

ESSAYS ON CONTEMPORARY MATCHING AND
ALLOCATION PROBLEMS

Musa Eren Çeldir

A Dissertation

Submitted to the Tepper School of Business

In Partial Fulfilment of the Requirements for the Degree of
Doctor of Philosophy in Operations Management

Dissertation Committee:

Soo-Haeng Cho (Chair)

Sridhar Tayur

Selman Erol

Alan Scheller-Wolf

James Best

May 2023

Abstract

In this dissertation, I investigate contemporary matching and allocation problems in various settings. In the first chapter, I provide remedies to overcome size-based disparities in the current liver allocation system. In the second chapter, I study the existence of popularity bias in the recommendation systems of online dating platforms. In the third chapter, I investigate how introducing protective measures to combat infectious diseases contribute to an increase in infection rates.

The first chapter studies the problem of achieving a fairer liver allocation system where there are disparities in organ access based on transplant patients' height and gender. Shorter candidates and women have higher average waiting times and mortality rates while waiting for a liver transplant because they can receive transplants from a smaller pool of available deceased donors than other candidates (e.g., tall patients, men). In order to address this problem, I model the liver transplant waitlist as a multiclass, heterogeneous multiserver overcrowded queueing system and use its first-order fluid approximation to solve for the objective of minimizing pre-transplant mortality over a finite time horizon with explicit equity constraints for all static patient classes, i.e., height. I solve the resulting optimal control problem to obtain the optimal policy of allocating deceased-donor livers, and I solve its discretized version with parameter estimates. I show that the optimal policy, the Equity Adjusted Mortality Risk policy, ranks patients according to their static patient classes and medical urgency to allocate deceased-donor livers. The shadow prices of the constraints in the optimal control problem can be mapped into the Model for End-Stage Liver Disease (MELD) exception points widely used in practice by the United Network for Organ Sharing System (UNOS). I provide these exception points to disadvantaged patient classes dynamically to move them to higher positions in the liver transplant waitlist to increase their chances of receiving a transplant. My simulations show that disadvantaged patient groups can significantly benefit from receiving

MELD exception points without decreasing the efficiency of the current liver allocation system. My work provides a remedy to reduce disparities in access to liver transplantation within the MELD-based allocation; hence, my framework of providing exception points can be used by policymakers to address inequity in practice.

The second chapter studies the existence of popularity bias in an online dating platform's recommendations and its consequences on users' likelihood of finding dating partners. Generating recommendations of compatible dating partners is challenging for online dating platforms because uncovering users' idiosyncratic preferences is complex. Thus, platforms tend to recommend popular users to others more frequently than unpopular users. From the data collected from a major online dating platform, I empirically find that a user's chance of being recommended by the platform's algorithm increases significantly with the user's popularity. Motivated by this empirical evidence, I study an online dating platform's incentive that generates popularity bias by modeling the platform's recommendations and users' subsequent interactions with a three-stage matching game. I also build a machine learning-based predictive model that estimates users' behavior and run simulations of the platform to validate our theoretical results. My analysis shows that the recommendations that maximize the platform's revenue and those that maximize the number of successful matches between users are not necessarily at odds, even though the former leads to a higher bias against unpopular users. Unbiased recommendations result in significantly lower revenue for the platform and fewer matches when users' implicit cost of evaluating incoming messages is low. Popular users help the platform generate more revenue and a higher number of successful matches as long as these popular users do not become "out of reach." My result indicates that an online dating platform can increase revenue and users' chances of finding dating partners simultaneously with a certain degree of bias against unpopular users. Online dating platforms can use my theoretical results to understand user behavior and my predictive model to improve their recommendation systems (e.g., by selecting a set of users leading to the highest probabilities of matching or other revenue-generating interactions).

The third chapter investigates the unintended consequences of introducing protective measures to combat infectious diseases. Higher availability and efficacy of protective measures against infectious diseases, such as vaccines, increase individuals' propensity to socialize. Consequently, the

number of visits to central points of interest (e.g., schools, gyms, grocery stores) and the rate of interactions with the agents employed therein (e.g., teachers, trainers, cashiers) increase. This opens more channels for the virus to transmit through the central agent or location, unlike in standard SIR models. This leads to a manifestation of network hazard. The infection rates can increase as protective measures become more effective and available. I confirm the testable predictions of the theory with the foot traffic data from 2019-2022 and historical COVID-19 vaccination and community transmission rates.

Acknowledgments

First and foremost, I would like to express my most sincere gratitude and appreciation to my advisor and dissertation chair, Prof. Soo-Haeng Cho, for his invaluable guidance, support, and encouragement throughout my doctoral journey. Your mentorship has been instrumental in helping me navigate the challenges of my Ph.D., and I am deeply grateful for your dedication to my success. I would also like to thank my dissertation committee members and co-authors, Prof. Mustafa Akan, Prof. Sridhar Tayur, and Prof. Selman Erol, whose expertise, insights, and suggestions were indispensable in the completion of my dissertation. I feel very fortunate to have had the opportunity to learn from you. Further thanks to my dissertation committee members, Prof. Alan Scheller-Wolf and Prof. James Best, for their insightful comments and discussions to improve my research.

My heartfelt gratitude goes to each and every one of my friends at Tepper who have been with me throughout my Ph.D. journey. I would like to thank Neda Mirzaeian for being a mentor and a supportive friend in challenging times. I also wish to thank Franco Berbeglia, Siddharth Singh, and Mehmet Aydemir for their guidance and honest feedback on my research. Whether it was lending a listening ear during moments of stress, providing a much-needed distraction from academic work, or simply being there to share in the joys and challenges of Ph.D. life, Sagnik Das, Violet Chen, Yuyan Wang, Zahra Ebrahimi, Martin Michelini, Sae-Seul Park, Kevin Mott, Savannah Tang, Neha, you have all played an integral role in my journey. I am grateful for the privilege of calling you my friends and colleagues. Many thanks to Lawrence Rapp and Laila Lee for their tireless efforts behind the scenes. Last but not least, my Turkish friends, Ozgun Elci, Ali Polat, Hakki Ozdenoren, Cagla Akin, Nilsu Uzunlar, and Orhun Gun, thank you for your unwavering support throughout these years.

To all my lifelong friends who have supported me through the ups and downs of my Ph.D. journey, thank you. I know I couldn't have made it through without you. My forever roommate and confidant, Melda Korkut, thank you for the support, the laughter, the tears, and all the late-night conversations that kept me grounded during this journey. You have always been there for me. The best friends a person can ask for, Dogan Ozen, Taha Erdogan, Gizem Caylak, and Utku Demiroz, thank you for everything you have done to support me, both personally and academically. I am grateful for your friendship.

This dissertation is dedicated to my mom, Aliye, whose love, sacrifices, vision, and encouragement have been an endless source of motivation in my life. Thank you for putting my and my sister's future above all, and thank you for believing in me during the most difficult times. I am also thankful for my dad, Ali. Thank you for expanding my horizon and pushing me to become a better version of myself. My lovely sister, Zehra, you always bring so much joy and comfort to my life. Even though we were apart during this journey, your love and understanding have made all the difference in the world. Thank you for being you. Finally, my one and only, Özlem, you joined me by the end of my Ph.D. journey but you cannot know how important for me to have your love and support. Thank you for your endless encouragement, guidance, and kindness. I'm looking forward to the next chapter of our lives together.

List of Coauthors

Chapter 1:

Mustafa Akan

Associate Professor of Operations Management, Tepper School of Business

Carnegie Mellon University

Sridhar Tayur

Ford Distinguished Research Chair

University Professor of Operations Management, Tepper School of Business

Carnegie Mellon University

Chapter 2:

Soo-Haeng Cho

IBM Professor of Operations Management and Strategy, Tepper School of Business

Carnegie Mellon University

Hyeunjung (Elina) Hwang

Associate Professor of Information Systems, Foster School of Business

University of Washington

Chapter 3:

Selman Erol

Assistant Professor of Economics, Tepper School of Business

Carnegie Mellon University

Contents

Abstract	ii
Acknowledgments	v
List of Coauthors	vii
0 Introduction	1
1 Dynamic Exception Points for Fair Liver Allocation	5
1.1 Introduction	5
1.2 Related Literature	9
1.3 A Fluid Model and Analysis	11
1.3.1 Model	11
1.3.2 Equity Adjusted Mortality Risk Policy	15
1.4 Providing MELD Exception Points	17
1.5 Simulation Results	19
1.6 Conclusion	26
2 Popularity Bias in Online Dating Platforms	29
2.1 Introduction	29
2.2 Related Literature	32
2.3 Empirical Setting and Data	35
2.4 A Two-Sided Matching Game	38
2.4.1 Model	39

2.4.2	The Platform’s Problem	45
2.4.2.1	Revenue-Maximizing Platform	45
2.4.2.2	Match-Maximizing Platform	47
2.5	Predictive Model	50
2.5.1	Simulations of the Platform	53
2.6	Extension: Two-Period Model	56
2.7	Conclusion	59
3	Network Hazard and Superspreaders	61
3.1	Introduction	61
3.2	Related Literature	63
3.3	Model	64
3.4	Data and Empirical Analysis	71
3.5	Conclusion	75
4	Conclusions	77
A	Additional Material for Chapter 1	80
A.1	Estimation of ESLD Patients’ 90-Day Mortality Risk	80
A.2	Summary of Notation	80
A.3	Proofs	81
A.4	Dynamics of the Liver Allocation System	88
A.5	Sized-Based Compatibility Analysis	90
B	Additional Material for Chapter 2	96
B.1	Description of User Attributes on the Platform	96
B.2	Summary Statistics of User Attributes	97
B.3	Correlation Matrices for User Attributes	98
B.4	Additional Discussion of Platform’s Recommendations and Regression Analysis on User Behavior	98
B.5	Summary of Notation	101
B.6	Proofs	101

B.7	Technical Details of Our Predictive Model	106
B.8	Impact of User Attributes on Users' Open, Send, and Accept Decisions	110
B.9	Extension: Match-Maximizing without Double Counting	115

List of Tables

1.1	Historical Data on Disparities in Access to Liver Transplant	6
1.2	Simulation Results on Equity	24
1.3	Simulation Results on Efficiency	25
1.4	Simulation Results on Recipients' Gender and Race	26
2.1	Summary of User Interaction Data	38
2.2	Regression Result on Platform's Recommendations	39
2.3	Accuracy of Prediction	52
2.4	Platform's Simulated Recommendations	54
3.1	Infection Probabilities of Agents	66
3.2	Summary of Cases of Equilibrium	68
3.3	Annual visits to points of interest	73
A.1	Logistic regression coefficients for predicting 90-day mortality risk	80
A.3	Size-Based Compatibility Matrix (%)	90
B.1	Regression Result on Users' Message Acceptance Rate	99

List of Figures

1.1	90-Day Mortality Risk of ESLD Patients	7
1.2	MELD Exception Point on Laboratory MELD - 90-day Mortality Risk Curve	18
1.3	Estimated Mortality Risk of Transplant Patients in Region 5	20
1.4	Optimization Model and Simulation Diagram	22
2.1	Sequence of Decisions in the Online Dating Platform	37
2.2	Numerical Experiment Results of Platform's Recommendations	50
2.3	Change in Platform's First-Period Recommendations with Growth Rate	58
3.1	Internal and Total Infections in Efficacy of Protection	69
3.2	Internal and Total Infections in Availability of Protection	70
3.3	Infection Transmission Level vs. Vaccination Rate in Allegheny County	72
3.4	Fast food restaurant visits	74
3.5	Gas station visits	74
3.6	Coffee shop visits	75
3.7	Gym visits	75
3.8	Airport visits	76
A.1	Diagram of the Liver Allocation System	89
A.2	Dynamics of class ij patients	90
A.3	BSA Histogram of ≤ 150 cm Donors	91
A.4	BSA Histogram of 151-156 cm Donors	91
A.5	BSA Histogram of 157-165 cm Donors	92
A.6	BSA Histogram of 166-175 cm Donors	92

A.7	BSA Histogram of ≥ 176 cm Donors	93
A.8	BSA Histogram of ≤ 150 cm Patients	93
A.9	BSA Histogram of 151-156 cm Patients	94
A.10	BSA Histogram of 157-165 cm Patients	94
A.11	BSA Histogram of 166-175 cm Patients	95
A.12	BSA Histogram of ≥ 176 cm Patients	95
B.1	Correlation Matrices for User Attributes	98
B.2	Percentage of Messages Sent over Users' Attractiveness Difference	100
B.3	User Interaction Data Tensor and Its Mode-1 Unfolding	107
B.4	Summary of Average SHAP Importance Scores Across All Six Decisions	111
B.5	(Color in PDF File) Individual SHAP Importance Scores of the Top 10 User At- tributes for Each of the Six Decisions	113

Chapter 0

Introduction

Matching and allocation problems have gained significant attention in the fields of operations management, economics, and computer science over the past few decades, as they deal with the allocation of scarce resources and the efficient and fair distribution of goods and services. These problems arise in various contexts, including labor markets, school choice, organ donation, and many others. This dissertation aims to contribute to the existing literature by studying the problems of fair allocation of deceased-donor livers, the existence and effects of popularity bias in online dating platforms' recommendation systems, and the network effects of introducing protective measures to combat infectious diseases.

Chapter 1 investigates the problem of achieving a fair liver allocation system where there are disparities in access to transplantation based on patients' height and gender. The historical data shows that shorter patients have a lower probability of receiving a deceased-donor liver transplant, and women, as they are of smaller stature in general, experience longer waiting times until transplantation. The reason is that disadvantaged patients can receive liver transplants from a smaller pool of available organs due to patient-donor size incompatibility. Currently, transplant patients are ranked based on medical urgency with respect to their Model for End-Stage Liver Disease (MELD) scores, which overlooks the aforementioned compatibility factor. We provide MELD score exception points to disadvantaged patients in order to move them to higher positions in the transplant waitlists to improve equity in the current liver allocation system.

We model the liver transplant waitlist as a multiclass, i.e., patient groups based on height,

overcrowded, i.e., the number of patients is significantly higher than the number of donors, queueing system with heterogenous servers, i.e., deceased donors with different liver sizes. In order to maintain tractability, we take its first-order fluid approximation and solve the resulting optimal control problem with the objective of minimizing pre-transplant mortality of patients with an explicit fairness constraint to ensure equity in access to transplantation. We show that the optimal policy, named Equity Adjusted Mortality Risk Policy, adjusts the MELD-based ranking in favor of the disadvantaged patient groups so that all patient groups have equal access to transplantation.

Implementing the Equity Adjusted Mortality Risk Policy in the current liver allocation system brings practical challenges. In order to address this challenge, we provide a computational framework to provide MELD score exception points for disadvantaged patient groups. Using these exception scores, disadvantaged patients' likelihood of receiving a deceased-donor liver transplant increase, and short-term mortality risk decrease. We test our exception points on a simulation of the national liver allocation system to examine how they affect various fairness and efficiency metrics of different patient groups as well as the overall system. Our simulations show that shorter transplant candidates' quality-adjusted life years and the likelihood of receiving a transplant increase while their mortality risk and the average waiting time until transplantation decrease. We also show that the overall efficiency of the current liver allocation system is not affected by the introduction of MELD exception points to shorter patients.

Chapter 2 studies the incentives for an online dating platform to employ recommendation algorithms favoring popular users and examines whether removing such popularity bias in the platform's recommendations would improve users' chances of finding compatible dating partners. Since generating recommendations of compatible dating partners is challenging for online dating platforms, they tend to recommend popular users more frequently than unpopular users, i.e., their algorithms are prone to popularity-based bias. Our empirical analysis of a major online dating platform's recommendations reveals that the users' chances of being recommended by the platform's algorithm increase significantly with their popularity among other users. This popularity bias results in congestion for popular users because popular users receive an excessively high number of messages, and it decreases unpopular users' chances of finding dating partners.

In order to address this problem, we model the platform's decision of recommending users and users' subsequent interactions with a three-stage matching game where users are divided into two

tiers, namely, popular and unpopular. We solve the platform's problem of maximizing expected revenue with recommendations and maximizing the expected number of matches between users. We use an unbiased platform as our benchmark for fairness to compare popular and unpopular users' matching probabilities, the platform's expected revenues, and the expected number of matches for three platforms. We empirically validate our model's assumptions using the extensive data collected from a major online dating platform in South Korea. We also utilize this data to predict users' future behavior on the platform which allows us to run simulations of the revenue-maximizing, match-maximizing, and unbiased platforms to test our theoretical results.

We show that the platform's revenue-maximizing and match-maximizing recommendations are not necessarily at odds when popular users are not selective in accepting messages. As popular users become more selective, the degree of bias in the match-maximizing platform's recommendations decreases significantly whereas it stays high in the revenue-maximizing platform's recommendations. Another factor that affects the degree of bias in a platform's recommendations is the congestion effect. When the congestion effect is high, both the revenue-maximizing and the match-maximizing platform's recommendations include popular and unpopular users more equally. Our simulations show that the match-maximizing recommendations are less biased against unpopular female users than the revenue-maximizing recommendations; however, both lead to a similar degree of bias against unpopular male users. Finally, we show that a platform can increase its revenue and users' chances of finding dating partners simultaneously with a certain degree of bias against unpopular users.

Chapter 3 describes a precautionary scenario on how unregulated levels of social activity respond to the high availability and efficacy of protective measures in combating infectious diseases. While protective measures such as vaccines are important tools in controlling the spread of infectious diseases such as COVID-19, they do not eliminate the risk of transmission entirely. However, individuals may feel more comfortable engaging in social activities such as going to restaurants, gyms, and grocery stores as they employ protective measures. This increase in social activity increases the risk of transmission in closed spaces. Also, individuals who work in central locations such as teachers, cashiers, and trainers can get infected and spread the infection to many visitors during their asymptomatic period. This is a manifestation of network hazard as the network of individuals becomes more concentrated for the virus to spread through the centers.

We consider a central location or agent such as a restaurant or a gym as the central node in a network and allow other agents to decide to form connections with the center to obtain certain benefits. As protective measures become more available and effective in reducing the transmission risk, more agents form connections with the center because the net utility of forming a connection, i.e., the difference between the utility of connection and the disutility of being infected, for each agent increases. As a direct result of this increased concentration in the network, we show that largely available and more effective protective measures have the direct effect of reducing aggregate infections; however, the indirect effect via increased contact through the center can potentially offset the direct positive effect in aggregate. The correlated nature of infections can also cause a disproportionately large number of simultaneous infections. Such superspreader events are particularly important to understand given the limited hospital capacities. Using monthly time-series data of visits to points of interest such as grocery stores, restaurants, and coffee shops, and the publicly available COVID-19 data, we find that the predictions of our model are observed in the data.

Chapter 4 concludes this dissertation, with a summary of our contributions and ideas for future research directions.

Chapter 1

Dynamic Exception Points for Fair Liver Allocation

1.1 Introduction

Liver transplantation is the only viable treatment for End-Stage Liver Disease (ESLD) and acute liver failure. Causes of ESLD include viral hepatitis, cirrhosis, non-alcoholic fatty liver disease, and hepatocellular carcinoma (HCC); it is the 12th leading cause of death in the United States (Cox-North et al. 2013). Patients with ESLD and acute liver failure join the transplant waiting list managed by the United Network for Organ Sharing System (UNOS) since the number of patients highly exceeds the number of available organs for transplantation. As of November 16, 2021, 11,864 patients are waiting for a liver transplant in the United States; 13,019 patients joined the liver transplant waiting list, but only 9,701 livers were donated in 2020. Due to the shortage of organ supply, the median waiting time until receiving a liver transplant is over three years for an adult, consequently, more than 40,000 patients died while waiting for a liver offer during 1995-2020.¹

Given the severity of ESLD and the patients' long waiting times until receiving a liver transplant, fairness in the allocation of the limited supply of organs becomes an important issue to be addressed by policymakers. In their three general principles of organ allocation, UNOS places fairness along with efficiency and respect for patients' autonomy in making their decisions of accepting/rejecting

¹Please see <https://unos.org/data/transplant-trends> for the recent transplant trends in the US.

Table 1.1: Historical Data on Disparities in Access to Liver Transplant

Height (cm)	Likelihood of Transplant (%)	Gender	Median Time Until Transplantation (months)
≤ 150	81	Men	3.6
151 - 165	89	Women	4.1
166 - 185	93		

(a) Based on patients' height

(b) Based on patients' gender

Notes. The data from the UCSF Liver Center is presented. **Likelihood of Transplant (%)** column reports the patients' likelihood of receiving a transplant relative to the patients with height ≥ 185 cm.

organs.² Within this context, we observe from the historical data that there are disparities in access to transplantation based on patients' height and gender. The data from the University of California San Francisco (UCSF) Liver Center, see Table 1.1(a) and 1.1(b), shows that shorter patients have a lower likelihood of receiving a liver transplant, and as expected, women experience longer waiting times until transplantation because they are of smaller stature in general.

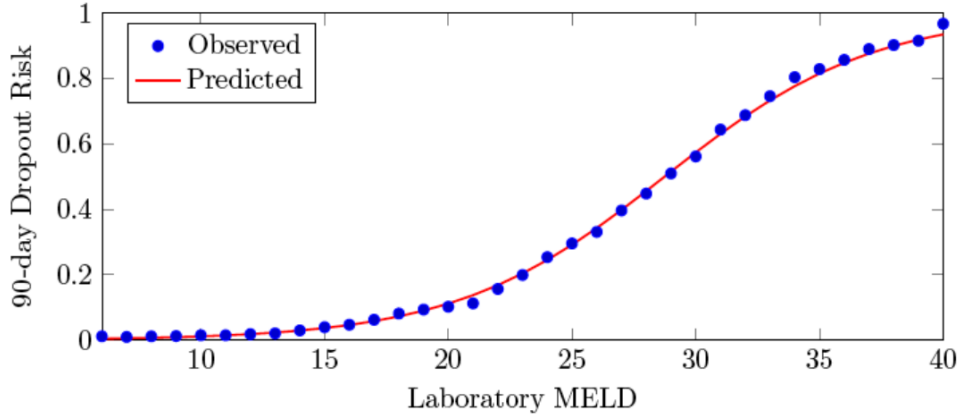
These disparities in organ access stem from the fact that disadvantaged patients can receive liver transplants from a smaller pool of available organs due to organ size incompatibility. Implantation of a large liver into a small recipient brings surgical difficulties (Reddy et al. 2013), and unmatched metabolic demand of recipient as well as physiologic mismatch aggravates the damage to liver graft, inevitably leading to graft failure (Fukazawa and Nishida 2016). The pool of adult donor livers has relatively few small livers nationally because most deceased-donor livers come from men; therefore smaller stature transplant candidates (e.g., women and Hispanic patients) are disadvantaged on the transplant waiting lists (Bernards et al. 2022).

In the current paradigm, the aforementioned donor-recipient compatibility factors are overlooked in ranking patients in the transplant waitlists because deceased-donor livers are allocated on the basis of medical urgency. A transplant patient has a Model for End-Stage Liver Disease (MELD) score that estimates a patient's chance of surviving their liver disease during the next three months. Figure 1.1 shows the relationship between ESLD patients' laboratory MELD score and their 90-day dropout risk. The MELD score of a patient is solely based on a patient's results from four blood tests³; it is updated more frequently as a patient's disease progresses (ranges from once

²"Ethical Principles in the Allocation of Human Organs" can be accessed at <https://optn.transplant.hrsa.gov/resources/ethics/ethical-principles-in-the-allocation-of-human-organs/>.

³MELD = $3.78 \times \ln[\text{serum bilirubin (mg/dL)}] + 11.2 \times \ln[\text{INR}] + 9.57 \times \ln[\text{serum creatinine (mg/dL)}]$

Figure 1.1: 90-Day Mortality Risk of ESLD Patients



Notes. The logistic regression coefficients can be found in Appendix A.1.

a week to once a year), and a higher MELD score indicates that a patient needs a liver transplant more urgently. When a deceased-donor liver becomes available, UNOS sequentially offers this liver to compatible patients who are ranked by their MELD scores after considering their geographical proximity to the donor. The number of offers is limited due to the cold-ischemia time of a liver (i.e., the time before a liver loses its functionality); the liver is discarded if a patient (or a surgeon) does not accept the organ on time.

The MELD score is an excellent predictor of survival for more than 70% patients on the transplant waitlists (Godfrey et al. 2019); however, some patients' severity of illness or risk of complications are not captured by their laboratory MELD scores. Transplant candidates whose MELD scores under-predict their short-term mortality risk apply for a MELD score exception to be placed in a higher position on the transplant waitlist. These exception scores are widely used in MELD-based allocation of deceased-donor livers; Hepatocellular Carcinoma (HCC) is the most common reason for MELD score exceptions along with 17 groups of diagnoses (Asrani and Kamath 2015). Despite their practical use by policymakers in order to overcome disparities in access to transplantation, the number of exception points allowed for diagnoses other than HCC has not been clearly defined (Massie et al. 2011).

In this chapter, we address the inequity in access to transplantation in the current liver allocation system where disadvantaged patient groups (e.g., short candidates, women) experience a longer time until transplantation and have a lower probability of receiving a liver transplant. For this,

we model the liver transplant waitlist as a multiclass overcrowded queueing system (patient groups based on height) with heterogeneous servers (deceased donors with different liver sizes) and take its first-order fluid approximation for tractability. We solve the resulting optimal control problem of minimizing pre-transplant mortality of patients over a finite time horizon with an explicit fairness constraint that equalizes the likelihood of receiving a transplant for all patient groups. We show that the optimal policy, Equity Adjusted Mortality Risk Policy, ranks patients with respect to their medical urgency, as in the current MELD-based allocation, but adjusts this ranking in favor of the disadvantaged patient groups so that all patient groups have equal access to transplantation.

In addition to our theoretical results, we provide a computational framework to provide MELD score exceptions for disadvantaged patient groups to achieve a fairer MELD-based liver allocation system in practice. For this purpose, we show that the shadow prices of the optimal control problem can be mapped into the MELD score exceptions for disadvantaged patient groups. Unlike the non-HCC MELD score exceptions used in practice, our proposed exception points are closely tied to the patients' short-term mortality risk while waiting for a transplant. Using these exception scores, we move disadvantaged patients to higher positions on the transplant waitlist to increase their likelihood of receiving a liver transplant. We test our exception points on a simulation of the national liver allocation system to analyze their performance on various fairness (e.g., average time until receiving a transplant, likelihood of transplantation) and efficiency (e.g., pre-transplant mortality, number of wasted organs, quality-adjusted life years) metrics over different patient groups as well as the entire system.

Our simulations show that ≤ 156 cm transplant candidates' (mostly female and Hispanic patients) likelihood of receiving a liver transplant improves significantly. As a result, their likelihood of death while waiting for a liver transplant decreases. Our computational framework ensures that all patient groups' access to transplantation converges to each other. We also show that ≤ 156 cm transplant candidates' quality-adjusted life years increase with our MELD exception points while their number of death on the transplant waiting list and the average waiting time until transplantation decrease. Overall, disadvantaged patient groups can greatly benefit from receiving MELD exception points without decreasing the efficiency of the current liver allocation system.

The rest of this chapter is organized as follows. In §1.2, we review the related literature and discuss our contributions. §1.3 presents the fluid model, the resulting optimal control problem,

and the optimal policy of allocating deceased-donor livers. Motivated by our optimal policy, we show how we provide MELD exception points to disadvantaged patient groups in §1.4. In §1.5, we present the results from our simulation study to discuss how equity and efficiency metrics are affected by introducing MELD exception points. We conclude this chapter and discuss potential research directions in §1.6.

1.2 Related Literature

This paper is broadly related to the literature on two-sided matching queues, where heterogeneous supply types (donor organs with different blood types and sizes) are matched with only a subset of demand types (patients with different blood types and heights). Under such compatibility constraints, Gurvich and Ward (2015) study the problem of minimizing finite-horizon cumulative holding costs of items, and Ashlagi et al. (2019) study the waiting time performance of various dynamic control policies in two-sided matching queues. Similarly, Nazari and Stolyar (2019) propose an optimal matching policy that maximizes the long-run average matching rewards while keeping queues stable, and Hu and Zhou (2022) design the optimal matching policies that maximize total discounted rewards. Afeche et al. (2022) find the optimal design of service compatibility topologies given the trade-off between customers' waiting time delays and maximizing match rewards. The main difference in transplant queueing systems is that the patients change classes while waiting for a transplant (e.g., health deterioration), leading us to use the first-order fluid approximation of the queueing system to minimize pre-transplant mortality in finite horizon under equity constraints.

The organ (e.g., kidney, liver, heart) allocation systems have been studied extensively by economists and operations researchers. In their early work, David and Yechiali (1985) consider a patient's problem of accepting a kidney offer as a time-dependent stopping problem to maximize their expected reward from transplantation. Righter (1989) models the kidney allocation process as a stochastic assignment problem with the objective of maximizing the total expected reward. Similarly, Ahn and Hornberger (1996) and Howard (2002) solve a transplant patient's problem of accepting/rejecting a kidney offer, and Alagoz et al. (2004, 2007) and Sandıkçı et al. (2008) consider a patient's problem for a liver offer. In a series of papers, Su and Zenios (2004, 2005, 2006) study how patient choice impacts the kidney allocation mechanisms. From a geographical perspective,

Kong et al. (2010) solve the problem of maximizing the efficiency of the liver allocation system, and Ata et al. (2017) address the region-based inequity in access to kidney transplantation. Bertsimas et al. (2013) propose a data-driven method for designing national policies for kidney allocation with the objectives of fairness and efficiency. Dai et al. (2020) analyze the welfare consequences of introducing the donor-priority rule, which grants registered organ donors priority in receiving organs in case they need transplants in the future.

Within this extensive literature, our paper contributes to the stream that focuses on the optimal design of organ allocation policies using the first-order fluid approximation of the transplant system. Zenios et al. (2000) find the best kidney allocation policy with the trade-off between clinical efficiency (i.e., QALY) and equity in access to transplantation, Akan et al. (2012) design optimal liver allocation policies where the trade-off is between medical urgency (i.e., the total number of patient deaths) and efficiency, and Hasankhani and Khademi (2021) propose optimal policies of allocating hearts with the trade-off between efficiency and equity. Given the prevalence of medical urgency in the current liver allocation system (i.e., MELD-based allocation), we restrict our focus to improving equity in access to liver transplantation while keeping medical urgency as our objective. In addition to proposing the optimal policy of allocating deceased-donor livers, differently from this stream of literature, we show that the shadow prices of the optimal control problem can be used to estimate transplant patients' short-term mortality risk; hence, we utilize our fluid model to introduce MELD exception points that can be used by policymakers to improve equity within the current MELD-based system.

Medical scientists study the problem of providing novel MELD exception points to patients with HCC whose mortality risks are under-predicted by their MELD scores. Toso et al. (2012) use a proportional hazard model, Vitale et al. (2014) run multivariable regressions, and Marvin et al. (2015) use a Cox regression model to estimate the short-term mortality risk of patients with HCC to provide MELD exception points. While this literature lacks studies to provide model-based MELD exception points to non-HCC patients (e.g., patients with cystic fibrosis, hepatopulmonary syndrome, etc.), our computational framework of providing MELD exception points can be generalized with non-HCC patients who apply for MELD exception points. Bernards et al. (2022) provide static MELD exception points to short patients regardless of their laboratory MELD scores. Our model takes the dynamics of the liver transplant waitlist (e.g., patients' health evolution, patient/donor

arrivals, mortality, etc.) into account, endogenously calculates transplant patients’ short-term mortality risk that can be directly mapped into MELD score exceptions, and provide different MELD exception points to disadvantaged patients based on candidates’ laboratory MELD scores.

Using the publicly available data from the Scientific Registry of Transplant Recipients (SRTR), we test our proposed MELD exception points in the simulation of the national liver allocation system to evaluate their impact on various efficiency and equity metrics for different patient groups as well as the overall system. Ruth et al. (1985) and Pritsker et al. (1995) are among the early papers that analyze the kidney and liver allocation systems via simulation, respectively. Kreke et al. (2002) and Shechter et al. (2005) incorporate the patients’ disease evolution, and Kim et al. (2015) develop a machine learning-based model to incorporate transplant patients’ accept/reject decisions into the simulations of the liver allocation system. Davis et al. (2013) and Sandıkçı et al. (2019) develop discrete-event simulations of the national kidney allocation system to evaluate potential policy changes in allocating kidneys. We incorporate the discretized version of our fluid model with Liver Simulated Allocation Model (LSAM) developed by SRTR to provide MELD exception points to disadvantaged patient groups in order to increase their access to deceased-donor livers.

1.3 A Fluid Model and Analysis

We model the liver transplant waitlist as a multiclass overcrowded queueing system with heterogeneous servers. Given the complex dynamics of this problem (patients’ health evolution, mortality, etc.), we use the first-order fluid approximation of the queueing system to solve the resulting optimal control problem in the finite horizon. We introduce our fluid model with the objective of minimizing patients’ pre-transplant mortality in the system in §1.3.1, propose our optimal policy, and discuss the interpretation of the shadow prices of the optimal control problem in §1.3.2. Table A.2 summarizes our notation. Proofs of our results can be found in Appendix A.3.

1.3.1 Model

We construct a stylized fluid model to characterize the liver allocation process and to track the dynamics of the system. An overview of this section is as follows. We first describe ESLD patients and their dynamics while waiting for a liver transplant. Next, we describe the deceased-donor livers

being harvested for the transplant patients, control variables that correspond to the allocation of donor livers to patients and the state of the system. Finally, we formulate our objective function as minimizing pre-transplant mortality in finite horizon with fairness constraints for different patient classes to ensure equity in access to transplantation.

We divide the ESLD patients who wait for a liver transplant into different classes along two dimensions: static patient characteristics such as height, and dynamic patient characteristics that represent their health status, i.e., laboratory MELD score. The former dimension is denoted by $i \in \mathcal{I} := \{1, 2, \dots, I\}$ where I is the total number of patient groups, and the latter dimension is denoted by $j \in \mathcal{J} := \{1, 2, \dots, J\}$ where J is the total number of laboratory MELD scores.⁴ MELD score of a patient may change over time, implying that the patients' dynamic class might change in our model. The diagram of the class structure of the ESLD patients can be found in Appendix A.4.

Patients of class ij arrive at the liver transplant waitlist with rate $\lambda_{ij}(t)$ for $t \geq 0$, and the number of class ij patients waiting for transplantation at time t is denoted by $x_{ij}(t)$; initially, there are $x_{ij}(0)$ patients in class ij . As we mentioned earlier, patients' health status (i.e., MELD score) changes while waiting for a transplant. In many cases, a patient's health condition deteriorates, leading to an increase in their MELD scores; however, it is also possible for some patients (e.g., patients with primary biliary cirrhosis) to experience a temporary recovery when they first join the waitlist leading to a decrease in their MELD scores. To be specific, we denote the rate at which a patient's MELD score changes from j to j' with $\alpha_{jj'}$ for $(j, j') \in \mathcal{J} \times \mathcal{J}$ without any structural assumptions. The rate at which a patient with MELD score j dies while waiting for a transplant is denoted by d_j . The patients with higher MELD scores are more likely to die; therefore, we assume that $d_j > d_{j'}$ for $j > j'$.

A type k deceased-donor liver arrives at the liver transplant system with rate $\mu_k(t)$ for $k \in \mathcal{K} := \{1, \dots, K\}$ and $t \geq 0$. The type of a liver is defined by its blood type and size. We denote the rate at which type k livers are dynamically allocated to class ij patients by $u_{ijk}(t)$ for $i \in \mathcal{I}$, $j \in \mathcal{J}$, $k \in \mathcal{K}$ and $t \geq 0$. The static type of a patient (i.e., height), denoted by i , must be compatible with the liver type k (i.e., size) so that a type k deceased-donor liver can be offered to class ij patients for

⁴In practice, there are 35 dynamic patient classes, i.e., $J = 35$, because MELD score takes integer values between 6 and 40.

all $j \in \mathcal{J}$. We ensure this with an incompatibility constraint on the control variables, specifically, $u_{ijk}(t) = 0$ for $(i, k) \in INF$, all $j \in \mathcal{J}$ and $t \geq 0$ where INF contains the incompatible patient type - donor liver type pairs.

Patients have the option to reject offered livers due to the expectation of receiving a better organ offer in the future. We denote the probability of a class ij patient accepting a type k liver by p_{ijk} . In our model, a deceased-donor liver can be offered to multiple patients on the waitlist. If a type k liver is offered to n patients of class ij , the probability of the liver being rejected by all patients, i.e., the organ is wasted, becomes $(1 - p_{ijk})^n$. As a result, the probability of a type k liver being transplanted when it is offered to n class ij patients, π_{ijk}^n , becomes $\pi_{ijk}^n = 1 - (1 - p_{ijk})^n$.

The state of the system is denoted by $x(t)$ that keeps track of the number of patients in each patient class at time t , i.e., $x(t) = (x_{11}(t), \dots, x_{ij}(t), \dots, x_{IJ}(t))^T$ for $t \geq 0$. Similarly, we denote the IJ -dimensional vector of control variables for each liver type k by $u^k(t)$ where $u^k(t) = (u_{11k}(t), \dots, u_{ijk}(t), \dots, u_{IJk}(t))^T$ for $k \in \mathcal{K}$ and $t \geq 0$. A feasible control $u(t)$ needs to satisfy three sets of constraints: (i) the total allocation of type k livers cannot exceed the supply of livers of the same type, (ii) the allocation of deceased-donor livers for the incompatible patient - donor type pairs must be zero, and (iii) the allocation of livers for the compatible patient - donor type pairs must be non-negative. Therefore, we define the set of feasible controls, $\Phi(t)$, as follows:

$$u(t) \in \Phi(t) := \{u(t) : e \cdot u^k(t) \leq \mu^k(t); u_{ijk}(t) = 0 \forall j \in \mathcal{J}, (i, k) \in INF; u^k(t) \geq 0\} \quad (1.1)$$

where e is an IJ -dimensional vector of ones.

Given a feasible control u , the system state evolves as follows:

$$\dot{x}(t) = \lambda(t) - \sum_{k=1}^K P^k u^k(t) - (d + \beta - \gamma)x(t), \quad t \geq 0, \quad (1.2)$$

where P^k is an $IJ \times IJ$ dimensional diagonal matrix with entries π_{ijk} for $i \in \mathcal{I}$, $j \in \mathcal{J}$ and each liver type $k \in \mathcal{K}$. $\lambda(t)$ is the IJ -dimensional vector of arrival rates of patients, $\lambda_{ij}(t)$, at time t . The square matrix of d has shape $IJ \times IJ$, and includes the death rate of patients for each MELD score, d_j , in its diagonal entries. Similarly, the square matrices β and γ , obtained from matrix α , include the health transition of patients at each MELD score. The former, β , contains the health

transition rate of patients out of each MELD score to other MELD scores, i.e., it has $\sum_{j \neq j'} \alpha_{jj'}$ and $\alpha_{jj} = 0$ in its diagonal entries for each $j \in \mathcal{J}$. The latter, γ has a shape $IJ \times IJ$, includes the health transition rate of patients into each MELD score from other MELD scores, i.e., it has $\alpha_{j'j}$ for each $j \in \mathcal{J}$ in block diagonal matrices of shape $J \times J$. Finally, we require that the number of patients in each class must be non-negative, that is,

$$x(t) \geq 0 \text{ for } t \geq 0. \quad (1.3)$$

Our aim is to come up with MELD exception points for disadvantaged patient groups to ensure their equal access to transplantation. For this reason, we focus on equalizing the likelihood of transplantation measure across all patient classes with respect to their static characteristics. The following constraint ensures that the ratio of the total amount of allocated deceased-donor livers to the total arrival rate of each static patient class must be the same:

$$\int_0^T \sum_{k=1}^K e \cdot u_i^k(t) dt = \frac{1}{\rho} \lambda_i T \text{ for } i \in \mathcal{I}, \quad (1.4)$$

where $1/\rho$ is the average likelihood of transplantation, $\lambda_i T$ is the total arrival rate of patients of class i , equals to $\int_0^T \sum_j \lambda_{ij}(t)$ over the finite time horizon, and $\int_0^T \sum_{k=1}^K e \cdot u_i^k(t) dt$ gives the total amount of allocated livers to class i patients. In reality, disadvantaged patient classes experience a lower likelihood of transplantation, i.e., the value of $1/\rho_i$ for disadvantaged patient class i is lower than the value of $1/\rho_{i'}$ for $i' \in \mathcal{I}$, so we ensure that the likelihood of transplantation is equal for all patients by enforcing the same value, $1/\rho$, independent of static patient classes.

We reflect the current liver allocation policy that prioritizes medical urgency (MELD-based allocation) with our objective function of minimizing pre-transplant mortality of patients while waiting for a liver transplant. Since patients with higher MELD scores have a higher mortality rate, this objective ensures that they are prioritized over patients with lower MELD scores in receiving liver transplant offers. As a result, the problem of minimizing pre-transplant mortality of patients while waiting for a transplant with the equity constraint becomes choosing an organ

allocation policy u so as to

$$\text{minimize } \int_0^T (e \cdot d) \cdot x(t) dt \quad \text{subject to (1.1) - (1.4)} \quad (\text{P})$$

with the initial state of the system $x(0) = x_0$.

1.3.2 Equity Adjusted Mortality Risk Policy

In this section, we first present the optimal policy of allocating deceased-donor livers to ESLD patients under fairness constraints by introducing the shadow prices from the dual problem formulation (D) of our optimal control problem (P). Next, we describe the implementation of the optimal policy, Equity Adjusted Mortality Risk Policy, from a policymaker's perspective. Finally, we provide the interpretation of the shadow prices from the dual problem that lays the groundwork for our computational framework of providing MELD exception points to disadvantaged patient groups.

The likelihood of transplant constraint (1.4) in the primal problem (P) requires the integration of the control variable $u(t)$, which makes the constraint nonlinear in $u(t)$. We linearize this constraint by introducing another state variable $w(t)$ where $w_i^k(t) = \int_0^t u_i^k(\tau) d\tau$, $\dot{w}_i^k(t) = u_i^k(t) dt$, $w_i^k(0) = 0$ for $i \in \mathcal{I}$, $k \in \mathcal{K}$ and $t \geq 0$. As a result, the nonlinear constraint (1.4) is replaced by a terminal condition as:

$$\sum_{k=1}^K e \cdot w_i^k(T) = \frac{1}{\rho} \lambda_i T \quad \text{for } i \in \mathcal{I}. \quad (1.5)$$

We adopt the general method introduced by Rockafellar and Wets (2009) to derive the dual control problem (D). In the dual formulation, the state vector $y_{ij}(t)$ is the shadow price that corresponds to the ij^{th} system evolution constraint (1.2) in the primal problem (P), $z_i^k(t)$ is the ik^{th} shadow price corresponding to the evolution of the control variable u , i.e., $\dot{w}_i^k(t) = u_i^k(t) dt$, and q_i is the i^{th} shadow price corresponding to the likelihood of transplant constraint for patient class i . Given the dual state vectors of (D), the structure of the optimal policy is presented in Theorem 1.

Theorem 1. *The primal problem (P) of minimizing pre-transplant mortality and the dual problem (D) have the same objective value. Furthermore, letting u and (y, z) pair be a feasible organ allocation policies for (P) and (D), the primal control u and the dual control (y, z) are optimal for (P) and*

(D), respectively, if and only if they satisfy the coextremality conditions given below. For $i \in \mathcal{I}$, $j \in \mathcal{J}$ and $t \in [0, T]$,

$$\dot{y}_{ij}(t) = d_{ij} + [y(t)(d + \beta - \gamma)]_{ij} \text{ if } x_{ij}(t) > 0, \quad (1.6)$$

$$\dot{z}(t) = 0, z(T) = \frac{q}{\rho}\lambda T, \quad (1.7)$$

$$u^k(t) \in \arg \min_{v \in \Phi(t)} \{(P^k \cdot y(t) - z(t))v^k\} \text{ for } k \in \mathcal{K}. \quad (1.8)$$

Equation (1.8) characterizes the optimal allocation of deceased-donor livers to minimize pre-transplant mortality in finite horizon. When a liver of type k arrives at time t , a policymaker ranks patients in class ij with respect to the quantity $y_{ij}(t)\pi_{ijk}^n - z_i(t)$ where the number of parallel organ offers is n and class ij patients' probability of accepting the liver offer is π_{ijk}^n . With the objective function of minimizing pre-transplant mortality, the dual state variable $y_{ij}(t)$ gives the potential increase in the objective function if we were to increase the number of patients in class ij by one, in other words, it gives us the mortality risk of an additional patient of class ij at the end of the time horizon. The dual state variable $z_i(t)$ gives the potential decrease in the objective function if we were to increase the likelihood of transplant of class i patients by 1%, i.e., it gives us the mortality risk of class i patients that can be avoided by increasing their access to transplantation. Therefore, the optimal policy, Equity Adjusted Mortality Risk Policy, non-decreasingly orders patients in terms of their adjusted mortality risk to allocate deceased-donor livers to the ESLD patients.

In the current liver allocation system, transplant patients are prioritized with respect to their medical urgency, i.e., MELD score, and their access to transplantation is not considered a factor while allocating deceased-donor livers. This is reflected in the first term of the optimal policy, $y_{ij}(t)\pi_{ijk}^n$, in the absence of fairness constraints for a fixed patient class i . Since patients with higher MELD scores have a higher mortality risk, i.e., $y_{ij}(t) < y_{ij'}(t)$ for $j < j'$ and $t \in [0, T]$, and have a higher probability of accepting incoming liver offers, i.e., $\pi_{ijk}^n < \pi_{ij'k}^n$ for $j < j'$, a policy maker offers a transplant organ to patients of class i starting from the patients with the highest MELD score. Adding the fairness constraints to the primal problem (P) brings a new term, $z_i(t)$, that allows for the case where a disadvantaged patient with relatively better health status, i.e., lower laboratory MELD score, might be prioritized to increase their access to transplantation.

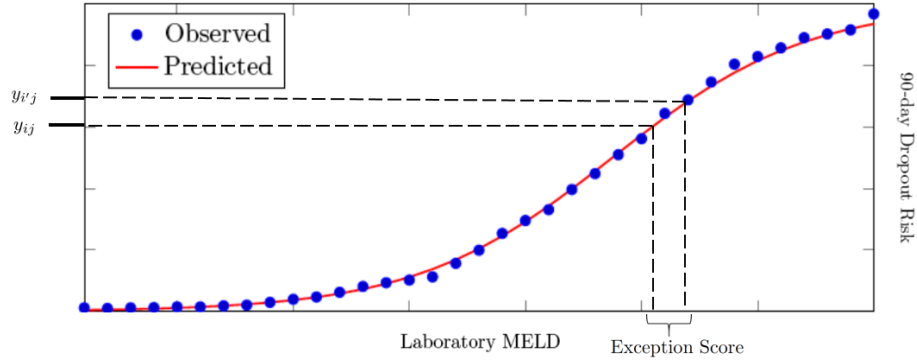
1.4 Providing MELD Exception Points

Inspired by the optimal policy of allocating deceased-donor livers to transplant patients under fairness constraints, we provide the framework of providing MELD score exception points to disadvantaged patient groups in this section. First, we establish the connection between the dual state variables of the optimal control problem and MELD exception points. Then, we provide an easy-to-implement algorithm to provide exception points by using parameter estimates. Finally, we demonstrate a numerical example to explain how MELD score exceptions would be provided to disadvantaged patient groups by the policymakers in practice.

For ease of discussion, we consider patients in two static classes where i represents the regular class and i' represents the disadvantaged patient class with a lower likelihood of receiving a liver transplant. As we have seen in the previous section, the dual state variable that corresponds to the ij th system evolution constraint (1.2) in the primal problem (P), $y_{ij}(t)$, gives us the mortality risk of a class ij patient at the end of the time horizon for $i \in \mathcal{I}$, $j \in \mathcal{J}$ and $t \in [0, T]$. We solve (P) by replacing $1/\rho$ with $1/\rho_i$ for class i patients and $1/\rho_{i'}$ for class i' patients in constraint (4). This gives us $y_{ij}(t)$ for a patient who belongs to the regular class with MELD score j and $y_{i'j}(t)$ for a patient who belongs to the disadvantaged class with the same MELD score where $y_{i'j}(t) < y_{ij}(t)$ because the likelihood of receiving a transplant for a class i patient, $1/\rho_i$, is greater than for a class i' patient, $1/\rho_{i'}$. The closest integer to the inverse of the difference between $y_{ij}(t)$ and $y_{i'j}(t)$ on the MELD score - short-term mortality risk curve becomes the MELD score exception point for the class $i'j$ patient. Figure 1.2 visualizes the MELD exception points for the disadvantaged patients in class i' at MELD score j .

The primal optimal control problem (P) is linear with respect to the state variable x and the control variable u ; therefore, the discretized version of it turns out to be a linear program that can be efficiently solved using parameter estimates. The dual state variable y can also be easily extracted from the primal problem letting us provide MELD exception points to disadvantaged patient groups. Given the parameter estimates of the primal problem, Algorithm 1 is easy to use by a policymaker who aims to increase disadvantaged patient groups' access to transplantation. The parameter estimates include the initial state of the system, i.e., the number of patients in each class x_0 , the average arrival rate of patients into the transplant waitlist, λ , the average arrival rate

Figure 1.2: MELD Exception Point on Laboratory MELD - 90-day Mortality Risk Curve



of deceased-donor livers, μ , the transplant acceptance probability of patients in each class, P , the rate of health transitions of patients in each MELD score, α , the mortality risk of patients, d , the feasible set of organ allocations, Φ , and the average likelihood of transplantation of each patient class, $1/\rho$. We note that the historical data or the forecasts for a pre-specified time horizon (3 months, 1 year) can be used to estimate parameters such as average arrival rates and the likelihood of receiving a liver transplant. After the linear program is solved using the estimated parameters, one can check the dual values of each patient class corresponding to the system dynamics constraints at each MELD score and give the exception points to the disadvantaged patient groups whose dual values appear to be higher than other patient groups.

For the numerical demonstration, we use the data from UNOS Region 5, which includes the states of Arizona, California, Nevada, New Mexico, and Utah, for the years between 2012 and 2017 to estimate transplant patient and deceased-donor arrival rates. We abstract away the blood type matching and focus only on the size matching to classify patients into different groups with respect to their heights. As we mentioned earlier, shorter patients can receive transplantation from a smaller pool of available donors, i.e., only small size deceased-donor livers, compared to medium-height and tall patients. As in Lai et al. (2010), we consider patients taller than 180 cm as tall, between 165-180 cm as medium-height, and shorter than 165 cm as short patients. The primal optimization problem (P) is solved for varying time horizons, and the results from the 1-year horizon are presented for brevity. Figure 1.3 shows the dual values that are endogenously calculated from the primal problem, i.e., tall, medium, and short patients' mortality risk at each MELD score, and the differences between dual values of all patient class pairs.

Algorithm 1 Providing MELD Exception Points

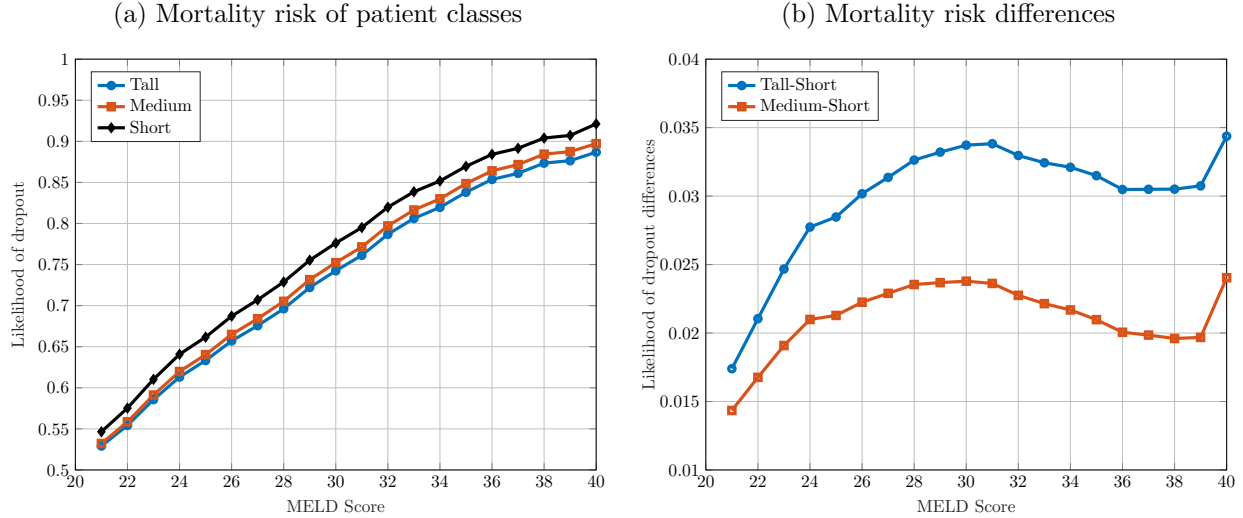
Require: $x_0, \lambda, \mu, P, \alpha, d, \Phi, 1/\rho$
discretize the time horizon;
solve the resulting Linear Program (LP);
for Patient class pairs (i,i') **do**
 for MELD score j **do**
 if $y_{i'j} > y_{ij}$ **then**
 $j' \leftarrow$ inverse of $d_j + y_{i'j} - y_{ij}$ on MELD - 90-day mortality risk curve
 Exception Point $_{i'j} \leftarrow \lfloor j' - j \rfloor$
 end if
 end for
end for

Figure 1.3(a) shows us that the mortality risk of all patients for low MELD scores, i.e., < 20 , are very close to each other regardless of their height because ESLD patients rarely receive transplantation when their health condition is relatively well. We do not provide exception points to short patients in this interval. The patients on the transplant waitlist receive deceased-donor liver offers for MELD scores higher than 20; therefore, we observe the discrepancy between the short, medium-height, and tall patients' mortality risks for higher MELD scores due to the differences in their access to transplantation. Using the differences in the mortality risk of patients in Figure 1.3(b), we provide MELD score exception points to short patients. For MELD scores between 21 and 34, the difference between the mortality risk of tall and short patients grants +1 exception point to short patients to artificially move them to higher positions in the transplant waitlist. For MELD scores between 35 and 38, short patients are granted +2 exception points because the difference between the mortality risks corresponds to a higher MELD score difference on the laboratory MELD score - short-term mortality risk curve. Medium-height patients do not qualify for MELD score exceptions because the difference between their mortality risk and tall patients' is not high enough for them to be granted. These exception points improve short patients' access to transplantation to decrease their mortality risk while waiting for a liver transplant.

1.5 Simulation Results

In this section, we present the results from the simulations of the national liver allocation system to study how various efficiency and equity metrics are affected by our proposed MELD score exception

Figure 1.3: Estimated Mortality Risk of Transplant Patients in Region 5



points to disadvantaged patient groups. First, we introduce the simulation model and describe the system dynamics that are updated over the course of the simulation. Second, we explain how our computational procedure is incorporated into the simulations to compute MELD exception points. Then, we present the simulation results on equity with our proposed exception points and compare the potential improvement over the current policy and static exception points. Improving the equity of a system comes at the cost of losing efficiency in general; therefore, we present the simulation results on efficiency metrics for each patient group as well as the overall system. Finally, we discuss how our proposed MELD exception points affect the trade-off between efficiency and equity.

We use the Liver Simulated Allocation Model (LSAM), a computer simulation program developed by the Scientific Registry of Transplant Recipients (SRTR), to simulate the allocation of livers to candidates on the Organ Procurement and Transplantation Network (OPTN) waiting list. A more detailed description of the simulation model and its validation can be found in SRTR (2019). The liver allocation system comprises 11 regions and 58 donor service areas (DSA). When transplant candidates arrive at a particular DSA, they are assigned a laboratory MELD score, blood type, Status 1 exception (for critically ill patients), and HCC exception. In addition to the aforementioned exception points that LSAM provides, we provide MELD exception points to patients based on their height and laboratory MELD score. During the simulation, transplant candidates' laboratory MELD scores are updated, reflecting the changes in their health status, and they are

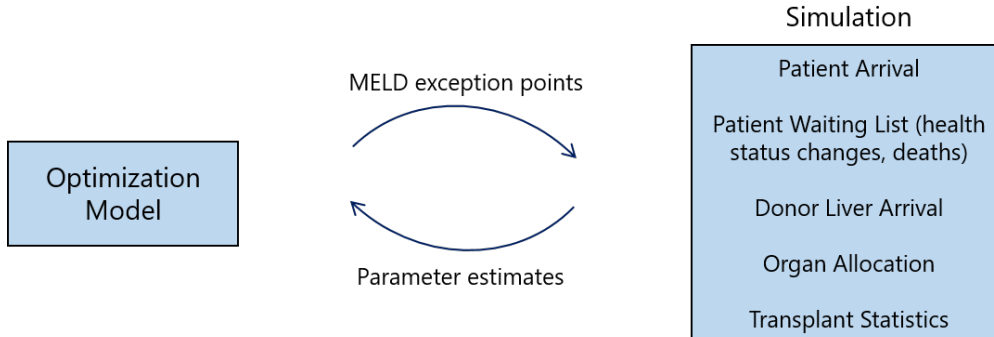
removed from the waitlist due to death or other reasons (e.g., being medically unsuitable for transplantation). After a deceased donor arrives, the liver is assigned a blood type with the donor’s height information, and it is offered to blood type and size-compatible candidates in accordance with the current allocation policies. The patient preferences module computes the acceptance probability of each transplant candidate when a deceased-donor liver becomes available, and the organ is discarded if it is rejected by all candidates who receive an offer. Lastly, the transplant statistics module computes the related performance metrics of the overall system that we discuss later in this section.

We run LSAM simulations of the current policy and our proposed policy for a 5-year horizon with three replications. In our simulation study, we focus on providing MELD score exception points based on transplant recipients’ height. Extending the example based on three height groups in §1.4, we divide transplant patients into five height tiers (≤ 150 cm, 151-156 cm, 157-165 cm, 166-175 cm, ≥ 176 cm) that align with the medical literature on liver transplantation (Bernards et al. 2022). In order to capture more granular differences in the most disadvantaged patients, the shortest patients are split into smaller tiers compared to tall patients. The data of transplant candidates and donors are collected from LSAM input files for the study period of July 1, 2011, to June 30, 2016. The average results are presented in this section for brevity.

The interaction between the simulation of the national liver allocation system and our optimization model is visualized in Figure 1.4. For a 1-year period, our simulation runs with the UNOS’ current policy of allocating livers based on medical urgency to estimate each patient group’s likelihood of receiving a liver transplant. Along with the estimated rates of patient and donor liver arrivals, and waiting list dynamics, these parameters are used in our optimization model to endogenously compute the mortality risk of patients that result in MELD score exception points for short patients in Algorithm 1. These exception points are then fed into the simulation to evaluate the changes in the performance metrics. Suppose there are still discrepancies between different patient classes regarding their likelihood of transplantation. In that case, we repeat the steps we have described so far until each patient group’s likelihood of transplantation converges to each other close enough.

We define size-based compatibility in deceased-donor liver transplantation with respect to the transplant recipients’ and donors’ heights by using the Body Surface Area Index (BSA_i) thresholds

Figure 1.4: Optimization Model and Simulation Diagram



in Fukazawa and Nishida (2016). Body Surface Area Index (BSA_i) of a patient-donor pair is the ratio of a donor’s Body Surface Area (BSA)⁵, correlated to the liver size, to a recipient’s BSA , correlated to the available abdominal volume. A donor liver is considered ”small-for-size” where $BSA_i < 0.78$, and ”large-for-size” where $BSA_i > 1.24$. In order to calculate an average compatibility metric based on available data, we analyze the BSA distribution of deceased donors and transplant recipients. Given $BSA_{\text{Recipient}}$, we first calculate what percent of BSA_{Donor} falls into the compatibility interval, i.e., $0.78 * BSA_{\text{Recipient}} < BSA_{\text{Donor}} < 1.24 * BSA_{\text{Recipient}}$. For each height tier of recipients, we replace $BSA_{\text{Recipient}}$ with $BSA_{\text{Recipient LB}}$ ($BSA_{\text{Recipient UB}}$) in the lower (upper) bound of the interval where $BSA_{\text{Recipient LB}}$ ($BSA_{\text{Recipient UB}}$) is the 5th (95th) percentile of the BSA of transplant patients. Since using $BSA_{\text{Recipient LB}}$ and $BSA_{\text{Recipient UB}}$ gives optimistic estimates for compatibility, we also calculate what percent of $BSA_{\text{Recipient}}$ falls into the compatibility interval given BSA_{Donor} , i.e., $0.78 / BSA_{\text{Donor}} < 1 / BSA_{\text{Recipient}} < 1.24 / BSA_{\text{Donor}}$. This time, we replace BSA_{Donor} with $BSA_{\text{Donor UB}}$ ($BSA_{\text{Donor LB}}$) in the lower (upper) bound of the interval where $BSA_{\text{Donor UB}}$ ($BSA_{\text{Donor LB}}$) is the 95th (5th) percentile of the BSA of donors. The minimum compatibility percentage obtained from the first set of intervals and the second set of intervals is taken as the size-based compatibility metric in the optimization model. The resulting compatibility percentages and the histogram of donor and recipient $BSAs$ can be found in Appendix A.5.

In order to examine the effect of MELD score exception points on the equity of liver transplantation, we measure the likelihood of receiving a transplant, the likelihood of death while waiting for a transplant, and the ratio of the likelihood of transplant to the sum of the likelihood of death and transplant metrics of each patient group. We compare the performance of our policy with three

⁵ $BSA = 0.007184 * \text{Height}^{0.725} * \text{Weight}^{0.425}$

benchmarks; current UNOS policy without exception points, (1,1,0) policy that grants +1 point to ≤ 150 cm patients, and +1 point to 151-156 cm patients with no exception points to remaining patients, and (2,1,0) policy that grants +2 points to ≤ 150 cm patients and +1 point to 151-156 cm patients with no exception points to remaining patients. We consider (1,1,0) and (2,1,0) policies as static policies because the exception points are granted to disadvantaged patients regardless of their laboratory MELD score, and they do not change during the simulation. As we have seen as an example in §1.4, our exception points depend on patients' MELD scores that are updated regularly, and they may change during the simulation if disadvantaged patients' likelihood of receiving a liver transplant changes.

The results on equity with the current UNOS policy, (1,1,0) and (2,1,0) static policies, and our policy are presented in Table 1.2. As discussed in §1.1, ≤ 150 cm and 151-156 cm candidates have a significantly lower likelihood of receiving a deceased-donor liver transplant (35.9% and 37.5%, respectively) compared to 157-165 cm (39.2%), 166-175 cm (40.1%), and ≥ 176 cm (41.0%) candidates with the current policy. Also, the likelihood of death while waiting for a liver transplant is higher for ≤ 150 cm and 151-156 candidates (10.7% and 10.6%, respectively) compared to 157-165 cm (9.6%), 166-175 cm (9.2%), and ≥ 176 cm (8.6%) candidates. Consequently, ≤ 150 cm and 151-156 cm candidates have lower transplant over death plus transplant percentage (77.0 % and 77.9%, respectively) compared to 157-165 cm (80.3%), 166-175 cm (81.3%), and ≥ 176 cm (82.7%) candidates.

(1,1,0) and (2,1,0) static policies improve ≤ 150 cm and 151-156 cm candidates' access to liver transplantation; ≤ 150 cm and 151-156 cm candidates' likelihood of transplant increases to 38.1% and 39.5% with (1,1,0) policy, and to 39.5% and 39.4% with (2,1,0) policy, respectively. These two policies result in a decrease in ≤ 150 cm and 151-156 cm candidates' likelihood of death while waiting for a transplant; ≤ 150 cm and 151-156 cm candidates' likelihood of death drops to 10.5% and 10.1% with (1,1,0) policy, and to 10.2% and 10.1% with (2,1,0) policy, respectively. Our dynamic policy further improves ≤ 150 cm and 151-156 cm candidates' access to transplantation compared to (1,1,0) and (2,1,0) policies; Table 1.2 shows that ≤ 150 cm and 151-156 cm patients' likelihood of transplant increases to 39.7% and 39.7%, respectively. We also observe a further decrease in 151-156 cm candidates' likelihood of death (9.9%) with our policy compared to (1,1,0) and (2,1,0) policies. Overall, all candidate groups' likelihood of transplant percentages converges

Table 1.2: Simulation Results on Equity

Height (cm)	Likelihood of Transplant (%)			
	Current Policy	(1,1,0) Policy	(2,1,0) Policy	Our Policy
≤ 150	35.9	38.1	39.3	39.7
151-156	37.5	39.5	39.4	39.7
157-165	39.2	38.9	39.3	39.6
166-175	40.1	39.7	39.7	39.7
≥ 176	41.0	40.4	40.2	40.0
Likelihood of Death (%)				
≤ 150	10.7	10.5	10.2	10.2
151-156	10.6	10.1	10.1	9.9
157-165	9.6	9.8	9.8	9.8
166-175	9.2	9.1	9.2	9.2
≥ 176	8.6	8.8	8.8	8.9
Transplant/(Death + Transplant) (%)				
≤ 150	77.0	78.4	79.4	79.7
151-156	77.9	79.6	79.6	80.0
157-165	80.3	79.9	80.0	80.2
166-175	81.3	81.4	81.2	81.2
≥ 176	82.7	82.1	82.0	81.8

to each other with our policy ensuring equal access to deceased-donor liver transplantation.

The MELD score exception points for disadvantaged patient groups improve equity in access to liver transplantation. Since improving the equity of a system comes at the cost of losing efficiency in general, we use various efficiency metrics to assess the effect of proposed MELD score exception points on the performance of the liver allocation system. In particular, we use the expected quality-adjusted life years of each patient group (QALY), the number of wasted livers from each donor group (NWL), the number of patients died while waiting for a liver transplant (NPD), and the average waiting time of each patient group until receiving a transplant (AWT).

The percentage improvements in the efficiency metrics over the current UNOS policy are presented in Table 1.3. As expected, ≤ 150 cm and 151-156 cm candidates benefit from receiving MELD exception points because all performance metrics improve for them. In particular, QALY and AWT objectives for ≤ 150 cm (16.8% and 11.2%, respectively) and 151-156 cm candidates (7.1% and 3.4%, respectively) improve substantially. The decrease in these two objectives is low

Table 1.3: Simulation Results on Efficiency

Height (cm)	% improvement over the current allocation			
	QALY	NPD	AWT	NWL
≤ 150	16.8	5.9	11.2	4.1
151-156	7.1	2.1	3.4	1.2
157-165	-1.1	-1.0	-0.1	0.2
166-175	-0.7	-2.2	-0.3	-0.4
≥ 176	-1.9	-3.1	-3.9	-1.1
Total	0.2	-0.1	0.9	-0.6

Notes. QALY: quality-adjusted life years, NPD: number of patients died while waiting for a transplant, AWT: average waiting time until transplant, NWL: number of wasted livers. % improvement takes a positive value when the QALY objective increases, NPD, AWT, and NWL objectives decrease.

for 157-165 cm (-1.1% and -0.1%, respectively) and 166-175 cm (-0.7% and -0.3, respectively) candidates compared to ≥ 176 cm candidates (-1.9% and -3.9%, respectively). Similarly, the NPD objective for ≤ 150 cm (5.9%) and 151-156 cm (2.1%) candidates improve even though this improvement is not as big as QALY and AWT objectives. This objective becomes worse for 157-165 cm (-1.0%), 166-175 cm (-2.2%), and ≥ 176 cm (-3.1%) candidates. On the supply side, NWL objective improves for ≤ 150 cm (4.1%), 151-156 cm (1.2%), and 157-165 cm (0.2%) donors with a decrease in 166-175 cm (-0.4%) and ≥ 176 cm (-1.1%) donors. Overall, our simulations show that the performance of the liver allocation system improves for QALY and AWT objectives (0.2% and 0.9%, respectively). In contrast, with our policy, NPD and NWL objectives worsen slightly (-0.1 and -0.6%, respectively). These results suggest that we can improve equity by introducing MELD exception points to disadvantaged candidates without sacrificing the efficiency of the liver transplant system.

Finally, Table 1.4 presents the change in various equity and efficiency metrics with respect to the patients' gender and race. The shortest stature candidates (≤ 150 cm and 151-156 cm) who receive MELD exception points in our policy represent a disproportionately female and Hispanic proportion of the liver transplant candidate population. Female candidates have a lower probability of receiving a liver transplant and a higher likelihood of death while waiting for a transplant (38.9% and 9.6%, respectively) compared to male candidates (40.3% and 8.8%, respectively) with the current policy. Hispanic candidates also have lower rates of liver transplant (38.3%), longer waiting times until

Table 1.4: Simulation Results on Recipients' Gender and Race

Gender	Current Policy			Our Policy		
	LT (%)	AWT	Death (%)	LT (%)	AWT	Death (%)
Women	38.9	257	9.6	39.8	256	9.4
Men	40.3	277	8.8	39.9	274	8.9
Race/Ethnicity						
Caucasian	39.4	271	9.1	39.2	275	9.2
Hispanic	38.3	296	10.9	38.7	289	10.6
African American	46.4	216	7.6	46.3	218	7.5
Asian	40.1	288	5.7	40.6	269	5.9

receiving a transplant (296 days), and higher rates of death (10.9%) in comparison to non-Hispanic candidates (e.g., Caucasian candidates have a transplant rate of 39.4%, an average waiting time of 271 days, and a dropout rate of 9.1%). Our policy almost equalizes female and male candidates' likelihood of liver transplantation (39.8% and 39.9%, respectively) while lowering female candidates' average waiting time until transplantation (from 257 to 256 days), and the likelihood of death (from 9.6% to 9.4%). Hispanic (38.7%) and Asian candidates' (40.6%) transplant rate increases, average waiting time until transplantation decreases (289 and 269 days, respectively), and Hispanic candidates' likelihood of death (10.6%) decreases with our MELD exception points. Altogether, our simulations demonstrate that disadvantaged candidates (female and Hispanic) greatly benefit from our policy, and more equal rates of liver transplantation and death across the entire transplant candidate population are obtained with MELD exception points.

1.6 Conclusion

The shortage of donor liver supply results in long waiting times for ESLD patients, which raises concerns over fairness in different patient groups' access to transplantation. In this regard, we study the problem of achieving a fairer liver allocation system where there are disparities in access to transplantation based on patients' height and gender. Shorter patients and women have higher average waiting times and mortality rates compared to other patient groups because they can receive liver transplants from a smaller pool of available deceased donors. To address this problem, (i) we develop a fluid model of the liver transplant system with fairness constraints, (ii)

derive the optimal policy of allocating deceased-donor livers to transplant patients, (iii) provide a computational framework to provide MELD score exception points to disadvantaged patient groups to increase their access to transplantation, and (iv) run simulations of the national liver allocation system to assess the effect of our proposed MELD exception points on the efficiency and equity of the current policy of allocating livers which is based on transplant patients' medical urgency.

Our analysis shows that the Equity Adjusted Mortality Risk policy, which ranks transplant patients in terms of their medical urgency but adjusts this ranking to ensure all patient groups have equal access to transplantation, is optimal in allocating deceased-donor livers. Without the fairness constraints, we show that the optimal policy coincides with the UNOS' current policy of allocating donor livers based on ESKD patients' medical urgency, i.e., ranking patients with respect to their laboratory MELD scores. The dual state variables are endogenously calculated while solving the primal optimal control problem of minimizing pre-transplant mortality. We show that these dual variables are a proxy for transplant patients' short-term mortality risk, which can be mapped into laboratory MELD scores. With an easy-to-implement algorithm, we utilize them to provide MELD exception points to disadvantaged patient groups because their short-term mortality risks are higher than other patient groups. These exception points move disadvantaged patients to higher positions in the transplant waitlist to improve their access to transplantation.

We run simulations of the national liver allocation system to test the effect of introducing MELD exception points over various efficiency and equity objectives. Our simulations show that disadvantaged patients can greatly benefit from receiving MELD score exceptions without decreasing the efficiency of the overall liver transplant system. While prior research focuses on optimizing the organ allocation rules that require fundamental policy changes, our work provides a remedy within the current liver allocation system where transplant patients are prioritized based on medical urgency. In addition to the static patient characteristics we have discussed in our work (height, gender, race), our methodology can be generalized with any factor that creates discrepancies in organ access. Also, our shadow price approach can be used to compare medical urgency across organs for patients who need dual organ transplants (a new liver and another new organ during the same surgical procedure) since these patients are listed on both organ waiting lists.

We discuss the potential considerations that are beyond the scope of this chapter. First, we solely focus on the allocation of deceased-donor livers; therefore, we do not consider living donor

liver transplantation in our theoretical and computational analysis. The reasons for this are: (i) living donor liver transplants constitute a small portion (5.5% in 2020) of liver transplants in the US, and (ii) most living donors donate a portion of their healthy livers to their family members or friends without participating in the national liver allocation system. Second, our model and analysis do not incorporate split liver transplants (SLT), which can potentially reduce disparities in organ access due to the size mismatch between the donor and the recipient. The practice of splitting deceased-donor livers provides liver transplants for two recipients (in general, one adult and one pediatric patient); however, only 3.8% of all deceased-donor livers are used for SLTs from 2010 to 2015 (Tang et al. 2021). Given their rising trends in the last few years, how to incorporate living donor liver transplants and SLTs into the MELD-based liver allocation system remains an interesting research question to study further in the future.

Chapter 2

Popularity Bias in Online Dating Platforms

2.1 Introduction

Since the launch of commercial dating sites such as Match.com in the 1990s, online dating has become a popular way of meeting and developing relationships with platforms serving billions of people around the world (Vogels 2020). The rise of online dating services gives people access to more potential dating partners than they could meet in traditional ways (e.g., through friends or other social activities). In the United States, three out of ten adults have used a dating site or app at some point, and over half of Americans say relationships that begin on a dating site or app are just as successful as those that begin in person (Anderson et al. 2020). Recently, the use of dating apps has increased substantially as people have sought to connect virtually during the Covid-19 pandemic (Wu 2021).

A great variety of online dating platforms exist today (see, e.g., the comparison of online dating services on Wikipedia). Early online dating services allowed users to browse, search, and contact other users freely, but those decentralized systems have been gradually replaced by online dating services that provide *recommendations* of compatible dating partners to their users. Matchmaking technologies have become an important value proposition for these new service platforms. For instance, the matchmaking website eHarmony claims to have a “scientific approach to matching

highly compatible singles” based on “29 dimensions of compatibility,” and OkCupid states that they “do a lot of crazy math stuff to help people connect faster” (Wu et al. 2018).

However, generating recommendations of compatible dating partners is challenging for online dating platforms because uncovering users’ idiosyncratic preferences is difficult. Thus, platforms tend to recommend popular users (those who are attractive, well-educated, well-employed, etc.) to others more frequently than unpopular users, even though users might be interested in them. In other words, their algorithms are prone to *popularity-based bias* (or popularity bias for short). For example, using data we collected over a 15-month period from a major online dating platform in South Korea, we find that a user’s chance of being recommended by the platform’s algorithm increases significantly with the increase in the user’s average attractiveness score as generated by other users. This popularity bias in a platform’s recommendations results in congestion for popular users due to the excessively high number of messages received, and it lowers unpopular users’ chances of finding compatible dating partners.

The challenge of building a high-performing recommendation algorithm is not the sole reason why we observe popularity bias in online dating platforms. The vast majority of platforms earn revenue not when users find a successful match but when users are in the process of searching for dating partners. For instance, some platforms generate revenues from showing ads to users while they view recommended profiles (e.g., Tinder, Bumble), some charge their users weekly/monthly subscription fees for the privileged opportunity to show interest to other users (e.g., eHarmony, Match.com), and others require their users to make in-app purchases to initiate contact (e.g., Coffee Meets Bagel, Noondate). As a result, a platform’s objective is to keep users engaged on the platform rather than maximize users’ chances of finding compatible dating partners. Recommending popular users more frequently can potentially increase user engagement on the platform (i.e., more likes and messages sent); however, it does not ensure a higher matching probability for users.

Motivated by the empirical evidence of popularity bias on a major online dating platform, this paper studies the incentives for an online dating platform to adopt recommendation algorithms favoring popular users and examines whether removing such popularity bias in the platform’s recommendations would improve match outcomes for its users. We model the platform’s decision and users’ subsequent interactions as a three-stage matching game, and we consider a platform with two tiers of users, namely, popular and unpopular. Popular users are more selective in sending and

accepting messages because they are more likely to have better options both on the platform and outside of it. We solve the platform’s problem of maximizing revenue and maximizing the number of matches. We use an unbiased platform (in which popular and unpopular users find equal chances to be recommended to others) as our benchmark for fairness to compare popular and unpopular users’ matching probabilities, and we evaluate the expected revenues and expected number of matches for three platforms (revenue-maximizing, match-maximizing, and unbiased).

We empirically validate our theoretical model’s assumptions using the extensive data collected from a major online dating platform. The data consists of individual attributes of approximately 243,000 users and over 30 million interactions among them over 15 months. The user attributes include detailed information about users’ demographics and preferences, and the interaction data capture the following user decisions: (1) seeking more information about other users recommended by the platform (referred to as “open”), (2) expressing interest to the recommended users by sending messages (referred to as “send”), and (3) accepting the messages received from other users (referred to as “accept”). In addition to validating our model assumptions, we utilize this data to predict users’ future behavior on the platform, i.e., we estimate the probability of interaction between any pair of users. Our predictions are free of assumptions that we use in our theoretical model; therefore, the estimated probabilities allow us to run simulations of the revenue-maximizing, match-maximizing, and unbiased recommendations to test our theoretical results.

Our theoretical analysis shows that the degree of popularity bias in the platform’s recommendations is affected by how selective popular users are in sending and accepting messages compared to unpopular users. We find that the revenue-maximizing recommendations and the match-maximizing recommendations are not necessarily at odds when popular users are not too selective in sending and accepting messages. Similar to the revenue-maximizing recommendations, the match-maximizing recommendations need to ensure a high number of messages are sent in order to maximize the number of successful matches between users. As popular users become more selective on the platform, the bias in the match-maximizing recommendations decreases significantly, whereas the bias in the revenue-maximizing recommendations remains high. The user interaction data indicates that popular female users are much more selective than unpopular female users, whereas popular and unpopular male users are similarly selective in sending and accepting messages. As a result, our simulations show that the match-maximizing recommendations are less

biased against unpopular female users than the revenue-maximizing recommendations; however, both the match-maximizing recommendations and the revenue-maximizing recommendations lead to a similar degree of bias against unpopular male users.

We note here that the degree of heterogeneity in users’ popularity is a double-edged sword for the platform’s success, both in terms of revenue and the number of successful matches. Even though popular users might help the platform increase revenue and the number of successful matches, they might also become too selective and may be perceived by others as “out of reach” so that other users are reluctant to initiate contact.

Another factor that affects the degree of bias in the platform’s recommendations is the congestion effect related to users’ incoming messages. We find that when the congestion effect is low, unbiased recommendations result in a significantly lower number of messages and successful matches on the platform than the revenue-maximizing and match-maximizing recommendations. As the congestion effect increases, both the revenue-maximizing and match-maximizing recommendations include popular and unpopular users more equally. Our simulations indicate that the platform can increase its revenue and users’ chances of finding dating partners simultaneously with a certain degree of bias against unpopular users.

The rest of this chapter is organized as follows. In §2.2, we review the related literature and discuss our contributions. In §2.3, we describe our empirical setting and analysis, and in §2.4, we present our theoretical model and analysis. In §2.5, we describe how we estimate users’ interaction decisions, discuss our method’s predictive performance, and present our simulation results. In §2.6, we present our two-period model and conclude this chapter in §2.7.

2.2 Related Literature

This paper is broadly related to the literature on online matching platforms. While there are a variety of online matching platforms (e.g., ride-sharing platforms such as Uber and Lyft, house rental platforms such as Airbnb, and freelancing platforms such as TaskRabbit and UpWork), our focus in this paper is on online dating platforms, which have received relatively less attention from operations researchers (see, e.g., Chen et al. (2020), Benjaafar and Hu (2020), Qi et al. (2020) and references therein). Chen et al. (2020) emphasize the importance of incorporating subjective,

idiosyncratic, and personalized user preferences to develop recommendation systems and identify research opportunities to investigate how online dating platforms should design their communication and recommendation mechanisms to improve matches. While incorporating users' idiosyncratic preferences, our work identifies why we observe popularity bias in a revenue-maximizing online dating platform's recommendation algorithm and how such bias affects users' chances of finding dating partners.

Online dating has been studied by researchers in diverse disciplines including economics, information systems, marketing, operations, sociology, and strategy. The literature related to online dating goes back to the extensive body of work that studies matching and sorting patterns in marriage markets, starting with Gale and Shapley (1962). Adachi (2003) shows that the stable equilibrium matching predicted by Gale and Shapley can be seen as the limit outcome of a two-sided search and matching model with negligible search costs. Noting that the environment Adachi (2003) studies is akin to that of online dating, Hitsch et al. (2010*b*) show that the observed correlations in user attributes in online dating largely agree with the correlations in the Gale-Shapley predicted stable matches in marriages. They use a discrete choice model in analyzing data from a major online dating service provider and estimate the probability that a user will contact another user after browsing their profile. In another study, Hitsch et al. (2010*a*) further examine how various user attributes, such as race, physical attributes, and education levels, affect the estimated probabilities. Their analysis is further extended by several researchers: Lin and Lundquist (2013) examine how race, gender, and education affect initial messages exchanged by any two users; Ong and Wang (2015) measure gender differences in preferences for mate income; Bruch et al. (2016) examine how age, height, and body mass index affect the probabilities of users' initial browsing and later contact decisions; and Lee (2016) studies the effect of online dating services on marital sorting.

Another stream of empirical research in online dating draws causal inference regarding the impact of a particular platform feature (or function) on users' decisions or matching outcomes. This line of work uses randomized field experiments, and researchers have investigated an anonymous search feature (Bapna et al. 2016), member-get-member referral programs (Belo and Li 2018), and optional phone verification of users (Shi and Viswanathan 2023), among other features. In a similar vein, Fong (2020) uses causal inference to study how market size and competition size affect users'

selectivity in searching and matching with other users.

While the above stream of research studies online dating services in which users can browse, search, and interact with other users (e.g., Match.com), a growing number of new online dating platforms (e.g., Bumble, Coffee Meets Bagel, eHarmony, OkCupid, Tinder) play a more active role by recommending users to each other (Wu and Padmanabhan 2019). This increasingly popular platform role has been recognized by a few recent analytical papers. For instance, Halaburda et al. (2018) develop a stylized model of online dating and explain how a recommendation-based platform can successfully compete with search-based platforms by limiting the number of choices it offers users while charging higher prices. Wu et al. (2018) note a potential conflict of interest between online dating platforms and their users since users may leave the platforms after finding their matches, reducing the platforms' revenue. They study this conflict in a game-theoretic model and offer insights into how various factors (e.g., competition, consumer patience) may undermine a platform's incentive to offer the best recommendations to its users. Kanoria and Saban (2021) model a dynamic two-sided matching market in which each user has an opportunity to see a random candidate (from one of two vertical quality tiers) who arrives according to a Poisson process. They find that simple platform interventions can improve matching outcomes, such as hiding quality information or preventing users on one side of the market from sending a message. Our study adopts similar assumptions; the utility that users receive from their matches includes a common evaluation by all users (average attractiveness) and an idiosyncratic element specific to each user pair, and one side of the platform includes two tiers of users (high-tier and low-tier). Rios et al. (2022) propose a family of heuristics to improve match rates in online dating platforms through assortment optimization. Finally, Bojd and Yoganarasimhan (2022) empirically examine how users' popularity information affects other users' decisions to initiate contact in a live dating game where a mobile dating app assigns four men (women) to four women (men), users rank-order their preferences, and the app provides matches to them with a stable match algorithm.

In line with this recent stream of research, we study online dating platforms in which users' interactions follow platforms' recommendations. Yet diverging from previous studies, we examine how a platform's objective and its users' behavior shape the platform's recommendations. We consider the platform, acting through its recommendation system, as the primary decision-maker that influences which users can interact with each other. Given the platform's recommendations,

users make interaction decisions, anticipating the probability of their interest being reciprocated by others. We assume that popular users are more attractive to others but more selective in sending and accepting messages and that users, regardless of popularity, become more selective in accepting messages if they receive too many of them due to the time and effort necessary to evaluate each potential dating partner (Hauser and Wernerfelt 1990, Iyengar and Lepper 2000, Pronk and Denissen 2020). We show that the degree of popularity bias in an online dating platform’s recommendations can vastly increase when the platform focuses on maximizing (short-term) revenue rather than users’ chances of finding successful matches (which may help build the platform’s reputation and increase long-term revenue). We run simulations of the platform to validate our theoretical results by developing a novel predictive model that estimates users’ likelihood of interaction.

Finally, our paper is also related to prior work that studies the connection between a platform’s objective and its bias in recommendations in non-dating contexts. For example, Lambrecht and Tucker (2019) show that an algorithm that optimizes cost-effectiveness in ad delivery promotes job opportunities in STEM (science, technology, engineering, and math) fields to men more than women because it is more expensive to show ads to younger women. Ciampaglia et al. (2018) find that popularity bias hinders the quality of content recommendations in cultural markets (e.g., social media, music streaming services) because popularity metrics lead users to make cognitively cheaper but less accurate choices in terms of quality. We contribute to this body of literature by showing the existence of popularity bias in a two-sided matching platform’s recommendations and its consequences for users’ likelihood of finding successful matches.

2.3 Empirical Setting and Data

In this section, we describe the data we collected from a major online dating platform in South Korea (subsequently referred to as “the platform”) and run a regression model to empirically check the presence of popularity bias in the platform’s recommendation algorithm. The data from the platform contains all active users during the fifteen months of our study period, amounting to 198,857 male users and 39,939 female users. The platform provided completely anonymized user data to protect users’ privacy. This data can be divided into two components: *user attributes* and *interactions*.

User attributes, summarized in Appendix B.1, consist of information on each user, including demographic information (age, education, religion), physical characteristics (height, weight), profile information (profile completion score, Facebook connected, whether or not a real name is verified), attractiveness score, and user-stated search filter settings for age, drinking, height, weight, religion, and smoking. The profile completion score is a number the platform calculates based on the degree of completion of a user’s profile and survey questions. Each user also has an attractiveness score which is the average value of the scores (between 1 and 5) the user received from the other users recommended by the platform in the past. This information provides a richer empirical setting than that studied by Hitsch et al. (2010a), who used attractiveness scores of users’ photos that were constructed based on evaluations provided by 100 university students in a lab environment. We refer to male and female users whose attractiveness scores are higher than the average attractiveness score for their gender as *popular* users and the remaining users as *unpopular* users. In addition to these individual user attributes, we construct *pairwise* attributes that apply to pairs of users, such as height difference, age difference, and religion match. Appendix B.2 provides summary statistics for each of the user attributes.

Unlike the user attributes, the interaction data relates to pairs of users and is generated by the users’ decisions on the platform. The platform’s operating sequence is shown in Figure 2.1. Every day the platform recommends two female (male) users to each male (female) user. The recommendation is given in the form of a *card* that contains only the recommended user’s basic information, such as age, area of residence, and one profile photo. After receiving the cards from the platform, users may interact with each other in the following three ways – for illustration, suppose the platform sends Bob’s card to Alice.

1. *Open*: Alice chooses whether or not to *open* Bob’s card based on Bob’s basic information shown on the card. When she opens the card, Alice can see Bob’s detailed profile.
2. *Send*: Assuming Alice opened Bob’s card, Alice chooses whether or not to *send* a message to Bob.
3. *Accept*: Assuming Alice sent a message to Bob, Bob chooses whether or not to *accept* Alice’s message. Once Bob accepts Alice’s message, the pair can exchange messages or chat freely.

Figure 2.1: Sequence of Decisions in the Online Dating Platform

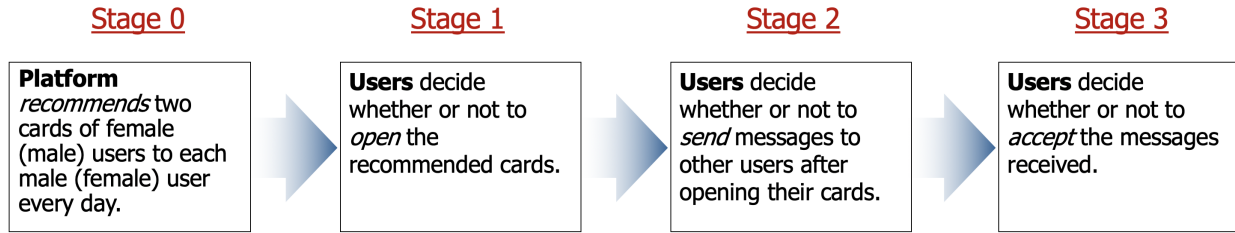


Table 2.1 summarizes the user interaction data. As one might expect, male users tend to express interest more actively than female users: male users’ likelihood of opening cards, 30.0%, is substantially higher than that of female users, 8.4%. We observe similar patterns for sending and accepting messages. Users’ likelihood of sending messages after opening cards is significantly lower than that of opening cards for both male and female users. This may be due to the additional information users obtain after opening recommended users’ cards. Interestingly, we observe that for both male and female users, the likelihood of accepting messages is higher than that of opening cards or sending messages. This means that users are more likely to be receptive to the interest shown by other users than they are to seek more information or initiate interactions by sending messages. Table 2.1 also informs us of how popular and unpopular male and female users behave differently when opening cards, sending, and accepting messages. Popular male users are more likely to open cards and send messages than unpopular male users, whereas popular and unpopular female users’ likelihoods are similar. We observe a different pattern for accepting messages. Popular male and female users’ likelihood of accepting messages is significantly lower than those of unpopular male and female users.

To empirically check the presence of popularity bias in the platform’s recommendations, we run a negative binomial regression. The regression model evaluates the effect of various user attributes on the platform’s recommendations (referred to as “number of recommendations”) over a 15-month period. The results in Table 2.2 indicate that a user’s average attractiveness score is a statistically significant factor in how many times the platform exposes a user’s profile to other users. The incidence rate ratio (IRR) in the second column of Table 2.2 gives us the percentage change in the dependent variable (number of recommendations) for a unit increase in an independent variable for numerical factors (e.g., average attractiveness, age). For instance, the IRR of average

Table 2.1: Summary of User Interaction Data

	Decision	Size	Likelihood (%)		Likelihood (%)		Likelihood (%)
Male Users	Open	18,288,720	30.0	Popular	32.7	Unpopular	28.9
	Send	5,483,418	7.9		10.8		6.7
	Accept	6,085	40.6		38.7		45.2
Female Users	Open	4,149,030	8.4	Popular	8.6	Unpopular	8.1
	Send	349,481	1.7		1.6		1.9
	Accept	435,180	9.2		7.8		13

Notes. The **Size** column reports the number of observed transactions, and the **Likelihood (%)** columns report the percentage of the interaction being observed.

attractiveness indicates that a user’s chance of being recommended by the platform’s algorithm increases by 64% for every unit increase in a user’s average attractiveness score. Because a higher attractiveness score implies greater popularity among users, this result suggests the presence of popularity bias in the platform’s recommendations. Combined with popular users’ lower likelihood of accepting messages, this popularity bias potentially decreases the number of successful matches on the platform. Additional comments on the platform’s recommendations are presented in Appendix B.4.

We also find that other individual user attributes presented in Table 2.2 are statistically significant in predicting the number of times a user’s profile is recommended to the other users on the platform. For instance, female users are recommended more (gender = 1 for female) simply because of the imbalance between the number of male and female users, and the platform recommends users more frequently if their names are verified (verified = 1). For categorical and binary variables (e.g., gender, verified, region), the IRR indicates the percentage change in the number of recommendations compared to their base level (e.g., gender = 0). For example, the platform recommends smoking users 18.5% less than non-smoking users and verified users 26.6% more than unverified users.

2.4 A Two-Sided Matching Game

Motivated by the empirical evidence presented in §2.3, in this section, we develop a theoretical model to examine the underlying mechanism that leads to popularity bias in the platform’s recommendations. Our model is a two-sided matching game with three stages. In the first stage, the

Table 2.2: Regression Result on Platform’s Recommendations

	Coefficients	Incidence Rate Ratio
(Intercept)	2.245*** (0.016)	
Average Attractiveness	0.494*** (0.003)	1.638
Gender	2.325*** (0.009)	10.232
Verified	0.236*** (0.005)	1.266
Age	-0.0144*** (0.001)	1.014
Platform Tenure	0.091*** (0.002)	0.913
Region	0.0103*** (0.001)	1.010
Religion	0.026*** (0.002)	1.026
Smoking	-0.205*** (0.006)	0.815

Notes. Dependent variable: Number of recommendations. Number of observations = 238,796. Standard errors are in parentheses. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. Incidence Rate Ratio = $\log((\mu_0 + 1)/\mu_0)$ where μ_0 denotes the mean value of the independent variable.

platform recommends a set of male (female) users to female (male) users. In the second stage, upon receiving the recommendations from the platform, users decide whether to send a message showing interest to the recommended user. In the third stage, users decide whether to accept the messages they received in the second stage. For instance, if the platform recommends user A to user B, a *match* between users A and B occurs if user B sends a message to user A in the second stage and user A accepts the message from user B in the third stage. These three stages are consistent with the sequence of decisions depicted in Figure 2.1, except that our model abstracts away from users’ open decisions for traceability. We describe the details of our model in §2.4.1 and solve for the platform’s optimal recommendations in §2.4.2. Appendix B.5 summarizes our notation. Proofs of our results can be found in Appendix B.6.

2.4.1 Model

We refer to the users on the platform as men (denoted by m) and women (denoted by w). Women on the platform are divided into two tiers: low (denoted by l) and high (denoted by h). The high

tier represents attractive women with high average attractiveness scores due to good appearance, education, employment, etc. The low tier represents relatively less attractive women with lower average attractiveness scores. We choose to construct only two tiers on one side of the platform in order to maintain tractability in characterizing users' behavior and the platform's strategy, and this approach suffices to show the existence of popularity bias. The matching game is played only once, so we fix the number of users in each tier. Let N^m, N^h , and N^l denote the number of men, high-tier women, and low-tier women, respectively, on the platform. The number of women on the platform is denoted by $N^w = N^h + N^l$. For simplicity, we assume an equal number of high-tier and low-tier women on the platform, i.e., $N^h = N^l$, and an equal number of men and women, i.e., $N^m = 2N^h = 2N^l$. The high and low tiers are determined by specifying a threshold for the overall average attractiveness score; scores above the threshold constitute the high tier, and scores below constitute the low tier. We thus divide the users into two tiers of the same size.

In the first stage, the platform recommends n users from one side to each user on the other side, where n is fixed and much smaller than the number of users on the platform; that is, $n \ll N^m, N^h, N^l$. The platform chooses a proportion of high-tier or low-tier women in men's recommendations. Specifically, the platform determines x_h and x_l where x_h (x_l) denotes the proportion of high-tier (low-tier) women in men's recommendations. This gives us $x_h \in [0, 1]$ where $x_h + x_l = 1$. For example, suppose that $n = 3$ and the platform chooses two high-tier women and one low-tier woman to recommend to men. This gives us $x_h = 2/3$ and $x_l = 1/3$. Given the platform's decision, each woman is equally likely to be recommended within their tier. Our model's primary variables of interest are x_h and x_l ; we characterize the optimal recommendation ratios given the platform's objective and investigate how changes in these decision variables influence users' behavior and expected utility on the platform. Furthermore, the platform's decision variables x_h and x_l help us examine whether popularity bias exists against a certain user tier on the platform. We say that the platform is *unbiased* when $x_h = x_l = 1/2$ because the platform has an equal chance of recommending any user in either tier (high or low). There is bias against high-tier women if the platform includes a higher proportion of low-tier women than high-tier women in men's recommendations, i.e., $x_h < 1/2$. Similarly, there exists bias against low-tier women if $x_l < 1/2$. As a measure of how strong bias is, we define the degree of bias as the Euclidean distance between the platform's chosen recommendations (x_h and x_l) and its unbiased recommendations.

In the second stage of the game, users receive recommendations from the platform and decide whether to send a message to each recommended user. A user sends a message to a recommended user when his or her expected match value exceeds the sum of the utility of remaining single (or the utility of outside options) and the cost of sending a message. We explain each of these model elements next.

First, if a pair of users match at the end of the game, each user earns a match value. Let v_{ij} denote the match value user i earns when matched with user j . This match value depends on two components: the *average attractiveness* of the tier that user j belongs to, q_j , and *idiosyncratic utility*, u_{ij} , specific to pair (i, j) , so that $v_{ij} = q_j + u_{ij}$. If user j is a woman from the high tier, then $q_j = q \in (0, 1)$. We normalize q_j to be 0 for low-tier women and all men on the platform. Idiosyncratic utility u_{ij} is distributed independently and identically with Uniform(0,1). This ensures that $v_{ij} \in [q, 1 + q]$ if user j is a high-tier woman and that $v_{ij} \in [0, 1]$ if user j is a low-tier woman or any man on the platform. Although a woman in the high tier is more likely to provide a higher match value than a woman in the low tier, this is not always true due to idiosyncratic utility. The distribution of idiosyncratic utility is known to all users on the platform, but the value of u_{ij} , along with q_j , can only be learned privately by user i when the platform recommends user j to user i . A similar assumption is made by Kanoria and Saban (2021). We note that idiosyncratic utility is not necessarily symmetric, i.e., u_{ij} may not equal u_{ji} for i and j . This captures the fact that two users may receive different utilities from the same match.

Second, users have options to find dating partners outside of the online dating platform (e.g., through family and friends or social activities). Such outside options would affect how *selective* users are on the platform: users with better outside options tend to apply higher standards in sending and accepting messages. Let θ_i denote the value of user i 's outside options, which can also be interpreted as user i 's utility from remaining single on the platform; a higher value of θ_i means user i is more selective in sending or accepting messages. We model heterogeneity of users' outside options by assuming θ_i is a random variable with Uniform($\alpha q_i, 1$) where $\alpha \in [0, 1]$ is a constant that captures high-tier women's degree of selectivity in sending and accepting messages. When $\alpha = 0$, θ_i is distributed with Uniform(0,1), suggesting that both high-tier and low-tier women are equally selective. When $\alpha > 0$, high-tier women tend to have better outside options and to be more selective in sending and accepting messages than low-tier women. By modeling random θ_i ,

our model allows low-tier users to be more selective than high-tier users, and it allows users in the same tier to have different values of outside options.

Finally, a user incurs a cost, denoted by $r \in [0, 1]$, when sending a message to a recommended user, and this cost is collected as revenue by the platform. This cost element is included to reflect most online dating platforms' revenue structure (e.g., users make in-app purchases to send likes to other users on platforms such as Coffee Meets Bagel and Noondate). Taken together, user i sends a message to user j , who was recommended by the platform in the first stage, if and only if

$$\mathbb{E}[v_{ij}] = \mathbb{P}[\text{user } j \text{ accepting user } i\text{'s message in the third stage}](q_j + u_{ij}) \geq \theta_i + r. \quad (2.1)$$

We denote the probability of a man sending a message to a high-tier (low-tier) woman as p_h^m (p_l^m) and the probability of a high-tier (low-tier) woman sending a message to a man as p_m^h (p_m^l).

In the third stage of the game, users see incoming messages and decide whether to accept or reject them. As is common in practice, we assume that a user does not incur any explicit cost in accepting a message and users can accept multiple messages (e.g., Bumble, Coffee Meets Bagel, Noondate); our model can be easily extended to the case where accepting a message has an explicit cost for users. If user i sends a message to user j in the second stage of the game, user j earns match value v_{ji} when she accepts user i 's message in the third stage. For user j to accept user i 's message, v_{ji} must be greater than or equal to the utility of remaining single, θ_j , and the bar of accepting a message, denoted by γ_j , on the platform. It is well established in the literature that a user's probability of rejecting a message increases as the user receives more messages (Iyengar and Lepper 2000, Pronk and Denissen 2020). We also verify this assumption in our data (see Appendix B.4). This increasing probability is due to the increase in the user's effort, which includes the effort to screen all incoming messages and decide whether to accept or reject each of them (Hauser and Wernerfelt 1990). We assume γ_j is a linear function of the total number of messages user j receives in the third stage, so for a high-tier woman j , $\gamma_j = \beta x_h p_h^m$, where $\beta \in [0, 1]$ is the *congestion effect*. (Recall that x_h determines the number of men to whom the platform recommends user j and p_h^m is the probability of a high-tier woman receiving messages from men.) Similarly, $\gamma_j = \beta x_l p_l^m$ for a low-tier woman j , and $\gamma_j = \beta/2(p_m^h + p_m^l)$ for man j . $\beta = 0$ corresponds to the case where there is no congestion effect, and as β increases, a user's bar of accepting a message increases. The value

of β is common to all users, and for the sake of analytical traceability, we assume it is known to all users and to the platform. Taken together, user j accepts user i 's message sent in the second stage if and only if

$$v_{ji} \geq \theta_j + \gamma_j. \quad (2.2)$$

We denote the probability of a high-tier (low-tier) woman accepting a man's message as a_m^h (a_m^l) and the probability of a man accepting a high-tier (low-tier) woman's message as a_h^m (a_l^m). When user j accepts user i 's message, a match occurs, and user i (user j) receives match value v_{ij} (v_{ji}).

Given the platform's strategy, denoted by x_h and x_l , users' decisions in the second and third stages of the game are interrelated. Suppose many users decide to send messages to the same subset of users in the second stage of the game. In that case, each user's matching probability decreases because the users who receive a high number of messages in the third stage face a high congestion effect. The same holds for high-tier women's degree of selectivity: high-tier women tend to provide a higher match value than low-tier women, so men are more likely to send messages to high-tier women. However, high-tier women are likely to have better outside options; hence, they are more selective in accepting messages in the third stage. On a dating platform where users receive only a few recommendations from the platform and the messages are costly signals, users are well aware of these two effects in sending messages (Fong 2020, Bojd and Yoganarasimhan 2022). We provide empirical evidence for this assumption using the data from a major online dating platform in Appendix B.4.

To analyze our three-stage matching game, we start by computing users' probabilities of accepting messages in the third stage. The following lemma gives these probabilities.

Lemma 2.1. *Given the platform's decision on x_h and x_l in the first stage and the users' probabilities p_k^m and p_m^k of sending messages in the second stage (where $k \in \{h, l\}$), users' probabilities of accepting messages in the third stage are $a_m^m = \max \left\{ 0, \frac{1 + 2q_k - \beta \sum_k p_m^k}{2} \right\}$ and $a_m^k = \max \left\{ 0, \frac{1 - \alpha q_k - 2\beta x_k p_k^m}{2} \right\}$.*

Lemma 2.1 shows that the platform's decision to recommend high-tier and low-tier women, x_k , affects women's acceptance probabilities, a_m^k . This is because the platform's decision determines how many men receive a woman's profile, which affects the number of men who send messages in the

second stage of the game. As the user's probabilities p_k^m and p_m^k of sending messages in the second stage increase, the user's probabilities of accepting messages, a_k^m and a_m^k , decrease in the presence of a positive congestion effect, β . High-tier women's average attractiveness term, q_k , increases the chance of their messages being accepted by men, a_h^m , but it decreases the probability that high-tier women will accept men's messages, a_m^h because high-tier women are more selective than low-tier women. Higher α increases high-tier women's selectivity, thus reducing their probability of accepting messages, a_m^h . Higher β increases the congestion effect for all users, reducing all users' probabilities of accepting messages (i.e., both a_k^m and a_m^k).

We next analyze the game's second stage and compute the users' probability of sending messages. The following lemma gives these probabilities.

Lemma 2.2. *Given the platform's decision on x_h and x_l in the first stage, the probabilities of users sending messages in the second stage, p_k^m and p_m^k for $k \in \{h, l\}$, are $p_h^m = \frac{(2q+1)(1-\alpha q) - 4r}{2(2q+1)\beta x_h + 4}$, $p_l^m = \frac{1-4r}{2(2q+1)\beta x_l + 4}$, $p_m^h = \frac{(4+\beta)(2q) + \beta(1-4r)}{8(2+\beta)}$ and $p_m^l = \frac{(1-4r)(8(2+\beta) - \beta^2) - \beta(4+\beta)2q}{8(4+\beta)(2+\beta)}$ for sufficiently small r .*

In Lemma 2.2, we focus on the case where all users have positive probabilities of sending messages, which happens for a sufficiently small cost of sending a message, r . Upon receiving the recommendations from the platform, men observe the number of high-tier and low-tier women in their recommendations, hence x_h and x_l , respectively. Recall from Lemma 2.1 that both high-tier and low-tier women become less likely to accept messages in the third stage as they receive more messages from men in the second stage of the game when there is a positive congestion effect, β . A higher value of x_h (x_l) implies that each high-tier (low-tier) woman is recommended to more men, intensifying competition for high-tier (low-tier) women in the third stage. This competition leads men to be more reluctant to send messages in the second stage of the game, i.e., p_h^m (p_l^m) decreases in x_h (x_l) and β . Similarly, Lemma 2.1 shows that high-tier women are more likely to reject messages when they are more selective, $\alpha > 0$ because they have better outside options. Thus, men's probability of sending messages to high-tier women, p_h^m , is decreasing in high-tier women's selectivity, α .

2.4.2 The Platform's Problem

Having characterized users' decisions to send and accept messages in the second and third stages of the game, we now turn to the platform's decisions on x_h and x_l . We consider two different cases: the platform that maximizes its revenue in §2.4.2.1 and the platform that maximizes the expected number of matches in §2.4.2.2.

2.4.2.1 Revenue-Maximizing Platform

As described in §2.3, the platform collects revenue from the messages that users send to others in the second stage of the game. To maximize the expected revenue, the platform solves the following problem:

$$\begin{aligned} \max_{x_h, x_l} \quad & nrN^m(p_h^m x_h + p_l^m x_l) + nrN^h p_m^h + nrN^l p_m^l \\ & x_h + x_l = 1 \\ & 0 \leq x_h, x_l \leq 1, \end{aligned} \tag{R}$$

where the first term in the objective function, $nrN^m(p_h^m x_h + p_l^m x_l)$, is the platform's expected revenue from the messages that men send to high-tier and low-tier women, and the second and third terms, $nrN^h p_m^h$ and $nrN^l p_m^l$, are the platform's expected revenue from the messages that high-tier and low-tier women, respectively, send to men.

The following lemma gives the revenue-maximizing platform's optimal proportion of high-tier and low-tier women in men's recommendations.

Lemma 2.3. *The revenue-maximizing platform constructs the proportion of high-tier women in men's recommendations as $x_h^* = \frac{1}{1 + \frac{\sqrt{h_2}}{\sqrt{h_3}}} + \frac{4(\sqrt{h_2} - \sqrt{h_3})}{h_1(\sqrt{h_2} + \sqrt{h_3})}$ and the proportion of low-tier women in men's recommendations as $x_l^* = 1 - x_h^*$, where $h_1 = 2(2q + 1)\beta$, $h_2 = (2q + 1)(1 - \alpha q) - 4r$ and $h_3 = 1 - 4r$.*

Lemma 2.3 shows that the proportion of high-tier women in the revenue-maximizing platform's recommendations, x_h^* , increases in h_2 while decreasing in h_3 . From Lemma 2.2, we observe that a man's probability of sending a message to a high-tier (low-tier) woman, p_h^m (p_l^m), is proportional to h_2 (h_3). Thus, as expected, the revenue-maximizing platform recommends high-tier (low-tier) women more when men are more likely to send a message to a high-tier (low-tier) woman. More

interestingly, we note that a higher value of h_1 reduces the impact of the difference between h_2 and h_3 in x_h^* (as captured by $\sqrt{h_2} - \sqrt{h_3}$). Since h_1 is proportional to the congestion effect, β , it shows the role of the congestion effect in the degree of popularity bias, as discussed next.

The following proposition gives the condition under which the revenue-maximizing platform becomes biased against high-tier or low-tier women in its recommendations.

Proposition 2.1. (a) *The revenue-maximizing platform is biased against low-tier women when $\alpha < \frac{2}{2q+1}$, against high-tier women when $\alpha > \frac{2}{2q+1}$, and unbiased when $\alpha = \frac{2}{2q+1}$. (b) *The congestion effect, β , reduces the bias (if bias exists) in the revenue-maximizing platform's recommendations.**

Proposition 2.1(a) shows that the bias in the revenue-maximizing platform's recommendations depends on the average attractiveness difference between high-tier and low-tier women, q , and the degree of high-tier women's selectivity in accepting messages, α . Given the average attractiveness difference q , when high-tier women are not too selective in accepting messages (i.e., α is low), the revenue-maximizing platform recommends high-tier women more than low-tier women. The reason is as follows. Recall from Lemma 2.2 that as high-tier women become more selective in accepting messages (i.e., α increases), men find it more difficult to match with high-tier women and become less likely to send messages (i.e., p_h^m decreases in α). When high-tier women become too selective in accepting messages (i.e., $\alpha > 2/(2q+1)$), men's probability of sending messages to low-tier women becomes higher than their probability of sending messages to high-tier women (i.e., $p_h^m < p_l^m$); this leads the revenue-maximizing platform to recommend low-tier women more than high-tier women. We also note from the condition that the revenue-maximizing platform is less likely to be biased against low-tier women when the average attractiveness difference, q , is larger. This is because men become discouraged from sending messages to high-tier women when high-tier women are too attractive (i.e., $2/(2q+1)$ is decreasing in q). This leads the revenue-maximizing platform to recommend high-tier users less, reducing bias against low-tier women.

Proposition 2.1(b) shows that the bias in the revenue-maximizing platform's recommendations decreases as the congestion effect, β , increases. This decrease in bias occurs both when the platform is biased against low-tier women, i.e., $\alpha < 2/(2q+1)$, and when the platform is biased against high-tier women, i.e., $\alpha > 2/(2q+1)$. When high-tier women are not too selective in accepting messages,

i.e., $\alpha < 2/(2q + 1)$, we know that men are more likely to send messages to high-tier women, $p_h^m > p_l^m$. An increase in the congestion effect, β , decreases men's probability of sending messages to high-tier women, p_h^m , more than their probability of sending messages to low-tier women, p_l^m . As a result, the platform chooses to recommend low-tier women more, i.e., x_l increases and the bias against low-tier women decreases. We can explain the case when $\alpha > 2/(2q + 1)$ similarly. Therefore, the high congestion effect, β , reduces popularity bias by preventing one tier of women (high- or low-tier) from receiving an overabundance of messages.

2.4.2.2 Match-Maximizing Platform

In this section, we consider the platform that aims to maximize the expected number of successful matches. The match-maximizing platform solves the following problem:

$$\begin{aligned}
\max_{x_h, x_l} \quad & N^m n(x_h p_h^m a_m^h + x_l p_l^m a_m^l) + N^h n p_m^h a_h^m + N^l n p_m^l a_l^m \\
& x_h + x_l = 1 \\
& 0 \leq x_h, x_l \leq 1,
\end{aligned} \tag{M}$$

where the first term, $N^m n(x_h p_h^m a_m^h + x_l p_l^m a_m^l)$, gives the expected number of matches when men send messages to high-tier and low-tier women in the second stage and when these high-tier and low-tier women accept the men's messages in the third stage. The second and third terms in the objective function, $N^h n p_m^h a_h^m$ and $N^l n p_m^l a_l^m$, give the expected number of matches when high-tier and low-tier women send messages to men in the second stage and the expected number of matches when men accept high-tier and low-tier women's messages in the third stage of the game, respectively.

The match-maximizing platform faces a more complex problem than the revenue-maximizing platform because the messages sent in the second stage of the game need to be reciprocated by their recipients in order to yield successful matches at the end of the game. If the platform's recommendations induce users to send messages to only a subset of users (which might maximize the number of messages sent), many messages may not be reciprocated.

The following proposition characterizes the condition under which the match-maximizing platform chooses biased recommendations against low-tier or high-tier women.

Proposition 2.2. *The match-maximizing platform is biased against low-tier women when $(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) > 0$, against high-tier women when $(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) < 0$, and is unbiased when $(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) = 0$, where h_1 , h_2 , and h_3 are defined in Lemma 2.3.*

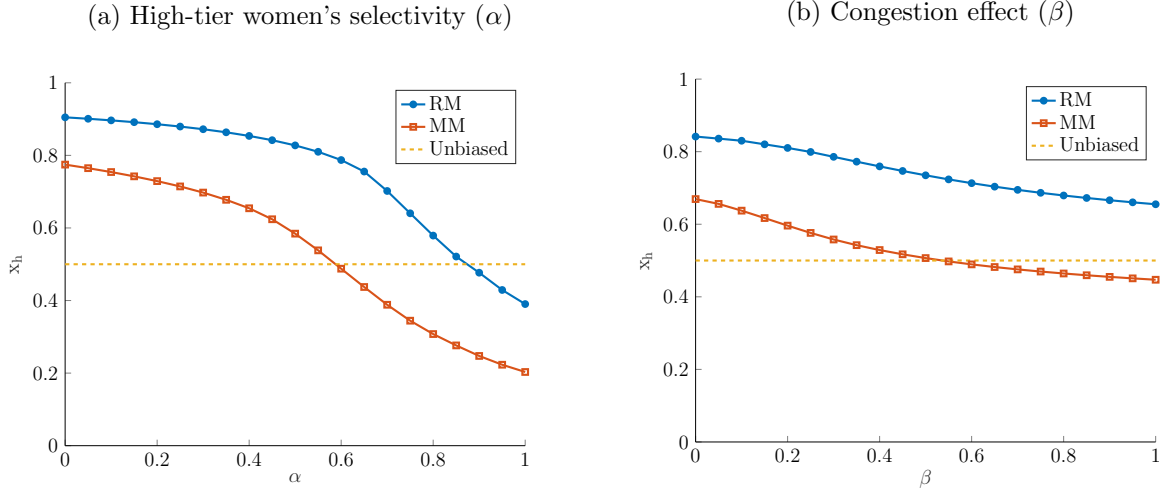
The condition given in Proposition 2.2 is intricate, so we look at each term closely. The first term, $(h_1 + 8)(h_2(1 - \alpha q) - h_3)$, has three parts: $(h_1 + 8)$, $h_2(1 - \alpha q)$, and h_3 . We note from Lemma 2.3 that h_1 is positive and increases in q and β ; h_2 is proportional to men's probability of sending messages to high-tier women, p_h^m , and h_3 is proportional to men's probability of sending messages to low-tier women, p_l^m . The term $(1 - \alpha q)$ captures high-tier women's selectivity in accepting messages, as high-tier women's probability of accepting messages, a_m^h , shown in Lemma 2.1, includes this term. Putting all of this together, we can interpret the first term of the condition as the difference between the probabilities of men's messages being reciprocated by high-tier women and by low-tier women in the third stage of the game. When this difference is positive (negative), i.e., $h_2(1 - \alpha q) - h_3 > 0$ ($h_2(1 - \alpha q) - h_3 < 0$), men are more likely to match with high-tier (low-tier) women; therefore, the match-maximizing platform has the incentive to recommend high-tier (low-tier) women more in order to generate a higher number of successful matches.

The second term of the condition given in Proposition 2.2, $4\beta(h_2^2 - h_3^2)$, ensures that women in one tier do not receive a significantly higher number of messages than those in the other tier for a positive congestion effect, β ; this prevents congestion in the third stage of the game. We know from Lemma 2.1 that a high-tier (low-tier) woman's probability of accepting a message, a_m^h (a_m^l), decreases by $\beta x_h p_h^m$ ($\beta x_l p_l^m$) as the number of messages high-tier (low-tier) women receive, $x_h p_h^m$ ($x_l p_l^m$), increases for $\beta > 0$. On top of that, Lemma 2.2 shows that a man's probability of sending messages to high-tier (low-tier) women, p_h^m (p_l^m), decreases as x_h (x_l) increases. Since men's probability of matching with high-tier (low-tier) women is a multiplication of a_m^h (a_m^l) and p_h^m (p_l^m), and h_2 (h_3) is proportional to p_h^m (p_l^m), we obtain the quadratic difference, $(h_2^2 - h_3^2)$, in the second term of the condition. This term means that the match-maximizing platform does not recommend high-tier (low-tier) women excessively for $\beta > 0$, even if men are more likely to send a message to high-tier (low-tier) women, i.e., $h_2 > h_3$ ($h_3 > h_2$), in order not to overwhelm them with an influx of messages.

Combining the two terms of the condition given in Proposition 2.2, we observe that the match-maximizing platform is biased against low-tier (high-tier) women when men are more likely to match with high-tier (low-tier) women while ensuring that high-tier (low-tier) women do not receive a significantly higher number of messages than low-tier (high-tier) women to avoid congestion in the third stage of the game.

Due to the complexity of the model, there is no closed-form solution for the match-maximizing platform's optimal proportion of high-tier and low-tier women in men's recommendations. Thus, we conduct an extensive numerical experiment to generate further insights into the effect of high-tier women's selectivity (α) and the congestion (β) due to the platform's recommendations. First, observe from Figure 2.2 that the optimal proportion x_h^* of high-tier women in men's recommendations is decreasing in both α and β for the match-maximizing platform. This is consistent with the revenue-maximizing platform discussed above. Second, as expected, Figure 2.2 shows that the revenue-maximizing platform recommends high-tier users more than the match-maximizing platform for all values of α and β , causing a higher bias against low-tier women. When α is low (e.g., $\alpha \in [0, 0.2]$), we observe from Figure 2.2(a) that both types of platforms recommend high-tier women much more than low-tier women. This is because men are more likely to send messages to high-tier women, and high-tier women are not very selective in accepting messages when α is low; therefore, recommending high-tier women more frequently results in a higher number of messages sent and successful matches between users. As values of α increase (e.g., $\alpha \in [0.4, 0.8]$), Figure 2.2(a) shows that the bias against low-tier women in the match-maximizing platform's recommendations decreases more than the bias in the revenue-maximizing platform's recommendations. As high-tier users become highly selective with a high value of α , they reject more messages; thus, the match-maximizing platform recommends high-tier users less than the revenue-maximizing platform, which is solely interested in the number of messages sent. Lastly, we observe from Figure 2.2(b) that both types of platforms are biased against low-tier women when the users' congestion effect, β , is low (e.g., $\beta \in [0, 0.5]$). In this region, the platform increases both the number of messages and the number of successful matches by recommending high-tier women more.

Figure 2.2: Numerical Experiment Results of Platform’s Recommendations



Notes. RM represents the revenue-maximizing platform, and MM represents the match-maximizing platform. In (a), each point on the plot for $\alpha \in \{0, 0.05, 0.1, \dots, 0.95, 1\}$ is the average of the optimal x_h over 1,331 scenarios with $\beta, q, r \in \{0, 0.1, \dots, 0.9, 1\}$. In (b), each point on the plot for $\beta \in \{0, 0.05, 0.1, \dots, 0.95, 1\}$ is the average of the optimal x_h over 1,331 scenarios with $\alpha, q, r \in \{0, 0.1, \dots, 0.9, 1\}$.

2.5 Predictive Model

We utilize different types of interaction data (i.e., users may open cards, send messages, and accept messages) to estimate each user interaction type. The platform’s interaction data can be viewed through the lens of matrix recovery, where a matrix contains one type of interaction (i.e., open, send, or accept) between users. In each matrix, any given user has interacted with only a small subset of other users, and we seek to predict the user’s likelihood of interacting with any other user. Collaborative filtering (CF) models are the most commonly used technique for this purpose. Their underlying principle is that people like items (i.e., users, in the context of online dating) that are similar to other items they like, and they like items that other people with similar tastes like. However, instead of recovering each matrix individually with CF, we seek to recover a three-dimensional tensor, i.e., we seek to fill in missing entries in such a tensor. A three-dimensional tensor in our model contains three types of interactions between users (open, send, and accept), where each slice of the tensor arises from one type of interaction data.

Our estimation procedure has two steps. In the first step, we adapt conventional matrix factorization (e.g., see Koren et al. 2009, Alhejaili and Fatima 2020) to the interaction data tensor and identify latent features, referred to as *CF features*, that are attributed to each user. Assuming that each user is associated with some unknown vector of latent features, we can ascribe to each male and female user a set of latent (i.e., unknown) features. While the first step estimates CF features using all slices of the interaction data tensor, the second step estimates each type of interaction based on the corresponding interaction data slice, the CF features identified in the first step, and the user attribute data (including both individual and pairwise attributes). Note that the CF features and the user attributes are common for all interaction types. For this second step, we employ a random forest model while using the model by Farias and Li (2019) and a neural network model as our benchmarks. We have chosen the random forest model because it allows us to utilize pairwise user attributes in addition to CF features and individual user attributes in making predictions. It also allows us to interpret the importance of input features in prediction results. We present our method’s technical details in Appendix B.7.

To test the accuracy of our model in predicting interactions, we randomly partition the interaction data: 70% is used for training the model and 30% for testing. Since the interaction data is binary, we evaluate performance vis-à-vis a binary classification task, i.e., the task of classifying entries as being equal to 0 or 1. We report the area under the receiver operating characteristic (ROC) curve as a summary statistic in Table 2.3. The ROC curve of a continuous-valued classifier represents the classifier’s trade-off between the true positive rate (defined as the number of true positives divided by the sum of the number of true positives and the number of false negatives) and the true negative rate (defined as the number of true negatives divided by the sum of the number of false positives and the number of true negatives). The area under the ROC curve (AUC) is the most commonly used measure of overall accuracy for prediction models. Its value is always between 0 and 1, and a higher value signifies greater accuracy. In general, an AUC of 0.5 suggests no discrimination, 0.7 to 0.8 is considered acceptable, and 0.8 to 0.9 is considered excellent; in practice, it is very unusual to observe an AUC greater than 0.9 (e.g., Hosmer et al. 2013). The AUC also has an important statistical property: it is equivalent to the probability that the classifier will rank a randomly selected positive sample higher than a randomly selected negative sample (e.g., Fawcett 2006).

Table 2.3: Accuracy of Prediction

		Area under the ROC Curve (AUC)				
Decision		(1)	(2)	(3)	(4)	(5)
		Bilinear	Random Forest	Random Forest	Random Forest	Neural Network
		(CF)	(CF)	(CF+I)	(CF+I+P)	(CF+I+P)
Male Users	Open	0.79	0.85	0.85	0.85	0.86
	Send	0.75	0.88	0.89	0.90	0.88
	Accept	0.60	0.70	0.70	0.72	0.71
Female Users	Open	0.74	0.90	0.91	0.91	0.91
	Send	0.62	0.77	0.83	0.83	0.81
	Accept	0.77	0.81	0.85	0.85	0.84

Notes. “Bilinear” refers to the (bilinear) model of Farias and Li (2019), which can utilize only CF features; “(CF)” means only CF features are used; “(CF+I)” means both CF features and individual attributes are used; and “(CF+I+P)” means CF features, individual attributes, and pairwise attributes are used for predictions.

We observe the following from the results reported in Table 2.3:

1. Comparing the AUC results of columns (1) and (2), we observe that the random forest model outperforms the bilinear model of Farias and Li (2019) on all three types of interaction decisions for both male and female users. The difference is an AUC of 0.11 on average across the six decisions, and it is often substantial: e.g., 0.16 for male users’ decisions to open.
2. Comparing the AUC results of columns (2) and (3) against column (4), we see incremental improvements in using individual and pairwise user attributes in the random forest model. These values are related to the amount of interaction data available. For example, the difference is small in male users’ decisions to open and send and in female users’ decisions to open, for which we observed a large number of interactions (see Table 2.1), whereas the difference is relatively larger in other decisions due to the smaller number of observed interactions.
3. Column (4) shows that the random forest model with all data yields high values of AUC for all three types of interaction decisions for both male and female users. All AUCs exceed 0.8 except male users’ decisions to accept, for which only a small amount of interaction data is available (see Table 2.1). As mentioned above, an AUC above 0.8 is considered excellent in practice.
4. Comparing the AUC results of columns (4) and (5), we observe that the random forest model

performs slightly better than the neural network model, except in the case of male users' opening decisions. Neural networks are known to be a superior model for speech recognition and image and text classification, but their superiority is, in general, not observed in binary classification tasks like ours (Fernández-Delgado et al. 2014). Male users' opening decisions have by far the highest number of observations (see Table 2.1) among all user decisions, helping the neural network achieve a higher AUC score than the random forest model.

Based on the results above, we use the predicted probabilities from the random forest with all data (i.e., CF+I+P in Table 2.3) for our simulations. The simulation results are presented in the next section. In addition, in Appendix B.8, we analyze how various user attributes affect users' interaction decisions based on the predicted probabilities.

2.5.1 Simulations of the Platform

Using the estimated probabilities of users' decisions from our predictive model, we run simulations of the revenue-maximizing platform, the match-maximizing platform, and the unbiased platform over a 15-month period. Our goal is to validate our theoretical findings and gain further insights into how the platform's objective affects users' chances of finding successful matches.

We simulate the sequence of decisions described in Figure 2.1 closely. The platform recommends two male (female) users to each female (male) user every day. The platform determines these recommendations based on its objective. Since the platform earns revenue when users open cards and send messages, the revenue-maximizing platform ranks users based on the probability of open \times send and recommends two users with the highest probability of prompting these actions. The match-maximizing platform aims to maximize the total number of successful matches between users; thus, it ranks users with respect to the probability of open \times send \times accept and recommends two users with the highest probability of prompting these actions. The unbiased platform recommends two users randomly. After simulating the platform's recommendations, we simulate users' open, send, and accept decisions using random draws based on the estimated probabilities from our predictive model. We randomly select a subset of users and then run simulations of the activity on the platform for 15 months.

Table 2.4: Platform’s Simulated Recommendations

	Female to Male Recommendations			Male to Female Recommendations		
	RM	MM	Unbiased	RM	MM	Unbiased
Number of Successful Interactions						
Open	1,900,235	1,828,208	1,680,854	1,066,400	1,062,056	740,816
Send	561,462	453,686	242,881	99,870	97,799	36,089
Accept	81,295	111,845	35,887	44,793	45,443	16,316
Likelihood of Successful Interactions (%)						
Open	47.5	45.7	42.0	26.7	26.6	18.5
Send	14.0	11.3	6.0	2.5	2.4	0.9
Accept	44.9	46.5	45.2	14.5	24.6	14.8
Percentage of Users in Recommendations						
High-Tier	0.69	0.52	0.45	0.68	0.66	0.43
Low-Tier	0.31	0.48	0.55	0.32	0.34	0.57

Notes. **Female to male recommendations (male to female recommendations)** means that the platform recommends female (male) users to male (female) users, male (female) users decide whether to open cards and send messages, and male (female) users decide whether to accept messages if female (male) users decide to open cards and send messages. **RM** represents the revenue-maximizing platform, and **MM** represents the match-maximizing platform. **Unbiased** represents the platform that recommends users randomly.

We measure the bias in recommendations by examining the percentage of high-tier users recommended by the platform. Among female (male) active users, 45% (43%) have attractiveness scores higher than the average attractiveness scores of female (male) users on the platform and are hence labeled as high-tier, and the remaining users are labeled as low-tier. If the platform’s recommendations consist of more than 45% (43%) high-tier female (male) users, we conclude that the platform is biased against low-tier female (male) users.

Table 2.4 presents our simulation results. Comparing the “female to male recommendations” columns with the “male to female recommendations” columns in Table 2.4, we find that the former generates a significantly higher number of successful interactions because male users are more active in opening cards and sending messages. This is expected and consistent with the actual user interaction patterns of the platform, as shown in Table 2.1.

Obtaining users' likelihood of successful interactions, as shown in Table 2.4, allows us to evaluate the performance of the platform's recommendations under different objectives (i.e., revenue-maximizing, match-maximizing, or unbiased). By comparing the recommendations of the revenue-maximizing platform with those of the match-maximizing platform, we can verify the results from the numerical experiment shown in Figure 2.2 that the recommendations of these two platforms are not necessarily at odds. For instance, in the case of male-to-female recommendations, we observe that both the revenue-maximizing recommendations (68%) and the match-maximizing recommendations (66%) consist of significantly more popular male users than the unbiased recommendations (43%). The reason why both objectives lead to a bias against unpopular users is that the revenue-maximizing recommendations maximize the number of messages sent in the second stage of the game, and the match-maximizing recommendations need to ensure a high number of messages sent to maximize the number of successful matches in the third stage of the game. From Figure 2.2(a), we know that this holds when popular users are not highly selective in sending and accepting messages compared to unpopular users, i.e., when α is low. The platform's data suggest that the value of α for popular male users is fairly low: popular male users are only 19% more selective than unpopular male users in sending and accepting messages (i.e., the chance of popular male users sending and accepting messages is only 19% less than that of unpopular users).

In contrast, the data show that the value of α for popular female users is high: popular female users are 70% more selective than unpopular female users in sending and accepting messages, so they are much more selective than popular male users. According to Figure 2.2(a), as popular users become more selective in sending and accepting messages (i.e., as α increases), the bias in the match-maximizing recommendations decreases significantly, whereas the bias in revenue-maximizing recommendations remains high. Consistent with our theory, we observe from Table 2.4 that the revenue-maximizing recommendations (69%) consist of popular female users considerably more than the match-maximizing recommendations (52%). This is because the revenue-maximizing recommendations continue to recommend popular users more in order to maximize the number of messages sent in the second stage of the game. When popular users are significantly more selective than unpopular users, popular users reject many of these messages, resulting in a low number of successful matches on the platform.

The unbiased recommendations result in a significantly lower number of both messages and successful matches than the revenue-maximizing and match-maximizing recommendations, especially when the congestion effect, β , is low. We observe from Figure 2.2(b) that when β is low, the platform is better off providing biased recommendations against low-tier women in order to maximize revenue (from users' messages sent) or the number of successful matches. Our simulation demonstrates that the unbiased recommendations result in 60% and 55% fewer messages sent by users and 59% and 66% fewer successful matches than the revenue-maximizing recommendations and the match-maximizing recommendations, respectively. This suggests that online dating platforms can increase their revenue and users' chances of finding partners simultaneously if they employ a certain degree of bias against unpopular users.

2.6 Extension: Two-Period Model

We present the two-period matching game in this section. In our two-period model, users who match in the first period leave the platform, and users who do not match in the first period stay on the platform for the second period. In addition to the users who stay on the platform, new users join the platform in the second period depending on the match rate observed in the first period. If the platform achieves a high match rate in the first period, more users decide to join the platform in the second period due to the platform's strong reputation in successfully matching its users. We call this the *reputation effect*.

In order to maintain tractability, we make certain simplifications in the model while keeping the main dynamics of our matching game intact. In our two-period model, only men receive recommendations from the platform and decide whether to send messages to the recommended women. After men's decisions, women receive messages from men and decide whether to reciprocate men's messages to match with them successfully. This simplification aligns with our data which show that more than 82% of conversations are initiated by men on the platform (Bapna et al. (2016) show that men are 75% more likely to initiate contact on online dating platforms).

Similar to our one-period model, there are N men and N women on the platform in the first period. The platform recommends one woman to each man. This recommendation is a high-tier (low-tier) woman with probability $x_{h(1)}$ ($x_{l(1)}$). Users' decisions of sending and accepting messages

on both the first and the second periods are the same as in our one-period model. Men on the platform receive a high-tier (low-tier) woman's profile with probability $x_{h(1)}$ ($x_{l(1)}$), send a message to the recommended woman with probability p_h^m (p_l^m), and their message is reciprocated by the high-tier woman with probability a_m^h (a_m^l). As a result, the expected match rate on the platform in the first period becomes $M = x_{h(1)}p_h^ma_m^h + x_{l(1)}p_l^ma_m^l$.

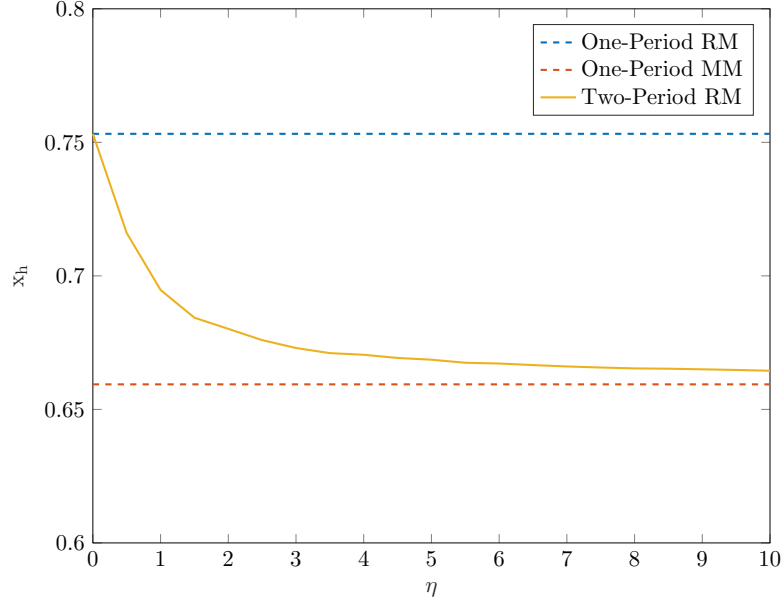
In the second period, new users join the platform depending on the match rate observed in the first period. If the platform is more successful in matching its users, more users decide to join due to the platform's strong reputation. We capture this effect as follows: ζN number of men and women join the platform in the second period where ζ is a function of the match rate, M , observed in the first period. We assume that this function is of linear form, i.e., $\zeta(M) = \eta M$, where η captures the platform's potential *growth rate* in the second period. As the value of η increases, the platform can expand its user base more in the second period if it provides a higher match rate in the first period.

Given the match rate, M , in the first period, the number of men on the platform in the second period, $N_{(2)}^m$, becomes the sum of the number of men who left unmatched in the first period, $N(1 - M)$, and the number of men who join the platform in the second period, $N\zeta(M) = N\eta M$. The platform then decides the probability of men receiving the recommendation of a high-tier (low-tier) woman $x_{h(2)}$ ($x_{l(2)}$) in the second period. The platform's second-period decision is not necessarily the same as in the first period. For instance, we can consider the first period as the platform's early stage, during which the platform aims to grow, and the second period as the platform's mature stage, where the platform already reaches a certain user base. Depending on the platform's objective, the platform would choose different recommendations, which eventually affect users' chances of finding suitable partners.

In the first period, the platform solves the problem of maximizing the number of messages sent in two periods given that the platform chooses the optimal recommendations, $x_{h(2)}^*$ and $x_{l(2)}^*$, in the second period. The first-period problem for the platform is as follows:

$$\begin{aligned}
\max_{x_{h(1)}, x_{l(1)}} \quad & x_{h(1)}Np_h^m(x_{h(1)}) + x_{l(1)}Np_l^m(x_{l(1)}) + x_{h(2)}^*N_{(2)}^mp_h^m(x_{h(2)}^*) + x_{l(2)}^*N_{(2)}^mp_l^m(x_{l(2)}^*) \\
& x_{h(1)} + x_{l(1)} = 1 \\
& 0 \leq x_{h(1)}, x_{l(1)} \leq 1,
\end{aligned} \tag{R_1}$$

Figure 2.3: Change in Platform’s First-Period Recommendations with Growth Rate



Notes. One-Period RM and MM represent the revenue-maximizing and match-maximizing platform’s recommendations in our base model, respectively. Two-Period RM represents the revenue-maximizing platform’s first-period recommendations in our two-period model. For $\eta \in \{0, 0.5, 1, \dots, 9.5, 10\}$, each point on the plot is the average of the optimal x_h over 14,641 scenarios with $\alpha, \beta, q, r \in \{0, 0.1, \dots, 0.9, 1\}$.

Lemma 2.4. *The objective function of R_1 is concave in $x_{h(1)}$ and there is a unique optimal solution $x_{h(1)}^*$ and $x_{l(1)}^*$ such that $x_{h(1)}^* + x_{l(1)}^* = 1$.*

Lemma 2.4 shows that the platform’s first-period recommendation decision takes unique values of $x_{h(1)}$ and $x_{l(1)}$ for each growth rate, η , in the second period. These values are depicted in Figure 2.3. In our two-period model, the revenue-maximizing platform’s first-period decision is the same as in the one-period model when the platform does not grow in the second period, i.e., $\eta = 0$. The reason is that there is no incentive for the platform to provide a high match rate in the first period when there will not be new users who join the platform in the second period to increase the platform’s revenue in two periods. This leads to a higher popularity bias against unpopular users. As η increases, more users join the platform in the second period if the platform provides a good match rate in the first period. Therefore, the revenue-maximizing platform’s first-period recommendations become closer and eventually converge to the match-maximizing platform’s recommendations in the one-period model. This decreases the popularity bias against unpopular

users. These results show that the platform prioritizes its short-term revenue if it already is in a mature stage with a large user base; however, it chooses to provide recommendations that lead to successful matches when there is room for growth in subsequent periods.

2.7 Conclusion

Online dating has experienced rapid growth in recent years. Even though online dating platforms provide users with a tremendous opportunity to connect with others, questions regarding fairness in their recommendation algorithms remain. To address these questions, we study popularity bias in an online dating platform’s recommendations and examine how such bias affects users’ likelihood of finding compatible dating partners and how bias is related to a platform’s ability to generate revenue. While prior research focuses on how to design a platform to maximize users’ chances of finding successful matches, we show that the platforms’ and users’ objectives are not in perfect alignment. The data from a major online dating platform demonstrates that the platform’s algorithm recommends popular users more despite their relatively lower likelihood of accepting messages. To address this problem, we model the platform and the users’ decisions with a three-stage matching game, utilize the data to develop a novel predictive model to estimate users’ decisions, and use the estimated decisions to run simulations of the platform. Our predictive model achieves high accuracy in predicting users’ decisions; therefore, it has great potential to increase the platform’s revenue and the number of successful matches between users.

Our analysis shows that a recommendation approach maximizing the platform’s revenue and an approach maximizing the number of successful matches are not necessarily at odds, even though the revenue-maximizing recommendations lead to greater bias against unpopular users. We show that as popular users become more selective in accepting messages, the revenue-maximizing recommendations remain biased against unpopular users, whereas the match-maximizing recommendations become more inclusive of unpopular users. Thus, popular users help the platform increase revenue and the number of successful matches as long as they do not become “out of reach” (i.e., as long as they do not become too selective in accepting messages, discouraging other users from initiating contact).

Our empirical analysis shows the implicit cost of congestion in users’ decision to accept messages; users tend to become more selective in accepting messages as they receive more messages from other users. We demonstrate that if this implicit cost becomes significantly high, both the revenue-maximizing and the match-maximizing recommendations then include popular and unpopular users more equally. However, our simulations suggest that when this cost is low, unbiased recommendations generate significantly less revenue for the platform and a lower number of successful matches.

This paper contributes to the literature on online matching platforms by studying fairness and bias in recommendation systems, building a novel predictive model to estimate user decisions, and bringing our theoretical and empirical models together to validate our results. Although we focus on a specific dating platform, our model and analysis can be applied to other matching platforms (e.g., freelancing platforms) where the platform makes recommendations to its users and users are heterogeneous (e.g., experienced/inexperienced workers).

Lastly, we discuss the potential considerations that are beyond the scope of this paper. First, various platforms in this business employ different revenue mechanisms. Our industry partner collects fees when users open recommended users’ cards and when they send those users messages, but not when users accept incoming messages. Thus, we believe that the platform can potentially increase total fees by recommending users that are most likely to prompt recipients to open the cards and send messages. In contrast, some platforms such as Tinder and Bumble generate revenue by showing users ads while they engage on the platform, while eHarmony and Match.com collect a fixed subscription fee per period. These platforms might consider recommending users with the highest match probabilities as the best strategy; however, this strategy does not align with the primary objective of such platforms, which is to keep users engaged. As noted by Wu et al. (2018), the main challenge for online dating platforms is that users may leave the platform after finding their matches, reducing the platform’s revenue. Second, a platform needs to consider various other strategic motives (e.g., engaging newly registered or inactive users by recommending the most attractive users to them). The question of how a platform should set its objective and constraints considering various trade-offs is an interesting topic for future study. The results presented by this paper as well as our predictive model, which accurately estimates users’ interaction decisions, would likely benefit such future research endeavors.

Chapter 3

Network Hazard and Superspreaders

3.1 Introduction

The outbreak of COVID-19 has had a significant impact on global health and the economy. The virus, caused by the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), was first identified in Wuhan, China in December 2019 and quickly spread to become a pandemic. The World Health Organization (WHO) declared COVID-19 a Public Health Emergency of International Concern in January 2020 and a pandemic in March 2020. As of 2022, the virus has infected over 600 million people and caused over 6 million deaths worldwide.¹ The rapid spread of COVID-19 and mixed public reactions to advised protective measures highlighted a need for better preparedness for and control of epidemics and pandemics. In particular, behavioral factors are not widely understood in the spread of coronavirus and may play a greater role in efforts to control epidemic and pandemic levels (Gupta et al. 2022).²

Within this context, our theory describes a precautionary scenario on how the unregulated levels of social activity respond to the high availability and efficacy of preventative measures. While vaccines are an important tool in controlling the spread of COVID-19, it is essential to recognize that they do not eliminate the risk of transmission entirely (Polack et al. 2020). As more individuals become vaccinated, they may feel more comfortable interacting with others in high-risk

¹See <https://covid19.who.int/> for WHO Coronavirus (COVID-19) Dashboard.

²Some examples of behavioral factors include how individuals interpret and respond to public health messages, how social norms and cultural practices shape behavior, or how economic and political factors influence the implementation of control measures (Berg and Lin 2020).

settings, such as restaurants, schools, gyms, and grocery stores, or attending family gatherings (Usherwood et al. 2021). The increased number of visitors and rates of activity can potentially push the viral load in closed spaces above threshold levels conducive to infections (Henriques et al. 2021). Additionally, agents employed in these central locations, such as waiters, teachers, trainers, or cashiers can get infected and spread the virus during their asymptomatic period to many visitors. We consider a central location or agent as a central node in a network and allow other agents to choose to form connections with the center to obtain certain benefits. Although largely available and more effective vaccines have the direct effect of reducing aggregate infections, the indirect effect via increased contact with centers and infections *through* the center can potentially offset the direct positive effect, in aggregate. We show that the increased interaction with connectors leads to a manifestation of network hazard (Erol 2019), as the network becomes more concentrated for the virus to spread *through* the centers. On top of increasing aggregate infection rates, the correlated nature of infections by virtue of being transmitted by the center causes a disproportionately large number of simultaneous infections. Such *superspreader* events are particularly important to understand in the face of limited hospital capacity and ventilators.

We use monthly time-series data of visits to various types of businesses such as grocery stores, restaurants, coffee shops, etc. over the period of 2019-2022 from the geospatial data company SafeGraph. This monthly foot traffic data is combined with the publicly available COVID-19 vaccination rate and community transmission rate datasets by the Centers for Disease Control and Prevention (CDC) for the same time period and geographical areas. We find that the testable predictions our model are observed in the data. As vaccination rates increase, there are more visits to the points of interest as well as higher rates of infections.

The rest of this chapter is organized as follows. In §3.2, we review the related literature and discuss our contributions. In §3.3, we present our network model and analysis. We describe our data and empirical results in §3.4 and conclude this chapter in §3.5.

3.2 Related Literature

The spread of COVID-19 in closed environments, such as households and workplaces, and the effectiveness of preventative measures, such as mask use, social distancing and ventilation, in controlling its transmission has been studied in epidemiology, mathematics, and physics.³ Abo and Smith (2020) compare the efficacy of various protective measures including physical distancing and vaccination. This literature does not factor in society’s endogenous behavior. We utilize this literature in specifying functional forms.

The effect of COVID-19 on retail operations, supply chain management, and public health policy design has been studied in operations research. Delasay et al. (2022) and Shumsky et al. (2021) study the impact of social distancing and other measures on consumer behavior and foot traffic.⁴ Additionally, research has played an important role in shaping public health policy by developing models to help with designing infection control policies (Kaplan 2020), forecasting local outbreaks (Chang and Kaplan 2023), and evaluating transmission risks in service facilities (Kang et al. 2022).

While the early literature on COVID-19 explored the mechanics of transmissions, the urgent need for incorporating human behaviors and social processes into mathematical epidemiological models gave rise to a growing literature. El Ouardighi et al. (2022) investigate the role of popular discontent and growing social fatigue in policymaker’s non-therapeutic interventions (e.g., mobility restrictions, securing social interactions) during a pandemic. Wu (2021) model an individual’s decision of whether to engage in social distancing as a social dilemma game played against his/her population. Usherwood et al. (2021) predict COVID-19 trends in the United States accounting for the population’s level of caution and sense of safety, which increases as more individuals take the vaccination (Liu and Wu 2022). In economics, Kaplan et al. (2020), Acemoglu et al. (2021),

³See Buonanno et al. (2020), Bazant et al. (2021), Bazant and Bush (2021), Henriques et al. (2021), Ooi et al. (2021), Salmenjoki et al. (2021), Shang et al. (2022). These works emphasize the importance of limiting cumulative exposure time which is the product of the number of occupants and their time in an enclosed space. This quantity depends on the type of respiratory activity (e.g., singing, talking, etc.) and the infectiousness of the respiratory aerosols, and it increases as the rate of ventilation, air filtration, size of the room, and face mask use increase. Kapoor et al. (2022) estimate the transmission probability of COVID-19 in enclosed spaces using an artificial neural network with real-time collected data.

⁴On the supply chain front, Han et al. (2022) investigate the impact of the pandemic on e-commerce operations, Khan et al. (2021) discuss its impact on medical supply chains, Mak et al. (2022) model and analyze two-dose vaccine distribution, Nikolopoulos et al. (2021) provide predictive analytics tools for forecasting and planning during a pandemic.

Fernández-Villaverde and Jones (2022) incorporate policy analysis into detailed SIR models using the expected arrival time of a vaccine. Our work contributes to this literature on two dimensions. First, we factor in agents that choose socialization rates in response to vaccine availability and efficacy in our theoretical and empirical analysis. Second, we bring the novel notion of network hazard, that is originally introduced in the context of financial networks, to the literature on epidemics to shed some light on superspreaders.

We use foot traffic data from cell phone records provided by SafeGraph. Cronin and Evans (2020) examine the role of state and local restrictions on foot traffic in different essential (e.g., retail) and nonessential (e.g., entertainment) industries. Goolsbee and Syverson (2021) measure how much of the decrease in economic activity resulted from government-imposed restrictions versus people voluntarily staying home to avoid infection. Our primary episode of focus is the recovery of foot traffic starting with the availability of vaccines.

3.3 Model

There is a unit mass of *agents* and one *center*. Each agent can choose to *connect* with the center to obtain some benefits. Agents who connect with the center are called *connected* agents. The center accepts all connections. The center can be a person or a group of people either with a high value from connections or with a commitment to meet all demands from the connections. Examples are teachers in schools, doctors in hospitals, cashiers in grocery stores, trainers in gyms, etc. The center can also be the physical space where people gather such as schools, hospitals, grocery stores, gyms, bars, and restaurants, etc.

There is an infectious *disease*. Agents can *contact* the disease and get *infected*. The contacts can happen exogenously, called *external contact*, at a given rate. The corresponding infections are called external infections. Agents who have not contacted the disease externally can still contact the disease if they are connected and some internally infected agents are also connected. Such contact is called internal contact and the corresponding infections are internal infections.⁵ When a

⁵Note that agents who contacted the disease externally but did not get infected are assumed to not get infected internally. For example, an agent who contacted the virus exogenously but did not get infected builds immunity and does not get infected out of endogenous contact either. Alternatively, all interactions happen repeatedly in a short timespan wherein infected agents are asymptomatic and agents' infection statuses are determined nearly simultaneously with their contact statuses. Since such agents do not get infected internally, we assume without loss of generality that they do not contact the disease externally either.

center is a person, infected connected agents infect the center, which can then create contact with other connected agents. When the center is a physical space, the mass of infected connected agents determine the viral density in the air at the center, which determines the contact probability of other connected agents (Henriques et al. (2021)). Overall, the mass of connected infected agents increases the internal contact probability. There is also a protective measure against infections, such as a vaccine, which we simply call *protection*.⁶ Using protection decreases the infection probability by a factor. We call an agent *protected* if the agent uses protection.

Formally, connecting to the center grants value $v < 1$ to each connected agent. If an agent is infected, it incurs a cost of 1. Each agent has an exogenous contact probability κ . Denoting χ the (endogenous) mass of connected infected agents, a connected agent who has not been contacted exogenously has probability $\Phi(\chi)$ of contacting the virus endogenously, where Φ is an increasing function. An agent who contacts the virus gets infected with probability ι , which is scaled down by $e < 1$ if the agent is protected, down to $e\iota$. We call e^{-1} is the *efficacy* of protection. For simplicity, we take $\iota = 1$.

We denote p the mass of protected agents. Using protection can be a choice or it can be mandated depending on the specific case at hand. We take p to be exogenous and assume that each agent has p probability of being protected. This way, we aim to capture the gradual availability of vaccines in the US during the COVID-19 pandemic. Agents know their protection status when making their connection choice, but contact and infections are unobservable. This is, the interactions happen during the asymptomatic interval of the disease.

Equilibrium Agents, when making their connection decisions, compare the value of the connection to the center with the internal contact probability and the cost of associated potential infection. In particular, a protected agent compares v and $e(1 - \kappa)\Phi(\chi)$ whereas an unprotected agent compares v and $(1 - \kappa)\Phi(\chi)$. This implies that protected agents have a higher expected value from connecting. Then, infection probabilities are given by Table 3.1.

The mass of agents who get infected through external contact is $\theta \equiv (pe + 1 - p)\kappa$. Denote $\mu_p \leq p$ the endogenous mass of connected protected agents and $\mu_u \leq 1 - p$ the endogenous mass of connected unprotected agents. The following cases characterize equilibria.

⁶The model can be generalized to include masks in case of airborne diseases. In the case of STIs, protection can also be condoms.

Table 3.1: Infection Probabilities of Agents

	Not Connected	Connected
Protected	$e\kappa$	$e(\kappa + (1 - \kappa)\Phi(\chi))$
Not Protected	κ	$\kappa + (1 - \kappa)\Phi(\chi)$

1. No agent is connected: $\mu_p = \mu_u = 0$. This case is characterized by $\chi = 0$, $e(1 - \kappa)\Phi(\chi) \geq v$. This is, even the protected agents prefer not to connect even if no other agent is connected. In this case, the mass of internal infections is zero as there are no connections.

2. Some protected agents are connected, no unprotected agents are connected: $\mu_p \in (0, p)$, $\mu_u = 0$. This case is characterized by $\chi = \mu_p e\kappa$, $e(1 - \kappa)\Phi(\chi) = v$. This is, protected agents connect to the center up to the point of being indifferent, at which point unprotected agents prefer not to connect. In this case, the mass of endogenously infected agents is $\mu_p(1 - \kappa)e\Phi(\chi)$. Denoting $\Psi(x) \equiv x\Phi(x)$, which is strictly increasing, the mass is given by $\mu_p(1 - \kappa)e\Phi(\chi) = \frac{1 - \kappa}{\kappa}\Psi(\chi) = \frac{1 - \kappa}{\kappa}\Psi(\Phi^{-1}(\frac{v}{e(1 - \kappa)}))$.

The mass of internal infections is constant in p , as the marginal connected agent has fixed internal infection probability. More importantly, internal infections are *increasing* in e^{-1} . This is, for more effective protection, there are more internal infections. This is an instance of network hazard. Agents do not internalize the infection probability and comfortably connect more when protection is better. In equilibrium, total infections increase.

3. All protected agents are connected, no unprotected agents are connected: $\mu_p = p$, $\mu_u = 0$. This case is characterized by $\chi = e\kappa p$, $(1 - \kappa)\Phi(\chi) \geq v \geq e(1 - \kappa)\Phi(\chi)$. This is, protected agents connect to the center up to the point of being indifferent, at which point unprotected agents prefer not to connect. In this case, the mass of endogenously infected agents is $p(1 - \kappa)e\Phi(\chi) = \frac{1 - \kappa}{\kappa}\Psi(\chi) = \frac{1 - \kappa}{\kappa}\Psi(e\kappa p)$.

The mass of internal infections is decreasing in e^{-1} , but, importantly, *increasing* in p . This is, for more widespread protection, there are more internal infections. This is also an instance of network hazard. When protection is good enough that protected agents prefer to connect, if protection gets more widespread, the mass of connected agents increases, increasing the mass of internal infections.

4. All protected agents are connected, some unprotected agents are connected: $\mu_p = p$, $\mu_u \in (0, 1 - p)$. This case is characterized by $\chi = e\kappa p + \kappa\mu_c$, $v = (1 - \kappa)\Phi(\chi)$. This is, unprotected agents connect to the center up to the point of being indifferent, at which point protected agents prefer to connect. In this case, the mass of endogenously infected agents is $p(1 - \kappa)e\Phi(\chi) + \mu_c(1 - \kappa)\Phi(\chi) = \frac{1-\kappa}{\kappa}\Psi(\chi) = \frac{1-\kappa}{\kappa}\Psi(\Phi^{-1}(\frac{v}{1-\kappa}))$.

This is constant in both p and e^{-1} . The marginal connected agent is unprotected; hence, neither the extent of protection in the society, p , nor the efficacy of protection, e , matters for the infection probability of the marginal agents.

5. All agents are connected: $\mu_p = p$, $\mu_u = 1 - p$. This case is characterized by $\chi = p\kappa + (1 - p)\kappa = \theta$, $v \geq (1 - \kappa)\Phi(\chi)$. This is, even the unprotected agents prefer to connect despite all agents being connected. In this case, the mass of endogenously infected agents is $p(1 - \kappa)e\Phi(\chi) + (1 - p)(1 - \kappa)\Phi(\chi) = \frac{1-\kappa}{\kappa}\Psi(\chi) = \frac{1-\kappa}{\kappa}\Psi(\theta)$.

As θ is decreasing in both p and e^{-1} , so is the mass of internally infected agents.

Note that internal infections are given by $\frac{1-\kappa}{\kappa}\Psi(\chi)$ which is isomorphic to χ . We summarize these cases in the following table of connected external infections χ .

These cases are parametrically exhaustive and hence characterize the equilibrium. For a fixed p , starting with $e = 1$, meaning completely ineffective protection, and gradually increasing the efficacy e^{-1} , we see that internal infections start increasing after the cutoff $e^{-1} = \frac{(1-\kappa)\Phi(0)}{v}$ between case 1 and case 2. During the phase of case 2, more protected agents connect as a function of the efficacy of protection. At the cutoff $e^{-1} = \kappa p \Psi^{-1}(\frac{v\kappa p}{1-\kappa})^{-1}$ between case 2 and case 3, all protected agents are connected and further improvements in efficacy decrease internal infections. This is portrayed in Figure 3.1(a).

Notice the dilemma here. As long as protected agents are not fully connected (case 2), the efficacy of protection increases internal infections. This is network hazard. Only after all protected agents are connected, the efficacy of protection starts reducing internal infections (case 3). However, when all protected agents and some unprotected agents are connected (case 4), the efficacy of protection does not affect internal infections. Therefore, network hazard hurts connected protected agents (case 2) but not the connected unprotected agents (case 4) although, in some sense, better protection is supposed to benefit protected agents compared to unprotected agents.

Table 3.2: Summary of Cases of Equilibrium

Case	Parametric Condition	Connected External Infections	Network Hazard
1	$e(1 - \kappa)\Phi(0) \geq v$	0	
2	$e(1 - \kappa)\Phi(e\kappa p) > v > e(1 - \kappa)\Phi(0)$	$\Phi^{-1}\left(\frac{v}{e(1-\kappa)}\right)$	Increasing in e^{-1}
3	$(1 - \kappa)\Phi(e\kappa p) \geq v \geq e(1 - \kappa)\Phi(e\kappa p)$	$p e \kappa$	Increasing in p
4	$(1 - \kappa)\Phi(\kappa(ep + 1 - p)) > v > (1 - \kappa)\Phi(e\kappa p)$	$\Phi^{-1}\left(\frac{v}{1-\kappa}\right)$	
5	$v \geq (1 - \kappa)\Phi(\kappa(ep + 1 - p))$	$\kappa(ep + 1 - p)$	


Next, consider a fixed e^{-1} . As we start with $p = 0$ meaning no availability of protection, and gradually increase p , we see that internal infections start increasing after the cutoff $p = \frac{1}{e\kappa}\Phi^{-1}\left(\frac{v}{e(1-\kappa)}\right)$ between case 2 and case 3. During the phase of case 3, all protected agents choose to be connected and no unprotected agents do so. In some sense, a protected population who are all exposed to each other through their connection to the center is being scaled up, and so the internal infections increase. At the cutoff $p = \frac{1}{e\kappa}\Phi^{-1}\left(\frac{v}{1-\kappa}\right)$ between case 3 and case 4, some unprotected agents finally find it optimal to connect, and the mass of unprotected agents further increases in protection availability. The marginal unprotected agent has a constant probability of internal infection so the internal infections are constant after this point. This is portrayed in Figure 3.2(a).

A similar dilemma appears here. Whenever there are non-trivial rates of internal infection across protected agents, larger availability of protection increases the mass of internal infections. Only after protection rate is big enough that it becomes optimal for all protected agents to connect and unprotected agents start connecting, the internal infection rate stops increasing. Therefore, larger availability of protection hurts the protected population due to externalities, which is another instance of network hazard.

Finally, note the complementarity between e^{-1} and p . The network hazard region for e^{-1} , i.e., case 2, is between $\frac{(1-\kappa)\Phi(0)}{v}$ and $\kappa p \Psi^{-1}\left(\frac{v\kappa p}{1-\kappa}\right)^{-1}$ which is wider for larger p . This is, if protection is more widespread, the adverse consequences of more effective protection are more prevalent. Similarly, the network hazard region for p , i.e., case 3, is between $\frac{1}{e\kappa}\Phi^{-1}\left(\frac{v}{e(1-\kappa)}\right)$ and $\frac{1}{e\kappa}\Phi^{-1}\left(\frac{v}{1-\kappa}\right)$. If the protection is more effective, the adverse consequences of more widespread protection manifest during higher availability.

The total mass of infections, $\tau \equiv \theta + \frac{1-\kappa}{\kappa}\Psi(\chi)$, is of interest as well. After all, higher efficacy

Figure 3.1: Internal and Total Infections in Efficacy of Protection



Notes. The parameter values are $v = 0.1$, $\kappa = 0.1$, $\alpha = 50$, $p = 0.5$.

e^{-1} and more availability p reduces external infections θ . Comparative statics in this case requires specifying a form for the transmission function Φ to pin down tradeoffs between internal and external infections. Referring to Henriques et al. (2021), Φ is described by $\Phi(\chi) = 1 - e^{-\alpha\chi}$ where $\alpha > 0$ is a constant that depends on a host of exogenous factors.

The adverse consequences of increased p and e^{-1} on internal infections appear for e^{-1} in case 2 and for p in case 3. In the other cases, internal infections are decreasing in e^{-1} and p , so the total infections also decrease. Therefore, we focus on e^{-1} in case 2 and p in case 3 to study total infections. Case 2 is given by $e\kappa p > \chi = -\alpha^{-1} \ln\left(1 - \frac{v}{e(1-\kappa)}\right)$ and some algebra yields


$$\frac{d\tau}{de^{-1}} = -\kappa p e^2 + \frac{v}{\alpha\kappa} \left(\frac{v}{e(1-\kappa) - v} + \chi\alpha \right)$$

Notice that p can be as small as $\frac{\chi}{e\kappa}$ under case 2 so $\frac{d\tau}{de^{-1}}$ can be as large as $\chi\left(\frac{v}{\alpha} - e\right) + \frac{v^2}{\alpha\kappa(e(1-\kappa) - v)}$. Thus, for relatively large e^{-1} , in particular $v > \alpha e$,⁷ we have $\frac{d\tau}{de^{-1}} > 0$. This is, the total infections can increase as protection efficacy e^{-1} increases, particularly if protection is not too widespread but it is highly effective. This is portrayed in Figure 3.1(b).

Next, consider p in case 3. Case 3 is given by $\frac{v}{e(1-\kappa)} > \Phi(\chi) = \Phi(e\kappa p) > \frac{v}{1-\kappa}$, and some algebra

⁷This does not contradict parametric specification of case 2.

Figure 3.2: Internal and Total Infections in Availability of Protection



Notes. The parameter values are $v = 0.1$, $\kappa = 0.1$, $\alpha = 50$, $1 - e = 0.85$.

yields

$$\frac{d\tau}{dp} = (1 - \kappa)e\Phi(e\kappa p) - \kappa(1 - e) + (1 - \kappa)\kappa e^2 p \Phi'(e\kappa p)$$

Similarly, p can be chosen to make $\Phi(e\kappa p)$ arbitrarily close to $\frac{v}{e(1-\kappa)}$, in which case v can be chosen arbitrarily large. In this case, $\frac{d\tau}{dp} > v - \kappa(1 - e)$ which is positive. Hence, the total infections can increase as protection availability p increases, particularly when protection is already widespread and the value of connections is large. This case is particularly relevant to our empirical analysis as we discuss in the next section. This is portrayed in Figure 3.2(b).

In summary, our analysis indicates that there are two cases of network hazard. First, when some protected individuals are connected but unprotected individuals are not connected with the center, the efficacy of protection increases the endogenous infection probability in the network because more protected individuals decide to form connections with the center. This increases the number of internal infections through the center due to negative externalities. The total number of infections, including both internal and external, might increase when the protection is highly effective but not widespread. Second, when all protected individuals are connected but unprotected individuals are not connected, the efficacy of protection decreases the infection probability expectedly. However, the availability of protective measures increases the infection probability because more protected individuals prefer to connect when the protection is good enough, which in turn, increases the


internal infections. Similarly, the total number of infections might increase when protection is already widespread which we observe in our empirical analysis. Only after the efficacy and the availability of protection are sufficiently large that all protected individuals connect, the internal infections decrease because the marginal connected agents' infection probability does not depend on protection efficacy.

3.4 Data and Empirical Analysis

We first present the historical changes in the COVID-19 community transmission rate and the vaccination rate in Allegheny County in Figure 3.3. After the first cases were recorded in March/April 2020, the first serious community transmission in Allegheny County occurred in December 2020/January 2021. With the increasing rate of vaccinations and the ability to socialize in outdoor settings, the number of cases decreased significantly in June/July 2021 to start increasing again in September 2021. The most dramatic increase in the number of cases for the complete pandemic timeline occurred in December 2021/January 2022 when the vaccination rate almost reached 70% in the county. Compared to the same time period previous year when the vaccinations had just started, the number of new cases per 100,000 people quadrupled during this time. This is a manifestation of network hazard where individuals forgo social distancing and engage in high-risk activities more with the comfort of being vaccinated, which in turn results in a significantly higher number of cases in the county.

The other testable prediction of our theory is the rate of social activities. As more doses of the vaccine become available, the number of visits to various central points of interest should increase as individuals are less concerned with infection. To test this, we construct a monthly, county-level time series data of visits to various types of businesses using cell phone location data from the geospatial data company SafeGraph. Couture et al. (2022) show that smartphone data cover a significant fraction of the US population and are broadly representative of the general population in terms of residential characteristics and movement patterns. Our data ranges from March 2019 to February 2022, divided into three-year-long episodes which allows us to net out seasonal effects. The *pre-pandemic* episode is the one-year episode from March 2019 to February 2020 serving as our benchmark. March 2020 is when the first cases of COVID-19 in the US were observed and

Figure 3.3: Infection Transmission Level vs. Vaccination Rate in Allegheny County



Notes. **New cases per 100,000 persons in the past 7 days** is calculated by adding the number of new cases in the last 7 days divided by the population in Allegheny county and multiplying by 100,000. **Vaccination Rate (%)** represents the percentage of people who have completed a primary series (have a second dose of a two-dose vaccine or one dose of a single-dose vaccine) in Allegheny County. Vaccination rates are by definition increased over time. The drop at the end of 2021 is likely a glitch in data collection.

WHO declared COVID-19 to be a pandemic. The *pre-vaccine* episode is the one-year period from March 2020 to February 2021.⁸ The *post-vaccine* episode is the one-year period from March 2021 to February 2022 covering the gradually increasing availability of vaccines up to the point of the Russian invasion of Ukraine. From March 2022 onwards there have been major macroeconomic changes and we observe significant declines in the number of visits to points of interest from March 2022 to December 2022, likely unrelated to COVID-19.

The points of interest we consider are restaurants, gas stations, big retail stores, grocery stores, coffee shops, gyms, and airports.⁹ We first present annual visits to the points of interest in Table 3.3. According to our theory, the visits to points of interest should decline from pre-pandemic episode to pre-vaccine episode, and increase from pre-vaccine episode to post-vaccine episode. This holds for restaurants, gas stations, coffee shops, gyms, and airports. These points of interest provide services or experiences that can not be completely replicated at home. In particular, visits to restaurants and gas stations come back to pre-pandemic levels. Visits to coffee shops, gyms, and airports increase but do not reach their pre-pandemic levels. This can be related to various factors

⁸The first administration of a vaccine in PA was in December 2020. The vaccination rate in March 2021 was around %15.

⁹Restaurants: McDonald's and Wendy's. Gas stations: GetGo, Sunoco, Sheetz. Big retailer stores: Target, Walmart, Costco. Grocery stores: Giant Eagle, ALDI, Market District. Coffee shops: Starbucks. Gyms: Planet Fitness, LA Fitness, Ascend. Airports: Pittsburgh International Airport.

Table 3.3: Annual visits to points of interest

	No. of locations	2019 - 2020		2020 - 2021		2021 - 2022	
		Mean	SD	Mean	SD	Mean	SD
Restaurants	22	11091	865	8518	1334	10448	549
Gas Stations	39	20020	1747	15221	2396	19701	1305
Big Retailers	9	16522	2782	15210	2132	14730	2343
Grocery Stores	21	10586	1477	8671	879	8027	464
Coffee Shops	19	16347	5641	8720	2308	10809	1481
Gyms	7	2441	191	1057	393	1449	151
Airports	1	24688	2938	8235	2916	13444	3541
All Places	118	205355	27835	132596	20353	156133	11650

Notes. Visits are counted starting from March of each year to February of following year.

that have altered consumption habits during the pandemic such as moving to suburbs or buying exercise equipment. The number of visits to big retailers and grocery stores keeps declining from pre-vaccine episode to post-vaccine episode. This could be attributed to a shift towards online shopping after most companies adjusted their infrastructure to accommodate deliveries during the pandemic. We believe this is beyond the scope of our paper and should be addressed in separate work. Accordingly, we focus on restaurants, gas stations, coffee shops, gyms, and airports.


Figures 3.4, 3.5, 3.6, 3.7, and 3.8 present the number of visits to these points of interest. Figures on the left columns overlap pre-pandemic, pre-vaccine, and post-vaccine episodes in annual plots to highlight seasonal changes in behavior. Our theory predicts an upward shift in visits from the pre-vaccine episode (orange lines) to the post-vaccine episode (green lines). Such a shift is evident in figures confirming our prediction.

Figures on the right columns display the entire time series spanning three years. Two points are marked. The black vertical dashed lines correspond to visits in January 2021. The first dose of a vaccine was administered on December 15 in Pennsylvania so it is more suitable to start analyzing the impact of vaccines on visits in the next month. Our theory predicts that visits should gradually increase starting with the black line. The red vertical dashed lines correspond to visits in January 2022. The evident spike in infection rates in Figure 3.3 spans the month of January 2022 from start to finish. Our theory predicts that the number of visits should increase up to January 2022. Such upward trends can be observed in the right columns of Figures 3.4, 3.5, 3.6, 3.7, and 3.8. The corresponding upward trend in infection rates can be seen in the right column of Figure 3.3 in the orange line.

Figure 3.4: Fast food restaurant visits



Figure 3.5: Gas station visits



In summary, our empirical analysis shows that the number of new COVID-19 cases in Allegheny County quadrupled in January 2022 when the vaccination rate reached 70% compared to January 2021 when only 2% of the population was vaccinated. This is the aforementioned network hazard in which individuals connect with the centers (e.g., restaurants, coffee shops) more with the comfort of being vaccinated, which increases the infection rate in the county. This hypothesis holds when we look at the monthly foot traffic data to the places of interest. Compared to the pre-pandemic period, the rate of social activities decreases significantly after the start of the pandemic before individuals started getting vaccinated. With mass vaccinations, we observe an increase in social activities coupled with a sharp increase in the number of COVID-19 cases which is possibly a manifestation of network hazard.

Figure 3.6: Coffee shop visits





Figure 3.7: Gym visits



3.5 Conclusion

The COVID-19 pandemic revealed the importance of incorporating patterns of social interaction and society’s behavioral characteristics in building mathematical epidemiological models. Within this context, our work brings the novel idea of network hazard, originally developed to understand and analyze financial networks, to highlight the potential downside of higher efficacy and availability of protective measures in controlling the spread of the virus if contact rates are left unregulated. While protective measures reduce the risk of transmission, higher availability and efficacy of protective measures potentially make individuals more comfortable in interacting with central agents or locations, which in turn, opens more channels for transmission. As imperfect protective measures are more available and effective, the increased number of contact with central agents and location

Figure 3.8: Airport visits



increases the risk of transmission *through* the center, which can offset the direct benefit of improved protection, in aggregate. Additionally, the correlation of exposures through the central agent or location can create a fat-tailed infection distribution, causing more frequent superspreaders which constrict the healthcare system and cause high fatality. We confirm the testable predictions of our theory by using two data sets regarding Allegheny County. First, we use CDC’s publicly available data on vaccination rates and infection rates. Second, we measure the changes in the number of visits to central points of interest using foot traffic data covering a three-year-long episode from pre-pandemic to post-pandemic episodes.

Our work contributes to the literature on epidemiological models by incorporating a model of individuals’ social behavior and shedding light on superspreaders and massive infections. Future work can incorporate a combination of protective measures and policies, such as masks and social distancing mandates. The use of protective measures is also a choice variable in various settings, which is an important avenue for future work in the light of polarized views of the public in the US. In a dynamic version of our framework, several other questions can be addressed. The evolution of the virus and its several variants would generate complex patterns of infection rates. Finally, agents would react to the news of updated infection rates which can add a robustness check to our broader theory.

Chapter 4

Conclusions

In this dissertation, we study three contemporary matching and allocation problems that arise in organ allocation, two-sided matching platforms, and combating infectious diseases. We provide theoretical models and analysis as well as empirical analysis and insights in order to shed light on these topics. Our work emphasizes the importance of fairness in designing matching and allocation mechanisms and understanding the behavioral factors in the spread of infectious diseases. We provide recommendations to policymakers to promote equity and efficiency in allocating scarce resources and enhance the benefits of current mechanisms for society.

In chapter 1, we study the problem of achieving a fairer liver allocation system where there are disparities in access to transplantation based on patients' height and gender. To address this, we develop a fluid model of the current liver transplant system with fairness constraints, present the optimal policy of allocating deceased-donor livers to transplant patients, provide a computational framework to provide MELD exceptions points to disadvantaged to increase their access to transplantation, and assess the effect of proposed exception points on the efficiency and equity of the current liver allocation system.

Our analysis shows that the Equity Adjusted Mortality Risk policy, which adjusts patients' medical-urgency-based rankings to ensure equity in access to transplantation, is optimal in allocating deceased-donor livers. We show that the dual of the optimal control problem provides a proxy for estimating patients' short-term mortality risk which can be mapped into their laboratory MELD scores. With an easy-to-implement algorithm, we utilize this to provide MELD exception points to disadvantaged patients. Our simulations show that disadvantaged patients can greatly

benefit from MELD score exceptions without decreasing the overall efficiency of the current liver allocation system.

We discuss the potential research directions that can be derived from this chapter. First, this chapter solely focuses on the allocation of deceased-donor livers; therefore, despite their small portion in the total number of liver transplants in the US, incorporating living donor liver transplants into our model and simulations can potentially provide further insights. Second, split liver transplantation (SLT) is a rising trend in the transplant community. How to incorporate SLTs into the MELD-based liver allocation system remains an interesting research question to study further in the future.

In chapter 2, we study the existence of popularity bias in an online dating platform’s recommendations and examine how such bias affects users’ chances of finding compatible dating partners on the platform. We model the platform and the users’ decisions with a three-stage matching game, utilize the data from a major online dating platform to develop a novel predictive model to estimate users’ decisions on the platform and use the estimated decisions to run simulations of the platform.

Our analysis shows that the platform’s revenue-maximizing and match-maximizing recommendations are not necessarily at odds, even though the former leads to greater bias against unpopular users. We show that the bias in the platform’s recommendations is affected by how selective popular users are on the platform and the congestion effect. As popular users become more selective in accepting messages, the degree of bias in the match-maximizing platform’s recommendations decreases whereas it stays high in the revenue-maximizing platform’s recommendations. As the congestion effect increases, both revenue-maximizing and match-maximizing platforms’ recommendations include popular and unpopular users more equally. Additionally, our novel predictive model achieves very high accuracy in estimating users’ decisions; therefore, it allows us to test and verify our theoretical results on the user interaction data.

There are interesting research directions to follow from this chapter. First, online dating platforms employ different revenue models such as fee-for-service and fixed subscription fees per period. A platform’s revenue model potentially affects its strategy in providing recommendations to its users. Therefore, one can investigate the effect of the platform’s revenue model on its bias in recommendations. Second, a platform needs to consider various strategic motives to expand its user base and keep existing users engaged on the platform. The question of how a platform should set

its objective considering various trade-offs can provide further insights into understanding current online dating platforms.

In chapter 3, we highlight the unintended consequences of higher efficacy and availability of protective measures in combating infectious diseases with the novel idea of network hazard. While protective measures reduce the risk of transmission, individuals feel more comfortable engaging in high-risk activities with higher availability and efficacy of protective measures. This potentially opens more channels for transmission and increases the probability of having superspreader events. We model this problem as a network of agents with a central location where agents choose to form a connection with the center to obtain benefits. Higher availability and efficacy of protective measures decrease each agent's risk of infection, leading more agents to form connections with the center. However, a higher density of connections with the center increases the risk of transmission through the center. This is a manifestation of network hazard.

Our analysis shows that there are two cases of network hazard. First is when some individuals who employ protective measures connect but unprotected individuals do not connect with the center. In this case, the infection probability through the center increases with the efficacy of protection because more protected individuals decide to form connections with the center. Second is when all protected individuals are connected but unprotected individuals are not connected. The availability of protective measures increases internal infection probability because more protected individuals choose to connect with the center when the efficacy of protection is good enough. Our empirical analysis shows that the number of COVID-19 infections quadrupled in January 2022 when the vaccination rate reached 70% compared to January 2021 when this number was only 2%. This is a case of network hazard where monthly foot traffic data show that significantly higher number of individuals engaged in social activities in late 2021 and early 2022 compared to late 2020 and early 2021.

We discuss the potential future research directions that can be followed from this chapter. First, we can investigate the effect of the combination of protective measures and mandates enforced by the local authorities in the spread of the virus. Second, individuals react to the news of updated infection rates differently which can be incorporated into the dynamic version of our model. Finally, the evolution of the virus and the different infection patterns of new virus variants can provide further insights as to how policymakers should approach the next pandemics and epidemics.

Appendix A

Additional Material for Chapter 1

A.1 Estimation of ESLD Patients' 90-Day Mortality Risk

The logistic regression coefficients used for predicting the 90-day mortality risk of ESLD patients can be seen in Table A.1.

Table A.1: Logistic regression coefficients for predicting 90-day mortality risk

	Coefficient	Standard Error	z value	$P(> z)$
Intercept	-6.817	0.024	-280.1	$< 10^{-6}$
Laboratory MELD	0.237	0.001	221.5	$< 10^{-6}$

A.2 Summary of Notation

Symbol	Definition
i	Index for the static class of the patient
j	Index for the health status of the patient, i.e., MELD score
k	Index for the type of the donor liver
t	Time index
n	Number of organ offers to patients

Continued on next page

Table A.2 – continued from previous page

Index	Definition
\mathcal{I}	Set of static classes of patients
\mathcal{J}	Set of the health status of patients, i.e., MELD scores
\mathcal{K}	Set of liver types
Φ	Feasible set of allocations
β	Matrix of patients' health transition rates out of MELD scores
γ	Matrix of patients' health transition rates into MELD scores
ρ	Inverse of the average likelihood of transplantation
λ_{ij}	Arrival rate of class ij patients
μ_k	Arrival rate of type k livers
P^k	Matrix of patients' probability of accepting liver offers
x_{ij}	Number of class ij patients in the system
$\alpha_{jj'}$	Transition rate of a patient's health status from j to j'
d_j	Mortality rate of a patient in health status j
u_{ijk}	The rate of allocating type k livers to class ij patients
INF	Set of incompatible patient type - liver type pairs
p_{ijk}	Probability of a class ij patient accepting an offered liver type of k
π_{ijk}^n	Probability of a type k liver being transplanted to a class ij patient when offered to n patients

A.3 Proofs

Proof of Theorem 1. The primal optimal control problem (P) includes a set of integral constraints (4) that ensures fairness with respect to the likelihood of transplantation. We introduce a new state variable $w_i^k(t)$ where $w_i^k(t) = \int_0^t u_i^k(\tau)d\tau$, $w_i^k(0) = 0$ for $i \in \mathcal{I}$, $k \in \mathcal{K}$ and $t \geq 0$. As a result, we

obtain $\dot{w}_i^k(t) = u_i^k(t)dt$ and constraint (4) in the primal problem becomes a terminal condition as:

$$\sum_{k=1}^K e \cdot w_i^k(T) = \frac{1}{\rho} \lambda_i T \text{ for } i \in \mathcal{I}. \quad (\text{A.1})$$

We follow the road map provided by Rockafellar and Wets (2009) to derive the dual problem of control associated with (P). We can write our convex integrand L on $[0, T] \times \mathbb{R}^{IJ} \times \mathbb{R}^{IJ} \times \mathbb{R}^{IJ} \times \mathbb{R}^{IJ}$ as follows (for the sake of notation, t is dropped for time dependent variables):

$$\begin{aligned} L(t, x, \dot{x}, w, \dot{w}) &= (e \cdot d) \cdot x + \chi_{\mathbb{R}_+^{IJ}}(x) + \chi_{\mathbb{R}_+^{IJ}}(w) + \sum_{k=1}^K \chi_{\mathbb{R}_+^{IJ}}(u^k) + \sum_{k=1}^K \chi_{\mathbb{R}_-}(e \cdot u^k - \mu_k(t)) \\ &+ \sum_{(i,k) \in INF} \chi_{\mathbb{R}_-}(e \cdot u_i^k - \epsilon) \end{aligned}$$

if $\dot{x}(t) = \lambda(t) - \sum_{k=1}^K P^k u^k(t) - (d + \beta - \gamma)x(t)$ and $\dot{w}(t) = u(t)$, otherwise $L(t, x, \dot{x}, w, \dot{w}) = \infty$.

This way, we append the hard constraints of (P) as penalty expressions to the objective function.

The expressions $\chi_{\mathbb{R}_+^{IJ}}(x)$ and $\chi_{\mathbb{R}_+^{IJ}}(w)$ ensure the non-negativity of the state variables x and w , and

the expression $\sum_{k=1}^K \chi_{\mathbb{R}_+^{IJ}}(u^k)$ ensures the non-negativity of the control variable u . The constraint

related to the allocation of organs not exceeding the supply is expressed by the penalty term

$\sum_{k=1}^K \chi_{\mathbb{R}_-}(e \cdot u^k - \mu_k(t))$ and the infeasible allocations is expressed by $\sum_{(i,k) \in INF} \chi_{\mathbb{R}_-}(e \cdot u_i^k - \epsilon)$.

We note that the infeasible allocations are restricted within ϵ . The system dynamics equations are

incorporated in L by requiring $\dot{x}(t)$ to be equal to $\lambda(t) - \sum_{k=1}^K P^k u^k(t) - (d + \beta - \gamma)x(t)$ and $\dot{w}(t)$

to be equal to $u(t)$.

Next, we define the functional l on $\mathbb{R}^{IJ} \times \mathbb{R}^{IJ} \times \mathbb{R}^{IJ}$ taking values on $\mathbb{R} \cup \{\infty\}$ for the initial

state of the problem and terminal conditions. Initially, there are $x_{ij}(0)$ patients in class ij and

$w(0)$ is equal to 0 because we have not allocated any organs yet. As the terminal condition, $x(T)$ is

not restricted because we minimize pre-transplant mortality on $[0, T]$ and we have $e \cdot w(T) = \lambda T / \rho$

to ensure that all patient classes have an equal likelihood of receiving an organ transplant. The

functional l is defined as $l(x_0, w_0, w_T) = l_0(x_0, w_0) + l_T(w_T)$ where $l_0(x_0, w_0) = \chi_{\{(x(0), 0)\}}(x_0, w_0)$ and

$l_T(w_T) = \chi_{\{\lambda T / \rho\}}(w_T)$. The functional l_0 and l_T dictate the initial and terminal state conditions,

respectively. As a result, the primal problem (P) becomes a problem of minimizing

$$\int_0^T L(t, x(t), \dot{x}(t), w(t), \dot{w}(t)) dt + l(x_0, w_0, w_T).$$

In our next step, we compute the conjugates to the functions L and l . Let L^* denote the conjugate to L . To be specific,

$$L^*(t, s, p, r, q) = \sup_{z \in \mathbb{R}^{IJ}, y \in \mathbb{R}^{IJ}, v \in \mathbb{R}^{IJ}, m \in \mathbb{R}^{IJ}} \{z \cdot s + y \cdot p + v \cdot r + m \cdot q - L(t, z, y, v, m)\}$$

for $s, p, r, q \in \mathbb{R}^{IJ}$. We can express L^* more explicitly as follows. Note that $L(t, z, y, v, k) < \infty$ only if $z \geq 0$, $v \geq 0$ and there exists some $u^k \in \mathbb{R}_+^{IJ}$ such that $y = \lambda(t) - \sum_{k=1}^K P^k u^k - (d + \beta - \gamma)z$, $m^k = u^k$, $e \cdot u^k \leq \mu_k(t)$, $u_i^k \leq \epsilon$ for $(i, k) \in INF$ and $u^k \geq 0$ for $k \in \mathcal{K}$. For $s, p, r, q \in \mathbb{R}^{IJ}$, we write L^* as

$$L^*(t, s, p, r, q) = \sup_{z \in \mathbb{R}_+^{IJ}, u(t) \in \Phi(t), v \in \mathbb{R}^{IJ}} \left\{ z \cdot s + p \cdot \left(\lambda(t) - \sum_{k=1}^K P^k u^k - (d + \beta - \gamma)z \right) + v \cdot r \right. \\ \left. + q \cdot u^k - (e \cdot d) \cdot z \right\}$$

by replacing y with $\lambda(t) - \sum_{k=1}^K P^k u^k - (d + \beta - \gamma)z$, m^k with u^k for feasible u^k and noting that $L(t, z, y, v, k) = (e \cdot d) \cdot z$. We rearrange the terms as follows

$$L^*(t, s, p, r, q) = \sup_{z \in \mathbb{R}_+^{IJ}} \{z \cdot (s - p(d + \beta - \gamma) - e \cdot d)\} + p \cdot \lambda(t) + \sup_{v \in \mathbb{R}^{IJ}} \{v \cdot r\} \\ + \sup_{u(t) \in \Phi(t)} \sum_{k=1}^K \{-p \cdot P^k u^k + q \cdot u^k\}$$

because we can take the supremum for z , v and u^k separately for each k . We have $\sup_{z \in \mathbb{R}_+^{IJ}} \{z \cdot (s - p(d + \beta - \gamma) - e \cdot d)\} = \chi_{\mathbb{R}_-^{IJ}} \{s - p(d + \beta - \gamma) - e \cdot d\}$ since $\sup_{z \in \mathbb{R}_+^{IJ}} \{z \cdot (s - p(d + \beta - \gamma) - e \cdot d)\}$ becomes ∞ if $(s - p(d + \beta - \gamma) - e \cdot d)_{ij} > 0$ for $i \in \mathcal{I}$ and $j \in \mathcal{J}$. Also, we obtain $\sup_{u(t) \in \Phi(t)} \sum_{k=1}^K \{-p \cdot P^k u^k + q \cdot u^k\} = \inf_{u(t) \in \Phi(t)} \sum_{k=1}^K \{(p \cdot P^k - q)u^k\}$. Therefore, L^* can be written as follows

$$L^*(t, s, p, r, q) = \chi_{\mathbb{R}_-^{IJ}} \{s - p(d + \beta - \gamma) - e \cdot d\} + p \cdot \lambda(t) + \sup_{v \in \mathbb{R}^{IJ}} \{v \cdot r\} + \inf_{u(t) \in \Phi(t)} \sum_{k=1}^K \{(p \cdot P^k - q)u^k\}.$$

Using the conjugate L^* of the primal integrand L , we calculate the dual integrand M . For $t \in [0, T]$ and $s, p, r, q \in \mathbb{R}^{IJ}$, the dual integrand M is given by $M(t, p, s, q, r) = L^*(t, s, p, r, q)$. That is, for $t \in [0, T]$, we have

$$\begin{aligned} M(t, y(t), \dot{y}(t), z(t), \dot{z}(t)) &= L^*(t, \dot{y}(t), y(t), \dot{z}(t), z(t)) \\ &= \chi_{\mathbb{R}^{IJ}} \{ \dot{y}(t) - y(t)(d + \beta - \gamma) - e \cdot d \} + y(t) \cdot \lambda(t) + \sup_{v \in \mathbb{R}^{IJ}} \{ v \cdot \dot{z}(t) \} \\ &\quad + \inf_{u(t) \in \Phi(t)} \sum_{k=1}^K \{ (y(t) \cdot P^k - z(t)) u^k \}, \end{aligned}$$

where $\chi_{\mathbb{R}^{IJ}} \{ \dot{y}(t) - y(t)(d + \beta - \gamma) - e \cdot d \}$ ensures that $\dot{y}(t) \leq y(t)(d + \beta - \gamma) + d$ for $t \in [0, T]$. Finally, we need to derive the terminal conditions associated with the dual problem. For this, we define the functional m on $\mathbb{R}^{IJ} \times \mathbb{R}^{IJ} \times \mathbb{R}^{IJ}$ as follows:

$$m(y_0, z_0, z_T) = l_0^*(y_0, z_0) + l_T^*(-z_T)$$

where l_0^* and l_T^* are the conjugates of l_0 and l_T . We calculate l_0^* as follows: $l_0^*(y_0, z_0) = \sup_{x, w} \{ y \cdot x + z \cdot w - l_0(x, w) \} = \sup_{x \in \{x(0)\}, w=0} \{ y \cdot x \} = x(0) \cdot y$. Similarly, $l_T^*(-z_T) = \sup_w \{ w \cdot z \} = \sup_{w=\lambda T/\rho} \{ -w \cdot z \} = -z \lambda T / \rho$. The dual problem of control become to minimize

$$\int_0^T M(t, y(t), \dot{y}(t), z(t), \dot{z}(t)) dt + m(y_0, z_0, z_T)$$

that is equivalent to minimizing

$$\int_0^T [y(t)\lambda(t) + f(t, y(t), z(t))] dt + x(0) \cdot y(0) - \frac{z(T)\lambda T}{\rho}$$

subject to

$$y(t) = y(0) + \int_0^t \dot{y}(s) ds \tag{D}$$

$$z(t) = \int_0^t \dot{z}(s) ds$$

$$\dot{z}(t) = 0$$

$$\dot{y}(t) \leq y(t)(d + \beta - \gamma) + d$$

where $f(t, y(t), z(t)) = \inf \sum_{k=1}^K \{(y(t) \cdot P^k - z(t))u^k : u(t) \in \Phi(t)\}$. To conclude that the objective function values of the primal problem (P) and the dual problem (D) are equal to each other by using Theorem 4 of Rockafellar and Wets (2009), we need to show that the primal problem (P) is feasible and bounded. It is bounded because $e \cdot u^k(t) \leq \mu_k(t)$ for $k \in \mathcal{K}$ and $t \geq 0$. Given that the primal problem is bounded, $\dot{z}(t) = 0$ to ensure that the dual problem is bounded as well. To show that there is a feasible u satisfying the constraints of P, we need an additional assumption on the likelihood of transplant constraint. The average likelihood of transplantation, $1/\rho$, must be small enough so that there exists u such that $e \cdot u^k(t) \leq \mu_k(t)$, $w(t) = \int_0^t u(\tau)d\tau$ and $w(T) = \lambda T/\rho$. With this additional assumption, we conclude that the primal problem (P) is also feasible ensuring that the objective function values of (P) and (D) are equal to each other.

We need to derive the coextremality conditions for the primal - dual problem pair to complete the proof. By Theorem 5 of Rockafellar and Wets (2009), let u be a feasible organ allocation for (P) with the corresponding state trajectories x and w , and let y and z be a feasible control for (D), the control u is optimal for (P) and y and z are optimal for (D), if and only if they satisfy the following coextremality conditions:

$$(y(0), z(0), -z(T)) \in \partial l(x(0), w(0), w(T)) \text{ and}$$

$$(\dot{y}(t), y(t), \dot{z}(t), z(t)) \in \partial L(t, x(t), \dot{x}(t), w(t), \dot{w}(t)) \text{ for almost every } t \in (0, T)$$

where ∂L and ∂l are the subgradients of the convex integrand L and the functional l , defined above. First, we calculate the subgradient of L from its epigraphical normals. For $h : \mathbb{R}^n \rightarrow [-\infty, +\infty]$ and any point \bar{x} at which h is finite, we have $\partial h(\bar{x}) = \{v : (v, -1) \in N_{\text{epi } h}(\bar{x}, h(\bar{x}))\}$ where $\text{epi } h$ denotes the epigraph of h defined as $\text{epi } h := \{(x, \gamma) \in \mathbb{R}^n \times \mathbb{R} : \gamma \geq h(x)\}$, and $N_{\text{epi } h}(\bar{x}, h(\bar{x}))$ is the set of vectors to the set $\text{epi } h$ at $(\bar{x}, h(\bar{x}))$ in the general sense as in Definition 6.3 of Rockafellar & Wets (1997). For $t \in [0, T]$, the epigraph of the integrand L is defined as follows: $\text{epi } L(t)$ consists of points $(x, \dot{x}, w, \dot{w}, \gamma) \in \mathbb{R}^{4I+1}$ such that

$$\dot{x} = \lambda(t) - \sum_{k=1}^K P^k u^k - (d + \beta - \gamma)x, \quad \dot{w}^k = u^k, \quad x \geq 0, \quad w \geq 0, \quad \gamma \geq (e \cdot d) \cdot x,$$

$$e \cdot u^k(t), u^k \geq 0 \text{ for } k \in \mathcal{K}, u_i^k(t) \leq \epsilon \text{ for } (i, k) \in INF,$$

since the points $(x, \dot{x}, w, \dot{w}) \in \mathbb{R}^{4IJ}$ where $L(t, x, \dot{x}, w, \dot{w}) = \infty$ are such that the vertical line $(x, \dot{x}, w, \dot{w}) \times \mathbb{R}$ misses epi $L(t)$. Then, we can write

$$\partial L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}) = \left\{ (v^1, v^2, v^3, v^4) \in \mathbb{R}^{4IJ} : (v^1, v^2, v^3, v^4, -1) \in N_{\text{epi } L(t)}(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}})) \right\}.$$

First, note that for $t \in [0, T]$, epi $L(t)$ is a convex set and the point $(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}))$ is an element of epi $L(t)$ for $(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}) \in \mathbb{R}^{4IJ}$. Let \mathbf{v} denote an arbitrary element of \mathbb{R}^{4IJ+1} where the first IJ components of \mathbf{v} is denoted as v^1 , the subsequent IJ components by v^2, v^3 and v^4 , and the last component by v^γ . That is, $\mathbf{v} = [v^1, v^2, v^3, v^4]^T$ where $v^1, v^2, v^3, v^4 \in \mathbb{R}^{IJ}$ and $v^\gamma \in \mathbb{R}$. Then, we use Theorem 6.9 of Rockafellar and Wets (2009) which gives

$$\begin{aligned} & N_{\text{epi } L(t)}(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}})) \\ &= \left\{ \mathbf{v} \in \mathbb{R}^{4IJ+1} : \left[(x, \dot{x}, w, \dot{w}, \gamma) - (\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}})) \right] \cdot \mathbf{v} \leq 0, \forall (x, \dot{x}, w, \dot{w}, \gamma) \in \text{epi } L(t) \right\}. \end{aligned} \tag{A.2}$$

We next establish the following properties of $N_{\text{epi } L(t)}(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}))$ for $t \in [0, T]$ which will assist us in finding the subgradients of L .

Property 1. For $t \in [0, T]$, if $\mathbf{v} = (v^1, v^2, v^3, v^4, v^\gamma)^T \in N_{\text{epi } L(t)}(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}))$, then $v^1 \leq v^2(d + \beta - \gamma) - v^\gamma d$. Moreover, $v_{ij}^1 = -v^\gamma d_{ij} + [v^2(d + \beta - \gamma)]_{ij}$ when $\bar{x}_{ij} > 0$.

To verify Property 1, we first show that any $\mathbf{v} = (v^1, v^2, v^3, v^4, v^\gamma)^T$ such that $v_{ij}^1 > -v^\gamma d_{ij} + [v^2(d + \beta - \gamma)]_{ij}$ for some ij cannot be in $N_{\text{epi } L(t)}(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}))$. Suppose not. Then, we could find an element $(\tilde{x}, \tilde{\dot{x}}, \tilde{w}, \tilde{\dot{w}}, \tilde{\gamma})$ of epi $L(t)$ such that it is equal to $(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}) + d_{ij}(\tilde{x}_{ij} - \bar{x}_{ij}))$ except for $\tilde{x}_{ij} > \bar{x}_{ij}$ and $\tilde{\dot{x}} = \lambda(t) - \sum_{k=1}^K P^k u^k - (d + \beta - \gamma)\tilde{x}$. However, in that case, we obtain

$$\begin{aligned} & \left[(\tilde{x}, \tilde{\dot{x}}, \tilde{w}, \tilde{\dot{w}}, \tilde{\gamma}) - (\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}})) \right] \cdot \mathbf{v} = v_{ij}^1(\tilde{x}_{ij} - \bar{x}_{ij}) + v^2 \cdot (\tilde{\dot{x}} - \bar{\dot{x}}) + v^\gamma d_{ij}(\tilde{x}_{ij} - \bar{x}_{ij}) \\ &= (\tilde{x}_{ij} - \bar{x}_{ij})(v_{ij}^1 + v^\gamma d_{ij}) + v^2(d + \beta - \gamma)(\tilde{\dot{x}} - \bar{\dot{x}}) \\ &= \left(v_{ij}^1 + v^\gamma d_{ij} - [v^2(d + \beta - \gamma)]_{ij} \right) (\tilde{x}_{ij} - \bar{x}_{ij}) \\ &> 0, \end{aligned}$$

contradicting that $(v^1, v^2, v^3, v^4, v^\gamma)^T \in N_{\text{epi } L(t)}(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$. Similarly, we can show that if $\bar{x}_{ij} > 0$, then any $\mathbf{v} = (v^1, v^2, v^3, v^4, v^\gamma)^T$ such that $v_{ij}^1 \neq -v^\gamma d_{ij} + [v^2(d + \beta - \gamma)]_{ij}$ for some ij cannot be in $N_{\text{epi } L(t)}(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$. Therefore, Property 1 proves the coextremality condition that for $t \in [0, T]$, $\dot{y}(t) \leq d + y(t)(d + \beta - \gamma)$ and whenever $x_{ij}(t) > 0$, it must be that $\dot{y}_{ij}(t) = d_{ij} + [y(t)(d + \beta - \gamma)]_{ij}$.

Property 2. For $t \in [0, T]$, if $\mathbf{v} = (v^1, v^2, v^3, v^4, v^\gamma)^T \in N_{\text{epi } L(t)}(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$, then $v^3 \leq 0$. Similar to Property 1, we show that any $\mathbf{v} = (v^1, v^2, v^3, v^4, v^\gamma)^T$ such that $v_{ij}^3 > 0$ for some ij cannot be in $N_{\text{epi } L(t)}(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$. Suppose not. Then, we could find an element $(\tilde{x}, \tilde{x}, \tilde{w}, \tilde{w}, \tilde{\gamma})$ of $\text{epi } L(t)$ such that it is equal to $(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$ except for $\tilde{w}_{ij} > \bar{w}_{ij}$. In that case, we obtain

$$\begin{aligned} \left[(\tilde{x}, \tilde{x}, \tilde{w}, \tilde{w}, \tilde{\gamma}) - (\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w})) \right] \cdot \mathbf{v} &= v_{ij}^3 (\tilde{w}_{ij} - \bar{w}_{ij}) \\ &> 0, \end{aligned}$$

contradicting that $(v^1, v^2, v^3, v^4, v^\gamma)^T \in N_{\text{epi } L(t)}(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$. Therefore, Property 2 along with the fact that the primal problem is bounded proves the coextremality condition that for $t \in [0, T]$, $\dot{z}(t) = 0$.

Property 3. For $t \in [0, T]$ and $k \in \mathcal{K}$, if $\bar{x} = \lambda(t) - \sum_{k=1}^K P^k \bar{u}^k - (d + \beta - \gamma)\bar{x}$ and $\bar{w} = \bar{u}$ for \bar{u}^k such that $\bar{u}^k \geq 0$, $e \cdot \bar{u}^k \leq \mu_k(t)$, $\bar{u}_i^k \leq \epsilon$ for $(i, k) \in INF$ and $\mathbf{v} = (v^1, v^2, v^3, v^4, v^\gamma) \in N_{\text{epi } L(t)}(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$, then, $\bar{u}^k \in \arg \min_{z \in \Phi(t)} \left\{ (P^k \cdot v^2 - v^4) \cdot z^k \right\}$.

To establish Property 3, we first recall that for any $(x, \dot{x}, w, \dot{w}, \gamma) \in \text{epi } L(t)$, there exists some $u^k \in \mathbb{R}^{IJ}$ for $k \in \mathcal{K}$ such that

$$\dot{x} = \lambda(t) - \sum_{k=1}^K P^k u^k - (d + \beta - \gamma)x, \quad \dot{w}^k = u^k, \quad x \geq 0, \quad w \geq 0, \quad u^k(t) \in \Phi(t), \quad \text{and } \gamma \geq (e \cdot d) \cdot x$$

For an arbitrary $k' \in \mathcal{K}$, consider now an element $(\bar{x}, \dot{x}, \bar{w}, \dot{w}, (e \cdot d) \cdot x) \in \text{epi } L(t)$ where $\dot{x} = \lambda(t) - (d + \beta - \gamma)\bar{x} - \sum_{k \neq k'}^K P^k \bar{u}^k - P^{k'} u^{k'}$, $\dot{w}^k = \bar{u}^k$ for $k \neq k'$ and $\dot{w}^{k'} = u^{k'}$. Then, the following holds for $\mathbf{v} = (v^1, v^2, v^3, v^4, v^\gamma) \in N_{\text{epi } L(t)}(\bar{x}, \bar{x}, \bar{w}, \bar{w}, L(t, \bar{x}, \bar{x}, \bar{w}, \bar{w}))$:

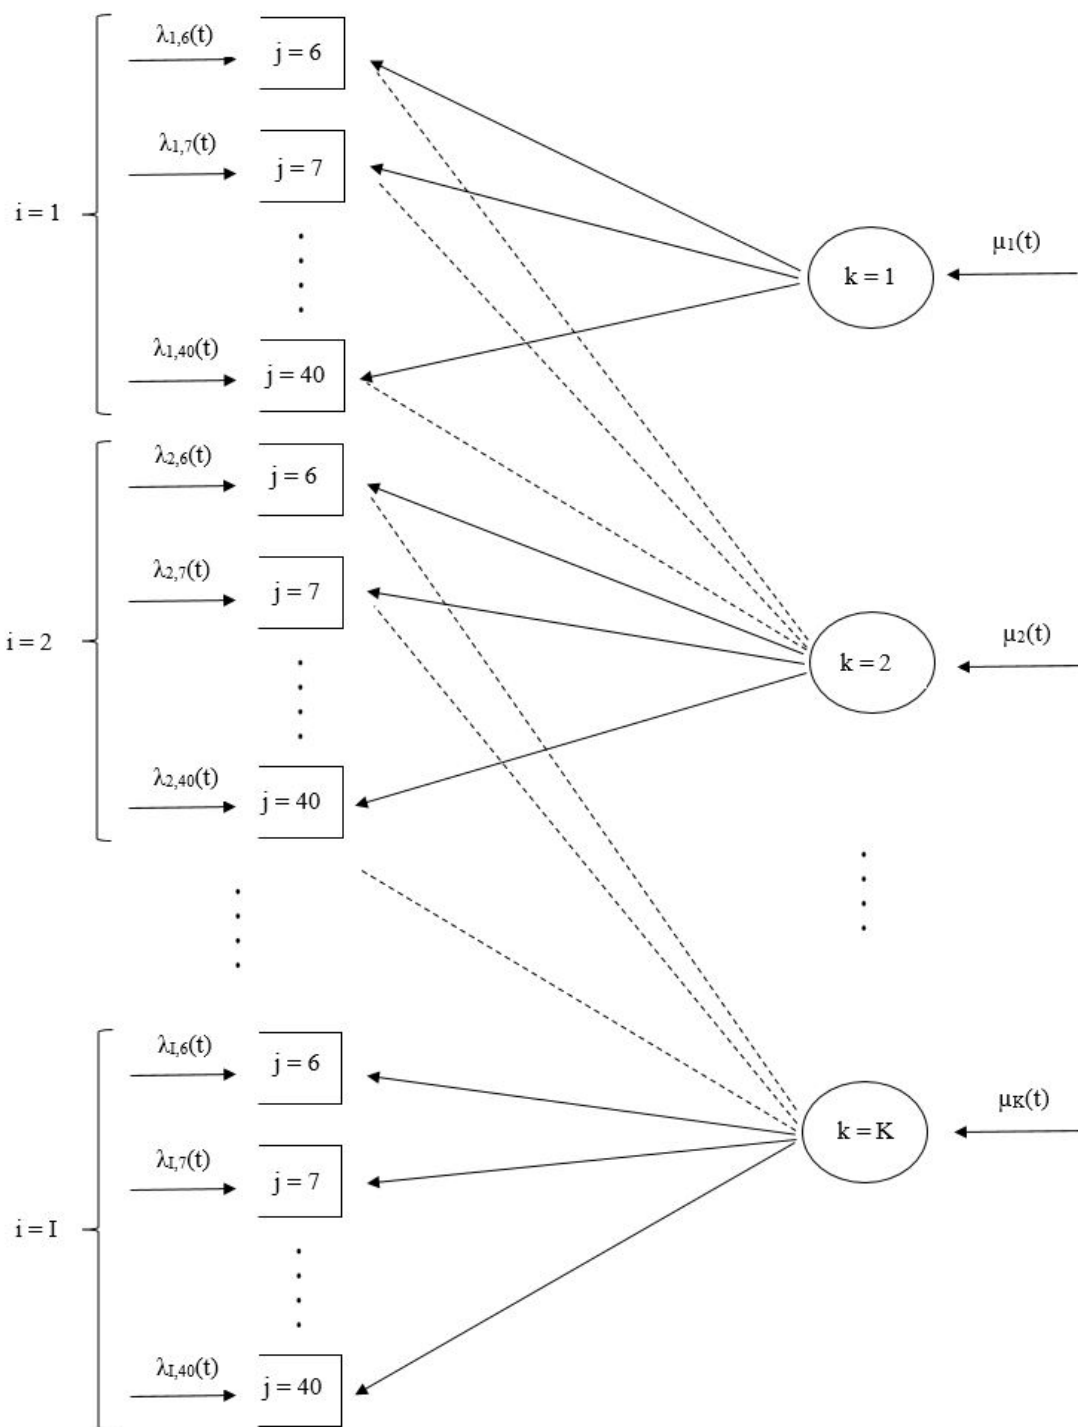
$$\begin{aligned}
& \left[\left(\bar{x}, \dot{x}, \bar{w}, \dot{w}, (e \cdot d) \cdot x \right) - \left(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}) \right) \right] \cdot \mathbf{v} \\
&= v^1 \cdot (\bar{x} - \bar{x}) + v^2 \cdot (\dot{x} - \bar{\dot{x}}) + v^3 \cdot (\bar{w} - \bar{w}) + v^4 \cdot (\dot{w} - \bar{\dot{w}}) + v^7 \cdot (-(e \cdot d) \cdot \bar{x} + (e \cdot d) \cdot \bar{x}) \\
&= v^2 \cdot (-P^{k'} u^{k'} + P^{k'} \bar{u}^{k'}) + v^4 \cdot (u^{k'} - \bar{u}^{k'}) \\
&= (P^{k'} \cdot v^2 - v^4)(\bar{u}^{k'} - u^{k'}).
\end{aligned}$$

Then, we have $(P^{k'} \cdot v^2 - v^4)(\bar{u}^{k'} - u^{k'}) \leq 0$, only if $(P^{k'} \cdot v^2 - v^4)u^{k'} \geq (P^{k'} \cdot v^2 - v^4)\bar{u}^{k'}$. From (6), since $(\bar{x}, \dot{x}, \bar{w}, \dot{w}, (e \cdot d) \cdot x)$ is an element of $\text{epi } L(t)$, this proves Property 3. Recall that the subgradient of L is related to the normal cone of its epigraph as $\partial L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}) = \left\{ (v^1, v^2, v^3, v^4) \in \mathbb{R}^{4IJ} : (v^1, v^2, v^3, v^4, -1) \in N_{\text{epi } L(t)}(\bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}}, L(t, \bar{x}, \bar{\dot{x}}, \bar{w}, \bar{\dot{w}})) \right\}$. The coextremality conditions state that for all $t \in [0, T]$, $(\dot{y}(t), y(t), \dot{z}(t), z(t)) \in \partial L(t, x(t), \dot{x}(t), w(t), \dot{w}(t))$. That is, for $t \in [0, T]$, $(\dot{y}(t), y(t), \dot{z}(t), z(t), -1) \in N_{\text{epi } L(t)}(x(t), \dot{x}(t), w(t), \dot{w}(t), L(t, x(t), \dot{x}(t), w(t), \dot{w}(t)))$. This implies that $u^k(t) \in \arg \min_{v \in \Phi(t)} \{(y(t) \cdot P^k - z(t))v\}$, which establishes the coextremality condition. This concludes the proof of Theorem 1.

A.4 Dynamics of the Liver Allocation System


The diagram of the liver allocation system is presented in Figure A.1. Figure A.2 shows the dynamic changes in class ij patients' health condition that is captured by their laboratory MELD score.

Figure A.1: Diagram of the Liver Allocation System



Notes. Static patient classes are denoted by $i \in \{1, 2, \dots, I\}$ and dynamic patient classes are denoted by $j \in \{6, \dots, 40\}$ corresponding to transplant candidates' laboratory MELD scores. The classes of deceased donor livers are denoted by $k \in \{1, 2, \dots, K\}$. The solid lines represent the identical donor-recipient matches and the dashed lines represent the other compatible matches.

Figure A.2: Dynamics of class ij patients



Notes. $\lambda_{ij}(t)$ denotes the arrival rate of class ij patients. d_j denotes the mortality rate of patients with MELD score j . $\alpha_{j'j}$ denotes patients' health transition rate from MELD score j' to j , and $\alpha_{jj'}$ denotes patients' health transition rate from MELD score j to j' .

A.5 Sized-Based Compatibility Analysis

As we discussed in §1.5, we created a compatibility matrix for each donor-recipient height class pair using the BSA analysis. The compatibility percentages can be seen in Table A.3.


Table A.3: Size-Based Compatibility Matrix (%)

		Recipient Height (cm)				
		≤ 150 cm	151-156 cm	157-165 cm	166-175 cm	≥ 176 cm
Donor Height (cm)	≤ 150 cm	97.8	97.4	97.4	93.2	67.2
	151-156 cm	98.2	99.2	98.4	98.4	83.4
	157-165 cm	93.6	1	1	1	97.8
	166-175 cm	85.9	99.3	1	1	1
	≥ 176 cm	40.5	90.7	97.3	1	1

Table A.3 shows that same height tier donor-recipient pairs are the most compatible as expected, i.e., the diagonal entries of the matrix are either or very close to one. The least compatible donor-recipient pairs are donors (recipients) with ≤ 150 cm and recipients (donors) with ≥ 176 cm. The remaining donor-recipient pairs show a high percentage of size-based compatibility ($\geq 80\%$).


BSA histograms and quantiles of each donor and recipient height tier can be found in figures below.

Figure A.3: BSA Histogram of ≤ 150 cm Donors




Quantiles (Definition 5)	
Level	Quantile
100% Max	2.101225
99%	1.886366
95%	1.651601
90%	1.610811
75% Q3	1.330189
50% Median	0.769907
25% Q1	0.512207
10%	0.361156
5%	0.276027
1%	0.253238
0% Min	0.223244

Figure A.4: BSA Histogram of 151-156 cm Donors




Quantiles (Definition 5)	
Level	Quantile
100% Max	2.64420
99%	2.34122
95%	2.00225
90%	1.86556
75% Q3	1.76759
50% Median	1.63977
25% Q1	1.49726
10%	1.40761
5%	1.37549
1%	1.22384
0% Min	1.10778

Figure A.5: BSA Histogram of 157-165 cm Donors



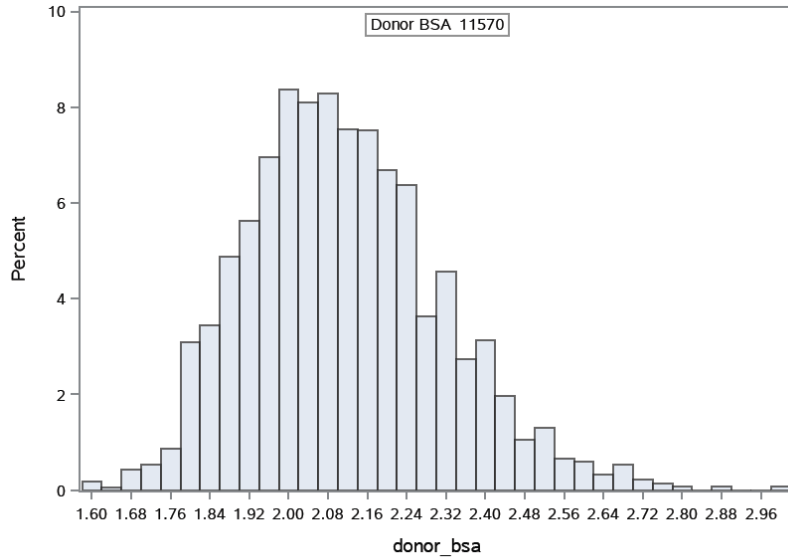
Quantiles (Definition 5)	
Level	Quantile
100% Max	2.33223
99%	2.22062
95%	2.12978
90%	2.01279
75% Q3	1.86448
50% Median	1.73205
25% Q1	1.61262
10%	1.52024
5%	1.47940
1%	1.37796
0% Min	1.30419

Figure A.6: BSA Histogram of 166-175 cm Donors




Quantiles (Definition 5)	
Level	Quantile
100% Max	2.64543
99%	2.35807
95%	2.19583
90%	2.10841
75% Q3	1.97668
50% Median	1.85206
25% Q1	1.75035
10%	1.67035
5%	1.61266
1%	1.53948
0% Min	1.39462

Figure A.7: BSA Histogram of ≥ 176 cm Donors




Quantiles (Definition 5)	
Level	Quantile
100% Max	3.00610
99%	2.67544
95%	2.46471
90%	2.38225
75% Q3	2.23854
50% Median	2.09470
25% Q1	1.97710
10%	1.86992
5%	1.81838
1%	1.72086
0% Min	1.58819

Figure A.8: BSA Histogram of ≤ 150 cm Patients




Quantiles (Definition 5)	
Level	Quantile
100% Max	2.150099
99%	1.812276
95%	1.686030
90%	1.607100
75% Q3	1.367628
50% Median	0.649442
25% Q1	0.376086
10%	0.300194
5%	0.273745
1%	0.229590
0% Min	0.159663

Figure A.9: BSA Histogram of 151-156 cm Patients



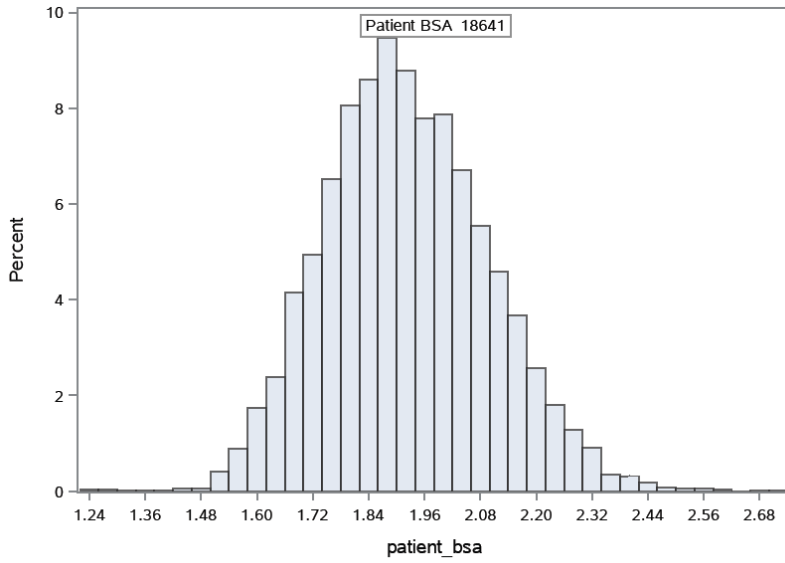
Quantiles (Definition 5)	
Level	Quantile
100% Max	2.31804
99%	2.03458
95%	1.92866
90%	1.87429
75% Q3	1.76302
50% Median	1.63632
25% Q1	1.53425
10%	1.45877
5%	1.40878
1%	1.32634
0% Min	1.03366

Figure A.10: BSA Histogram of 157-165 cm Patients



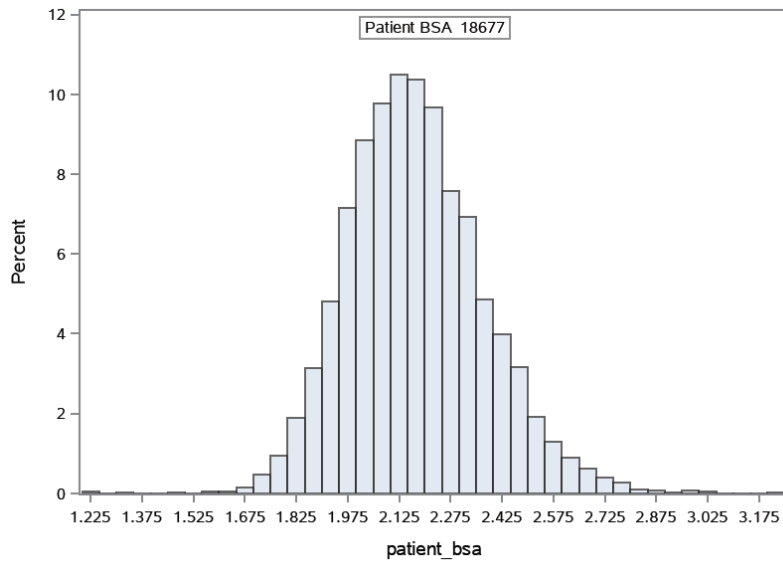
Quantiles (Definition 5)	
Level	Quantile
100% Max	2.57888
99%	2.20805
95%	2.05249
90%	1.98475
75% Q3	1.86796
50% Median	1.74776
25% Q1	1.63664
10%	1.55089
5%	1.49967
1%	1.41770
0% Min	1.16642

Figure A.11: BSA Histogram of 166-175 cm Patients



Quantiles (Definition 5)	
Level	Quantile
100% Max	2.72365
99%	2.34734
95%	2.22297
90%	2.15373
75% Q3	2.03878
50% Median	1.91099
25% Q1	1.79946
10%	1.70303
5%	1.64907
1%	1.55777
0% Min	1.25442

Figure A.12: BSA Histogram of ≥ 176 cm Patients



Quantiles (Definition 5)	
Level	Quantile
100% Max	3.24683
99%	2.69605
95%	2.51241
90%	2.43060
75% Q3	2.29516
50% Median	2.16134
25% Q1	2.03710
10%	1.93774
5%	1.87478
1%	1.76693
0% Min	1.22324

Appendix B

Additional Material for Chapter 2

B.1 Description of User Attributes on the Platform

User Attribute	Type	Additional Details
Age	Numerical	Age
Average Attractiveness	Numerical	Average attractiveness score received from other users
Height	Numerical	Height in centimeters
Num. Photos	Numerical	Number of profile photos that show a user's face
Num. Photobooks	Numerical	Number of non-profile photos included in a photobook
Profile Completion Score	Numerical	Score indicating the individual's degree of profile completion
Platform Tenure	Numerical	Time since a user joined the platform in days
Drinking	Ordinal	Frequency of drinking
Smoking	Ordinal	Frequency of smoking
Weight	Ordinal	Self-reported categories of body shape: skinny, little skinny, slim & sporty, average, muscular, little fat, fat
Gender	Binary	Male, Female
College	Binary	Whether or not user has filled out college info
Facebook Connected	Binary	Whether or not connected with Facebook
iPhone	Binary	Whether or not using iPhone when signing up (versus Android Phone or web)
Verified	Binary	Whether or not a real name is verified
Filtering/Search	Binary	Whether or not a user stated search filtering preferences for a partner's age, drinking, height, weight, religion, and smoking
Blood Type	Categorical	One of: O, A, B, AB, missing
Religion	Categorical	One of 13 religious designations
Region	Categorical	One of 8 provinces in South Korea

Notes. “Ordinal” attributes are categorical and ordered, whereas “categorical” attributes are categorical and unordered. Correlation matrices for user attributes can be found in Appendix B.3.

B.2 Summary Statistics of User Attributes

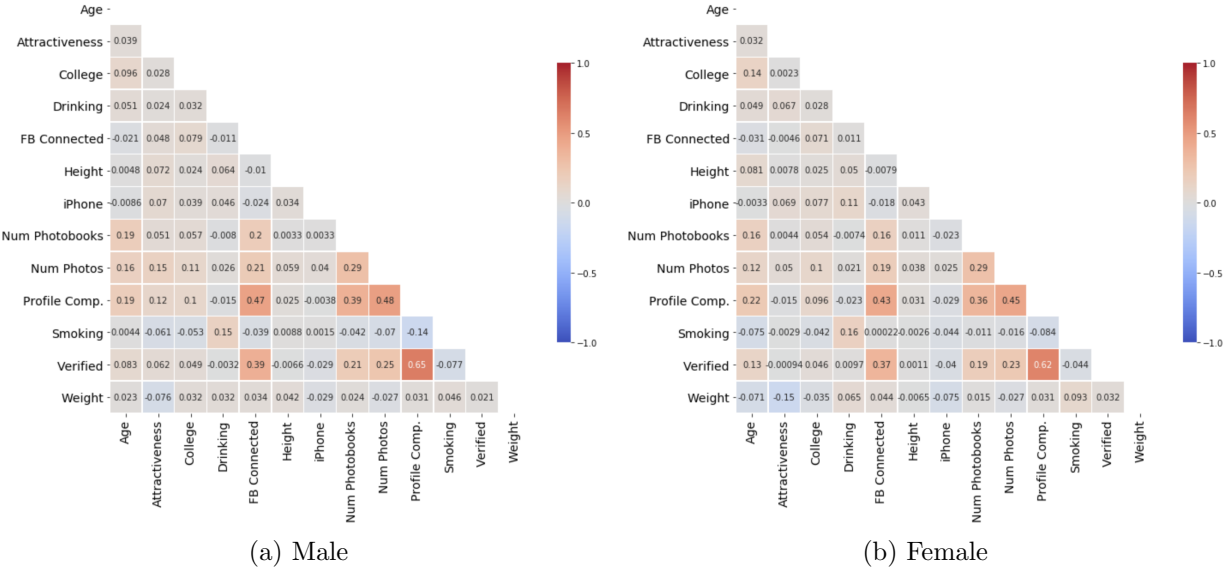
Type	Male							Female						
	Mean	St.Dev.	Min	1Q	2Q	3Q	Max	Mean	St.Dev.	Min	1Q	2Q	3Q	Max
Numerical														
Age	27.55	5.99	15.0	23.2	26.7	31.2	67.0	25.1	5.5	15.0	21.0	24.1	28.1	66.9
Attractiveness	1.93	0.87	0	1.66	2	2.46	5	3.69	0.99	0	3.50	3.90	4.20	5
Height (cm)	175.60	5.45	0	172	175	179	220	162.17	5.38	0	159	162	165	220
Num. Photobooks	1.19	2.88	0	0	0	1	283	0.84	2.59	0	0	0	1	223
Num. Photos	3.33	1.37	1	2	3	4	7	3.09	1.32	1	2	3	4	7
Platform Tenure	378.25	507.61	1	59	161	412	2165	215.90	299.75	0	51	132	219	2158
Profile Completion Score	60.37	24.33	10.5	36.8	59.6	79.5	123.8	52.4	22.5	13.5	32.5	45.9	69.6	123.8
Ordinal	Mean	St.Dev.	Min	1Q	2Q	3Q	Max	Mean	St.Dev.	Min	1Q	2Q	3Q	Max
Drinking	3.18	0.86	0	3	3	4	6	2.98	1.00	0	2	3	4	6
Smoking	1.31	0.48	0	1	1	2	2	1.10	0.31	0	1	1	1	2
Weight	3.85	1.46	0	3	3	5	7	4.12	1.79	0	3	3	7	7
Binary	Mean							Mean						
College	0.59							0.47						
Facebook Connected	0.31							0.23						
iPhone	0.35							0.45						
Verified	0.51							0.31						
Categorical			C1	C2	C3	C4	Rest			C1	C2	C3	C4	Rest
Blood Type			33.7	27.7	27.2	11.3	0.1			33.7	27.8	27.1	11.3	0.1
Religion			72.7	15.7	5.9	5.4	0.3			65.0	22.0	7.1	5.5	0.4
Region			27.0	7.4	6.3	5.9	53.3			33.8	7.5	6.7	5.7	46.3

Notes. For numerical and ordinal (i.e., ordered categorical) attributes, **1Q**, **2Q**, and **3Q** refer to the values of the first, second, and third quantiles, respectively. For categorical (i.e., unordered categorical) attributes, **C1** to **C4** refer to the percentage of users that make up the first to fourth most common categories, and **Rest** refers to the percentage making up the remainder.

B.3 Correlation Matrices for User Attributes

The correlation matrices for user attributes, divided by gender, are presented in Figure B.1.

Figure B.1: Correlation Matrices for User Attributes



B.4 Additional Discussion of Platform’s Recommendations and Regression Analysis on User Behavior

Our theoretical model assumes that a user’s probability of accepting messages on the platform decreases as the number of messages they receive increases. To verify this assumption, we run a regression model analyzing how the number of messages users receive affects their message acceptance rate. To eliminate sparsity in this analysis, we consider the subset of users who received and accepted at least one message. We compute the dependent variable, users’ message acceptance rate, as the ratio of the number of messages users accept out of the number of messages they receive from other users on the platform. We use a similar set of covariates capturing individual user characteristics as we used previously in Table 2.2. We exclude the variable *Gender* because of the multicollinearity issue: it has a variance inflation factor (VIF) greater than 5. This is rather expected because female users receive a significantly higher number of messages than male users, i.e., a male user receives 0.03 messages on average, whereas a female user receives 10.9.

Table B.1 presents the regression result. Our results indicate that the number of messages received has a statistically significant and negative coefficient, validating our assumption in the

Table B.1: Regression Result on Users' Message Acceptance Rate

	Accept Rate
(Intercept)	1.213*** (0.0284)
Message Received	-0.0023*** (0.0001)
Average Attractiveness	-0.3094*** (0.0017)
Verified	-0.0162*** (0.0042)
Age	-0.0037*** (0.0004)
Platform Tenure	-0.0037 (0.0017)
Region	0.0017*** (0.0002)
Religion	-0.0050 (0.0019)
Smoking	-0.0044*** (0.0007)


Notes. Dependent variable: Accept rate, Number of observations = 45907, Standard errors are in parentheses. *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

theoretical model. In addition, the coefficient of the variable *Average Attractiveness* indicates that higher average attractiveness decreases the message acceptance rate of users, demonstrating that popular users are more selective than unpopular users in accepting messages on the platform.

Figure B.2 shows the percentage of messages sent with respect to the average attractiveness difference between the focal user (e.g., a male user) and the user on the other side of the interaction (e.g., a female user). As we discussed earlier, when users receive only a few recommendations from the platform and there is a cost associated with showing interest, they take the probability of their interest being reciprocated into account while sending messages. Even though it is difficult to validate this assumption using the interaction data, we provide some descriptive evidence supporting the idea that users consider both their own attractiveness score and recommended users' attractiveness while sending messages.

To better explain the graph presented in Figure B.2, we consider the following example: Bob (with an average attractiveness score of 2.3) receives Alice's profile as a recommendation from the platform. Alice's average attractiveness score is 3.8. The difference between Alice's and Bob's average attractiveness scores is 1.5 ($= 3.8 - 2.3$). Because the average attractiveness scores of users range between 1 and 5, attractiveness difference can take values between -4 and 4. Negative (positive) values correspond to the cases where users with higher (lower) attractiveness scores send

Figure B.2: Percentage of Messages Sent over Users' Attractiveness Difference



messages to users with lower (higher) attractiveness scores. As expected, we find that users are very unlikely to send messages to others with lower attractiveness scores. In Figure B.2, the percentage of messages sent starts increasing when the attractiveness difference becomes positive, implying that users tend to initiate contact with users who are, on average, more attractive than them. However, this trend changes when the attractiveness score difference between users becomes considerably high (i.e., over 2). This pattern suggests that users are reluctant to send messages to users who are too attractive for them, i.e., users who are “out of their reach.” One potential explanation for this behavior is that users do not expect that very attractive users would reciprocate their interest, so they choose not to send messages to them.

B.5 Summary of Notation

Symbol	Definition
i, j	Index for users on the platform
k	Index for men, high-tier women, and low-tier women ($\in \{m, h, l\}$)
N^k	Number of k -tier users
n	Number of users recommended to each user
x_k	Proportion of women in tier k in men's recommendations
r	Cost of each message to users ($\in [0, 1]$)
v_{ij}	Match value that user i receives from matching with user j
q_i	Average attractiveness of the tier user i belongs to ($\in \{q, 0\}$)
u_{ij}	Idiosyncratic utility that user i receives from matching with user j
θ_i	User i 's utility of remaining single
γ_i	User i 's bar of accepting messages in the third stage
α	The magnitude of high-tier women's degree of selectivity ($\in [0, 1]$)
β	The magnitude of the cost of evaluating incoming messages ($\in [0, 1]$)
$p_{k'}^k$	Probability of user in tier k sending a message to a user in tier k' with the opposite gender
$a_{k'}^k$	Probability of user in tier k accepting a message from a user in tier k' with the opposite gender

B.6 Proofs

Proof of Lemma 2.1. We compute the probability that the match value user j receives from matching with user i , v_{ji} , exceeds the sum of user j 's utility from remaining single, θ_j , and the congestion effect that user j incurs from other users' messages in the third stage of the game, γ_j .

We first compute the probability of user j accepting user i 's message conditional on user j 's utility from remaining single, θ_j : $\mathbb{P}(v_{ji} \geq \theta_j + \gamma_j | \theta_j) = \mathbb{P}(q_i + u_{ji} \geq \theta_j + \gamma_j | \theta_j) = \mathbb{P}(u_{ji} \geq \theta_j + \gamma_j - q_i | \theta_j) = 1 - \theta_j - \gamma_j + q_i$ because $v_{ji} = q_i + u_{ji}$ and $u_{ji} \sim U[0, 1]$. We then compute the unconditional probability as $\int_{\alpha q_j}^1 (1 - \theta_j - \gamma_j + q_i) \frac{1}{1 - \alpha q_j} d\theta_j = (\theta_j - \frac{1}{2}\theta_j^2 - \theta_j\gamma_j + \theta_j q_i) \frac{1}{1 - \alpha q_j} \Big|_{\alpha q_j}^1 = 1 - \frac{1 + \alpha q_j}{2} - \beta x_j p_j^i + q_i$ because $\gamma_j = \beta x_j p_j^i$ and $\theta_j \sim U[\alpha q_j, 1]$. The probability of accepting messages should be nonnegative, so we rearrange the terms to obtain $a_j^i = \max \left\{ 0, \frac{1 - \alpha q_j - 2\beta x_j p_j^i + 2q_i}{2} \right\}$. We

obtain $a_m^k = \max \left\{ 0, \frac{1 - \alpha q_j - 2\beta x_j p_j^i}{2} \right\}$ because $q_i = 0$ for male users. Similarly, we obtain $a_k^m = \max \left\{ 0, \frac{1 + 2q_k - \beta(p_m^h + p_m^l)}{2} \right\}$. \square

Proof of Lemma 2.2. We compute the probability that the expected match value user i receives

from matching with user j , $\mathbb{E}[v_{ij}]$, exceeds user i 's utility from remaining single and the cost of sending a message, r . We compute the probability of a man i sending a message to a high-tier woman j , p_h^m , and derive the rest of the probabilities similarly. For man i and high-tier woman j , $\mathbb{E}[v_{ij}] = \mathbb{E}[a_m^h(q + u_{ij})] = a_m^h q + a_m^h \mathbb{E}[u_{ij}] = a_m^h(q + 1/2)$ since $u_{ij} \sim U[0, 1]$. Therefore, $p_h^m = \mathbb{P}(a_m^h(q + 1/2) \geq \theta_i + r) = \mathbb{P}(\theta_i \leq a_m^h(q + 1/2) - r) = a_m^h(q + 1/2) - r$ because $\theta_i \sim U[0, 1]$ for man i . $a_m^h = \frac{1 - \alpha q - 2\beta x_h p_h^m}{2}$ follows from Lemma 2.1. Rearranging the terms gives $p_h^m = \frac{1 - \alpha q - 2\beta x_h p_h^m}{2} \left(q + \frac{1}{2}\right) - r$. Solving this equation for p_h^m gives $p_h^m = \frac{(2q + 1)(1 - \alpha q) - 4r}{2(2q + 1)\beta x_h + 4}$. Similarly, we obtain $p_l^m = \frac{1 - 4r}{2(2q + 1)\beta x_l + 4}$. \square

Proof of Lemma 2.3. To characterize the revenue-maximizing platform's optimal recommendation ratios, we solve for the first order conditions of the revenue maximization problem denoted by R . Substituting $x_l = 1 - x_h$ and taking the first order condition of the objective function gives

$$\begin{aligned} \frac{dR}{dx_h} &= \frac{(2q + 1)(1 - \alpha q) - 4r}{2(2q + 1)\beta x_h + 4} - x_h \frac{2(2q + 1)\beta((2q + 1)(1 - \alpha q) - 4r)}{(2(2q + 1)\beta x_h + 4)^2} \\ &\quad - \frac{1 - 4r}{2(2q + 1)\beta(1 - x_h) + 4} + (1 - x_h) \frac{2(2q + 1)\beta(1 - 4r)}{(2(2q + 1)\beta(1 - x_h) + 4)^2} = 0. \end{aligned}$$

To simplify the algebra, we denote $h_1 = 2(2q + 1)\beta$, $h_2 = (2q + 1)(1 - \alpha q) - 4r$ and $h_3 = 1 - 4r$. The first order condition becomes $\frac{dR}{dx_h} = \frac{h_2}{h_1 x_h + 4} - x_h \frac{h_1 h_2}{(h_1 x_h + 4)^2} - \frac{h_3}{h_1(1 - x_h) + 4} + (1 - x_h) \frac{h_1 h_3}{(h_1(1 - x_h) + 4)^2} = 0$. Rearranging the terms gives us

$$\frac{h_2}{(h_1 x_h + 4)^2} = \frac{h_3}{(h_1(1 - x_h) + 4)^2}.$$

Solving this equation for x_h gives $x_h^* = \frac{h_1 \sqrt{h_2} + 4(\sqrt{h_2} - \sqrt{h_3})}{h_1 \sqrt{h_2} + h_1 \sqrt{h_3}}$, which can be further simplified to x_h^* given in the proposition.

To verify whether x_h^* is indeed the maximizer of the expected revenue, we check the second order condition of the objective function:

$$\frac{d^2 R}{dx_h^2} = -2h_1 h_2 (h_1 x_h + 4)^{-2} + 2h_1^2 h_2 x_h (h_1 x_h + 4)^{-3} - 2h_1 h_3 (h_1(1 - x_h) + 4)^{-2} + 2h_1^2 h_3 (1 - x_h)(h_1(1 - x_h) + 4)^{-3}$$

where the sum of the first two terms is negative because $h_1 x_h (h_1 x_h + 4)^{-1} < 1 \Rightarrow 2h_1^2 h_2 x_h (h_1 x_h + 4)^{-1} < 2h_1 h_2 \Rightarrow 2h_1^2 h_2 x_h (h_1 x_h + 4)^{-3} < 2h_1 h_2 (h_1 x_h + 4)^{-2} \Rightarrow -2h_1 h_2 (h_1 x_h + 4)^{-2} + 2h_1^2 h_2 x_h (h_1 x_h + 4)^{-3} < 0$. Similarly, $h_1(1 - x_h)(h_1(1 - x_h) + 4)^{-1} < 1 \Rightarrow 2h_1^2 h_3 (1 - x_h)(h_1(1 - x_h) + 4)^{-1} < 2h_1 h_3 \Rightarrow 2h_1^2 h_3 (1 - x_h)(h_1(1 - x_h) + 4)^{-3} < 2h_1 h_3 (h_1(1 - x_h) + 4)^{-2} \Rightarrow -2h_1 h_3 (h_1(1 - x_h) + 4)^{-2} + 2h_1^2 h_3 (1 -$

$x_h)(h_1(1 - x_h) + 4)^{-3} < 0$. Therefore, $d^2R/dx_h^2 < 0$ ensures that x_h^* is indeed the maximizer of the objective function. \square

Proof of Proposition 2.1. (a) In order to be unbiased and given an equal number of high-tier and low-tier women on the platform, the revenue-maximizing platform should recommend both tiers equally, i.e., x_h and x_l should be $1/2$. We know from Lemma 2.3 that the revenue-maximizing platform's recommendation ratios are $x_h^* = \frac{h_1\sqrt{h_2} + 4(\sqrt{h_2} - \sqrt{h_3})}{h_1\sqrt{h_2} + h_1\sqrt{h_3}}$ and $x_l^* = 1 - x_h^*$. Therefore, unbiased recommendations satisfy the following condition:

$$x_h^* = \frac{h_1\sqrt{h_2} + 4(\sqrt{h_2} - \sqrt{h_3})}{h_1\sqrt{h_2} + h_1\sqrt{h_3}} = \frac{1}{2}.$$

This equality gives us $h_2 = h_3$ where $h_2 = (2q + 1)(1 - \alpha q) - 4r$ and $h_3 = 1 - 4r$. Solving this equality gives us two roots: $q = 0$ and $\alpha = 2/(2q + 1)$. The first root implies that the platform uses unbiased recommendations when there is no difference between two tiers in terms of quality, so that there is only one tier of women on the platform. This is trivial and violates our assumption that $q > 0$. The second root gives us the condition given in Proposition 2.1.

Finally, we show that $h_2 = h_3$ is a necessary and sufficient condition for $x_h^* = 1/2$. Suppose $h_2 > h_3$. Then, we have $(h_1 + 8)(\sqrt{h_2} - \sqrt{h_3}) > 0$ since $h_1 > 0$. $(h_1 + 8)(\sqrt{h_2} - \sqrt{h_3}) > 0 \Rightarrow h_1(\sqrt{h_2} - \sqrt{h_3}) + 8(\sqrt{h_2} - \sqrt{h_3}) > 0 \Rightarrow h_1\sqrt{h_2} + 8(\sqrt{h_2} - \sqrt{h_3}) - h_1\sqrt{h_3} > 0 \Rightarrow h_1\sqrt{h_2} + 8(\sqrt{h_2} - \sqrt{h_3}) > h_1\sqrt{h_3} \Rightarrow 2h_1\sqrt{h_2} + 8(\sqrt{h_2} - \sqrt{h_3}) > h_1\sqrt{h_2} + h_1\sqrt{h_3} \Rightarrow 2(h_1\sqrt{h_2} + 4(\sqrt{h_2} - \sqrt{h_3})) > h_1\sqrt{h_2} + h_1\sqrt{h_3} \Rightarrow \frac{h_1\sqrt{h_2} + 4(\sqrt{h_2} - \sqrt{h_3})}{h_1\sqrt{h_2} + h_1\sqrt{h_3}} > \frac{1}{2}$ for $h_1 > 0$ and $h_2 > h_3$. Therefore, in this case $x_h > 1/2$. Similarly, we can show that if $h_2 < h_3$, then $x_h < 1/2$. Therefore, h_2 must be equal to h_3 to obtain $x_h^* = 1/2$.

(b) First, we look at the case where the revenue-maximizing platform is biased against low-tier women, $\alpha < 2/(2q + 1)$. This gives $h_2 > h_3$ and $x_h^* = \frac{\sqrt{h_2}}{\sqrt{h_2} + \sqrt{h_3}} + \frac{4(\sqrt{h_2} - \sqrt{h_3})}{h_1(\sqrt{h_2} + \sqrt{h_3})} > \frac{1}{2}$ where $h_1 = 2(2q + 1)\beta$, $h_2 = (2q + 1)(1 - \alpha q) - 4r$, and $h_3 = 1 - 4r$. As β increases, h_1 increases and x_h^* decreases. Since $x_h^* > 1/2$ in this case, this reduces the bias on the revenue-maximizing platform. Second, we look at the case where the revenue-maximizing platform is biased against high-tier women, $\alpha > 2/(2q + 1)$. This gives $h_2 < h_3$ and $x_h^* = \frac{\sqrt{h_2}}{\sqrt{h_2} + \sqrt{h_3}} + \frac{4(\sqrt{h_2} - \sqrt{h_3})}{h_1(\sqrt{h_2} + \sqrt{h_3})} < \frac{1}{2}$. As β increases, h_1 increases, so $\frac{4(\sqrt{h_2} - \sqrt{h_3})}{h_1(\sqrt{h_2} + \sqrt{h_3})} (< 0)$, hence x_h^* , increases. Since $x_h^* < 1/2$ in this case, this reduces the bias on the revenue-maximizing platform. \square

Proof of Proposition 2.2. Formulating the Lagrangian of the match-maximizing platform's optimization problem (M) gives $\mathcal{L}(x_h, x_l, \lambda) = x_h p_h^m a_m^h + x_l p_l^m a_m^l + \lambda(1 - x_h - x_l)$. We derive the

first order conditions:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x_h} &= p_h^m a_m^h + x_h \frac{\partial p_h^m}{\partial x_h} a_m^h + x_h p_h^m \frac{\partial a_m^h}{\partial x_h} - \lambda = 0, \\ \frac{\partial \mathcal{L}}{\partial x_l} &= p_l^m a_m^l + x_l \frac{\partial p_l^m}{\partial x_l} a_m^l + x_l p_l^m \frac{\partial a_m^l}{\partial x_l} - \lambda = 0, \text{ and} \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= 1 - x_h - x_l = 0,\end{aligned}$$

where $\frac{\partial p_h^m}{\partial x_h} = -\frac{h_1 h_2}{(h_1 x_h + 4)^2}$, $\frac{\partial p_l^m}{\partial x_l} = -\frac{h_1 h_3}{(h_1 x_l + 4)^2}$, $\frac{\partial a_m^h}{\partial x_h} = -\frac{4\beta h_2}{(h_1 x_h + 4)^2}$, and $\frac{\partial a_m^l}{\partial x_l} = -\frac{4\beta h_3}{(h_1 x_l + 4)^2}$ for $h_1 = 2(2q+1)\beta$, $h_2 = (2q+1)(1-\alpha q) - 4r$, and $h_3 = 1 - 4r$. Rewriting the first order conditions gives

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x_h} &= \frac{h_2(1-\alpha q)}{2(h_1 x_h + 4)} - \frac{(2\beta h_2^2 + (1-\alpha q)h_1 h_2)x_h}{2(h_1 x_h + 4)^2} + \frac{\beta h_1 h_2^2 x_h^2 - 4\beta h_2^2 x_h}{(h_1 x_h + 4)^3} - \lambda = 0, \\ \frac{\partial \mathcal{L}}{\partial x_l} &= \frac{h_3}{2(h_1 x_l + 4)} - \frac{(2\beta h_3^2 + h_1 h_3)x_l}{2(h_1 x_l + 4)^2} + \frac{\beta h_1 h_3^2 x_l^2 - 4\beta h_3^2 x_l}{(h_1 x_l + 4)^3} - \lambda = 0, \text{ and} \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= 1 - x_h - x_l = 0.\end{aligned}$$

At $x_h = x_h^*$ and $x_l = x_l^*$, we have $\partial \mathcal{L}/\partial x_h = \partial \mathcal{L}/\partial x_l$, which can be written as:

$$\begin{aligned}\frac{h_2(1-\alpha q)}{2(h_1 x_h^* + 4)} - \frac{(2\beta h_2^2 + (1-\alpha q)h_1 h_2)x_h^*}{2(h_1 x_h^* + 4)^2} + \frac{\beta h_1 h_2^2 x_h^{*2} - 4\beta h_2^2 x_h^*}{(h_1 x_h^* + 4)^3} &= \frac{h_3}{2(h_1 x_l^* + 4)} - \frac{(2\beta h_3^2 + h_1 h_3)x_l^*}{2(h_1 x_l^* + 4)^2} \\ &+ \frac{\beta h_1 h_3^2 x_l^{*2} - 4\beta h_3^2 x_l^*}{(h_1 x_l^* + 4)^3}.\end{aligned}\quad (\text{B.1})$$

To derive the condition under which the match-maximizing platform is unbiased, we examine the first order conditions of the match-maximizing platform's problem under which $\partial \mathcal{L}/\partial x_h = \partial \mathcal{L}/\partial x_l$ for $x_h^* = x_l^* = 1/2$. Substituting $x_h^* = x_l^* = 1/2$, we obtain

$$\begin{aligned}\frac{h_2(1-\alpha q)}{2(h_1/2 + 4)} - \frac{(2\beta h_2^2 + (1-\alpha q)h_1 h_2)/2}{2(h_1/2 + 4)^2} + \frac{\beta h_1 h_2^2/4 - 2\beta h_2^2}{(h_1/2 + 4)^3} &= \frac{h_3}{2(h_1/2 + 4)} - \frac{(2\beta h_3^2 + h_1 h_3)/2}{2(h_1/2 + 4)^2} \\ &+ \frac{\beta h_1 h_3^2/4 - 2\beta h_3^2}{(h_1/2 + 4)^3}.\end{aligned}\quad (\text{B.2})$$

After simplification, we obtain $(h_1 + 8)(h_2(1-\alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) = 0$.

Suppose the left-hand side of (B.6) is greater than the right-hand side at $x_h = x_l = 1/2$, i.e., $(h_1 + 8)(h_2(1-\alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) > 0$. In this case, we want to show that $x_h^* > 1/2$ and $x_l^* < 1/2$. To this end, we first show that $\partial^2 \mathcal{L}/\partial x_h^2 < 0$ for all x_h in the following:

$$\begin{aligned}
\frac{\partial^2 \mathcal{L}}{\partial x_h^2} &= -\frac{h_1 h_2 (1 - \alpha q)}{2(h_1 x_h + 4)^2} - \frac{2\beta h_2^2 + (1 - \alpha q) h_1 h_2}{2(h_1 x_h + 4)^2} + \frac{(2\beta h_2^2 + (1 - \alpha q) h_1 h_2) h_1 x_h}{(h_1 x_h + 4)^3} + \frac{2\beta h_1 h_2^2 x_h - 4\beta h_2^2}{(h_1 x_h + 4)^3} \\
&\quad - \frac{3\beta h_1^2 h_2^2 x_h^2 + 12\beta h_1 h_2^2 x_h}{(h_1 x_h + 4)^4} \\
&= -\frac{8h_1 h_2 (1 - \alpha q) + 16\beta h_2^2}{(h_1 x_h + 4)^3} < 0,
\end{aligned}$$

because $\alpha q < 1$, $\beta > 0$, $h_1 > 0$ and $h_2 > 0$. Similarly, $\partial^2 \mathcal{L} / \partial x_l^2 = -(16h_1 h_3 + 64\beta h_3^2) / (h_1 x_l + 4)^3 < 0$. The discriminant $D = (\partial^2 \mathcal{L} / \partial x_h^2)(\partial^2 \mathcal{L} / \partial x_l^2) - (\partial \mathcal{L} / \partial x_h x_l)^2 > 0$ because $\partial^2 \mathcal{L} / \partial x_h^2 < 0$, $\partial^2 \mathcal{L} / \partial x_l^2 < 0$, and $\partial \mathcal{L} / \partial x_h x_l = 0$. Since \mathcal{L} is strictly concave in x_h and x_l and $D > 0$, there is a unique optimal solution (x_h^*, x_l^*) (where $x_h^* + x_l^* = 1$) that satisfies $\partial \mathcal{L} / \partial x_h = \partial \mathcal{L} / \partial x_l = 0$. Equivalently, at this optimal solution, the left-hand side of (B.5) is equal to the right-hand side of (B.5) that is equal to λ . Let us denote the left-hand side of equation (B.5) with $f(x_h^*)$ and the right-hand side with $g(x_l^*)$. We claim that if $f(x_h) - g(x_l) > 0$, i.e., $(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) > 0$ at $x_h = x_l = 1/2$, then it must be that $f(x_h) > f(x_h^*) = g(x_l^*) = \lambda > g(x_l)$. Suppose this does not hold. Then, it must be that either $\lambda > f(x_h) > g(x_l)$ or $f(x_h) > g(x_l) > \lambda$ for $x_h + x_l = 1$. If $\lambda > f(x_h) > g(x_l)$ for $x_h + x_l = 1$, then the optimal $x_h^* > x_h$ and $x_l^* > x_l$ because $\partial^2 f / \partial x_h^2 < 0$ and $\partial^2 g / \partial x_l^2 < 0$. However, $x_h^* + x_l^* > x_h + x_l = 1$, which contradicts $x_h^* + x_l^* = 1$. Similarly, if $f(x_h) > g(x_l) > \lambda$ for $x_h + x_l = 1$, then both $x_h^* < x_h$ and $x_l^* < x_l$. This contradicts $x_h^* + x_l^* = 1$ because $x_h^* + x_l^* < x_h + x_l = 1$. Therefore, we conclude that $f(x_h) > \lambda > g(x_l)$ if $f(x_h) - g(x_l) > 0$. This implies that if $(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) > 0$ at $x_h = x_l = 1/2$, $f(1/2) > \lambda > g(1/2)$. Since $\partial^2 f / \partial x_h^2 < 0$ and $\partial^2 g / \partial x_l^2 < 0$, $f(x_h^*) = g(x_l^*) = \lambda$ where $x_h^* > 1/2$ and $x_l^* < 1/2$.

Similarly, if $(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4\beta(h_2^2 - h_3^2) < 0$, then $x_h^* < 1/2$ and $x_l^* > 1/2$. \square

Proof of Lemma 2.4. First, we note that $x_{h(2)}^*$ and $x_{l(2)}^*$ are characterized by Lemma 2.3 because the platform's second-period problem is the same as our one period model for the revenue-maximizing platform. Therefore, $x_{h(2)}^*$, $x_{l(2)}^*$, $p_h^m(x_{h(2)}^*)$ and $p_l^m(x_{l(2)}^*)$ are not functions of $x_{h(1)}$ and they can be treated as constants. As a result, we obtain the sum of the third and fourth terms of the objective function, $x_{h(2)}^* N_{(2)}^m p_h^m(x_{h(2)}^*) + x_{l(2)}^* N_{(2)}^m p_l^m(x_{l(2)}^*)$, is a linear function of $N_{(2)}^m$. $N_{(2)}^m$ is a function of $x_{h(1)}$ because $N_{(2)}^m = N(1 - M) + \eta NM$ where $M = x_{h(1)} p_h^m a_m^h + x_{l(1)} p_l^m a_m^l$. This gives $N_{(2)}^m$ as a linear function of M . The match rate obtained in the first period M is the same as the match-maximizing platform's objective function in our one-period model. We have shown in Proposition 2.2 that this objective function is concave in $x_{h(1)}$; therefore, M is a concave function of $x_{h(1)}$. The sum of the first two terms of the objective function, $x_{h(1)} N p_h^m(x_{h(1)}) + x_{l(1)} N p_l^m(x_{l(1)})$, is the same as the revenue-maximizing platform's problem in our one-period model. We have shown

in Lemma 2.3 that this function is concave in $x_{h(1)}$. The sum of two concave functions is concave; therefore, the objective function of R_1 is concave in $x_{h(1)}$. Similarly, we have shown in Lemma 2.3 and in Proposition 2.2 that the revenue-maximization problem (R) and match maximization problem (M) have unique optimal solutions for $x_h + x_l = 1$. Therefore, the sum of these two functions, R_1 , has a unique optimal solution such that $x_{h(1)}^* + x_{l(1)}^* = 1$. \square

B.7 Technical Details of Our Predictive Model

In this appendix, we provide technical details of our predictive model. We encode the data from each interaction type (open, send and accept) as a separate matrix, with male users on one axis and female users on the other axis. To introduce some notation, let us label male users as $i = 1, 2, \dots, m_1$ and female users as $j = 1, 2, \dots, m_2$ ($m_1 = 198,895$ and $m_2 = 43,907$ in our data). We also label the different interaction types as $k = 1, 2, 3$, where $k = 1$ denotes open, $k = 2$ denotes send, and $k = 3$ denotes accept. For each interaction type k , our interaction data is encoded in a matrix X^k which has m_1 rows and m_2 columns. Each entry of these matrices is then either unobserved (e.g., Bob did not receive Alice’s card) or takes on a binary value, with 1 indicating that the interaction occurred (e.g., Bob received Alice’s card and opened it) and 0 indicating that it did not occur (e.g., Bob received Alice’s card but did not open it).

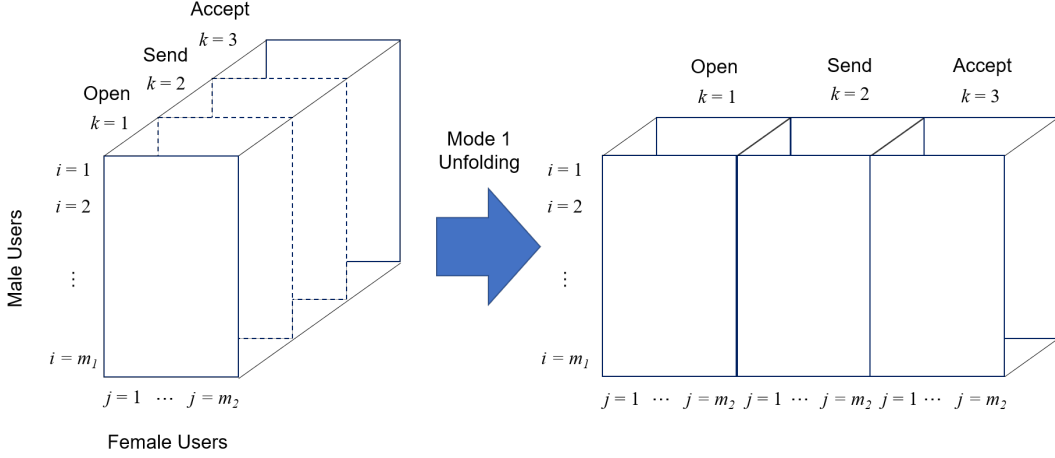
Our estimation procedure operates on three matrices (X^1, X^2 and X^3) jointly as a single tensor. A tensor is a higher-dimensional analogue of a matrix; for example, if we view a matrix as representing data along two dimensions (its rows and columns), then it is equally natural to represent data along more than two dimensions in suitable settings. In our setting, the interaction data is naturally represented as a three-dimensional tensor by stacking the matrices X^1, X^2 , and X^3 (see Figure B.3, left); we let X denote this tensor. We then construct mode-1 and mode-2 unfolding of X , i.e., flattening the tensor into a matrix. The mode-1 unfolding of X , denoted as $X_{(1)}$, is the $m_1 \times 3m_2$ matrix whose columns are the columns of X^1, X^2 , and X^3 (see Figure B.3, right). Similarly the mode-2 unfolding, denoted as $X_{(2)}$, is the $m_2 \times 3m_1$ matrix whose columns are the transposed rows of X^1, X^2 , and X^3 .

Given the interaction data X , our task is to estimate M , which is a fixed tensor of the same size as X , representing the probabilities of user interactions. Formally,

$$X = M + \epsilon,$$

where ϵ is a random tensor of the same size as X and represents noise. Let \hat{M} denote the obtained estimate of M . The performance of any estimation method is measured by how close (e.g., in terms

Figure B.3: User Interaction Data Tensor and Its Mode-1 Unfolding



of mean squared error) this estimate is to the true M .

As explained in §2.5, the estimation procedure proceeds in two steps. In the first step, we adapt conventional matrix factorization to the interaction data tensor. Matrix factorization is one of the most successful methods used for model-based collaborative filtering in the literature (e.g., see Koren et al. 2009, Alhejaili and Fatima 2020). To apply this method to the tensor, we construct the mode-1 unfolding $X_{(1)}$ and identify latent features of users via singular value decomposition (SVD). SVD is a well-established technique for identifying latent semantic factors in information retrieval (e.g., Su and Khoshgoftaar 2009). More precisely, assuming that $X_{(1)}$ admits the singular value decomposition $X_{(1)} = U_{(1)}\Sigma_{(1)}V_{(1)}^\top$, we set \hat{U} to be the columns of $U_{(1)}$ corresponding to the r_1 largest singular values (breaking ties arbitrarily). In this same first step, we also apply a similar procedure using the mode-2 unfolding $X_{(2)}$, computing a matrix \hat{V} as the first r_2 left singular vectors of $X_{(2)}$. Letting u_i and v_j , respectively, be the rows of \hat{U} and \hat{V} , we can interpret this decomposition nicely as follows. Assuming that each user is associated with some unknown vector of latent features, we can ascribe to each male user i a set of r_1 latent (i.e., unknown) features u_i and similarly to each female user j a set of r_2 latent features v_j .

In the second step, we estimate a user’s probability of interacting with the other users, i.e., $M_{i,j}^k$, which is the probability of user i interacting with user j in interaction type k . When using the CF features identified in the first step, we can write $M_{i,j}^k$ as follows:

$$M_{i,j}^k = f^k(u_i, v_j), \quad k = 1, 2, 3, \tag{B.3}$$

where f^k is a function to be estimated. For this estimation, we test three alternative models: the model of Farias and Li (2019), a neural network model, and a random forest model.

One of our benchmark models is the one developed by Farias and Li (2019), who assume that f^k has a bilinear form: $\hat{M}^k = \hat{U}S^k\hat{V}^\top$. They then solve a least-squares problem to obtain a closed-form solution for the estimate of S^k as $\hat{S}^k = \hat{U}^\top X^k \hat{V}$ and finally compute \hat{M}^k as:

$$\hat{M}^k = \hat{U}\hat{U}^\top X^k \hat{V}\hat{V}^\top, \quad k = 1, 2, 3. \tag{B.4}$$

Although Farias and Li (2019) theoretically derive upper and lower bounds of expected estimation error, their approach has potential limitations in prediction accuracy and interpretability. First, their model cannot fully utilize the user attribute data in prediction. One might consider expanding the CF feature space by appending the user attributes, but this is not straightforward because the values of the user attributes have different scales than the values of the CF features (e.g., numerical features such as height, binary features such as verified name, and categorical features such as region). In addition, their model does not allow us to utilize the pairwise user attributes such as age difference because pairwise user attributes cannot be used in a product form to make predictions with the bilinear model. Second, even though their model provides a convenient way to compute probabilities of interaction in a closed form, it requires the bilinear assumption in making predictions, and there is no guarantee that this assumption fits the underlying data structure. For instance, when predicting the probability of interaction between two users with respect to their ages, the bilinear assumption means that the probability of interaction increases with the product of two users' ages, which is unlikely to hold in reality. Last but not least, their model does not offer any insight into how the CF features and the user attributes affect users' decisions. In other words, prediction results have very limited interpretability.

For these reasons, we employ two alternative machine learning techniques, neural network and random forest, in the second step of our estimation. A neural network is based on a collection of connected units or nodes called artificial neurons aggregated in layers and connected by the edges to the neurons in other layers. We first preprocess the data to feed into the model to run the neural network. We normalize individual user attributes (e.g., average attractiveness, height, weight) and pairwise attributes (e.g., height difference, age difference), then scale them with respect to CF features to avoid overfitting. Our base model entails one input layer, one hidden layer, and one output layer that contains 16, 32, and 2 neurons, respectively. After normalizing the individual and pairwise user attributes, the first layer takes the CF features as well as the individual and the pairwise user attributes. It outputs to the hidden layer, which then outputs to the last layer. The last layer estimates the likelihood of a user interacting with the other users.

We build our neural network model using Keras in Tensorflow. We use the ReLu activation

function in the input and hidden layers and the Softmax activation function in the output layer to generate the probabilities of male and female users’ open, send and accept decisions. We use the Adam optimizer with a learning rate of 0.0001 and a “sparse_categorical_crossentropy” loss function. Different batch sizes (10, 20, 30) and epoch numbers (20, 30, 50, 100, 1000) are used throughout our analysis. We test various network structures with changing numbers of layers and activation functions, employing techniques to overcome over-fitting (regularization, dropouts, etc.). The simple network described above results in the highest accuracy and AUC score, as we discuss in §2.5. This model, however, also lacks interpretability because it is a black-box model that does not give any insights into the importance of input features.

To overcome this shortcoming of a neural network model, we employ a random forest model as our final model. Like a neural network model, a random forest model is a supervised learning algorithm, and it is one of the most popular classification algorithms used in practice. The random forest consists of many decision trees built upon a random subset of features. This model works best when features have some predictive power, as in our setting. It generates a score between 0 and 1 to estimate the likelihood of a user interacting with another user. For our random forest model, we use the RandomForestClassifier class from the ensemble package in the scikit-learn library. The number of estimators, `n_estimators`, is set to the default value of 100, the number of jobs to run in parallel, `n_jobs`, 8, and the random state, 0. To avoid over-fitting, we set the minimum number of samples required to be at a leaf node, `min_samples_leaf`, equal to 10.

Using random forests, we estimate users’ probability of interaction without making any structural assumption about the underlying data structure. Similar to the neural network model, this model also utilizes CF features as well as individual and pairwise user attributes. The main advantage of the tree model over the neural network model is its ability to interpret model predictions. Understanding why a machine learning model makes a certain prediction is as crucial as the prediction’s accuracy. We adopt the SHAP (SHapley Additive exPlanations) framework recently developed by Lundberg and Lee (2017), which is extensively used in numerous applications. Using this framework, we shed light on the impact of various user attributes on users’ decisions to interact with other users. We present the discussion on the effect of user attributes on users’ open, send, and accept decisions in Appendix B.8.

B.8 Impact of User Attributes on Users’ Open, Send, and Accept Decisions

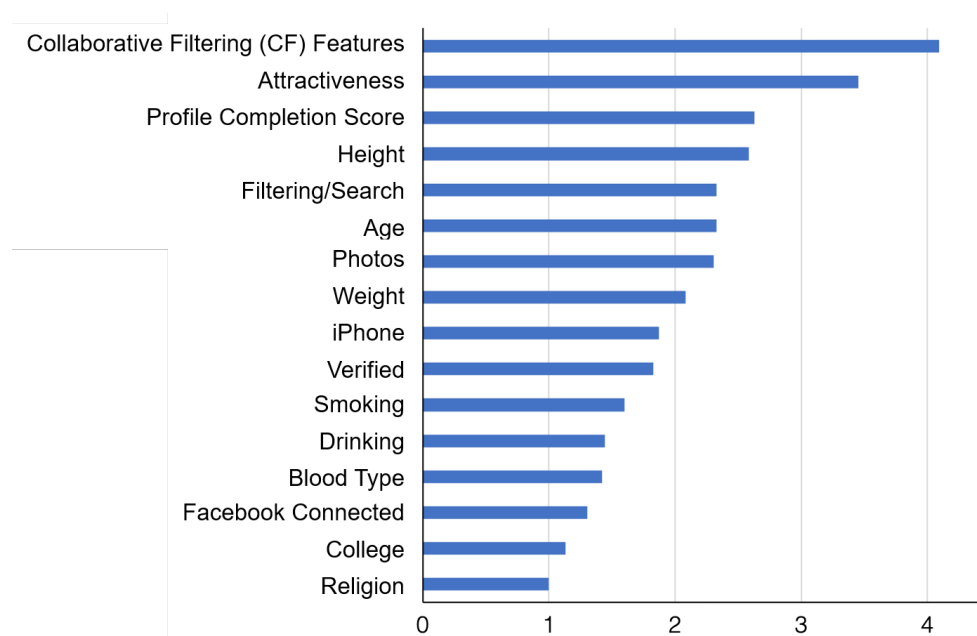
To interpret our prediction results, we compute and analyze importance scores for all user attributes (i.e., non-CF features) as well as the CF features used to predict users’ decisions. Among many approaches for calculating the importance of a feature, one recent family of approaches that have proven to be effective in terms of interpretability are those based on the concept of the Shapley value from game theory. The link between feature importance and cooperative games was originally noted by Lundberg and Lee (2017). Loosely, we can view the different features in a prediction task as cooperating to collect reward in the form of predictive accuracy – the game-theoretic question then is how to divide this reward in a “fair” manner, and the Shapley value is the unique solution that satisfies a set of reasonable preconditions for a fair division. Since calculating exact Shapley values is computationally intractable for even moderate-sized data, the body of work that has followed has focused on producing tractable approximations. The dominant approach for tree ensembles, such as our random forest model, is the *TreeSHAP* algorithm (Lundberg et al. 2020), which we use here.

SHAP importance scores can be calculated for each attribute, on each pair of users, for each of our six decisions that we predict (i.e., open, send, and accept decisions for male and female users). Broadly speaking, the SHAP score approximates the direction and magnitude of each attribute’s effect on the overall prediction for that pair of users. To summarize the *overall* importance of each attribute, we average the magnitudes of the SHAP scores for each attribute across all pairs of users and all decisions for a random selection of ten thousand pairs of users. The results are reported in Figure B.4.

Our analysis shows that the CF features have the highest average SHAP importance scores, about 2.7 greater than all other attributes combined. This result reiterates the importance of leveraging multiple types of interaction data among users, which has been largely ignored in prior literature on online dating.

For individual user attributes, as Figure B.4 shows, users’ attractiveness has the highest average SHAP score among all user attributes, by a wide margin (6.7 times greater than the next highest user attribute). This attribute represents a user’s popularity as perceived by other users (of the opposite gender). Obviously, this is an important attribute to predict other users’ interest in a focal user. What is interesting is that it dominates all other attributes by large margins. This means that how other users view a user is a determining factor in predicting users’ decisions – even more important than the user’s own preferences, captured by filtering/search settings (which is the

Figure B.4: Summary of Average SHAP Importance Scores Across All Six Decisions



Notes. The horizontal scale is logarithmic, so for example, the highest SHAP score (for the collaborative filtering features) is 10^4 times greater than the lowest SHAP score (for the religion attribute). Some attributes are aggregated: photos combines num. photobooks and num. photos, age includes age difference, height includes height difference, and religion includes religion match.

fifth most important attribute). Profile completion score is the third most important attribute, followed by attributes that represent physical attributes such as height, age, photos, and weight. This result indicates that beyond physical attributes users seek other types of information in order to find potentially compatible partners. Yet attributes that represent users' credibility, such as name verified and Facebook connected, turn out to be relatively less important than the other attributes we have discussed. Lastly, iPhone, smoking, drinking, blood type, college, and religion appear to have little influence on users' decisions.

It is worth noting that, whereas prior empirical studies primarily used physical attributes and education to make predictions, we find that attractiveness and profile completion score are the most important user attributes that have the highest average importance scores. Note that weight in our data indicates only body types, not exact weight or body mass index values. Even if the platform did collect numeric data about weight, there would be no way to verify whether self-reported weights were true – in fact, this is precisely why the platform does not ask users to enter numerical weight information. Dishonesty is a pervasive issue on online dating platforms (Anderson et al. 2020). Similarly, although users can enter the names of their colleges, 43% of users did not

provide this information. This is why college does not have a high importance score. Finally, we note that race is not an issue in our data, as most of the population in South Korea is ethnically homogeneous. Thus, the richness of our data allows us to identify many determinants of same-race preferences.

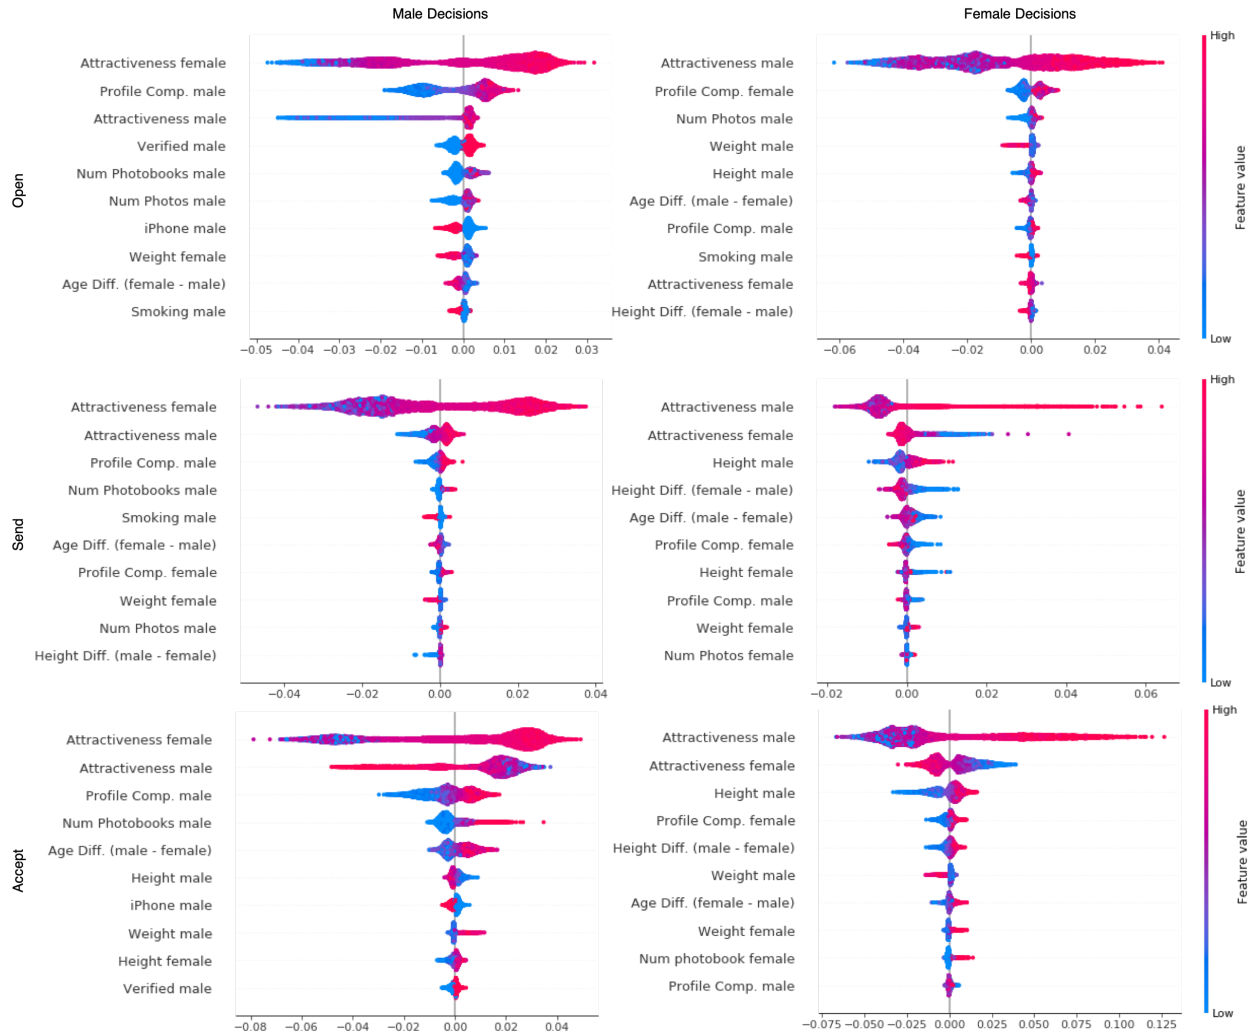
To dive deeper into the effect of individual attributes in predicting users’ interactions, we can analyze the SHAP scores for each of the six decisions (i.e., open, send and accept for male and female users) separately. Figure B.5 summarizes these SHAP scores for ten thousand randomly-selected pairs of users for the top ten user attributes (in average magnitude, as defined in the previous subsection) in each decision. The top ten attributes are ranked in descending order. For each of the six decisions shown in Figure B.5, the SHAP score of each attribute is plotted as a single point on the horizontal axis. The horizontal location of each point indicates whether the SHAP score is associated with a higher or lower prediction; specifically, a positive score means the attribute is pushing the prediction to 1, whereas a negative score means the attribute is pushing the prediction to 0. The figure also shows attribute values by color: red means that the value is high, and blue means that it is low.

Our first observation regarding Figure B.5 is that the attractiveness of the opposite gender is the most important attribute that determines users’ choices across all six decisions. In all three decisions of male users (the left column in Figure B.5), the attractiveness of female users is the most important user attribute overall, and we can observe that higher attractiveness (red) tends to yield higher predicted probabilities of all three decisions, as expected. The same effect is observed in all three decisions of female users (the right column in Figure B.5), where the attractiveness of male users is the most important user attribute, with higher attractiveness (red) leading to higher predicted probabilities.

Below the attractiveness of the opposite gender, three attributes appear in the top ten user attributes of all six decisions in Figure B.5: users’ own attractiveness, users’ own completion scores, and age difference. However, the importance ranks and their impacts on predicted probabilities are often different across gender and different types of interaction decisions. For example, we observe in male users’ decisions to open and send (top two subfigures in the left column of Figure B.5) that “attractiveness male” predominantly shows red data points in positive SHAP scores, whereas in male users’ decisions to accept (bottom subfigure in the left column of Figure B.5), “attractiveness male” shows plenty of both red and blue data points in positive SHAP scores. This means that male users with higher attractiveness scores tend to have higher probabilities of opening and sending, but they do not necessarily have higher probabilities of accepting.

This result indicates that more attractive male users tend to be more active in showing interest

Figure B.5: (Color in PDF File) Individual SHAP Importance Scores of the Top 10 User Attributes for Each of the Six Decisions



Notes. All individual attributes in Appendix B.1 are used separately for male and female users. In addition, three composite attributes (i.e., height difference, age difference, and religion match) are used. The filtering/search attribute is further dis-aggregated into specific criteria for age, drinking, height, etc., but none of these criteria are of top importance.

to female users, but they can be more selective in accepting messages from female users. The same pattern is not observed for female users: female users' own attractiveness is ranked much lower (9th) in their decisions to open than their decisions to send or accept (2nd), and female users with higher attractiveness scores tend to have lower probabilities of sending or accepting (i.e., in female decisions to send or accept in Figure B.5, blue dots are predominant in positive SHAP scores). This suggests that female users with higher attractiveness scores tend to be more reluctant to send messages and accept incoming messages. In almost all decisions (with the exception of female users' opening decisions), users with higher profile completion scores tend to have higher probabilities of opening, sending and accepting, indicating that they are more eager to interact with others. We also observe that male users prefer to interact with younger female users, while the impact of age difference on female users' decisions is not unambiguous.

In contrast with the attributes we have discussed so far, there are a few user attributes that are influential only in male users' decisions. The number of male users' photobooks is one of the top ten user attributes of all three decisions made by male users. In addition, various attributes (e.g., name verified, smoking, and iPhone) that characterize male users play important roles in predicting their own decisions. The effect of the number of photobooks and the verification of real names is similar to that of profile completion scores, as higher values in these attributes imply that male users are making efforts to provide more information about themselves. Female users' physical attributes, such as weight and height, also appear in the top ten attributes of some, but not all, male decisions. Female users' profile completion scores are important only in male users' decisions to send. This seems explicable, as male users who can see female users' detailed profile information might be more likely to send messages.

There are also a few distinct attributes that are implicated only in female users' decisions. Three user attributes, namely, height of male users, height difference, and profile completion score, are in the top ten user attributes for all three decisions of female users. This is in contrast to the earlier observation that female users' physical attributes and profile completion scores are in the top ten attributes for some decisions of male users, but not all three. As for the impact of height, we observe that taller male users are, on average, more appealing to female users. Interestingly, the amount of information male users provide can be a double-edged sword: higher profile completion scores of male users tend to increase the probability that female users will open their cards but decrease the probability that female users will send them messages. The latter observation is the opposite of the observation that in male decisions, higher profile completion scores of female users tend to increase the probability that male users will send them messages. Finally, compared to male users' own attributes, female users' own attributes seem less important: for example, for male

users' accept decisions, seven out of the top ten attributes are related to their own attributes, whereas for female users' accept decisions, only four out of the top ten attributes are related to their own attributes. Lastly, we note that our machine-learning model focuses on prediction rather than on testing causal hypotheses. In our discussion of directions (not magnitude) above, SHAP values represent correlation rather than causality, which is better identified by experimental designs or similar approaches.

B.9 Extension: Match-Maximizing without Double Counting

In the platform's match-maximizing objective, there are small chances that the same pairs of men and women are recommended to each other, and send messages to each other. This causes double counting in the objective function. In order to eliminate this, we revise the match-maximizing platform's problem as follows:

$$\begin{aligned}
\max_{x_h, x_l} \quad & N^m n(x_h p_h^m a_m^h + x_l p_l^m a_m^l) + N^h n p_m^h a_h^m + N^l n p_m^l a_l^m - n^2 x_h p_h^m p_m^h - n^2 x_l p_l^m p_m^l \\
& x_h + x_l = 1 \\
& 0 \leq x_h, x_l \leq 1,
\end{aligned} \tag{M-Ext}$$

where fourth ($n^2 x_h p_h^m p_m^h$) and fifth ($n^2 x_l p_l^m p_m^l$) terms represent the expected number of matches that result from the same high-tier woman-man and low-tier woman-man pairs messaging each other in the second stage of the game, respectively.

We note that the last two terms of the objective function in (M-Ext) due to double counting are negligibly small compared to the other terms in the objective function. This is due to the fact that the number of users on the platform, N^m , is significantly higher than the number of recommendations each user receives, n . We can illustrate this with a numerical example. For $N^m = 50,000$, $n = 2$, $\alpha = 0.4$, $\beta = 0.2$, $q = 0.2$, $x_h = 0.4$, and $x_l = 0.6$, we obtain $p_h^m = 0.21$, $p_l^m = 0.14$, $p_m^h = 0.11$, $p_m^l = 0.13$, $a_m^h = 0.44$, $a_m^l = 0.48$, $a_h^m = 0.62$, and $a_l^m = 0.55$. As a result, the expected number of matches between men and women is 3901, and the expected number of double-counted matches is 0.4, which is 0.01% of the objective function value.

The following corollary provides a revised condition given in Proposition 2.2, under which the match-maximizing platform is biased against popular or unpopular users.

Corollary B.1. *The match-maximizing platform is biased against low-tier women when $N^m n(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n\beta(h_2^2 - h_3^2) + 4n^2(h_3 p_m^l - h_2 p_m^h)(h_1 + 8) > 0$, against high-tier women when $N^m n(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n\beta(h_2^2 - h_3^2) + 4n^2(h_3 p_m^l - h_2 p_m^h)(h_1 + 8) < 0$, and is*

unbiased when $N^m n(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n\beta(h_2^2 - h_3^2) + 4n^2(h_3 p_m^l - h_2 p_m^h)(h_1 + 8) = 0$, where h_1 , h_2 , and h_3 are defined in Lemma 2.3.

Proof of Corollary B.1. Formulating the Lagrangian of the match-maximizing platform's optimization problem (M) gives $\mathcal{L}(x_h, x_l, \lambda) = N^m n x_h p_h^m a_m^h + N^m n x_l p_l^m a_m^l - n^2 x_h p_h^m p_m^h - n^2 x_l p_l^m p_m^l + \lambda(1 - x_h - x_l)$. We derive the first-order conditions:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x_h} &= N^m n p_h^m a_m^h + N^m n x_h \frac{\partial p_h^m}{\partial x_h} a_m^h + N^m n x_h p_h^m \frac{\partial a_m^h}{\partial x_h} - n^2 p_h^m p_m^h - n^2 x_h \frac{\partial p_h^m}{\partial x_h} p_m^h - \lambda = 0, \\ \frac{\partial \mathcal{L}}{\partial x_l} &= N^m n p_l^m a_m^l + N^m n x_l \frac{\partial p_l^m}{\partial x_l} a_m^l + N^m n x_l p_l^m \frac{\partial a_m^l}{\partial x_l} - n^2 p_l^m p_m^l - n^2 x_l \frac{\partial p_l^m}{\partial x_l} p_m^l - \lambda = 0, \text{ and} \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= 1 - x_h - x_l = 0,\end{aligned}$$

where $\frac{\partial p_h^m}{\partial x_h} = -\frac{h_1 h_2}{(h_1 x_h + 4)^2}$, $\frac{\partial p_l^m}{\partial x_l} = -\frac{h_1 h_3}{(h_1 x_l + 4)^2}$, $\frac{\partial a_m^h}{\partial x_h} = -\frac{4\beta h_2}{(h_1 x_h + 4)^2}$, and $\frac{\partial a_m^l}{\partial x_l} = -\frac{4\beta h_3}{(h_1 x_l + 4)^2}$ for $h_1 = 2(2q+1)\beta$, $h_2 = (2q+1)(1-\alpha q) - 4r$, and $h_3 = 1 - 4r$. Rewriting the first-order conditions gives

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x_h} &= \frac{N^m n h_2(1 - \alpha q)}{2(h_1 x_h + 4)} - \frac{N^m n(2\beta h_2^2 + (1 - \alpha q)h_1 h_2)x_h}{2(h_1 x_h + 4)^2} + \frac{N^m n(\beta h_1 h_2^2 x_h^2 - 4\beta h_2^2 x_h)}{(h_1 x_h + 4)^3} \\ &\quad - \frac{n^2 h_2 p_m^h}{h_1 x_h + 4} + \frac{n^2 x_h h_1 h_2 p_m^h}{(h_1 x_h + 4)^2} - \lambda = 0, \\ \frac{\partial \mathcal{L}}{\partial x_l} &= \frac{N^m n h_3}{2(h_1 x_l + 4)} - \frac{N^m n(2\beta h_3^2 + h_1 h_3)x_l}{2(h_1 x_l + 4)^2} + \frac{N^m n(\beta h_1 h_3^2 x_l^2 - 4\beta h_3^2 x_l)}{(h_1 x_l + 4)^3} \\ &\quad - \frac{n^2 h_3 p_m^l}{h_1 x_h + 4} + \frac{n^2 x_l h_1 h_3 p_m^l}{(h_1 x_h + 4)^2} - \lambda = 0, \text{ and} \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= 1 - x_h - x_l = 0.\end{aligned}$$

At $x_h = x_h^*$ and $x_l = x_l^*$, we have $\partial \mathcal{L}/x_h = \partial \mathcal{L}/x_l$, which can be written as:

$$\begin{aligned}&\frac{N^m n h_2(1 - \alpha q)}{2(h_1 x_h^* + 4)} - \frac{N^m n(2\beta h_2^2 + (1 - \alpha q)h_1 h_2)x_h^*}{2(h_1 x_h^* + 4)^2} + \frac{N^m n(\beta h_1 h_2^2 x_h^{*2} - 4\beta h_2^2 x_h^*)}{(h_1 x_h^* + 4)^3} - \frac{n^2 h_2 p_m^h}{h_1 x_h + 4} + \frac{n^2 x_h h_1 h_2 p_m^h}{(h_1 x_h + 4)^2} \\ &= \frac{N^m n h_3}{2(h_1 x_l^* + 4)} - \frac{N^m n(2\beta h_3^2 + h_1 h_3)x_l^*}{2(h_1 x_l^* + 4)^2} + \frac{N^m n(\beta h_1 h_3^2 x_l^{*2} - 4\beta h_3^2 x_l^*)}{(h_1 x_l^* + 4)^3} - \frac{n^2 h_3 p_m^l}{h_1 x_h + 4} + \frac{n^2 x_l h_1 h_3 p_m^l}{(h_1 x_h + 4)^2}. \quad (\text{B.5})\end{aligned}$$

To derive the condition under which the match-maximizing platform is unbiased, we examine the first order conditions of the match-maximizing platform's problem under which $\partial \mathcal{L}/\partial x_h = \partial \mathcal{L}/\partial x_l$

for $x_h^* = x_l^* = 1/2$. Substituting $x_h^* = x_l^* = 1/2$ into (B.5), we obtain

$$\begin{aligned} & \frac{N^m n h_2 (1 - \alpha q)}{2(h_1/2 + 4)} - \frac{N^m n (2\beta h_2^2 + (1 - \alpha q) h_1 h_2)/2}{2(h_1/2 + 4)^2} + \frac{N^m n (\beta h_1 h_2^2/4 - 2\beta h_2^2)}{(h_1/2 + 4)^3} - \frac{n^2 h_2 p_m^h}{h_1/2 + 4} + \frac{n^2 h_1 h_2 p_m^h/2}{(h_1/2 + 4)^2} \\ &= \frac{N^m n h_3}{2(h_1/2 + 4)} - \frac{N^m n (2\beta h_3^2 + h_1 h_3)/2}{2(h_1/2 + 4)^2} + \frac{N^m n (\beta h_1 h_3^2/4 - 2\beta h_3^2)}{(h_1/2 + 4)^3} - \frac{n^2 h_3 p_m^l}{h_1/2 + 4} + \frac{n^2 h_1 h_3 p_m^l/2}{(h_1/2 + 4)^2}. \end{aligned} \quad (\text{B.6})$$

Simplifying (B.6), we obtain $N^m n (h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n \beta (h_2^2 - h_3^2) + 4n^2 (h_3 p_m^l - h_2 p_m^h)(h_1 + 8) = 0$.

Suppose the left-hand side of (B.6) is greater than the right-hand side at $x_h = x_l = 1/2$, i.e., $N^m n (h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n \beta (h_2^2 - h_3^2) + 4n^2 (h_3 p_m^l - h_2 p_m^h)(h_1 + 8) > 0$. In this case, we want to show that $x_h^* > 1/2$ and $x_l^* < 1/2$. To this end, we first show that $\partial^2 \mathcal{L} / \partial x_h^2 < 0$ for all x_h in the following:

$$\begin{aligned} \frac{\partial^2 \mathcal{L}}{\partial x_h^2} &= -\frac{N^m n h_1 h_2 (1 - \alpha q)}{2(h_1 x_h + 4)^2} - \frac{N^m n (2\beta h_2^2 + (1 - \alpha q) h_1 h_2)}{2(h_1 x_h + 4)^2} + \frac{N^m n (2\beta h_2^2 + (1 - \alpha q) h_1 h_2) h_1 x_h}{(h_1 x_h + 4)^3} \\ &+ \frac{N^m n (2\beta h_1 h_2^2 x_h - 4\beta h_2^2)}{(h_1 x_h + 4)^3} - \frac{3N^m n (\beta h_1^2 h_2^2 x_h^2 + 4\beta h_1 h_2^2 x_h)}{(h_1 x_h + 4)^4} + \frac{2n^2 h_1 h_2 p_m^h}{(h_1 x_h + 4)^2} - \frac{2n^2 x_h h_1^2 h_2 p_m^h}{(h_1 x_h + 4)^3} \\ &= -\frac{8N^m n h_2 (h_1 (1 - \alpha q) + 2\beta h_2)}{(h_1 x_h + 4)^3} + \frac{2n^2 h_1 h_2 p_m^h (h_1 x_h + 4)(1 - x_h h_1)}{(h_1 x_h + 4)^3} < 0, \end{aligned}$$

where the first term is negative because $\alpha q < 1$, $\beta > 0$, $h_1 > 0$, $h_2 > 0$, and the second term can be negative or positive depending on whether $x_h h_1 > 1$ or $x_h h_1 < 1$. However, the overall expression is negative because $N^m \gg n$; therefore, the magnitude of the first term is greater than the second term. Similarly, $\partial^2 \mathcal{L} / \partial x_l^2 = -16N^m h_3 (h_1 + 4\beta h_3) / (h_1 x_l + 4)^3 + 2n^2 h_1 h_3 p_m^l (1 - x_l h_1) / (h_1 x_h + 4)^3 < 0$. The discriminant $D = (\partial^2 \mathcal{L} / \partial x_h^2)(\partial^2 \mathcal{L} / \partial x_l^2) - (\partial \mathcal{L} / \partial x_h x_l)^2 > 0$ because $\partial^2 \mathcal{L} / \partial x_h^2 < 0$, $\partial^2 \mathcal{L} / \partial x_l^2 < 0$, and $\partial \mathcal{L} / \partial x_h x_l = 0$. Since \mathcal{L} is strictly concave in x_h and x_l and $D > 0$, there is a unique optimal solution (x_h^*, x_l^*) (where $x_h^* + x_l^* = 1$) that satisfies $\partial \mathcal{L} / \partial x_h = \partial \mathcal{L} / \partial x_l = 0$. Equivalently, at this optimal solution, the left-hand side of (B.5) is equal to the right-hand side of (B.5) that is equal to λ . Let us denote the left-hand side of equation (B.5) with $f(x_h^*)$ and the right-hand side with $g(x_l^*)$. We claim that if $f(x_h) - g(x_l) > 0$, i.e., $N^m n (h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n \beta (h_2^2 - h_3^2) + 4n^2 (h_3 p_m^l - h_2 p_m^h)(h_1 + 8) > 0$ at $x_h = x_l = 1/2$, then it must be that $f(x_h) > f(x_h^*) = g(x_l^*) = \lambda > g(x_l)$. Suppose this does not hold. Then, it must be that either $\lambda > f(x_h) > g(x_l)$ or $f(x_h) > g(x_l) > \lambda$ for $x_h + x_l = 1$. If $\lambda > f(x_h) > g(x_l)$ for $x_h + x_l = 1$, then the optimal $x_h^* > x_h$ and $x_l^* > x_l$ because $\partial^2 f / \partial x_h^2 < 0$ and $\partial^2 g / \partial x_l^2 < 0$. However, $x_h^* + x_l^* > x_h + x_l = 1$, which contradicts $x_h^* + x_l^* = 1$. Similarly, if $f(x_h) > g(x_l) > \lambda$ for $x_h + x_l = 1$, then both $x_h^* < x_h$ and $x_l^* < x_l$. This contradicts $x_h^* + x_l^* = 1$ because $x_h^* + x_l^* < x_h + x_l = 1$. Therefore, we conclude that $f(x_h) > \lambda > g(x_l)$ if $f(x_h) - g(x_l) > 0$. This implies that if

$N^m n(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n\beta(h_2^2 - h_3^2) + 4n^2(h_3 p_m^l - h_2 p_m^h)(h_1 + 8) > 0$ at $x_h = x_l = 1/2$, $f(1/2) > \lambda > g(1/2)$. Since $\partial^2 f / \partial x_h^2 < 0$ and $\partial^2 g / \partial x_l^2 < 0$, $f(x_h^*) = g(x_l^*) = \lambda$ where $x_h^* > 1/2$ and $x_l^* < 1/2$.

Similarly, if $N^m n(h_1 + 8)(h_2(1 - \alpha q) - h_3) - 4N^m n\beta(h_2^2 - h_3^2) + 4n^2(h_3 p_m^l - h_2 p_m^h)(h_1 + 8) < 0$, then $x_h^* < 1/2$ and $x_l^* > 1/2$. □

The insights derived from Proposition 2.2 remain the same with the change in the objective function. We can see this by comparing the condition presented in Proposition 2.2 with the condition derived in Corollary B.1. The first two terms of the condition derived in Corollary B.1 are the same as in Proposition 2.2 multiplied by $N^m n$. The third term comes from the double-counted matches. Given that $N^m n \gg n^2$ because $N^m \gg n$, the first two terms determine the sign of the condition in Corollary B.1.

Bibliography

- Abo, S. M. and Smith, S. R. (2020), ‘Is a covid-19 vaccine likely to make things worse?’, *Vaccines* **8**(4), 761.
- Acemoglu, D., Chernozhukov, V., Werning, I. and Whinston, M. D. (2021), ‘Optimal targeted lockdowns in a multigroup sir model’, *American Economic Review: Insights* **3**(4), 487–502.
- Adachi, H. (2003), ‘A search model of two-sided matching under nontransferable utility’, *Journal of Economic Theory* **113**(2), 182–198.
- Afeche, P., Caldentey, R. and Gupta, V. (2022), ‘On the optimal design of a bipartite matching queueing system’, *Operations Research* **70**(1), 363–401.
- Ahn, J.-H. and Hornberger, J. C. (1996), ‘Involving patients in the cadaveric kidney transplant allocation process: A decision-theoretic perspective’, *Management Science* **42**(5), 629–641.
- Akan, M., Alagoz, O., Ata, B., Erenay, F. S. and Said, A. (2012), ‘A broader view of designing the liver allocation system’, *Operations research* **60**(4), 757–770.
- Alagoz, O., Maillart, L. M., Schaefer, A. J. and Roberts, M. S. (2004), ‘The optimal timing of living-donor liver transplantation’, *Management Science* **50**(10), 1420–1430.
- Alagoz, O., Maillart, L. M., Schaefer, A. J. and Roberts, M. S. (2007), ‘Choosing among living-donor and cadaveric livers’, *Management Science* **53**(11), 1702–1715.
- Alhejaili, A. and Fatima, S. (2020), Latent feature modelling for recommender systems, *in* ‘2020 IEEE 21st International Conference on Information Reuse and Integration for Data Science (IRI)’, IEEE, pp. 349–356.
- Anderson, M., Vogels, E. A. and Turner, E. (2020), ‘The virtues and downsides of online dating’.
- Ashlagi, I., Burq, M., Jaillet, P. and Manshadi, V. (2019), ‘On matching and thickness in heterogeneous dynamic markets’, *Operations Research* **67**(4), 927–949.

- Asrani, S. K. and Kamath, P. S. (2015), ‘Model for end-stage liver disease score and meld exceptions: 15 years later’, *Hepatology international* **9**(3), 346–354.
- Ata, B., Skaro, A. and Tayur, S. (2017), ‘Organjet: Overcoming geographical disparities in access to deceased donor kidneys in the united states’, *Management Science* **63**(9), 2776–2794.
- Bapna, R., Ramaprasad, J., Shmueli, G. and Umyarov, A. (2016), ‘One-way mirrors in online dating: A randomized field experiment’, *Management Science* **62**(11), 3100–3122.
- Bazant, M. Z. and Bush, J. W. (2021), ‘A guideline to limit indoor airborne transmission of covid-19’, *Proceedings of the National Academy of Sciences* **118**(17), e2018995118.
- Bazant, M. Z., Kodio, O., Cohen, A. E., Khan, K., Gu, Z. and Bush, J. W. (2021), ‘Monitoring carbon dioxide to quantify the risk of indoor airborne transmission of covid-19’, *Flow* **1**, E10.
- Belo, R. and Li, T. (2018), ‘Referral programs for platform growth: Evidence from a randomized field experiment’.
- Benjaafar, S. and Hu, M. (2020), ‘Operations management in the age of the sharing economy: What is old and what is new?’, *Manufacturing & Service Operations Management* **22**(1), 93–101.
- Berg, M. B. and Lin, L. (2020), ‘Prevalence and predictors of early covid-19 behavioral intentions in the united states’, *Translational Behavioral Medicine* **10**(4), 843–849.
- Bernards, S., Lee, E., Leung, N., Akan, M., Gan, K., Zhao, H., Sarkar, M., Tayur, S. and Mehta, N. (2022), ‘Awarding additional meld points to the shortest waitlist candidates improves sex disparity in access to liver transplant in the united states’, *American Journal of Transplantation* **22**(12), 2912–2920.
- Bertsimas, D., Farias, V. F. and Trichakis, N. (2013), ‘Fairness, efficiency, and flexibility in organ allocation for kidney transplantation’, *Operations Research* **61**(1), 73–87.
- Bojd, B. and Yoganarasimhan, H. (2022), ‘Star-cursed lovers: Role of popularity information in online dating’, *Marketing Science* **41**(1), 73–92.
- Bruch, E., Feinberg, F. and Lee, K. Y. (2016), ‘Extracting multistage screening rules from online dating activity data’, *Proceedings of the National Academy of Sciences* **113**(38), 10530–10535.
- Buonanno, G., Morawska, L. and Stabile, L. (2020), ‘Quantitative assessment of the risk of airborne transmission of sars-cov-2 infection: prospective and retrospective applications’, *Environment international* **145**, 106112.

- Chang, J. T. and Kaplan, E. H. (2023), ‘Modeling local coronavirus outbreaks’, *European journal of operational research* **304**(1), 57–68.
- Chen, Y.-J., Dai, T., Korpeoglu, C. G., Körpeoğlu, E., Sahin, O., Tang, C. S. and Xiao, S. (2020), ‘Om forum—innovative online platforms: Research opportunities’, *Manufacturing & Service Operations Management* **22**(3), 430–445.
- Ciampaglia, G. L., Nematzadeh, A., Menczer, F. and Flammini, A. (2018), ‘How algorithmic popularity bias hinders or promotes quality’, *Scientific reports* **8**(1), 15951.
- Couture, V., Dingel, J. I., Green, A., Handbury, J. and Williams, K. R. (2022), ‘Jue insight: Measuring movement and social contact with smartphone data: a real-time application to covid-19’, *Journal of Urban Economics* **127**, 103328.
- Cox-North, P., Doorenbos, A., Shannon, S. E., Scott, J. and Curtis, J. R. (2013), ‘The transition to end-of-life care in end-stage liver disease’, *Journal of Hospice & Palliative Nursing* **15**(4), 209–215.
- Cronin, C. J. and Evans, W. N. (2020), Private precaution and public restrictions: what drives social distancing and industry foot traffic in the covid-19 era?, Technical report, National Bureau of Economic Research.
- Dai, T., Zheng, R. and Sycara, K. (2020), ‘Jumping the line, charitably: Analysis and remedy of donor-priority rule’, *Management Science* **66**(2), 622–641.
- David, I. and Yechiali, U. (1985), ‘A time-dependent stopping problem with application to live organ transplants’, *Operations Research* **33**(3), 491–504.
- Davis, A., Mehrotra, S., Friedewald, J. and Ladner, D. (2013), Characteristics of a simulation model of the national kidney transplantation system, in ‘2013 Winter Simulations Conference (WSC)’, IEEE, pp. 2320–2329.
- Delasay, M., Jain, A. and Kumar, S. (2022), ‘Impacts of the covid-19 pandemic on grocery retail operations: An analytical model’, *Production and Operations Management* **31**(5), 2237–2255.
- El Ouardighi, F., Khmelnitsky, E. and Sethi, S. P. (2022), ‘Epidemic control with endogenous treatment capability under popular discontent and social fatigue’, *Production and Operations Management* **31**(4), 1734–1752.
- Erol, S. (2019), ‘Network hazard and bailouts’, *Available at SSRN 3034406* .

- Farias, V. F. and Li, A. A. (2019), ‘Learning preferences with side information’, *Management Science* **65**(7), 3131–3149.
- Fawcett, T. (2006), ‘An introduction to roc analysis’, *Pattern recognition letters* **27**(8), 861–874.
- Fernández-Villaverde, J. and Jones, C. I. (2022), ‘Estimating and simulating a sird model of covid-19 for many countries, states, and cities’, *Journal of Economic Dynamics and Control* **140**, 104318.
- Fong, J. (2020), ‘Search, selectivity, and market thickness in two-sided markets: Evidence from online dating’, *Selectivity, and Market Thickness in Two-Sided Markets: Evidence from Online Dating (December 19, 2020)* .
- Fukazawa, K. and Nishida, S. (2016), ‘Size mismatch in liver transplantation’, *Journal of Hepato-Biliary-Pancreatic Sciences* **23**(8), 457–466.
- Gale, D. and Shapley, L. S. (1962), ‘College admissions and the stability of marriage’, *The American Mathematical Monthly* **69**(1), 9–15.
- Godfrey, E. L., Malik, T. H., Lai, J. C., Mindikoglu, A. L., Galván, N. T. N., Cotton, R. T., O’Mahony, C. A., Goss, J. A. and Rana, A. (2019), ‘The decreasing predictive power of meld in an era of changing etiology of liver disease’, *American Journal of Transplantation* **19**(12), 3299–3307.
- Goolsbee, A. and Syverson, C. (2021), ‘Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020’, *Journal of public economics* **193**, 104311.
- Gupta, S., Starr, M. K., Farahani, R. Z. and Asgari, N. (2022), ‘Om forum-pandemics/epidemics: Challenges and opportunities for operations management research’, *Manufacturing & Service Operations Management* **24**(1), 1–23.
- Gurvich, I. and Ward, A. (2015), ‘On the dynamic control of matching queues’, *Stochastic Systems* **4**(2), 479–523.
- Halaburda, H., Jan Piskorski, M. and Yildirim, P. (2018), ‘Competing by restricting choice: The case of matching platforms’, *Management Science* **64**(8), 3574–3594.
- Han, B. R., Sun, T., Chu, L. Y. and Wu, L. (2022), ‘Covid-19 and e-commerce operations: evidence from alibaba’, *Manufacturing & Service Operations Management* **24**(3), 1388–1405.
- Hasankhani, F. and Khademi, A. (2021), ‘Is it time to include post-transplant survival in heart transplantation allocation rules?’, *Production and Operations Management* **30**(8), 2653–2671.

- Hauser, J. R. and Wernerfelt, B. (1990), ‘An evaluation cost model of consideration sets’, *Journal of consumer research* **16**(4), 393–408.
- Henriques, A., Azzopardi, G., Tarocco, N., Devine, J., Rognlien, M. K., Andreini, M., Elson, P. J. and Mounet, N. (2021), Modelling airborne transmission of sars-cov-2: Risk assessment for enclosed spaces, Technical report.
- Hitsch, G. J., Hortaçsu, A. and Ariely, D. (2010a), ‘What makes you click?—mate preferences in online dating’, *Quantitative marketing and Economics* **8**, 393–427.
- Hitsch, G. J., Hortaçsu, A. Ariely, D. (2010b), ‘Matching and sorting in online dating’, *American Economic Review* **100**(1), 130–163.
- Hosmer, D. W., Lemeshow, S. and Sturdivant, R. X. (2013), *Applied logistic regression*, Vol. 398, John Wiley & Sons.
- Howard, D. H. (2002), ‘Why do transplant surgeons turn down organs?: A model of the accept/reject decision’, *Journal of health economics* **21**(6), 957–969.
- Hu, M. and Zhou, Y. (2022), ‘Dynamic type matching’, *Manufacturing & Service Operations Management* **24**(1), 125–142.
- Iyengar, S. S. and Lepper, M. R. (2000), ‘When choice is demotivating: Can one desire too much of a good thing?’, *Journal of personality and social psychology* **79**(6), 995.
- Kang, K., Doroudi, S., Delasay, M. and Wickeham, A. (2022), ‘A queueing-theoretic framework for evaluating transmission risks in service facilities during a pandemic’, *Production and Operations Management* .
- Kanoria, Y. and Saban, D. (2021), ‘Facilitating the search for partners on matching platforms’, *Management Science* **67**(10), 5990–6029.
- Kaplan, E. H. (2020), ‘Om forum-covid-19 scratch models to support local decisions’, *Manufacturing & Service Operations Management* **22**(4), 645–655.
- Kaplan, G., Moll, B. and Violante, G. L. (2020), The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the us, Technical report, National Bureau of Economic Research.
- Kapoor, N. R., Kumar, A., Kumar, A., Kumar, A. and Kumar, K. (2022), ‘Transmission probability of sars-cov-2 in office environment using artificial neural network’, *Ieee Access* **10**, 121204–121229.

- Khan, S., Haleem, A., Deshmukh, S. and Javaid, M. (2021), ‘Exploring the impact of covid-19 pandemic on medical supply chain disruption’, *Journal of industrial integration and management* **6**(02), 235–255.
- Kim, S.-P., Gupta, D., Israni, A. K. and Kasiske, B. L. (2015), ‘Accept/decline decision module for the liver simulated allocation model’, *Health care management science* **18**, 35–57.
- Kong, N., Schaefer, A. J., Hunsaker, B. and Roberts, M. S. (2010), ‘Maximizing the efficiency of the us liver allocation system through region design’, *Management Science* **56**(12), 2111–2122.
- Koren, Y., Bell, R. and Volinsky, C. (2009), ‘Matrix factorization techniques for recommender systems’, *Computer* **42**(8), 30–37.
- Kreke, J., Schaefer, A. J., Angus, D. C., Bryce, C. L. and Roberts, M. S. (2002), Incorporating biology into discrete event simulation models of organ allocation, *in* ‘Proceedings of the Winter Simulation Conference’, Vol. 1, IEEE, pp. 532–536.
- Lai, J. C., Terrault, N. A., Vittinghoff, E. and Biggins, S. W. (2010), ‘Height contributes to the gender difference in wait-list mortality under the meld-based liver allocation system’, *American Journal of Transplantation* **10**(12), 2658–2664.
- Lambrecht, A. and Tucker, C. (2019), ‘Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads’, *Management science* **65**(7), 2966–2981.
- Lee, S. (2016), ‘Effect of online dating on assortative mating: Evidence from south korea’, *Journal of Applied Econometrics* **31**(6), 1120–1139.
- Lin, K.-H. and Lundquist, J. (2013), ‘Mate selection in cyberspace: The intersection of race, gender, and education’, *American Journal of Sociology* **119**(1), 183–215.
- Liu, Y. and Wu, B. (2022), ‘Coevolution of vaccination behavior and perceived vaccination risk can lead to a stag-hunt-like game’, *Physical Review E* **106**(3), 034308.
- Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N. and Lee, S.-I. (2020), ‘From local explanations to global understanding with explainable ai for trees’, *Nature machine intelligence* **2**(1), 56–67.
- Lundberg, S. M. and Lee, S.-I. (2017), ‘A unified approach to interpreting model predictions’, *Advances in neural information processing systems* **30**.

- Mak, H.-Y., Dai, T. and Tang, C. S. (2022), ‘Managing two-dose covid-19 vaccine rollouts with limited supply: Operations strategies for distributing time-sensitive resources’, *Production and Operations Management* **31**(12), 4424–4442.
- Marvin, M. R., Ferguson, N., Cannon, R. M., Jones, C. M. and Brock, G. N. (2015), ‘Meldeq: An alternative model for end-stage liver disease score for patients with hepatocellular carcinoma’, *Liver Transplantation* **21**(5), 612–622.
- Massie, A. B., Caffo, B., Gentry, S. E., Hall, E. C., Axelrod, D., Lentine, K. L., Schnitzler, M. A., Gheorghian, A., Salvalaggio, P. R. and Segev, D. L. (2011), ‘Meld exceptions and rates of waiting list outcomes’, *American Journal of Transplantation* **11**(11), 2362–2371.
- Nazari, M. and Stolyar, A. L. (2019), ‘Reward maximization in general dynamic matching systems’, *Queueing Systems* **91**, 143–170.
- Nikolopoulos, K., Punia, S., Schäfers, A., Tsinopoulos, C. and Vasilakis, C. (2021), ‘Forecasting and planning during a pandemic: Covid-19 growth rates, supply chain disruptions, and governmental decisions’, *European journal of operational research* **290**(1), 99–115.
- Ong, D. and Wang, J. (2015), ‘Income attraction: An online dating field experiment’, *Journal of Economic Behavior & Organization* **111**, 13–22.
- Ooi, C. C., Suwardi, A., Ou Yang, Z. L., Xu, G., Tan, C. K. I., Daniel, D., Li, H., Ge, Z., Leong, F. Y., Marimuthu, K. et al. (2021), ‘Risk assessment of airborne covid-19 exposure in social settings’, *Physics of Fluids* **33**(8), 087118.
- Polack, F. P., Thomas, S. J., Kitchin, N., Absalon, J., Gurtman, A., Lockhart, S., Perez, J. L., Pérez Marc, G., Moreira, E. D., Zerbini, C. et al. (2020), ‘Safety and efficacy of the bnt162b2 mrna covid-19 vaccine’, *New England journal of medicine* **383**(27), 2603–2615.
- Pritsker, A. A. B., Martin, D. L., Reust, J. S., Wagner, M. A., Daily, O. P., Harper, A. M., Edwards, E. B., Bennett, L. E., Wilson, J. R., Kuhl, M. E. et al. (1995), Organ transplantation policy evaluation, in ‘Proceedings of the 27th conference on winter simulation’, pp. 1314–1323.
- Pronk, T. M. and Denissen, J. J. (2020), ‘A rejection mind-set: Choice overload in online dating’, *Social Psychological and Personality Science* **11**(3), 388–396.
- Qi, A., Sethi, S., Wei, L. and Zhang, J. (2020), ‘Strategic overcapacity in live-streaming platform selling’, *Available at SSRN* .

- Reddy, M. S., Varghese, J., Venkataraman, J. and Rela, M. (2013), ‘Matching donor to recipient in liver transplantation: relevance in clinical practice’, *World journal of hepatology* **5**(11), 603.
- Righter, R. (1989), ‘A resource allocation problem in a random environment’, *Operations Research* **37**(2), 329–338.
- Rios, I., Saban, D. and Zheng, F. (2022), ‘Improving match rates in dating markets through assortment optimization’, *Manufacturing & Service Operations Management* .
- Rockafellar, R. T. and Wets, R. J.-B. (2009), *Variational analysis*, Vol. 317, Springer Science & Business Media.
- Ruth, R. J., Wyszewianski, L. and Herline, G. (1985), ‘Kidney transplantation: A simulation model for examining demand and supply’, *Management Science* **31**(5), 515–526.
- Salmenjoki, H., Korhonen, M., Puisto, A., Vuorinen, V. and Alava, M. J. (2021), ‘Modelling aerosol-based exposure to sars-cov-2 by an agent based monte carlo method: Risk estimates in a shop and bar’, *Plos one* **16**(11), e0260237.
- Sandıkçı, B., Maillart, L. M., Schaefer, A. J., Alagoz, O. and Roberts, M. S. (2008), ‘Estimating the patient’s price of privacy in liver transplantation’, *Operations Research* **56**(6), 1393–1410.
- Sandıkçı, B., Tunç, S. and Tanrıover, B. (2019), A new simulation model for kidney transplantation in the united states, *in* ‘2019 Winter Simulation Conference (WSC)’, IEEE, pp. 1079–1090.
- Shang, Y., Dong, J., Tian, L., He, F. and Tu, J. (2022), ‘An improved numerical model for epidemic transmission and infection risks assessment in indoor environment’, *Journal of Aerosol Science* **162**, 105943.
- Shechter, S. M., Bryce, C. L., Alagoz, O., Kreke, J. E., Stahl, J. E., Schaefer, A. J., Angus, D. C. and Roberts, M. S. (2005), ‘A clinically based discrete-event simulation of end-stage liver disease and the organ allocation process’, *Medical Decision Making* **25**(2), 199–209.
- Shi, L. and Viswanathan, S. (2023), ‘Optional verification and signaling in online matching markets: Evidence from a randomized field experiment’, *Information Systems Research* .
- Shumsky, R. A., Debo, L., Lebeaux, R. M., Nguyen, Q. P. and Hoen, A. G. (2021), ‘Retail store customer flow and covid-19 transmission’, *Proceedings of the National Academy of Sciences* **118**(11), e2019225118.

- Su, X. and Khoshgoftaar, T. M. (2009), ‘A survey of collaborative filtering techniques’, *Advances in artificial intelligence* **2009**.
- Su, X. and Zenios, S. (2004), ‘Patient choice in kidney allocation: The role of the queueing discipline’, *Manufacturing & Service Operations Management* **6**(4), 280–301.
- Su, X. and Zenios, S. A. (2005), ‘Patient choice in kidney allocation: A sequential stochastic assignment model’, *Operations research* **53**(3), 443–455.
- Su, X. and Zenios, S. A. (2006), ‘Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model’, *Management science* **52**(11), 1647–1660.
- Tang, Y. S., Scheller-Wolf, A. A. and Tayur, S. R. (2021), ‘Multi-armed bandits with endogenous learning and queueing: An application to split liver transplantation’, *Available at SSRN 3855206*.
- Toso, C., Dupuis-Lozeron, E., Majno, P., Berney, T., Kneteman, N. M., Perneger, T., Morel, P., Mentha, G. and Combescure, C. (2012), ‘A model for dropout assessment of candidates with or without hepatocellular carcinoma on a common liver transplant waiting list’, *Hepatology* **56**(1), 149–156.
- Usherwood, T., LaJoie, Z. and Srivastava, V. (2021), ‘A model and predictions for covid-19 considering population behavior and vaccination’, *Scientific reports* **11**(1), 1–11.
- Vitale, A., Volk, M. L., De Feo, T. M., Burra, P., Frigo, A. C., Morales, R. R., De Carlis, L., Belli, L., Colledan, M., Fagioli, S. et al. (2014), ‘A method for establishing allocation equity among patients with and without hepatocellular carcinoma on a common liver transplant waiting list’, *Journal of hepatology* **60**(2), 290–297.
- Vogels, E. A. (2020), ‘10 facts about americans and online dating’.
- Wu, Y. and Padmanabhan, V. (2019), ‘The strategy puzzle of subscription-based dating sites’.
URL: <https://hbr.org/2019/01/the-strategy-puzzle-of-subscription-based-dating-sites>
- Wu, Y., Zhang, K. and Padmanabhan, V. (2018), ‘Matchmaker competition and technology provision’, *Journal of Marketing Research* **55**(3), 396–413.
- Wu, Z. (2021), ‘Social distancing is a social dilemma game played by every individual against his/her population’, *Plos one* **16**(8), e0255543.
- Zenios, S. A., Chertow, G. M. and Wein, L. M. (2000), ‘Dynamic allocation of kidneys to candidates on the transplant waiting list’, *Operations Research* **48**(4), 549–569.