processing is performed in python (version 3.7), with driving distances calculated via the Open Street Maps API [16]. Where required, we use the implementation of Gale & Shapley’s student-proposing DA algorithm from python’s `matching` package [101].³ Formulation of MILO and our pre-solve algorithm are programmed in Julia (version 1.5.2), while the MILOs are solved using Gurobi’s Mixed-Integer Linear Programming optimizer (version 9.1) and a time limit of 2 hours unless otherwise stated.

5.4.2 Results

In the first set of experiments, we consider 500 random instances of the BPS matching problem using all 16,255 students and 92 schools. Each student ranks between 5 and 10 schools, resulting in problems that have on the order ≈ 135 K admissible pairs. The matching objective is to minimize the total distance between students and their assigned schools, i.e., $\sum_{ij} d_{ij}x_{ij}$, subject to full stability (formulation 5.1-5.9). We apply the pre-solve algorithm of Section 5.3.2 to help solve each instance.

Table 5.1 reports average statistics on matching problem size before and after pre-solve. We observe that the pre-solver fully determines the matchings for 97.7% of students on average, while the number of admissible pairs is reduced by 99.1%. From a computational perspective, the improvements are evident: the pre-solve algorithm terminates within 30 minutes for all instances, while the resulting MILOs are solved to provable optimality in seconds. We do not report solve times without pre-solve, as Gurobi fails to solve the root relaxation of any problem within 24 time limit.

We also emphasize here that the reduction in problem size (first four rows of Ta-

³Note that the package’s documentation refers to the college admissions problem as the hospital-resident assignment problem.

Table 5.1: Matching problem size before and after pre-solve. Values indicate the average over 500 full-sized BPS instances, with the standard deviation in parentheses.

	Before pre-solve	After pre-solve	Reduction
# Students	16,255 (0)	369 (252)	97.7% (1.5%)
# Schools	92 (0)	59 (20)	35.7% (21.6%)
# Admissible pairs	134,663 (1,566)	1,256 (851)	99.1% (0.6%)
# Potential blocking pairs	134,663 (1,566)	1,423 (1,030)	98.9% (0.8%)
# Binary variables	269,323 (3,137)	1,571 (1,114)	99.4% (0.4%)
# Constraints	135.3M (29.5M)	11,271 (30,557)	$\approx 100\%$ (0.0%)
Pre-solve time (sec)	-	927.8 (250.3)	-
MILP solve time (sec)	-	0.3 (0.7)	-

ble 5.1) is applicable beyond our proposed approach, and can improve computational performance for any matching algorithm that guarantees stability. In particular, other MILO formulations of stability can leverage the same inferences to reduce the number of variables and constraints.

In terms of the experimental objectives, the pre-solver results also provide some intuition as to the nature of the stable solution space. The substantial reduction in problem size implies that any variance in the optimization objective must come from the assignment of just 2.3% of students. In other words, the requirement for stability does not leave a lot of room to optimize transport distance.

In Figure 5-1 we plot a histogram of the total number of stable solutions in each of the 500 BPS instances. Here, we explicitly enumerate the stable matching space by solving a sequence of MILOs that use binary no-good constraints to eliminate previously-found solutions from the feasible set (see Appendix 5.6.1). In 158 of the 500 instances, we find that there is only a single stable solution, while another 199 have five or fewer solutions. Moreover, Figure 5-2 plots a histogram of the percentage gap between the minimum- and maximum-distance stable solutions, i.e., $\frac{\max - \min}{\min}$, for each instance. We observe that, even when there are multiple stable solutions, the difference in their objective values is no more than 0.13% in any instance. Therefore, in this case at least, it does not seem possible to optimize the alternative objective function significantly under the requirement of strict stability.




Figure 5-1: Distribution of the number of stable solutions over 500 full-sized BPS matching instances. The number of solutions was determined by solving a sequence of MILOs, each time excluding all previously found solution from the feasible set via binary no-good constraints (see Appendix 5.6.1).

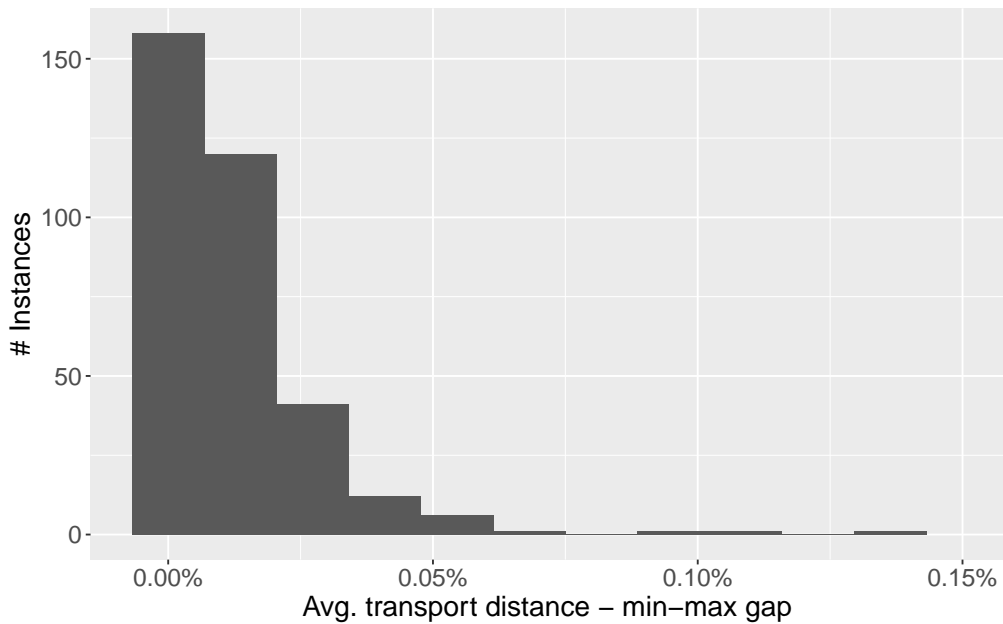


Figure 5-2: Distribution of the percentage difference in objective value for the minimum and maximum distance stable solutions over 342 full-sized BPS matching instances. Excluded are the 158 instances where the stable solution was unique.

These observations motivate a second set of experiments, where we relax stability and instead consider how much transport distance can be reduced if some blocking pairs are allowed in the matching. As noted in Section 5.3.2, here we cannot leverage the Rural Hospitals Theorem in order pre-solve assignments, and so we present proof-of-concept results on a set of smaller instances. These are generated as before, but are limited to a subset of 450 students and 10 schools chosen to be distributed as widely as possible across the Boston area (see Appendix 5.6.1).

We generate 250 such instances with varying preference profiles, and implement the relaxed stability formulation described in Section 5.3.1. We vary the level of stability γ and again set the objective to minimize total transport distance. Since the number of matched students can vary for different γ , we enforce via constraints that the solution matches at least as many students as a 1-stable solution, as computed by the DA algorithm. To compare across instances, we report the objective values of each solution as a percentage of the objective value of the 1-stable solution found by the DA algorithm in the same instance.⁴

Figure 5-3 plots the distribution of optimized transport distance over the 250 instances as a function of $1 - \gamma$. On the far left, solutions are constrained to have no blocking pairs, and therefore do not significantly improve in objective value over the DA solution (per our earlier results). As stability is relaxed, however, the optimization realizes gains in terms of reduced transport distance: with just 1% of pairs allowed to be blocking, we observe a $12.7\% \pm 4.2\%$ (avg \pm sd) reduction in objective value compared to the DA solution. At 5% of pairs, the reduction is $28.2\% \pm 8.6\%$. These results suggest that, while stability is a rather constraining property, there may be considerable gains in terms of alternative policy objectives if it is even slightly relaxed.

⁴Note that, since the set of 1-stable solutions is a subset of γ -stable solutions, the percentage must always be $\leq 100\%$.




Figure 5-3: Students’ average transport distance in the optimal γ -stable solution as a function of percentage of unstable pairs ($1 - \gamma$). The y-axis expressed as a percentage of the average transport distance in the 1-stable solution given by the student-proposing DA algorithm. The solid line, error bands, and dashed lines denote the median, 25th/75th percentiles, and minimum/maximum respectively over 250 medium-sized BPS instances.

5.5 Conclusions and future work

In this work, we have presented an empirical analysis of stable matching mechanisms in the context of school choice. While the notion of fairness that stability ensures is both intuitive and elegant, it is just one of many possible policy objectives that districts should consider when designing assignment mechanisms. Global optimization methods provide a natural language for combining multiple objectives and constraints (including stability) in the matching problem, and so form the basis for our approach.

Concretely, we develop a custom pre-solve algorithm for MILO and successfully scale global optimization in the fully stable case. Our experiments on large-scale synthetic problems suggest that this setting may be overly constraining when optimizing for an alternative objective function, e.g., student transportation cost. Motivated by this, we propose a natural way to relax stability so that it can be traded off against alternative objectives. Results on smaller synthetic problems suggest that districts can in fact realize significant gains in the alternative objective with only small sacrifices to stability.

While these proof-of-concept results are promising, some work remains in scaling the approach to larger problem sizes for the relaxed stability case. Though the Rural Hospitals Theorem no longer applies directly, constraint propagation rules such as in Section 5.3.2 are still valid for inferring assignments once certain variable values have been fixed. This suggests the possibility of a custom branch-and-bound approach to solve MILOs: after branching, one might use a modified version of Algorithm 2 to fix additional variable values or detect infeasibility, and so more efficiently fathom/prune nodes. We still expect large cascades of inferences, particularly at nodes where the budget of blocking pairs has been exhausted and the remaining assignments are highly constrained by the need for stability. The approach could be further augmented with custom branching rules, e.g., prioritizing assignments that are likely to cause many blocking pairs, or primal heuristics, e.g., running the DA algorithm on remaining assignments to produce a solution that is guaranteed to satisfy the stability constraint.

Beyond computational improvements, we also wish to extend our experiments to

non-synthetic datasets from real school districts. We expect that stability will prove to be a constraining property in these settings as well, while relaxing it should allow for matching mechanisms that address districts' other policy objectives. One might envision applying the "ethics-by-design" framework we developed in Chapters 3 and 4 to help evaluate global tradeoffs in school choice. Here, the "black-box" function would be the end-to-end MILO solving process, using a parameterized objective and constraints to control the utility/efficiency/fairness characteristics of the resulting matching. An interactive optimization tool could be used, wherein district officials specify their desired objectives in terms of assignment equity, transportation cost, level of stability, etc., and a surrogate "meta-optimization" designs a conforming assignment mechanism.

5.6 Appendix

5.6.1 Data and methods

Data pre-processing

As described in Section 5.4.1, our experiments use synthetic preference data generated for students and schools in the Boston Public Schools (BPS) district. The 2017 BPS Transportation Challenge dataset, which forms the basis of our data generation process, contains the geographic addresses of 22,420 hypothetical BPS students and the 134 schools they must be transported to. Before running our experiments, we augment the dataset as follows:

- For each school, we scrape their 2019 School Quality Framework (SQF) score [17] from www.bostonschoolfinder.org. We divide the scores by 100 to have them lie on a unit scale, and refer to school j 's normalized SQF score as SQF_j .
- There are 42 public and charter schools that do not have SQF scores, which we remove from the dataset along with their assigned students. This leaves 16,255 students and 92 schools.
- We compute the number of students needing transport to each school, denoted by N_j for school j , to be used as a proxy for school capacity.
- We compute the driving distance between each remaining student's home and each school, measured in kilometers. We use a graph representation of Boston's road network available on Open Street Maps [16], and compute pair-wise distances as the sum of: (i) the Euclidean distance between a student's location and the nearest network node; (ii) Euclidean distance between the school's location and the nearest network node; and (iii) shortest path network distance between those two nodes. We denote the distance between student i and school j by d_{ij} .

Data generating process

We next describe the random data generating process that creates an instance of the BPS school choice problem:

1. We assume that the dependence of students' preferences for schools on distance (if any) takes the following form: the student assigns a distance score that equals 1 if the school is within 2km, 0 if it is farther than 10km, and decreases linearly in between; that is, $f(d) := \max\{0, \min\{1, \frac{d-2}{8}\}\}$.
2. Student i 's utility for being assigned to school j is parameterized by a single value $\lambda \in [0, 1]$ as follows:

$$p_{ij}(\lambda) = \frac{\lambda}{2}f(d_{ij}) + \frac{\lambda}{2}\text{SQF}_j + (1 - \lambda)X_{ij}$$

where $X_{ij} \stackrel{iid}{\sim} \text{Uniform}(0, 1)$.

3. School j 's utility for being assigned student i is parameterized by a single value μ as follows:

$$q_{ij}(\mu) = \mu\widehat{d}_{ij} + (1 - \mu)Y_{ij}$$

where $Y_{ij} \stackrel{iid}{\sim} \text{Uniform}(0, 1)$. Here \widehat{d}_{ij} is the driving distance calculated in the previous section, but min-max normalized over the entire dataset (so that it lies on the unit scale).

4. Parameters P_{\min}/P_{\max} denote the minimum and maximum number of schools that any student can express preferences for. Each student expresses a preference for their top $P_i \stackrel{iid}{\sim} \text{DiscreteUniform}(P_{\min}, P_{\max} - 1)$ schools, ranked according to $p_{ij}(\lambda)$, as well as their originally assigned school.⁵ The remaining schools are considered inadmissible for the student.
5. Schools express preference for any student that expressed a preferences for them, and rank them according to $q_{ij}(\mu)$.

⁵Original assignments are included so that there always exists a matching that is capacity feasible.

6. The capacity inflation parameter π denotes the extra number of school positions to add to the instance, expressed as a percentage of the total number of students. Each extra position is opened at a school chosen uniformly-at-random and independently from other positions. Thus a school’s capacity C_j is given by N_j plus however many of the extra positions they were randomly assigned.

To summarise, our data generating process is parameterized by the student utility weight λ , school utility weight μ , minimum and maximum number of student preferences P_{\min} and P_{\max} , and capacity inflation factor π . Randomization over instances is given by the random factors in student and school utility functions, number of preferences expressed by each student, and which schools are assigned extra seats.

Experimental design

For the first set of experiments in Section 5.4, we use the full dataset of 16,255 students and 92 schools. We vary both λ and μ over a grid of $\{0, 0.25, 0.5, 0.75, 1\}$ and fix $(P_{\min}, P_{\max}, \pi) = (5, 10, 5\%)$, for a total of 25 parameter settings. We then generate 20 random instances for each parameter setting, for a total of 500 matching problems.

Each problem is processed by the pre-solver of Section 5.3.2, and the induced subproblem formulated as an MILO. In order to enumerate the stable solution space, we make extensive use of so-called “no-good” constraints. A no-good constraint is one that eliminates just a single solution from the feasible set. To illustrate, suppose that the values of the x_{ij} variables in the no-good solution are given by \bar{x}_{ij} . Since x_{ij} are binary and fully determine a matching, the constraint:

$$\sum_{(i,j):\bar{x}_{ij}=0} x_{ij} + \sum_{(i,j):\bar{x}_{ij}=1} (1 - x_{ij}) \geq 1$$

has the effect of eliminating just the given solution from the feasible set. Intuitively, it states that any feasible point must have x whose Hamming distance from \bar{x} is at least one.

Thus for each of the 500 problems, we solve a sequence of MILOs that minimize total matching distance. Each iteration a single no-good constraint is added to eliminate the most recently found solution from the feasible set. We repeat the process, keeping track of the computed solutions, until either the problem becomes infeasible—in which case the solution space has been enumerated—or we have found more than 20 feasible solutions. In the latter case, we solve one final MILO to maximize total matching distance, so that both the minimum- and maximum- distance solutions are in our set.

For the second set of experiments we cannot use the pre-solve algorithm as stated and so focus on a subset of schools and their assigned students. We formulate a facility location problem to select the K schools that minimize total distance between all 16,255 students and their nearest selected school. The goal is for these schools to span the entire geographic area of Boston, so that there is variance within each student’s distance-based preference scores. We further require that the total capacity among selected schools is bounded between N_{\min} and N_{\max} . The MILO facility-location formulation is as follows:

$$\begin{aligned}
\min_{x,z} \quad & \sum_{(i,j) \in \mathcal{I} \times \mathcal{J}} d_{ij} x_{ij} \\
\text{s.t.} \quad & \sum_{j \in \mathcal{J}} z_j = K \\
& N_{\min} \leq \sum_{j \in \mathcal{J}} N_j z_j \leq N_{\max} \\
& \sum_{j \in \mathcal{J}} x_{ij} = 1 \quad \forall i \in \mathcal{I} \\
& x_{ij} \leq z_j \quad \forall i \in \mathcal{I}, j \in \mathcal{J} \\
& x_{ij}, z_j \in \{0, 1\} \quad \forall i \in \mathcal{I}, j \in \mathcal{J}
\end{aligned}$$

where z_j encodes whether school j is selected, and x_{ij} whether school j is the closest selected school to student i .

We solve the facility location problem setting $(K, N_{\min}, N_{\max}) = (10, 400, 450)$, resulting in a reduced dataset of exactly 450 students and 10 schools. The data generation process is exactly as before, where we vary both λ and μ across a grid of $\{0, 0.25, 0.5, 0.75, 1\}$, and fix $(P_{\min}, P_{\max}, \pi) = (5, 10, 5\%)$. We generate 10 random instances for each parameters setting, for a total of 250 random instances. We then run the γ -stable MILO formulation minimizing total matching distance in each instance. To generate tradeoff curves, we vary γ between 0.8 and 0.9 in increments of 0.025, and between 0.91 and 1 in increments of 0.01, for a total of 3750 MILO solves.

All experiments are parallelized on a cluster of standard CPU machines with ≤ 32 GB of RAM. The pre-solve algorithm is written in the Julia programming language (version 1.5.2), and MILOs are solved using Gurobi's Mixed-Integer Linear Programming optimizer.

Chapter 6

Conclusions

This thesis develops multi-objective optimization techniques to support policymakers in designing more efficient, fair, and inclusive policies, with a focus on applications in transplantation and public education.

In Chapter 2, we consider a fundamental tension between efficiency and fairness in the geographic distribution of organs recovered for transplantation. We use multi-objective optimization to generate tradeoff curves comparing three different frameworks for incorporating geography into patients' allocation priority. We show that a continuous distribution concept allows for both the greatest reduction in patient mortality, and the most equitable geographic distribution, among all others considered. Following the completion of this work, the Organ Procurement & Transplantation Network (OPTN) Board of Directors issued a directive that all future policy development would focus on continuous distribution.

In Chapter 3, the OPTN's decision prompted us to develop a general framework for ethics-by-design in scarce resource allocation. We create an optimization-based interactive application that allows policymakers to specify their desired policy outcomes in terms of an objective and constraints, and our methodology computes a conforming policy in near real-time. This enables stakeholders, even those without technical expertise, to quickly iterate on different policy scenarios and refine their value judgments on relevant tradeoffs in an evidence-driven way.

Chapter 4 details our collaboration with the United Network for Organ Sharing

(UNOS) and OPTN Lung Transplantation Committee to help design a new national allocation policy for lungs. We applied our ethics-by-design framework towards the design of a continuous distribution allocation policy, and presented key tradeoff analyses on placement efficiency and pediatric priority to help guide the committee’s decision-making process. The committee’s proposal was unanimously approved by the OPTN Board of Directors in December of 2021. Starting in 2023, all deceased-donor lungs in the US will be allocated according to a policy that was significantly informed by this work, promising to improve patient welfare for years to come.

Finally, in Chapter 5 we present an empirical analysis of school assignment mechanisms using global optimization. We investigate tradeoffs between the popular notion of matching stability and alternative objective functions, e.g., bus transportation cost. In our experiments, we find that stability is a particularly constraining property that does not allow for significant optimization of transport cost. We propose instead a relaxed notion of stability that can be readily incorporated into a multi-objective optimization problem, and would allow school districts to balance stability with their other policy objectives.

Bibliography

- [1] Ross Anderson, Joey Huchette, Will Ma, Christian Tjandraatmadja, and Juan Pablo Vielma. Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming*, pages 1–37, 2020.
- [2] Itai Ashlagi and Peng Shi. Improving community cohesion in school choice via correlated-lottery implementation. *Operations Research*, 62(6):1247–1264, 2014.
- [3] David A Axelrod, Parsia A Vagefi, and John P Roberts. The evolution of organ allocation for liver transplantation: tackling geographic disparity through broader sharing. *Annals of surgery*, 262(2):224–227, 2015.
- [4] Mourad Baiou and Michel Balinski. The stable admissions polytope. *Mathematical programming*, 87(3):427–439, 2000.
- [5] Solon Barocas and Andrew D Selbst. Big data’s disparate impact. *Calif. L. Rev.*, 104:671, 2016.
- [6] Hamsa Bastani, Kimon Drakopoulos, Vishal Gupta, Ioannis Vlachogiannis, Christos Hadjicristodoulou, Pagona Lagiou, Gkikas Magiorkinis, Dimitrios Paraskevis, and Sotirios Tsiodras. Efficient and targeted covid-19 border testing via reinforcement learning. *Nature*, 599(7883):108–113, 2021.
- [7] Peter C Belford and H Donald Ratliff. A network-flow model for racially balancing schools. *Operations Research*, 20(3):619–628, 1972.
- [8] Dimitris Bertsimas, Vassilis Digalakis Jr, Alexander Jacquillat, Michael Lingzhi Li, and Alessandro Previero. Where to locate covid-19 mass vaccination facilities? *Naval Research Logistics (NRL)*, 69(2):179–200, 2022.
- [9] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.
- [10] Dimitris Bertsimas and Mohammad M Fazel-Zarandi. Prescriptive machine learning for public policy: The case of immigration enforcement. *Under review*, 2022.

- [11] Dimitris Bertsimas, Jerry Kung, Nikolaos Trichakis, Yuchen Wang, Ryutaro Hirose, and Parsia A Vagefi. Development and validation of an optimized prediction of mortality for candidates awaiting liver transplantation. *American Journal of Transplantation*, 19(4):1109–1118, 2019.
- [12] Dimitris Bertsimas, Theodore Papalexopoulos, Nikolaos Trichakis, Yuchen Wang, Ryutaro Hirose, and Parsia A Vagefi. Balancing efficiency and fairness in liver transplant access: tradeoff curves for the assessment of organ distribution policies. *Transplantation*, 104(5):981–987, 2020.
- [13] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.
- [14] Max Biggs, Rim Hariss, and Georgia Perakis. Optimizing objective functions determined from random forests. *Available at SSRN 2986630*, 2017.
- [15] Aaron L Bodoh-Creed. Optimizing for distributional goals in school choice problems. *Management Science*, 66(8):3657–3676, 2020.
- [16] Geoff Boeing. Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017.
- [17] BPS. School Quality Framework. <https://www.bostonpublicschools.org/schoolquality>, 2017. Boston Public Schools. Online; accessed 11 November 2021.
- [18] BPS. Transportation Challenge: Simulated "Fake" Student Address Dataset. <https://www.bostonpublicschools.org/transportationchallenge>, 2017. Boston Public Schools. Online; accessed 11 November 2021.
- [19] BPS. Boston public schools budget. [https://www.bostonpublicschools.org/cms/lib/MA01906464/Centricity/domain/184/budgetvisualization/index.html#/,](https://www.bostonpublicschools.org/cms/lib/MA01906464/Centricity/domain/184/budgetvisualization/index.html#/) 2021. Boston Public Schools. Online; accessed 16 March 2022.
- [20] Michael M Burgess. From ‘trust us’ to participatory governance: deliberative publics and science policy. *Public understanding of science*, 23(1):48–52, 2014.
- [21] Felipe Caro, Takeshi Shirabe, Monique Guignard, and Andrés Weintraub. School redistricting: Embedding gis tools with integer programming. *Journal of the Operational Research Society*, 55(8):836–849, 2004.
- [22] Rodolfo Carvajal, Miguel Constantino, Marcos Goycoolea, Juan Pablo Vielma, and Andrés Weintraub. Imposing connectivity constraints in forest planning models. *Operations Research*, 61(4):824–836, 2013.
- [23] Jonathan P Caulkins. The role of operations research in public policy. In *Operations Research Proceedings 2002*, pages 1–13. Springer, 2003.

- [24] Cruz v. DHHS. Cruz v. United States Department of Health and Human Services, 2018. S.D.N.Y. 18-CV-06371.
- [25] Maxence Delorme, Sergio García, Jacek Gondzio, Joerg Kalcsics, David Manlove, and William Pettersson. Mathematical models for stable matching problems with ties and incomplete lists. *European Journal of Operational Research*, 277(2):426–441, 2019.
- [26] Ranjit Deshpande, Ryutaro Hirose, and David Mulligan. Liver allocation and distribution: time for a change. *Current opinion in organ transplantation*, 22(2):162–168, 2017.
- [27] DHHS. Organ Procurement and Transplantation Network Final Rule, 2000. Electronic Code of Federal Regulations, Title 42–Public Health, Chapter I–Public Health Service, Department of Health and Human Services, Subchapter K–Health Resources Development, Part 121.
- [28] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3):1155–1170, 2018.
- [29] Gordon M Fisher. The development and history of the poverty thresholds. *Soc. Sec. Bull.*, 55:3, 1992.
- [30] John Forrest, Ted Ralphs, Stefan Vigerske, LouHafer, Bjarni Kristjansson, jpfasano, EdwinStraver, Miles Lubin, Haroldo Gambini Santos, rlougee, and Matthew Saltzman. coin-or/cbc: Version 2.9.9, July 2018.
- [31] Archon Fung. Recipes for public spheres: Eight institutional design choices and their consequences, 11 j. *Pol. Phil*, 338:345, 2003.
- [32] Archon Fung, Erik Olin Wright, et al. *Deepening democracy: Institutional innovations in empowered participatory governance*, volume 4. Verso, 2003.
- [33] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [34] Sommer E Gentry, Allan B Massie, Sidney W Cheek, Krista L Lentine, EH Chow, Corey E Wickliffe, Nino Dzebashvili, Paolo R Salvalaggio, Mark A Schnitzler, David A Axelrod, et al. Addressing geographic disparities in liver transplantation through redistricting. *American Journal of Transplantation*, 13(8):2052–2058, 2013.
- [35] Alexandra K Glazier. The lung lawsuit: A case study in organ allocation policy and administrative law. *Journal of Health & Biomedical Law*, XIV(1):139–148, 2018.

- [36] Aparna Goel, W Ray Kim, Joshua Pyke, David P Schladt, Bertram L Kasiske, Jon J Snyder, John R Lake, and Ajay K Israni. Liver simulated allocation modeling: were the predictions accurate for share 35? *Transplantation*, 102(5):769–774, 2018.
- [37] David S Goldberg, Matthew Levine, Seth Karp, Richard Gilroy, and Peter L Abt. Share 35 changes in center-level liver acceptance practices. *Liver Transplantation*, 23(5):604–613, 2017.
- [38] Jan-Henrik Haunert and Alexander Wolff. Area aggregation in map generalisation by mixed-integer programming. *International Journal of Geographical Information Science*, 24(12):1871–1897, 2010.
- [39] Sidney Wayne Hess, JB Weaver, HJ Siegfeldt, JN Whelan, and PA Zitlau. Non-partisan political redistricting by computer. *Operations Research*, 13(6):998–1006, 1965.
- [40] Caroline Minter Hoxby. School choice and school competition: Evidence from the united states. *Swedish Economic Policy Review*, 2003.
- [41] Ajay K Israni, Nicholas Salkowski, Sally Gustafson, Jon J Snyder, John J Friedewald, Richard N Formica, Xinyue Wang, Eugene Shteyn, Wida Cherikh, Darren Stewart, et al. New national allocation policy for deceased donor kidneys in the united states and possible effect on patient outcomes. *Journal of the American Society of Nephrology*, 25(8):1842–1848, 2014.
- [42] Muriel Jean-Jacques and Howard Bauchner. Vaccine distribution—equity left behind? *Jama*, 325(9):829–830, 2021.
- [43] Tim Jin. Why is california’s age-based covid-19 vaccine policy overlooking disabled people like me? <https://www.latimes.com/opinion/story/2021-01-29/covid-vaccine-disabled-people-priority>. Los Angeles Times. Online; published 29 January 2021, accessed 15 March 2021.
- [44] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein. Discrimination in the age of algorithms. *Journal of Legal Analysis*, 10:113–174, 2018.
- [45] Slawomir Koziel and Leifur Leifsson. *Surrogate-based modeling and optimization*. Springer, 2013.
- [46] Augustine Kwanashie and David F Manlove. An integer programming approach to the hospitals/residents problem with ties. In *Operations research proceedings 2013*, pages 263–269. Springer, 2014.
- [47] Daniela Lamas and Lisa Rosenbaum. Very complicated math—reconfiguring organ allocation. *New England Journal of Medicine*, 371(26):2447–2450, 2014.

- [48] Rebecca R Lehman and Kevin M Chan. Elimination of the donor service area (dsa) from lung allocation: no turning back, 2019.
- [49] Carli J Lehr, Melissa Skeans, and Maryam Valapour. Validating thoracic simulated allocation model predictions for impact of broader geographic sharing of donor lungs on transplant waitlist outcomes. *The Journal of Heart and Lung Transplantation*, 39(5):433–440, 2020.
- [50] Retsef Levi, Manoj Rajan, Somya Singhvi, and Yanchong Zheng. Improving farmers’ income on online agri-platforms: Theory and field implementation of a two-stage auction. *Available at SSRN 3486623*, 2020.
- [51] Michael Lingzhi Li, Hamza Tazi Bouardi, Omar Skali Lami, Thomas A Trikalinos, Nikolaos K Trichakis, and Dimitris Bertsimas. Forecasting covid-19 and analyzing the effect of government interventions. *MedRxiv*, pages 2020–06, 2021.
- [52] Robin Segerblom Liggett. The application of an implicit enumeration algorithm to the school desegregation problem. *Management Science*, 20(2):159–168, 1973.
- [53] LTC. Public Comment Proposal: Establish Continuous Distribution of Lungs. https://optn.transplant.hrsa.gov/media/3111/thoracic_publiccomment_201908.pdf, 2019. Organ Procurement & Transplantation Network, Thoracic Organ Transplantation Committee. Online; accessed 8 November 2021.
- [54] LTC. OPTN Lung Transplantation Committee Meeting Summary, March 25, 2021. https://optn.transplant.hrsa.gov/media/4567/20210325_lung_meeting_summary.pdf, 2021. Online; accessed 8 November 2021.
- [55] LTC. OPTN Lung Transplantation Committee Meeting Summary, March 31, 2021. https://optn.transplant.hrsa.gov/media/4579/20210331_lung-meeting-summary.pdf, 2021. Online; accessed 8 November 2021.
- [56] LTC. OPTN Lung Transplantation Committee Meeting Summary, October 22, 2021. https://optn.transplant.hrsa.gov/media/203b23pe/20211022_lung-committee-meeting-summary_draft.pdf, 2021. Online; accessed 8 November 2021.
- [57] LTC. Public Comment Proposal: Establish Continuous Distribution of Lungs. https://optn.transplant.hrsa.gov/media/4772/continuous_distribution_of_lungs-public_comment.pdf, 2021. Organ Procurement & Transplantation Network, Lung Transplantation Committee. Online; accessed 8 November 2021.
- [58] Vahideh Manshadi and Scott Rodilitz. Online policies for efficient volunteer crowdsourcing. *Management Science*, 2022.

- [59] Donato Maragno, Holly Wiberg, Dimitris Bertsimas, S. Ilker Birbil, Dick den Hertog, and Adejuyigbe Fajemisin. Mixed-integer optimization with constraint learning, 2021.
- [60] Allan B Massie and John Paul Roberts. Geographic disparity in liver allocation: time to act or have others act for us. *Transplantation*, 102(2):189–190, 2018.
- [61] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.
- [62] Anuj Mehrotra, Ellis L Johnson, and George L Nemhauser. An optimization based heuristic for political districting. *Management Science*, 44(8):1100–1114, 1998.
- [63] Sanjay Mehrotra, Vikram Kilambi, Kevin Bui, Richard Gilroy, Sophoclis P Alexopoulos, David S Goldberg, Daniela P Ladner, and Goran B Klintmalm. A concentric neighborhood solution to disparity in liver access that contains current unos districts. *Transplantation*, 102(2):255–278, 2018.
- [64] Velibor V Mišić. Optimization of tree ensembles. *Operations Research*, 68(5):1605–1624, 2020.
- [65] Logan Patrick Moore and David L Weimer. The geography of life and death: Evidence and values in the evolution of us liver transplant rules. *World Medical & Health Policy*, 2021.
- [66] National Academies of Sciences, Engineering, and Medicine. *Framework for Equitable Allocation of COVID-19 Vaccine*. The National Academies Press, Washington, DC, 2020.
- [67] NCES. Back-to-school statistics. <https://nces.ed.gov/fastfacts/display.asp?id=372>, 2021. National Center for Education Statistics. Online; accessed 16 March 2022.
- [68] OPTN. Broader sharing of adult donor lungs. https://optn.transplant.hrsa.gov/media/2314/broader_sharing_lungs_20171124.pdf, 2017. OPTN/UNOS Executive Committee. Online; accessed 25 February 2019.
- [69] OPTN. Executive Summary of the OPTN/UNOS Board of Directors Meeting, December 3-4, 2018. https://optn.transplant.hrsa.gov/media/2787/board_executivesummary_201812.pdf, 2018. Organ Procurement & Transplantation Network. Online; accessed 8 November 2021.
- [70] OPTN. Frameworks for organ distribution. https://optn.transplant.hrsa.gov/media/2762/geography_boardreport_201812.pdf, 2018. OPTN/UNOS Ad Hoc Geography Committee. Online; accessed 25 February 2019.

- [71] OPTN. Key Initiatives: Continuous distribution. <https://optn.transplant.hrsa.gov/governance/key-initiatives/continuous-distribution/>, 2018. Organ Procurement and Transplantation Network. Online; accessed 25 February 2019.
- [72] OPTN. Liver and intestine distribution using distance from donor hospital. https://optn.transplant.hrsa.gov/media/2766/liver_boardreport_201812.pdf, 2018. OPTN/UNOS Liver and Intestine Transplantation Committee. Online; accessed 25 February 2019.
- [73] OPTN. Organ Procurement & Transplantation Network Policies. https://optn.transplant.hrsa.gov/media/eavh5bf3/optn_policies.pdf, 2018. Online; accessed 23 November 2018.
- [74] OPTN. 10-Step Policy Development Process. <https://optn.transplant.hrsa.gov/media/3115/optn-policy-development-process-explanatory-document.pdf>, 2021. Organ Procurement & Transplantation Network. Online; accessed 8 November 2021.
- [75] OPTN. Executive Summary of the OPTN/UNOS Board of Directors Meeting, December 6, 2021. <https://optn.transplant.hrsa.gov/media/g23hdtxk/20211206-optn-bod-summary.pdf>, 2021. Organ Procurement & Transplantation Network. Online; accessed 12 January 2022.
- [76] OPTN. Organ Procurement & Transplantation Network Policies. https://optn.transplant.hrsa.gov/media/eavh5bf3/optn_policies.pdf, 2021. Online; accessed 8 November 2021.
- [77] Daniel T O'Brien, NE Hill, M Contreras, NE Phillips, and G Sidoni. An evaluation of equity in the boston public schools' home-based assignment policy. *Boston: Boston Area Research Initiative*, 2018.
- [78] Theodore Papalexopoulos, Dimitris Bertsimas, I Glenn Gohen, Rebecca Goff, Darren Stewart, and Nikolaos Trichakis. Ethics-by-design: Efficient, fair and inclusive resource allocation using machine learning. *Journal of Law and the Biosciences*, 2022. Forthcoming.
- [79] Theodore Papalexopoulos, Dimitris Bertsimas, Nikolaos Trichakis, James Alcorn, Rebecca Goff, and Darren Stewart. Reshaping organ allocation policy through multi-objective optimization. Working paper.
- [80] Philip A Perry and Timothy Hotze. Oregon's experiment with prioritizing public health care services. *AMA Journal of Ethics*, 13(3):241–247, 2011.
- [81] Arti Kaur Rai. Rationing through choice: A new approach to cost-effectiveness analysis in health care. *Ind. LJ*, 72:1015, 1996.

- [82] RFI. Kidney Allocation Concepts: Request for Information. <https://asts.org/docs/default-source/optn-unos/proposed-kidney-allocation-concepts-rfi-september-24-2008.pdf>, 2008. OPTN/UNOS Kidney Transplantation Committee. Online; accessed 8 November 2021.
- [83] Francesca Rossi, Peter Van Beek, and Toby Walsh. *Handbook of constraint programming*. Elsevier, 2006.
- [84] Alvin E Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of political Economy*, 92(6):991–1016, 1984.
- [85] Alvin E Roth. On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica: Journal of the Econometric Society*, pages 425–427, 1986.
- [86] Alvin E Roth and Marilda Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541, 1992.
- [87] Jay Sethuraman, Chung-Piaw Teo, and Liwen Qian. Many-to-one stable matching: geometry and fairness. *Mathematics of Operations Research*, 31(3):581–596, 2006.
- [88] Mohammad Shahabsafa, Tamás Terlaky, Naga Venkata Chaitanya Gudapati, Anshul Sharma, George R Wilson, Louis J Plebani, and Kristofer B Bucklen. The inmate assignment and scheduling problem and its application in the pennsylvania department of corrections. *Interfaces*, 48(5):467–483, 2018.
- [89] Peng Shi. Guiding school-choice reform through novel applications of operations research. *Interfaces*, 45(2):117–132, 2015.
- [90] Takeshi Shirabe. A model of contiguity for spatial unit allocation. *Geographical Analysis*, 37(1):2–16, 2005.
- [91] Noah A Smith and Roy W Tromble. Sampling uniformly from the unit simplex. *Johns Hopkins University, Tech. Rep*, 29, 2004.
- [92] Jon J Snyder, Nicholas Salkowski, Andrew Wey, Joshua Pyke, Ajay K Israni, and Bertram L Kasiske. Organ distribution without geographic boundaries: a possible framework for organ allocation. *American Journal of Transplantation*, 18(11):2635–2640, 2018.
- [93] SRTR. Continuous distribution simulations for lung transplant: Round 2. https://optn.transplant.hrsa.gov/media/4646/lu2021_01_cont_distn_report_final.pdf, 2021. Scientific Registry of Transplant Recipients. Online; accessed 8 November 2021.

- [94] SRTR. Simulated Allocation Models. <https://www.srtr.org/requesting-srtr-data/simulated-allocation-models/>, 2021. Scientific Registry of Transplant Recipients. Online; accessed 8 November 2021.
- [95] Mark D Stegall, Peter G Stock, Kenneth Andreoni, John J Friedewald, and Alan B Leichtman. Why do we have the kidney allocation system we have today? a history of the 2014 kidney allocation system. *Human immunology*, 78(1):4–8, 2017.
- [96] Darren E Stewart, Dallas W Wood, James B Alcorn, Erika D Lease, Michael Hayes, Brett Hauber, and Rebecca E Goff. A revealed preference analysis to develop composite scores approximating lung allocation policy in the us. *BMC Medical Informatics and Decision Making*, 21(1):1–11, 2021.
- [97] Harini Suresh and John Guttag. A framework for understanding sources of harm throughout the machine learning life cycle. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–9. Equity and Access in Algorithms, Mechanisms, and Optimization, 2021.
- [98] Parsia A Vagefi, Nancy L Ascher, Chris E Freise, Jennifer L Dodge, and John P Roberts. Use of living donor liver transplantation varies with the availability of deceased donor liver transplantation. *Liver Transplantation*, 18(2):160–165, 2012.
- [99] Parsia A Vagefi, Sandy Feng, Jennifer L Dodge, James F Markmann, and John P Roberts. Multiple listings as a reflection of geographic disparity in liver transplantation. *Journal of the American College of Surgeons*, 219(3):496–504, 2014.
- [100] David L Weimer. *Medical governance: Values, expertise, and interests in organ transplantation*. Georgetown University Press, 2010.
- [101] Henry Wilde, Vincent Knight, and Jonathan Gillard. Matching: A python library for solving matching games. *Journal of Open Source Software*, 5(48):2169, 2020.
- [102] Justin C Williams. A zero-one programming model for contiguous land acquisition. *Geographical Analysis*, 34(4):330–349, 2002.
- [103] Zhizhou Yang, William D Gerull, Jason M Gauthier, Bryan F Meyers, Benjamin D Kozower, G Alexander Patterson, Ruben G Nava, Ramsey R Hachem, Chad A Witt, Derek E Byers, et al. Shipping lungs greater distances increases costs without cutting waitlist mortality. *The Annals of Thoracic Surgery*, 110(5):1691–1697, 2020.
- [104] Zelda B. Zabinsky and Robert L. Smith. Hit-and-run methods. In Saul I. Gass and Michael C. Fu, editors, *Encyclopedia of Operations Research and Management Science*, pages 721–729. Springer US, Boston, MA, 2013.

- [105] Andris A Zoltners and Prabhakant Sinha. Sales territory alignment: A review and model. *Management Science*, 29(11):1237–1256, 1983.

Acronyms

σ **MMaT** standard deviation of MMaT across DSAs.

AC Acuity Circles.

AI Artificial Intelligence.

B2C Broader 2-Circle Distribution.

BPS Boston Public Schools.

CAS Composite Allocation Score.

CD Continuous Distribution.

DA Deferred Acceptance.

DCD Donation after Cardiac Death.

DHHS Department of Health and Human Services.

DSA Donor Service Area.

HHRI Hennepin Healthcare Research Institute.

HRSA Health Resources and Services Administration.

KAS Kidney Allocation System.

LAS Lung Allocation Score.

LSAM Liver Simulated Allocation Model.

LTC Lung Transplantation Committee.

MAE mean absolute error.

MELD Model for End-Stage Liver Disease.

MILO Mixed Integer Linear Optimization.

MIO Mixed Integer Optimization.

ML Machine Learning.

MLaT Median LAS score at transplant.

MMaT Median MELD at Transplant.

MMRF Minneapolis Medical Research Foundation.

OD Optimized Districts.

OPO Organ Procurement Organization.

OPOM Optimized Prediction of Mortality.

OPTN Organ Procurement & Transplantation Network.

OR Operations research.

PTAUC Post-Transplant Area Under the (survival) Curve.

RHT Rural Hospitals Theorem.

rMSE root mean squared error.

SAM Simulated Allocation Model.

SQF School Quality Framework.

SRTR Scientific Registry of Transplant Recipients.

TSAM Thoracic Simulated Allocation Model.

UNOS United Network for Organ Sharing.

WLAUC Waitlist mortality Area Under the (survival) Curve.

WMAD Weighted Mean Absolution Deviation.

We note that some of these terms are explained in more detail in the Glossary.