Winter 12-15-2018

Computational Explorations of Information and Mechanism Design in Markets

Zhuoshu Li
Washington University in St. Louis

Washington University in St. Louis

School of Engineering and Applied Science

Department of Computer Science and Engineering

Dissertation Examination Committee:
Sanmay Das, Chair
John P. Dickerson
Roman Garnett
Jason R. Wellen
William Yeoh

Computational Explorations of Information and Mechanism Design in Markets

by

Zhuoshu Li

A dissertation presented to
The Graduate School
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

December 2018
Saint Louis, Missouri

Contents

List of Tables

List of Figures

Acknowledgments

This thesis, or finishing my Ph.D., could not have happened without the constant support and guidance of my advisor, Sanmay Das. He is a great mentor who has always been unfailingly generous with his time and advice. As a researcher, who has been my role model, he really taught me about professionalism and continues to be resourceful in all aspects of academia. For the past five years, I have been bugging him for innumerable questions, asking for advice and suggestions. His office is always opened for me, whenever a help is needed.

I am deeply indebted to my dissertation committee of John Dickerson, Roman Garnett, Jason Wellen, and William Yeoh, who have offered much invaluable insight, guidance, and feedback throughout this process. I have also enjoyed my conversations with all of them and learned so much outside the scope of the dissertation.

I want to thank my senior lab-mates, Meenal Chhabra, Mithun Chakraborty, and Allen Lavoie, who would be my models for successful graduate students. They were always ready to help. I cannot image how struggle I would have been without them in my early Ph.D. life. I also want to say thank you to Hao Yan and Shali Jiang for being great lab mates and friends of my second half of the graduate school. They are always generous on time whenever I need a discussion in research, job-seeking preparations, and expanding my knowledge outside of my research area. Thanks are due to my publication co-authors outside of WUSTL, Sofia Carrillo, Neal Gupta, and Tuomas Sandholm. I have also had incredible colleagues and mentors in the industry. Zhitang Chen introduced me to a brand new area and walked me through all the kernel and reinforcement learning things that I can rely on in the future. Fei Huang shared all his expertise with me on infrastructure, algorithm implementations and model training from an industrial view, which greatly sharpened my skillset. Yanhui Geng has demonstrated me how to be a successful program manager and director, which inspired me to think about all the possibilities in my future career.

I would like to thank Eddie Xu, Wenlin Chen, Ming Zou, and Wei Qiao, for all kinds of help in my job search. I am thankful for my friends and gym-mates, Xiaoxin Liu, Wenlin Wang, Junjie Liu, Liang Zhou, Mingquan Yuan, Muhan Zhang, Liran Wang, Jinghan Yang, Haipeng Dai, and Wei Tang, for keeping my life joyful and healthy, and also feeding me with delicious hotpot.

Much gratitude goes to all members of the faculty and administrative staff at WUSTL, as well as all the students I have interacted with. They all contributed to the wonderful graduate school life that I had. In particular, I have greatly appreciated working with Chien-Ju Ho and REU students in our lab.

I am thankful for financial support from NSF IIS awards 1414452 and 1527037.

On a personal note, none of this could have happened without my parents, whose unconditional love made me move forward fearlessly. Last but not the least, my life cannot be as happy as I am now without my husband, Yang Zhang, who offers continuous and unlimited support, and made me a better and stronger person.

Zhuoshu Li

Washington University in Saint Louis
December 2018

ABSTRACT OF THE DISSERTATION

Computational Explorations of Information and Mechanism Design in Markets

by

Zhuoshu Li

Doctor of Philosophy in Computer Science

Washington University in St. Louis, December 2018

Research Advisor: Professor Sanmay Das

Markets or platforms assemble multiple selfishly-motivated and strategic agents. The outcomes of such agent interactions depend heavily on the rules, regulations, and norms of the platform, as well as the information available to agents. This thesis investigates the design and analysis of mechanisms and information structures through the "computational lens" in both static and dynamic settings. It both addresses the outcome of single platforms and fills a gap in the study of the dynamics of multiple platform interactions.

In static market settings, we are particularly interested in the role of information, because mechanisms are harder to change than the information available to participants. We approach information design through specific examples, i.e., matching markets and auction markets. First, in matching markets, we study the situation where the matching is preceded by a costly interviewing stage in which firms acquire information about the qualities of candidates. We focus on the impact of the signals of quality available prior to the interviewing stage. We show that more "commonality" in the quality of information can be harmful, yielding fewer matches. Second, in auction markets, we design an information environment

for revenue enhancement in a sealed-bid second price auction. Much of the previous literature has focused on signal design in settings where bidders are symmetrically informed, or on the design of optimal mechanisms under fixed information structures. Here, we provide new theoretical insights for complex situations like corporate mergers, where the sender of the signal has the opportunity to communicate in different ways to different receivers.

Next, in dynamic markets, we focus on two dimensions: (1) the effects of different market-clearing rules on market outcomes and (2) the dynamics of multiple platform interactions. Considering both dimensions, we investigate two important real-world dynamic markets: kidney exchange and financial markets. Specifically, in kidney exchange, we analyze the performance of different market-clearing algorithms and design a competing-market model to quantify the social welfare loss caused by market competition and exchange fragmentation. Here, we present the first analysis of equilibrium behavior in these dynamic competing matching market systems, from the viewpoints of both agents and markets. To improve the performance of kidney exchange in terms of both social welfare and individual utility, we analyze the benefit of convincing directed donation pairs to participate in paired kidney exchange, measured in terms of long-term graft survival. We provide the first empirical evidence that including compatible pairs dramatically benefits both social welfare and individual outcomes.

For financial markets, in the debate over high frequency trading, the frequent call (Call) mechanism has recently received considerable attention as a proposal for replacing the continuous double auction (CDA) mechanisms that currently run most financial markets. We examine agents' profit under CDA and frequent call auctions in a dynamic environment. We design an agent-based model to study the competition between these two market policies and show that CALL markets can drive trade away from CDAs. The results help to inform this very important debate.

Preface

Most of this thesis is based on works jointly published with others. In particular:

- Chapter 2 represents joint work with Sanmay Das [40], which appears in *Proceedings of the 10th International Conference on Web and Internet Economics (WINE 2014)*, under the title "The Role of Common and Private Signals in Two-Sided Matching with Interviews."

- Chapter 3 contains joint work with Sanmay Das [96], which was presented at the IJCAI Algorithmic Game Theory Workshop, 2017, under the title "Revenue Enhancement Via Asymmetric Signaling in Interdependent-Value Auctions."

- Chapter 4 is based on joint work with Sanmay Das, John P. Dickerson, and Tuomas Sandholm [37], which appears in *Proceedings of the 3rd Conference on Auctions, Market Mechanisms and Their Applications (AMMA 2015)*, under the title "Competing Dynamic Matching Markets."

- Chapter 5 is based on joint work with Neal Gupta, Sanmay Das, and John P. Dickerson [98], which appears in *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, under the title "Equilibrium Behavior in Competing Dynamic Matching Markets."

- Chapter 6 contains joint work with Sanmay Das, Sofia Carrillo, and Jason R. Wellen [97], which is presented at the AAMAS-IJCAI Workshop on Agents and Incentives in

Artificial Intelligence, 2017, under the title "Estimating the Benefits of Incorporating Compatible Pairs in Kidney Exchange."

- Chapter 7 is based on joint work with Sanmay Das [95], which appears in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS 2016)*, under the title "An Agent-Based Model of Competition Between Financial Exchanges: Can Frequent Call Mechanisms Drive Trade Away from CDAs?"

Chapter 1

Introduction

Markets or platforms assemble multiple selfishly-motivated, strategic agents. The outcomes of such agents' interactions depend heavily on how the market is designed, which involve the rules of the platform as well as the information available to the agents. Poorly designed platforms and information environments suffer from unexpected and undesirable results, for example, agent manipulation (the scandal of badminton at the 2012 Summer Olympics - Women's doubles [1], where eight players were found guilty of "not using best efforts" by playing to lose matches in order to manipulate the draw for the knockout stage), unnecessary expense of human effort and equipment, social welfare loss (high-frequency trading arms race [27]), and so on.

This thesis studies mechanisms and information structures of markets through the "computational lens." We investigate problems for both static settings and dynamic settings. In a static setting, agents and items arrive at the market and are matched at the same time, then the market disappears (examples include matching medical residents to hospitals and

[1] https://en.wikipedia.org/wiki/Badminton_at_the_2012_Summer_Olympics_%E2%80%93_Women%27s_doubles

matching students to schools). In a dynamic setting, agents (or both agents and items) arrive and depart over time in a persistent market (for example, kidney exchange or financial exchange (NYSE, Nasdaq)). Further, we address the design and analysis of mechanisms and information structures on the outcome of single platforms, as well as the dynamics of multiple platform interaction. Using multiple approaches, including computational game theory, multi-agent simulation, and empirical game analysis, this thesis provides insights for important real-world domains like kidney exchange and financial markets, and also methodological advances in modeling complex and dynamic agent interaction environments.

The rest of this chapter introduces the background of this thesis for each part and overviews the structure and high-level contributions of the thesis.

1.1 Static Markets and Information Design

Traditionally, researchers and policymakers have studied platforms where agents interact for economic or social purposes to inform mechanism and market design. One specific goal is to provide a descriptive model of how the rules that the platform imposes on intermediate interactions between the agents affect individual and social outcomes. However, one can also then use these models to guide the design of these platforms to achieve certain social or commercial goals. This is the focus of the field of market design, which can be thought of as "microeconomic engineering" and has clearly become one of the key areas where economics, computer science, and operations research intersect [82]. In most of the literature on mechanism design, the model assumes that agents' information is given, and then searches for rules of the game that yield desired outcomes. However, there has recently been considerable interest in the parallel problem of designing the *information environment* that agents will encounter [35, 77]. This paradigm is clearly applicable in many scenarios of interest to AI researchers, including online advertising and internet marketplaces. This line of research is motivated by asymmetric information, where some relevant information is available to only one side of the market. For example, the seminal job market signaling model of Michael Spence [132] considers how the employers try to infer the quality of candidates from observable characteristics. In systems with entrenched mechanisms that are unlikely to change, manipulating information available to agents (employers in Spence's model), and thus affecting the outcomes of the mechanism becomes more valuable, especially when the valuations or preferences of agents are unknown or noisy. In this thesis, we investigate the role and design of information in both matching markets (Chapter 2 [40]) and auction markets (Chapter 3 [96]) in static settings.

1.1.1 Matching Markets and the Role of Information

Matching markets have a long history of study in economics, operations research, and other areas, which focus on who gets what. In 1962, Gale and Shapley published their pioneering paper on college admissions and marriage [59], and since then a large theoretical literature has grown from this paper. Matching can be one-sided, i.e., allocating indivisible items among agents, where the items do not have preferences but the agents do. Alternatively, it can be two-sided, i.e., agents from two different sides, such as firms and workers, and they cannot just choose, but also have to be chosen.

In standard matching mechanisms, ordinal preference/priority rankings are submitted by the participants to represent their individual choice, and often this is sufficient to induce desirable or stable outcomes. While there has been much work on the theory and applications of matching, the literature has typically assumed that agents know their preferences before the mechanism is run.[2] Recently, there have been papers that try to relax this stringent assumption [30, 38, 90]. This work can be divided into two main categories: one-shot settings and repeated match settings. In one-shot settings, agents come into the matching setting with unknown (or partially known) true preferences, but can learn more through a costly information acquisition (*interviewing*) stage before the actual matching happens (for example, academic job markets). In repeated matching settings, the "match" is not final, but conveys information to participants on quality (for example, the matching between task requesters and contractors in crowdsourcing platforms [75], or potential mates in a dating market [38]). In this thesis, we focus on the former, one-shot settings.

We are motivated by labor markets. In most labor markets, employers interview potential employees before offering them positions. The interview is an information acquisition stage, where both employers and employees can learn more about their true preferences. Lee and Schwartz [91] proposed what may be the first model of matching with an interviewing stage, in which they ask about the employer's decision of whom to interview, given that interviews are costly and all employers and workers on either side of the market are ex-ante identical. Another recent piece of work on interviewing is that of Rastegari et al., who look at the problem of minimizing the number of interviews while guaranteeing stability and proposer-optimality; they assume that agents have correct partial orderings and use interviewing to refine and complete these partial orderings [121]. Our research is motivated by their models, but the main issue we investigate is different. We consider the role of information ahead of the interview process.

[2]There are many interesting variants where agents know their own preferences, but don't know the preferences of others [125], or where the mechanism does not wish to elicit complete preference information [49], that we will not consider in detail here.

Consider the matching process that academic departments go through when interviewing and hiring faculty candidates. Typically, departments have a budget, say they can interview three or four candidates for a position. They start off the process by receiving a noisy signal about their preferences over candidates – CVs, letters of recommendation, and word-of-mouth can yield much information about candidates, but not nearly as much as an in-person interview. Once they have received these noisy signals, each department chooses which candidates to interview to further form their true preferences. Following Lee and Schwartz, after all the interviews have taken place, we can model the matching process as Gale-Shapley matching with departments submitting ranked lists of the candidates they interviewed. While this ignores some frictions (like exploding offers [41, 99]) that can be important, those are likely to be a second-order effect compared with the choice of candidates to interview.

We study the effects of prior information signals that can be incorporated into firms' interviewing decisions and seek to illuminate how information influences the preference learning and matching process. We are interested in both the overall efficiency of market outcomes and distributional differences in expected outcomes. We show that more commonality in the quality signals can be harmful, yielding fewer matches as some firms make the same mistakes in choosing whom to interview. Relatively high and medium quality candidates are most likely to suffer lower match probabilities. The effect can be mitigated when firms use "more rational" interviewing strategies, or through the availability of private signals of candidate quality to the firms.

1.1.2 Auction Markets and Information Design

The other particular domain where manipulating information is interesting is in auctions with signaling, which have been studied extensively in both economics and computer science. Auction itself is one of the oldest ideas of selling and spans many different domains. For example, Christie's, founded in 1744 and Sotheby's, founded in 1766, use them to sell art. Governments use auctions to sell treasury bills, spectrum, or oil leases. Auction theory was initiated in the seminal 1961 article by William Vickrey [141], which is the first game-theoretic analysis of auctions, and then developed by researchers including Wilson, Clarke, Groves, Milgrom, Weber, Myerson, Maskin, and Riley [15]. There are many different ways of defining auctions, and they promote different kinds of behavior among bidders. Particularly, this thesis analyzes the type called sealed-bid auctions (i.e., static settings) for single items, where all bidders simultaneously submit sealed bids to the auctioneer, ideally without knowing their opponents' bids. The auctioneer unseals the bids and determines a winner, usually the highest bidder. The three most commonly studied sealed-bid formats are:

- *First-price sealed-bid auctions.* The terminology reflects the original format for such auctions, where the highest bidder wins the object and pays the value of her bid.
- *Second-price sealed-bid auctions,* also called *Vickrey auctions.* The highest bidder wins the object and pays the value of the second-highest bid.
- *All pay auctions.* The highest bidder will be awarded the item and every bidder pays her bid.

The underlying assumption of auction models is that each bidder has an intrinsic value for the item being auctioned and she is willing to purchase the item up to this value. The bidders can share a common value, which we refer it as *common value auctions*, but the bidders may have different information about the item's value; for example, the value of an oil-lease depends on how much oil is under the ground, but bidders may have their own experts to estimate the amount. In contrast, in *private value auctions*, each bidder's private valuation of the item is different and independent of peer's valuations. A key feature of auctions is the presence of asymmetric information, where one side, either the buyer or seller, may have better information than the other. Because both sellers and bidders are trying to maximize their utility, signaling begins to play an important role.

The theory of signaling begins with the groundbreaking work of Michael Spence [132] as we mentioned in the beginning of this section, which considers how the employers try to infer the quality of candidates from observable characteristics. Since then, auctions with signaling have been studied in several different contexts. Much of the literature assumes that agents are symmetric with respect to the information they receive about the value of the item, in the sense that the bidders' signals are drawn from the same distribution. For example, the seminal "Linkage Principal" of Milgrom and Weber [107] states that fully and publicly announcing all information available to the seller is the expected-revenue-maximizing policy in common value auctions. Somewhat less is known about auctions with asymmetrically informed bidders, and most of that literature has focused on understanding how information asymmetries affect revenue rather than on the design of the optimal signal structure. There has also been a line of work on so-called "deliberative auctions" [24, 88] in AI domain, where agents have the opportunity to acquire information about valuations before entering a bidding process. Most of this literature focuses on strategic choices by the bidders and how this affects equilibrium outcomes of the auction.

Recently, a popular strand of research has considered the power of signaling in the so-called persuasion model. Kamenica and Gentzkow [77] consider the problem of designing the optimal information environment for the case between one self-interested agent ("sender") and one decision-maker ("receiver"), where both of them are rational Bayesians. The sender can design the *information structure* or *signal structure* to release information about the state of the world to receiver before the receiver makes her choice.

In this thesis, we analyze the signal design of a sealed-bid second price auction as a persuasion game. As usual, the winner is the bidder who submitted the highest bid (with ties broken equiprobably in either direction), but pays to the seller the second highest bid. The bidders and seller share the same common prior on the underlying state of the item. Before the bidding stage, the seller can provide a (noisy) signal to each bidder based on the state of the world. She commits to a signaling strategy in advance, which can be asymmetric for each bidder, and the resultant structure becomes common knowledge. We explore the following two auction games: (1) a basic common-value auction model, where the value of the item is determined either by a single attribute or by two independent attributes when each bidder can receive information from exactly one of the attributes; (2) an interdependent-value auction, where the valuation for each bidder is decided by a common value attribute and a private attribute. We show that in the common-value auction settings, there is no benefit to the auctioneer in terms of expected revenue from sharing information with the bidders, although there are effects on the distribution of revenues. In an interdependent-value model with mixed private- and common-value components, however, we show that asymmetric, information-revealing signals can increase revenue.

Our model contributes to the growing literature on Bayesian persuasion with multiple receivers; this literature usually focuses on public signals [50, e.g.] or symmetric signal structures [39, e.g.]. Our model is applicable to complex situations where the sender of the signal has the opportunity to communicate in different ways to different receivers. This can happen in situations like corporate mergers [19, 123], where targets (sellers or signal senders in our case) have to communicate with potential acquirers (the signal receivers). It is known that targets often inflate their output [64] or themselves may not be aware of their value to an acquirer due to the complexity and intangible characteristics which cannot be easily observed [81].

1.2 Dynamic Markets and Platform Competition

Instead of static settings as we discussed above, many real-world market problems are dynamic, with agents (or both agents and items) arriving and departing over time in a persistent market. Dynamic markets in a *single* market have been explored in many domain-specific applications. We can still consider the categories of one-sided markets and two-sided markets as we did for the static settings. For example, famous public housing assignment problem, which tries to assign scarce public housing to low income households [1, 78, 79, 92]; cadaveric organ allocation, which matches cadaveric organs to patients based on their medical characteristics [20, 133, 150], and general barter markets like kidney exchange [48]. These examples are in the area of one-sided markets, where only one side (the *agents*) has preferences over the other (the *items*). There are also many applications in two-sided markets, and examples of such markets are myriad: online dating (e.g., Match.com and OkCupid) [21]; rideshare, like

Didi Chuxing and Uber [71]; and financial or commodity markets (e.g., NYSE and Nasdaq), where we have buy orders and sell orders.

In dynamic markets, the planner needs to select a subset of acceptable transactions at any point in time. Thus, the natural question is, what kind of algorithm to use for matching? There is greate debate about this in this area across different domains [4, 11, 13, 17, 142, 143]. In this part of our research, we investigate the role of matching algorithms or market-clearing rules in helping inform the debate; Particularly, we focus on two extreme but representative cases: *greedy policy*, which attempts to match each entering agent immediately; and *patient* or *batch policy*, which allows agents to accumulate in the market.

Furthermore, many dynamic applications, like rideshare services, universities, and organ exchanges, involve multiple clearing houses that compete to attract participants, and they may share overlapping pools of agents. These platforms may be self-interested: their eventual goal is to survive in the competitive environment and optimize their profit. The interactions of markets or platforms are understudied by literature, especially in the area of general barter exchange or in the context of matching markets. Accordingly, we focus on quantifying the social welfare and on capturing individual equilibrium behaviors of both agents and markets in the multi-market competition environment. Overall, this part of the thesis can be applied to many important real-world domains, two important examples are (1) kidney exchange and (2) financial markets.

1.2.1 Kidney Exchange

According to the National Kidney Foundation[3], in the last few years, more than 100,000 patients have been waiting for a kidney transplant in the US. In 2014, which is the latest year we have data, about only 17,000 transplants were conducted, and close to one-third of those were from living donors. One major issue with living-donor transplantation is that willing donors must be medically compatible with the patient. Unfortunately, due to ABO blood-type incompatibility and positive crossmatches, some pairs are incompatible. One idea proposed decades ago is to have incompatible pairs enter a kidney exchange [120].

A kidney exchange allows patients who suffer from terminal kidney failure, and have been lucky enough to find a willing but incompatible kidney donor, to swap donors. Apparently, a kidney exchange is a dynamic matching market, where the patient-donor pairs (i.e., agents) arrive gradually over time. They stay in the market until they find a compatible pair unless the patients' situation deteriorates so that kidney transplants are no longer feasible, in which case the agent leaves the market. Ünver [139] was the first to address dynamic kidney exchange, with recent follow-up work by Ashlagi et al. [13] and Anderson et al. [6]. All three papers look at matching policies that aim to maximize (discounted) social welfare (i.e., the sum of expected utility). Particularly relevant to real-world kidney exchanges are *batching* policies, where a market clearing occurs at a fixed interval; several theoretical and empirical explorations of this class of policy have been performed [6,8,13,16]. Learning policies are also designed based on different data distributions or use potential data distribution to inform myopia algorithms [12, 33, 45, 46, 48].

[3] https://www.kidney.org/news/newsroom/factsheets/Organ-Donation-and-Transplantation-Stats

Kidney Exchange and Market Fragmentation

While paired kidney donation of this kind has had success in the United States, a raft of coordination problems and exchange fragmentation has prevented it from accounting for a truly significant fraction of transplants. In practice, kidney exchange accounts or only $10\% - 12\%$ of living donations. The interaction of multiple competing kidney exchanges—a problem that is especially relevant in the US now, and, as kidney exchanges move to international swapping, will soon become relevant worldwide—is little reviewed in the literature, and we seek to fill the gap. In this part of the thesis, we make two contributions:

We address market fragmentation and quantify the social welfare loss directly in these competing markets (Chapter 4) [37]. In the United States, multiple fielded kidney exchanges exist, and patient-donor pairs are entered simultaneously into one or more of these markets, based on geographical location, travel preferences, home transplant center preferences, or other logistical reasons. Individual kidney exchange clearinghouses have the incentive to compete on the number of matches performed within their specific pools. We explore the effect of competition between exchanges with different matching policies on global social welfare in the context of the number of matched patients.

We formalize a two-market model where agents enter one market or both markets stochastically; they can then be matched to other agents who have joined the same market or both markets. The markets adhere to different matching policies, with one matching greedily (Greedy market) and the other building market thickness through a policy of patience (Patient market). From both theoretical support and experimental evidence, we show that market fragmentation caused by the competition leads to worse global loss than a single market.

We then provide the first analysis of equilibrium behavior in this competing market (Chapter 5) [98]. Akbarpour et al. [4] show that platforms may maximize the number of transplants achieved by being *patient* (Patient policy) instead of trying to match new pairs immediately (Greedy policy). The intuition is that waiting helps the market become thicker. Thus, Patient market is actually more socially preferable. However, any given individual almost certainly seeks to maximize her own utility instead of considering the social welfare of the market. Now we approach this market competition problem from a game-theoretic point of view—under what circumstance and which type of agents have the incentive to participate in Patient or Greedy market? Similarly, because markets seek to maximize their own utility at a potential cost to overall social welfare, how should they adapt their matching rates?

We utilize the above two-market model, and first allow agents to strategically choose a market, given the knowledge of their own criticality. Our model considers two types of agents in terms of criticality, *short-lived* and *long-lived*. An agent receives zero utility if she perishes. If she is matched, she receives a utility of 1, discounted at rate δ. Thus, the market choice is actually a tradeoff between matching probability and utility. That is, entering a Patient market gives an agent a higher matching probability but lower utility as the patient's situation deteriorates during waiting; in contrast, immediate matching from a Greedy market provides a higher utility but may lower the probability of matching since the market is not thick enough.

Second, we prescribe agency to the markets themselves, allowing them to choose overall matching policies (defined by the frequency at which they decide to match) strategically to maximize their overall utility. In this case, the agents are stochastic in their choice to join one or the other market or to enter both markets. We quantify via best response dynamics the social welfare loss of this competitive marketplace under a variety of initial conditions and compare that loss to the lower bound provided by a single market running an optimal matching policy.

Our work is among the first to study strategic issues in market/platform competition in the context of matching markets. Matching markets of the kind we study here are quite distinct from markets for securities or for other kinds of products, because in our setting the value of a matching is idiosyncratic (to the pair matched) and utility is nontransferable [28]. Therefore the standard price mechanism is unavailable as one of the levers available to the platforms or the agents to change outcomes, and the matching policy becomes of central importance. These restrictions (no money changing hands and nontransferable utility) are necessary for modeling domains like kidney exchange (where exchanging money is prohibited by law and utility can be a function of waiting time as well as kidney quality, although in this part we focus on a model for the former), or dating (where it would be considered problematic for a dating app to pay users to go on dates with certain other members).

Our work applies techniques from computational game theory to studying platform competition in the context of dynamic matching markets. An established model of matching due to Akbarpour et al. [4] has found that greedy matching (making matches as soon as they are possible) can lead to worse social outcomes. If one defines social loss in terms of the additive inverse of the waiting-time discounted number of matched pairs, then adopting a Greedy strategy can result in exponentially worse loss than a Patient strategy for sufficiently low discount rates. We extended that model to quantify the costs of market fragmentation when greedy and patient markets compete. The central insight in both cases is that thinner markets lead to fewer (or poorer quality) matches. In this thesis, we show that the thinness problem actually gets worse when either agents are strategic about market entry or markets are strategic about choosing matching frequency/cadence. The existence of greedy and patient alternatives can lead to patients with privation information about their type into separating equilibria that further fragment markets, leaving everyone worse off than if there were a single monopolist market (even if that market were greedy rather than patient!). The presence of another market that could be greedy can lead a market to choose a greedy matching policy, even when it would have been better off choosing a patient policy if it were a monopolist.

The above results are not merely theoretical concerns. In the United States, for example, two of the largest kidney exchanges are the National Kidney Registry (NKR) and the United Network for Organ Sharing (UNOS). NKR matches in an essentially greedy fashion. UNOS started by matching once per month, then moved to twice per month, then weekly, and now 2+ times per week, in part to reduce the "failure rate" caused by competition with the fast-matching NKR. We see this behavior replicated in our model, and can quantify social welfare loss as well. Combinations of analytic and simulation results of this nature have set policy in kidney exchanges before (e.g., [47] and [48] have set parts of UNOS policy), and our model could help inform this debate.

Utility Design and Incorporating Compatible Pairs in Kidney Exchange

To improve the performance of kidney exchange, instead of solving the market fragmentation problem directly, one proposal for transplanting more recipients from incompatible pairs has been to incorporate *compatible* pairs into exchanges. This idea was first proposed by Gentry *et al* [61], but it has not been studied much. Part of the reason is that it is very tough to come up a reason or a quantitive measure for the compatible pairs to participate in a kidney exchange instead of just transplanting directly with their compatible donors. The recent development of new metrics for the quality of a living donor transplant [104] presents an opportunity to reassess the possible benefits in the context of realistic models of compatible pair behavior, while also evaluating benefits in terms of both additional transplants made possible and improved outcomes from transplants. Further, it is reasonable to believe that compatible pairs may be more willing to enter exchanges if (1) their waiting times are kept low, and (2) they have a more precise idea of the potential benefit to doing so.

In this part of the thesis (Chapter 6), our main contributions are two-fold: (1) We first present a framework for studying kidney exchange in a weighted or cardinal utility setting, which can directly present how much benefit a compatible pair can receive in terms of long-time graft survival; (2) Then we use this framework to estimate the benefit of including compatible pairs in kidney exchange. Using data from Barnes Jewish Hospital in St. Louis, Missouri, we develop a novel simulator that generates realistic distributions of graft survival (based on the recent introduction of the Living Kidney Donor Profile Index [104]) and combine this with a well-known compatibility simulator [127] in a manner that is faithful to data on real arriving pairs. We use our simulator across different matching mechanisms to estimate both the increased numbers of transplants of incompatible pairs (almost doubling the number transplanted) and the improved match quality for recipients in compatible pairs (increasing expected graft survival by between 1 and 2 years). Our results are robust across several different exchange sizes in the static setting, in dynamic settings where compatible pairs must be immediately matched, and across assumptions about incompatible to compatible pair ratios. The results are also promising for hard-to-match subpopulations, including blood group O recipients and highly sensitized patients.

1.2.2 High-frequency Trading and Competition Between Financial Exchanges

Another interesting line of dynamic markets is financial markets. Most modern financial exchanges operate using the continuous double auction (CDA) mechanism, a greedy fashion mechanism, which in principle allows for trading in continuous time, at least to within our measurement and implementation capabilities [58]. In this kind of market, agents submit bids, or *limit orders*, which represent the maximum price at which they would like to buy, or the minimum price at which they would like to sell. Outstanding orders are maintained in two priority queues: one for bids (the buy orderbook) and one for asks (the sell orderbook). Bids and asks are prioritized first by price and second by time. When a new order comes in, it is added to the corresponding order book. A trade is executed immediately if the highest bid exceeds or is equal to the lowest ask. The execution involves the orders at the top of the bid and ask queues, at the price of the older of the two orders involved.

In the last two decades, The existence of this continuous time feature has led to the development of the phenomenon called High-frequency trading. Essentially people are competing to get to market a little bit quicker than someone else, because having small of time advantage can make a big difference. With companies keeping investing in faster infrastructure for trading, and events like the "flash crash" of May 2010, high frequency trading (HFT) has become an increasingly debated topic in both the media and policy spheres [93]. Proponents claim that high-frequency trading improves liquidity and price discovery. Improved liquidity means lower transaction costs for average investors, while better price discovery serves the social information aggregation and dissemination role of market prices [26, 106]. However, there is increasing evidence that at least one form of high frequency trading, namely latency arbitrage, has reached a point of socially diminishing returns. Budish et al demonstrate this both empirically and through a simple model: empirically, they show that correlations between virtually identical assets being traded in different markets break down at very small timescales, while they are essentially perfect at larger timescales [27]. Correlation breakdown can almost be thought of as a law of physics – there is no natural force tying the assets or markets together, so there is no way to make them actually move simultaneously. What is problematic is the "arms race" this creates to extract the maximum profit from squeezing this reaction time down as much as possible. Budish et al. show that this is not only socially inefficient, it can actually create thinner markets. Along similar lines, Wah and Wellman build a model where an asset is traded on two markets, and there is an infinitely fast latency arbitrageur present. They show how the presence of the arbitrageur can hurt social welfare [143]. Both sets of authors recommend frequent batch auctions as a market structure that could replace CDAs, since the minimum time period between trades is specified, and there is no benefit to being faster than that.

An important question for the possible use of frequent batch (or call) auctions is how they would work in the presence of existing CDA markets. Competition between exchanges or platforms that try to attract trade is a vast topic, and there is evidence in many domains that platforms with better welfare properties *assuming that there is only one platform is considered at a time* may not be able to capture enough of the market for these properties to become evident when they face competition from other platforms. For example, as what we showed in living-donor paired kidney exchange, even though exchanges that wait to build thickness may be socially preferable, exchanges that match greedily can make them non-viable. Therefore, even though they may have desirable welfare properties, could call auction based markets actually take volume away from CDA markets if both existed simultaneously? Wah et al have engaged this question using empirical game theoretic analysis [142]. They develop a model where the environment is populated by fast (HFT) and slow (non-HFT) traders. They argue that a frequent call market in the wild could attract sufficient volume for viability from two perspectives: first, in equilibrium, welfare of slow traders is generally higher in the call market, where they are relatively protected from sniping and adverse selection, and second, fast traders are willing to follow the slow traders to either market, including to the call market, so it could serve as a basin of attraction. Wah et al.'s model does not consider traders who have a preference for immediacy, and it also restrict traders to choose a single market and then do not allow traders to move. While their results are quite promising, we seek to build a richer model that combines aspects of classic financial market microstructure models and agent-based models that are known to replicate important properties of order books.

Another line of literature relates to the Trading Agent Competition Market Design Competition (CAT) [111]. In this competition, participants aim to design better mechanisms to maximize a score (a combination of profit, market share and transaction success rate) when traders are drawn from a known population of different types. CAT gives a general view of competition among different markets, but we focus on a comparison of two more specific market mechanisms and how they influence the social welfare of traders.

In this part of the thesis (Chapter 7 [95]), we contribute to this nascent literature by developing an agent-based model of competition between a Call market and a CDA market. Agent-based modeling seeks to fill the hole in simple stylized models which may not represent agent behavior in sufficiently complex manners to really capture the essence of the important phenomena. The last two decades have seen substantial work on agent-based modeling of financial markets, using both sophisticated [32,52,57,114] and simple [56,80,89] trader models in the population. Our model is as parsimonious as possible while attempting to capture the essential relevant behaviors that are important to understanding the behavior of these markets. As such, it follows the basic structure of classic models of market microstructure such as those of Glosten and Milgrom [66] and of Kyle [85]. In these models, there are informed traders, who possess superior information and trade in search of profit, liquidity (or background) traders, who trade for exogenous reasons (e.g. retirement funds that receive cash and need to track indices, or investors liquidating portfolios in retirement or in order to buy a house, say) and demand immediacy, and market makers, who may be employed in order to facilitate price discovery and trade execution.

We show that there is a strong tendency for the Call market to absorb a significant fraction of trade under most equilibrium and approximate-equilibrium conditions. These equilibria typically lead to significantly higher welfare for the background traders, an important measure of social value, than the operation of an isolated CDA market.

1.3 Contributions and Structure of the Thesis

Part I of the thesis studies static markets and information design, which consists of Chapter 2 and Chapter 3.

Chapter 2 addresses the role of information in matching markets. We study two-sided matching markets where the matching is preceded by a costly interviewing stage in which firms acquire information about the qualities of candidates. Our focus is on the impact of the signals of quality available prior to the interviewing stage. Equilibrium interviewing decisions are hard to characterize in complex models with differentiated quality, so we use a mixture of simulation, numerical, and empirical game theoretic analysis to analyze social outcomes. We show that more commonality in the quality signals can be harmful, yielding fewer matches as some firms make the same mistakes in choosing whom to interview. Relatively high and medium quality candidates are most likely to suffer lower match probabilities. The effect can be mitigated when firms use "more rational" interviewing strategies, or through the availability of private signals of candidate quality to the firms.

Chapter 3 investigates the information design in auction markets. We consider the problem of designing the information environment for revenue maximization in a sealed-bid second price auction with two bidders. Much of the prior literature has focused on signal design in settings where bidders are symmetrically informed, or on the design of optimal mechanisms under fixed information structures. We study common- and interdependent-value settings where the mechanism is fixed (a second-price auction), but the auctioneer controls the signal structure for sellers. We show that in a standard common-value auction setting, there is no benefit to the auctioneer in terms of expected revenue from sharing information with the bidders, although there are effects on the distribution of revenues. In an interdependent-value model with mixed private- and common-value components, however, we show that asymmetric, information revealing signals can increase revenue.

Part II of the thesis contributes to dynamic markets and platform interactions or competitions, and consists of Chapter 4, Chapter 5, Chapter 6, and Chapter 7. Chapter 4, Chapter 5 and Chapter 6 focus on kidney exchange, and 7 studies financial markets.

Chapter 4 and Chapter 5 investigate dynamic matching market competition. While dynamic matching markets are usually modeled in isolation, assuming that every agent to be matched enters that market, in many real-world settings there exist *rival* matching markets with overlapping pools of agents. We extend the framework of dynamic matching due to Akbarpour et al. [4] to characterize outcomes in cases where two such rival matching markets compete with each other. One market matches quickly, while the other builds market thickness by matching slowly. We give an analytic bound on the loss—the expected fraction of unmatched vertices—of this two-market environment relative to one in which all agents enter either one market or the other, and numerically quantify its exact loss, demonstrating that rival markets increase overall loss compared to a single market that builds thickness. We then look at two competing kidney exchanges, where patients with end-stage renal failure swap willing but incompatible donors, and show that matching with rival barter exchanges performs qualitatively the same as matching with rival matching markets—that is, rival markets increase the global loss. We also provide the first analysis of equilibrium behavior in dynamic competing matching market systems—first from the points of view of individual participants when market policies are fixed, and then from the points of view of markets when agents are stochastic.

To improve the performance of kidney exchange in terms of both social welfare and individual utility, Chapter 6 analyzes the benefit of convincing directed donation pairs to participate in paired kidney exchange. This possibility has been relatively understudied by literature. Possibly, incorporation of compatible pairs in exchanges has not taken off because the potential benefits to recipients in compatible pairs have been pooly qualified. The recent introduction of the Living Donor Kidney Profile Index (LKDPI), which can be transformed into an expected survival time for the graft, presents an opportunity to better estimate the potential benefits, and to present compatible pairs with a compelling medical reason to participate in an exchange rather than proceeding with a direct donation. Using data from Barnes Jewish Hospital, we develop a novel simulator for LKDPIs that generates realistic distributions of graft survival, and we combine this with a well-known compatibility simulator in a manner that is faithful to data on real arriving pairs. We use our simulator across different matching mechanisms to estimate both the increased numbers of transplants of incompatible pairs (almost doubling the number transplanted) as well as the improved match quality for recipients in compatible pairs (increasing expected graft survival by between 1 and 2 years). Our results are robust across several different exchange sizes in the static setting, dynamic settings where compatible pairs must be immediately matched, and across assumptions about incompatible to compatible pair ratios. The results are also promising for hard-to-match subpopulations, including blood group O recipients and highly sensitized patients.

Chapter 7 studies dynamic markets from the view of financial markets. In the debate over high-frequency trading, the frequent call (Call) mechanism has recently received considerable attention as a proposal for replacing the continuous double auction (CDA) mechanisms that currently run most financial markets. We examine agents' profit under CDA and frequent call auctions in a dynamic environment. Another natural question, which has begun to spur the development of new models, is the effect of competition between platforms that use these two different mechanisms when agents can strategize over platform choice. We contribute to this nascent literature by developing an agent-based model of competition between a Call market and a CDA market. Our model incorporates patient informed traders (both high-frequency and not) who are willing to wait for order execution at their preferred price and impatient background traders who demand immediate execution. We show that there is a strong tendency for the Call market to absorb a significant fraction of trade under most equilibrium and approximate-equilibrium conditions. These equilibria typically lead to significantly higher welfare for the background traders, an important measure of social value, than the operation of an isolated CDA market.

Chapter 2

Matching Markets and the Role of Information

In this chapter, we investigate the role of information in matching markets. Specifically, we focus on labor markets with interviewing. In most labor markets, employers interview potential employees before offering them positions. The interview is an information acquisition stage, where both employers and employees can learn more about their true preferences. Lee and Schwartz proposed what may be the first model of matching with an interviewing stage, where employers first simultaneously choose a subset of workers to interview, and then, in a second stage, submit preferences to a (Gale-Shapley) matching algorithm that then forms the matching [90]. An interview is a precondition for a possible matching to be formed between an employer and a worker. The basic question that Lee and Schwartz ask is about the employer's decision of whom to interview, given that interviews are costly and all employers and workers on either side of the market are ex ante identical. The main complexity is then that the marginal benefit of interviewing a worker goes down as her number of other interviews goes up. The major result is that in symmetric equilibria (where each employer and worker has the same number of interviews), the number of agents matched goes up in the overlap, a measure characterizing the number of common interview partners among agents.

In this chapter, we follow Lee and Schwartz's model but consider a different issue. we are interested in situations where firms and workers are of different qualities instead of ex ante identical, and some quality signals are available prior to the interviewing stage. We study the effects of prior information signals that can be incorporated into firms' interviewing decisions. We are interested in both the overall efficiency of market outcomes and distributional differences in expected outcomes.

In order to elucidate these issues, we look at a stylized model where there is a universally shared, common knowledge ranking of all firms, and there is a "true" universally shared ranking of all candidates, but this true ranking is not known – instead, firms receive different signals of candidates' rankings or qualities. If the true ranking were known to everyone, there would be only one stable matching, the assortative one, and any rational interviewing process would lead to the stable outcome in the matching stage. When signals of quality or ranking are noisy, firms must reason both about the true quality of candidates and about strategic issues in deciding whom to interview. This can lead to inefficiencies, where some candidates andfirms do not end up getting matched whereas they would have with better information; these inefficiencies may fall disproportionately on some portion of the population of candidates and firms.

We are particularly interested in the roles of *common* and *private* information on aggregate and distributional outcomes in such matching markets. Common signals are shared across firms – for example, the quality of a CV, number of publications, LinkedIn endorsements, or public contributions to open source projects, can all be thought of as common signals of varying precision. A private signal is, as the name suggests, private to a particular firm. Private signals can be generated through phone screens, preliminary interviews, etc. We assume that common and private signals are conditionally independent given the true ranking or value of the candidate. The central question of this chapter is the effect of the relative precision of common signals and private signals on market outcomes. While a perfect common signal would reduce the problem to one with known rankings of both firms and candidates (and lead to the assortative matching and no inefficiencies under any reasonable model), our main finding is that the presence of a strong, but imperfect, common signal in addition to existing private signals can actually have significant negative effects, with fewer matchings occurring than with a private signal alone. The burden of this is typically borne by the candidates who are ranked relatively high (but not in the highest echelon). The mechanism is interesting – when these candidates end up with a common signal that is "too high", they interview at firms that are ranked too high for their actual quality. The firms that are closer to their true range choose not to interview them, but when these candidates' true qualities are revealed, they often don't get offers from the places that did interview them.

In this chapter, we start from introducing the formal matching model and the various models of quality signals that we consider. All the interviewing strategies we develop are predicated on the posterior belief of each firm after receiving the common signal and the private signal of employee qualities, so in Section 2.2 we describe the inference procedures necessary to compute these posterior beliefs. In Section 2.3, we analyze matching outcomes when all firms use the heuristic interviewing strategy of interviewing applicants "around" their own rank. Section 2.4 develops more sophisticated strategies in a sequential setting, and examines the distributional and aggregate effects on matching outcomes when firms use these strategies.

2.1 Model

There are n workers and n firms, represented by the sets $W = \{w_1, ..., w_n\}$ and $F = \{f_1, ..., f_n\}$. The matching market operates in two stages, following the model of Lee and Schwarz [90]. In the first stage, each firm selects k workers (or candidates) to interview; this decision is made on the basis of information present in the signals received by firms (described below). During the interview process, the true ranking of the set of candidates that is interviewed is revealed to each firm. The second stage can then be thought of as a Gale-Shapley matching where each firm submits a ranked list of the candidates it interviewed (others are unacceptable), and each candidate submits a ranked list of firms.

2.1.1 Signals and Preferences

All workers know their preference rankings over employers with certainty. We assume that the workers all have exactly the same preferences over potential employers (for example, all workers rank departments solely on the basis of the US News and World Report program ranking).[4] Further, there exists a universal "true" ranking of all the workers as well, but this ranking is unobserved. Employers receive a private signal of their preferences as well as a common signal. In this chapter we consider two possibilities:

1. **Random-utility models:** w_i has a true value v_i (which is drawn from a normal distribution). f_j's private signal is a vector $\mathbf{s_j} = (s_1, s_2, \dots s_n)$. Each $s_i, 1 \leq i \leq n$ is a noisy realization of the true value of v_i. We consider two noise distributions, Gaussian and uniform. For Gaussian noise, $s_i \sim \mathcal{N}(v_i, \sigma_p)$, where $\mathcal{N}(v_i, \sigma_p)$ denotes the Gaussian

[4] Because of this, who proposes in the second stage becomes irrelevant for the rest of this chapter, since either side proposing would yield the same outcome.

density function with mean v_i and standard deviation σ_p. σ_p is constant for all i. For uniform noise, $s_i \sim \mathcal{U}(v_i - b_p, v_i + b_p)$, where $\mathcal{U}(v_i - b_p, v_i + b_p)$ denotes the uniform distribution with support on $[v_i - b_p, v_i + b_p]$. b_p is also constant for all i. The common signal, received by all employers, is a single vector $\mathbf{z_C} = (z_1, z_2, \ldots z_n)$. Similar to the realizations for the private signal, the realizations for the common signal are also noisy, with $z_i \sim \mathcal{N}(v_i, \sigma_C)$ in the case of Gaussian noise and $z_i \sim \mathcal{U}(v_i - b_C, v_i + b_C)$ in the case of uniform noise.

2. **Mallows model:** We can also directly consider signals over the ranking space, instead of the value space. The *Mallows model* [103] is a distance-based model, which defines the probability of a permutation according to its distance to a modal permutation [101]. Following Lu and Boutilier's description of its form, we say that each employer's private signal is a *ranking* Γ_j sampled from the distribution which assigns $P(\Gamma_j|\Gamma, \phi_p) = \frac{1}{Z}{\phi_p}^{d(\Gamma_j,\Gamma)}$, where Γ is the modal ranking (which in our case is the true ranking), $\phi_p \in (0, 1]$ is a dispersion parameter such that the smaller ϕ_p is, the more the distribution will be concentrated around the modal ranking, d is a distance function between rankings (in our case the classic Kendall tau distance which counts the pairwise disagreements between the two rankings), and Z is a normalizing factor. The common signal, Γ_C is sampled from a Mallows model with the same modal ranking Γ and a possibly different dispersion parameter ϕ_C.

In both cases, we assume common knowledge of all the relevant parameters of the distributions; the only unknowns are the true values or rankings. Note that in all cases, if the true rankings were known, the only stable matching is assortative, with the best worker getting matched to the top employer, the second-best to the second-best, and so on.

2.2 Inference

In order to describe interviewing strategies and outcomes, we first need to specify appropriate inference techniques for firms to compute posteriors given the private and common signals they receive. We denote posterior density functions on values of workers as $f_v(\cdot)$. In both the random utility and Mallows models, it is computationally difficult to perform full Bayesian reasoning over the whole space of possible posterior rankings, so we assume that firms compute the single most likely posterior ranking from the common and private signals (defined explicitly below in either case), which we denote as $\widetilde{\Gamma}_j$, and use this single ranking for interviewing decisions.

2.2.1 Inference in the random utility models

The main ideas for combining different signals in the random utility models follow from those developed by MacQueen [102]. Given f_j's private signal $\mathbf{s_j} = (s_1, s_2, \ldots, s_n)$, and common signal $\mathbf{z_C} = (z_1, z_2, \ldots z_n)$, the posterior on w_i's value v_i is given by Bayes' rule:

$$f_v(v_i|s_i, z_i) = \frac{f_v(s_i|v_i) f_v(z_i|v_i) f_v(v_i)}{\int_{-\infty}^{+\infty} f_v(s_i|v_i) f_v(z_i|v_i) f_v(v_i) dv_i}.$$

When the noise is Gaussian, $v_i \sim \mathcal{N}(\mu, \sigma)$, $s_i \sim \mathcal{N}(v_i, \sigma_p)$ and $z_i \sim \mathcal{N}(v_i, \sigma_C)$. Thus, the expected value of v_i is

$$E(v_i|s_i, z_i) = \frac{\frac{s_i}{\sigma_p{}^2} + \frac{z_i}{\sigma_C{}^2} + \frac{\mu}{\sigma^2}}{\frac{1}{\sigma_p{}^2} + \frac{1}{\sigma_C{}^2} + \frac{1}{\sigma^2}}.$$

When the noise is uniformly distributed, $v_i \sim \mathcal{N}(\mu, \sigma)$, $s_i \sim \mathcal{U}(v_i - b_p, v_i + b_p)$, and $z_i \sim \mathcal{U}(v_i - b_C, v_i + b_C)$. The expected value v_i is

$$E(v_i|s_i, z_i) = \int_D \frac{v_i f(v_i)}{\int_D f(v_i) dv_i} dv_i,$$

where D is the intersection of $[s_i - b_p, s_i + b_p]$ and $[z_i - b_C, z_i + b_C]$. In both cases, the posterior ranking $\widetilde{\Gamma}_j$ is found by sorting the $E(v_i|s_i, z_i)$ in descending order.

2.2.2 Inference in the Mallows model

For inference in the Mallows model, we use an algorithm based on the one devised by Qin *et al* [117] in the coset-permutation distance based stagewise (CPS) model (which is equivalent to the Mallows model using the Kendall tau distance). Here we describe their model in terms of our problem (our exposition below follows theirs, adapted to our domain). W is the set to be ranked. A ranking π is a bijection from W to itself; $\pi(i)$ denotes the rank of w_i and $\pi^{-1}(i)$ denotes the worker assigned to position i. The bracket alternative notation is also used to represent a permutation, i.e., $\pi = \langle \pi^{-1}(1), \ldots, \pi^{-1}(n) \rangle$. Let S_n denote the symmetric group of order n (a non-Abelian group under composition). The right coset $S_{n-k}\pi = \{r\pi | r \in S_{n-k}\}$ is a subset of permutations whose top-k objects are exactly the same as in π (here S_{n-k} denotes the subgroup of S_n consisting of all permutations whose first k positions are fixed: $S_{n-k} = \{\pi \in S_n | \pi(i) = w_i, \forall i = 1, ..., k\}$). The *coset-permutation distance* is a measure of the average distance between the permutations in the coset and the reference permutation. Given a permutation distance d (we use the Kendall tau distance), the coset-permutation distance $\hat{d}$ from a coset $S_{n-k}\pi$ to a target permutation r is $\hat{d}(S_{n-k}\pi, r) = \frac{1}{|S_{n-k}\pi|} \sum_{\tau \in S_{n-k}\pi} d(\tau, r)$, where $|S_{n-k}\pi|$ is the number of permutations in set $S_{n-k}\pi$. The CPS model defines the probability of a permutation π conditioned on a dispersion parameter $\phi \in (0,1]$ and a reference permutation r as,

$$P(\pi|r,\phi) = \prod_{k=1}^{n} \frac{\phi^{\hat{d}(S_{n-k}\pi,r)}}{\sum_{j=k}^{n} \phi^{\hat{d}(S_{n-k}(\pi,k,j),r)}},$$

where $S_{n-k}(\pi,k,j)$ denotes the right coset including all the permutations that rank workers $\pi^{-1}(1),...,\pi^{k-1}(k-1)$ and $\pi^{-1}(j)$ in the top k positions respectively. When the coset-permutation distance in the CPS model is induced by the Kendall tau distance, the CPS model is mathematically equivalent to the Mallows model defined with the Kendall tau distance.

We apply the *sequential inference* algorithm (shown as Algorithm 1) of [117] to get a single posterior ranking $\widetilde{\Gamma}_j$. This algorithm approximates the single highest probability posterior ranking conditioned on the input rankings. The algorithm decomposes the inference into n steps. At the kth step, it selects worker w_i who minimizes the coset-permutation distance,

$$\sum_m (-\ln(\phi_m))\hat{d}(S_{n-1}(\langle \widetilde{\Gamma}_j^{-1}(1),...,\widetilde{\Gamma}_j^{-1}(k-1),w_i\rangle),\Pi_m),$$

and puts this worker at the kth position.

2.3 Market Outcomes With a Simple Interviewing Strategy

We first examine outcomes in a market where firms all use the same simple and intuitive interviewing strategy. They each compute their posterior ranking based on the available

ALGORITHM 1: Sequential Inference for Posterior Ranking (from [117])

Input: : W, input rankings $\mathbf{\Pi}$ where Π_1 is the private signal, Π_2 is the common signal, and parameters $\boldsymbol{\phi}$ where ϕ_1 is the private signal dispersion parameter and ϕ_2 is the common signal dispersion parameter.

Output: the final ranking $\widetilde{\Gamma_j}$.

$\widetilde{\Gamma_j}^{-1}(1) = \arg\min_{w_i \in W} \sum_m (-\ln(\phi_m))\hat{d}(S_{n-1}(\langle w_i \rangle), \Pi_m)$;

Remove worker $\widetilde{\Gamma_j}^{-1}(1)$ from set W;

repeat

$\widetilde{\Gamma_j}^{-1}(k) = \arg\min_{w_i \in W} \sum_m (-\ln(\phi_m))\hat{d}(S_{n-1}(\langle \widetilde{\Gamma_j}^{-1}(1), ..., \widetilde{\Gamma_j}^{-1}(k-1), w_i \rangle), \Pi_m)$;

Remove worker $\widetilde{\Gamma_j}^{-1}(k)$ from set W;

until $W = \emptyset$;

signals, and then interview the k candidates who are ranked "around" the firms own ranking. So, suppose $k = 5$, then the firm ranked number 11 will interview the candidates it ranks in positions 9 through 13. Firms at the top and bottom of the firm rankings adjust their interview set downwards and upwards respectively (so, the top three ranked firms all interview candidates 1-5, and the bottom three all interview candidates 26-30, although the particular candidates occupying these ranking positions can be different for each firm, since they may be based on a posterior computed using private information).

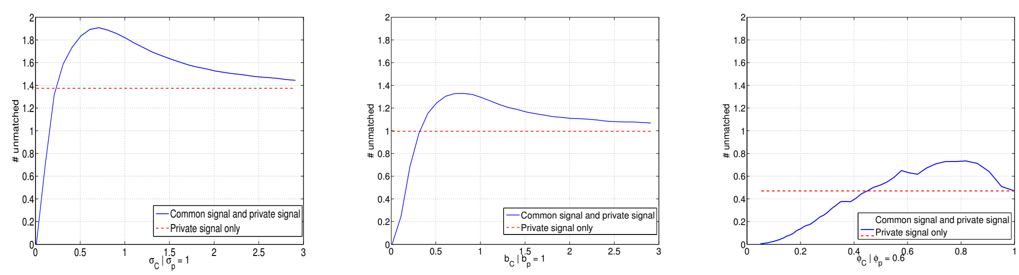

Figure 2.1: Simulation results with 30 employers, 30 workers and an interview budget of 5 for each firm. The graphs show the average number of agents left unmatched (Y axis) versus a decreasing function of the precision of the common signal (σ_C for Gaussian noise (**left**), b_C for uniform noise (**middle**), and ϕ_C for the Mallows model (**right**)), holding the precision of private signals fixed. The dashed line shows the number that are left unmatched when there is no common signal.

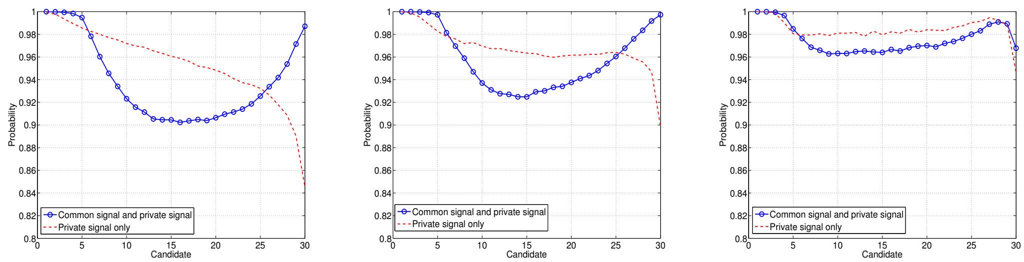

Figure 2.2: The probability that the candidate of a particular rank is matched when firms have access to both a common signal and a private signal. Left: Gaussian noise ($\sigma_C = 0.6, \sigma_p = 0.5$), Center: Uniform noise ($b_C = 0.6, b_p = 0.5$), Right: Mallows model ($\phi_C = 0.7, \phi_p = 0.6$).

In order to study market outcomes, we run 50000 simulations for each of the random utility and Mallows models; each simulation is of a market with 30 firms and 30 workers, each with interview budget 5. In each run, we hold the private signal parameters fixed, which are σ_p, b_p in the random utility models and ϕ_p in the Mallows model, and vary the common signal parameters, which are σ_C, b_C in the random utility models and ϕ_C in the Mallows model.

2.3.1 Analysis

Based on the observation that the only stable matching if true preferences were known is the assortative matching, and that adding a common signal gives everyone more information about the true ranking, one would assume that adding the common signal always leads to more agents being matched. At the extreme, this is obvious – suppose the common signal had no noise and contained perfect information. Then the rational inference is just to use that signal. In this case, the assortative match would occur for sure.

But it turns out that, as the signal becomes less precise, the number of unmatched agents goes up sharply, and quickly exceeds the expected number of unmatched agents when no

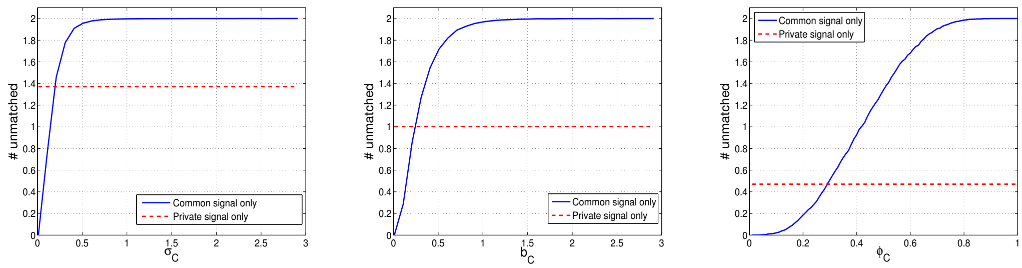

Figure 2.3: The average number of agents left unmatched when firms have access to only a single common signal. The X axis in each graph is a decreasing function of the precision of the common signal. Left: Gaussian noise, Center: Uniform noise, Right: Mallows model.

common signal is present! Surprisingly, on the candidates' side, the candidates who are less likely to get matched are actually the higher ranked ones (except for the very top ranked ones) (see Figures 2.1 and 2.2). The typical case for such a candidate being left unmatched is when the candidate gets a common signal that is too high. Then the candidate interviews at firms that are ranked too high for their actual quality. The firms that are closer to their true range choose not to interview them, but when these candidates' true qualities are revealed, they often don't get offers from the places that did interview them. The truth-revealing nature of the interview phase means that it can be disadvantageous to "place too high" in the first (interview selection) stage. These effects are even more extreme when there is only a single, common signal available to the firms for all candidates (see Figures 2.3 and 2.4). In order to gain a little more insight into this process, we look at a simpler example.

2.3.2 The Case of 4 Firms and 4 Workers

In order to investigate further, we look at a simpler case, with only four employers and four workers, and a slightly different interviewing strategy – in this case, employers interview the 3 candidates ranked from their rank down (so Employer 2 would interview the candidates it

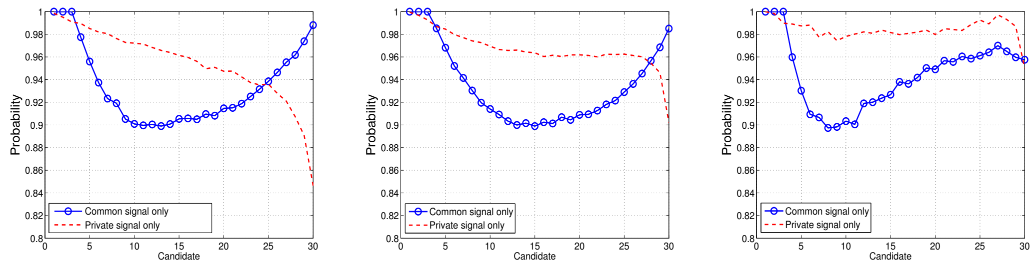

Figure 2.4: The probability that the candidate of a particular rank is matched when firms have access to only a single common signal. Left: Gaussian noise ($\sigma_C = 0.6$), Center: Uniform noise ($b_C = 1$), Right: Mallows model ($\phi_C = 0.6$).

ranks from 2-4). We focus in this part on only the Mallows model. These changes do not affect any of the results substantially, but allow us to obtain exact numerical results, rather than simulation results, and the simpler model yields insight into the basic properties of the larger markets above. With only 4 firms and 4 candidates, we can find the exact probability that each candidate, ranked from 1-4, is left unmatched by breaking the probability up into components. For example, for Candidate 2 to remain unmatched, it must be the case that (1) Candidate 1 was interviewed by Employer 1, and, (2) Candidate 2 was not interviewed by any of Employers 2-4. This is because, once a candidate is interviewed, their true ranking is revealed, so if Candidate 2 interviewed with any of Employers 2, 3, or 4 (and Candidate 1 went to Employer 1), then Candidate 2 would be the highest ranked for any of those Employers and would match with them. But for Candidate 2 to not be interviewed by Employers 2-4, they must *all* have ranked her as the top candidate based on the private and common signals they received.

These probabilities can all be efficiently exactly computed in the Mallows model. Figure 2.5 shows the probabilities that each of the candidates gets matched. Interestingly, the probability of candidate 2 being matched is lower than the probabilities of candidates 3 and 4 being matched. The effect is stronger with just a common signal (left graph), and the

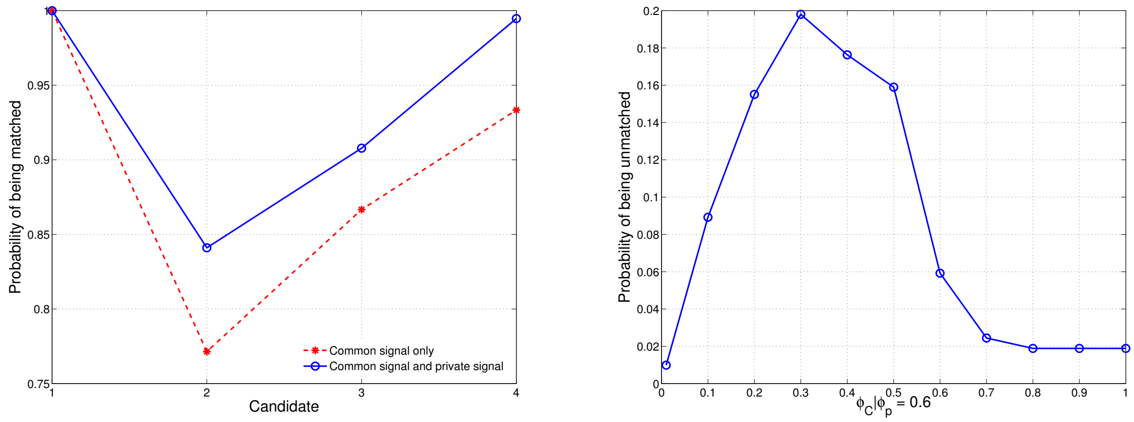

Figure 2.5: Left: The probabilities that each of candidates 1-4 is matched with a single common signal (dashed red line), and with both common and private signals (solid blue line). Right: The probability that candidate 2 is left unmatched as a function of ϕ_C (the common signal becomes less accurate as ϕ_C increases), holding ϕ_p fixed.

probability that candidate 2 remains unmatched decreases as the strength of the common signal declines with respect to the strength of the private signal (right graph). This helps us understand the mechanism at play. In a more coordinated environment, as created by a common signal, the correlation between employers' estimates of a workers desirability is higher. Thus, it is more likely that several employers *all* make the mistake of thinking a particular worker is too good or too bad for them.[5] In fact, when there is *only* a common signal, the probability of Candidate 2 being matched in the 4 agent model drops to less than 0.8 with $\phi_C = 0.5$. When opinions are more independent, as is the case when the private signal is stronger, it is less likely that someone will fall through the cracks in this manner. Therefore, more homogeneity of opinion, with even a little bit of noise, can create worse outcomes!

[5] For example, suppose a middle-ranked candidate gets early "buzz" on the job market, he may not get interviews from departments actually ranked in his vicinity because they think he is out of reach, but may not get offers once he is interviewed by higher ranked places and they realize he isn't quite at their level.

2.4 Alternative Interviewing Strategies

The results in the previous section apply to one simple strategy, albeit one, which, anecdotally, is often used in practice. A natural question is whether the inefficiencies we document are a result of irrational interviewing strategies, rather than an inherent feature of the types of signals available to agents. Therefore, in this section we analyze more sophisticated interviewing strategies, using the basic idea of empirical game theoretic analysis [76, 147].

The fundamental strategic decision faced by a firm is to choose a set of k candidates to interview. Game-theoretically, an (ex-ante) Bayes-Nash equilibrium would be one where each firm would not change the set of candidates it chose to interview, given the strategies of other firms, and the information available to them prior to the interview stage. Unfortunately, this game is very complex to analyze – even the set of strategies available to one firm is combinatorial ($\binom{n}{k}$). Therefore, we restrict our attention to a manageable set of strategies: each firm can decide on any set of k contiguously ranked candidates (in its posterior private ranking). So, say firm i has posterior ranking $\widetilde{\Gamma}_i$, then when it uses strategy T_i, uniquely identified by some integer m, it interviews the candidates ranked from m to $m+k-1$ in $\widetilde{\Gamma}_i$. This includes the "interview around my own rank" strategy discussed in Section 4 as a special case with $m = i-(k-1)/2$, and also allows firms to shoot higher or lower than their own ranking.

Which of these strategies is best? One simplification in analyzing which strategy is best for firm i to follow arises from the observation that firm i's best strategy depends only on the choices of the firms ranked above i. Since firm rankings are all common and common knowledge, firm i will always get its pick over any firm that is ranked lower. Since a Gale-Shapley mechanism is used in the second stage, firm i will get the candidate that it likes the most among the candidates it interviews who do not effectively receive an offer from a higher ranked firm.

The determination of firm strategies is iterative. Firm 1 should always interview the top k candidates in $\widetilde{\Gamma_1}$. Given the strategies being used by all firms ranked above it, firm i can run Monte Carlo simulations of outcomes for all strategies it can use, and pick the one that yields the highest utility or highest rank on average.[6] Importantly, note that the choice of firm i's strategy has no effect on the utilities being achieved by any of the strategies of the firms ranked above it, because those firms would always be able to get any candidate they interview before firm i could, by virtue of being higher ranked. Therefore, the choice of strategies by all firms is (approximately, because of the simulation) an equilibrium given the restricted range of strategies available to the firms.

2.4.1 Effects of common and private signals on strategies

First, we turn to understanding what kinds of interviewing behavior result as a choice of the intelligent strategy selection method above. Figure 2.6 shows what happens when there is only a common signal, and firms do not receive any private signals. The first row is for the case where the common signal is perfect information ($\sigma_C = 0$). The top k firms all interview the top k candidates, and every firm thereafter moves one candidate down. As σ_C increases, an increasingly zig-zag behavior in choice of strategy becomes apparent. As there is more uncertainty in the actual qualities of candidates, there can be more benefit to lower-ranked firms in interviewing higher-ranked sets of candidates, because the probability that they will not match with a higher-ranked firm increases. Of course, this tradeoff is governed by how

[6] Here, we focus on minimizing the average rank of the candidate the firm is matched to. The "average rank" when the firm is left unmatched can be important, but our results are qualitatively similar across many choices.

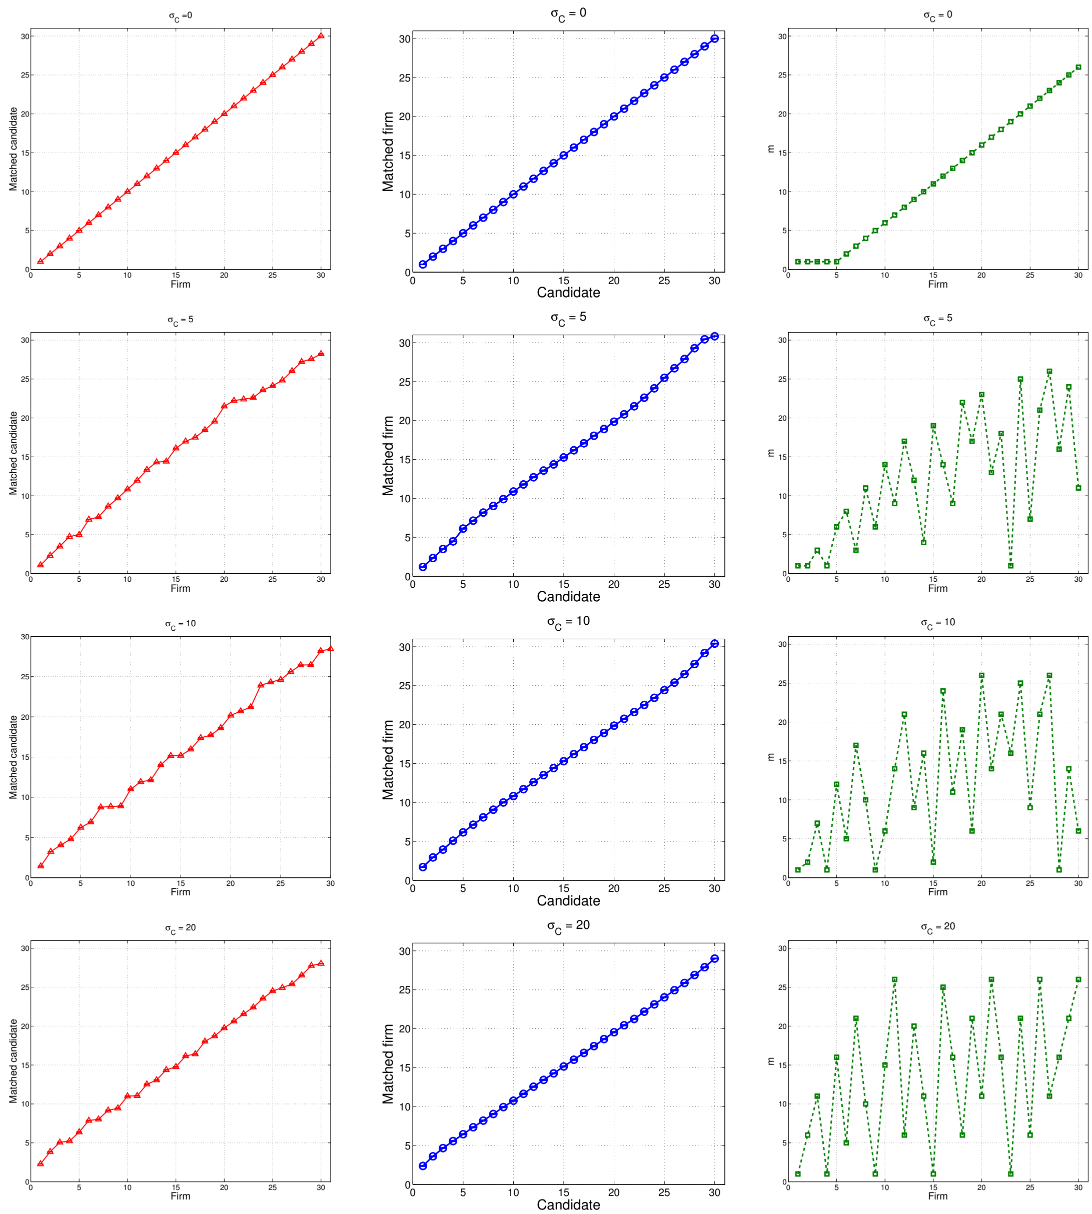

Figure 2.6: Strategies when firms have only a single common signal. Noise in the common signal is increasing as we move down the rows ($\sigma_C = 0, 5, 10, 20$ respectively). The left column shows the average rank of the candidate each firm is matched to, the center column shows the average rank of the firm each candidate is matched to, and the right column shows the strategy employed by each firm.

many interviews they already have, so once one firm jumps up (like firm 14 in the second row of the figure), there is less benefit to the next firm of also jumping that high, thus yielding the zig-zag behavior. Note that, despite this, the average ranks of candidates that firms are matched to are monotonically increasing in firm rank (and vice versa), as we would expect in equilibrium from the fact that firms all have to choose from the same set of strategies, so a lower-ranked firm cannot do better than a higher ranked one in expectation and still have an equilibrium.

The behavior with both common and private signals (Figure 2.7) is different. In some cases, the zig-zag behavior of the strategy still manifests itself, but less so as σ_C increases. With a high σ_C compared with σ_p, the difference in private signals is strong enough that interviewing mostly based on the private signal is the best strategy, providing enough differentiation in the sets being interviewed.

2.4.2 Effects of complex strategies on matching outcomes

Can "more rational" interviewing strategies resolve some of the inefficiency in terms of the number of participants left unmatched? Figure 2.8 shows the average number left unmatched as a function of the strength of the common signal. In addition to the "interview around my own rank" strategy of Section 4, we include the empirically determined strategies T_i, with three different values for the "penalty" rank assigned to a firm when it is unmatched: 31, 50, and 100. With only common signals, when the penalty is high enough, the better strategies, do, in fact, reduce the number left unmatched. However, with both common and private signals, the more complex strategies actually lead to a greater number of agents being left unmatched. This indicates that the greater complexity in strategy selection does not have

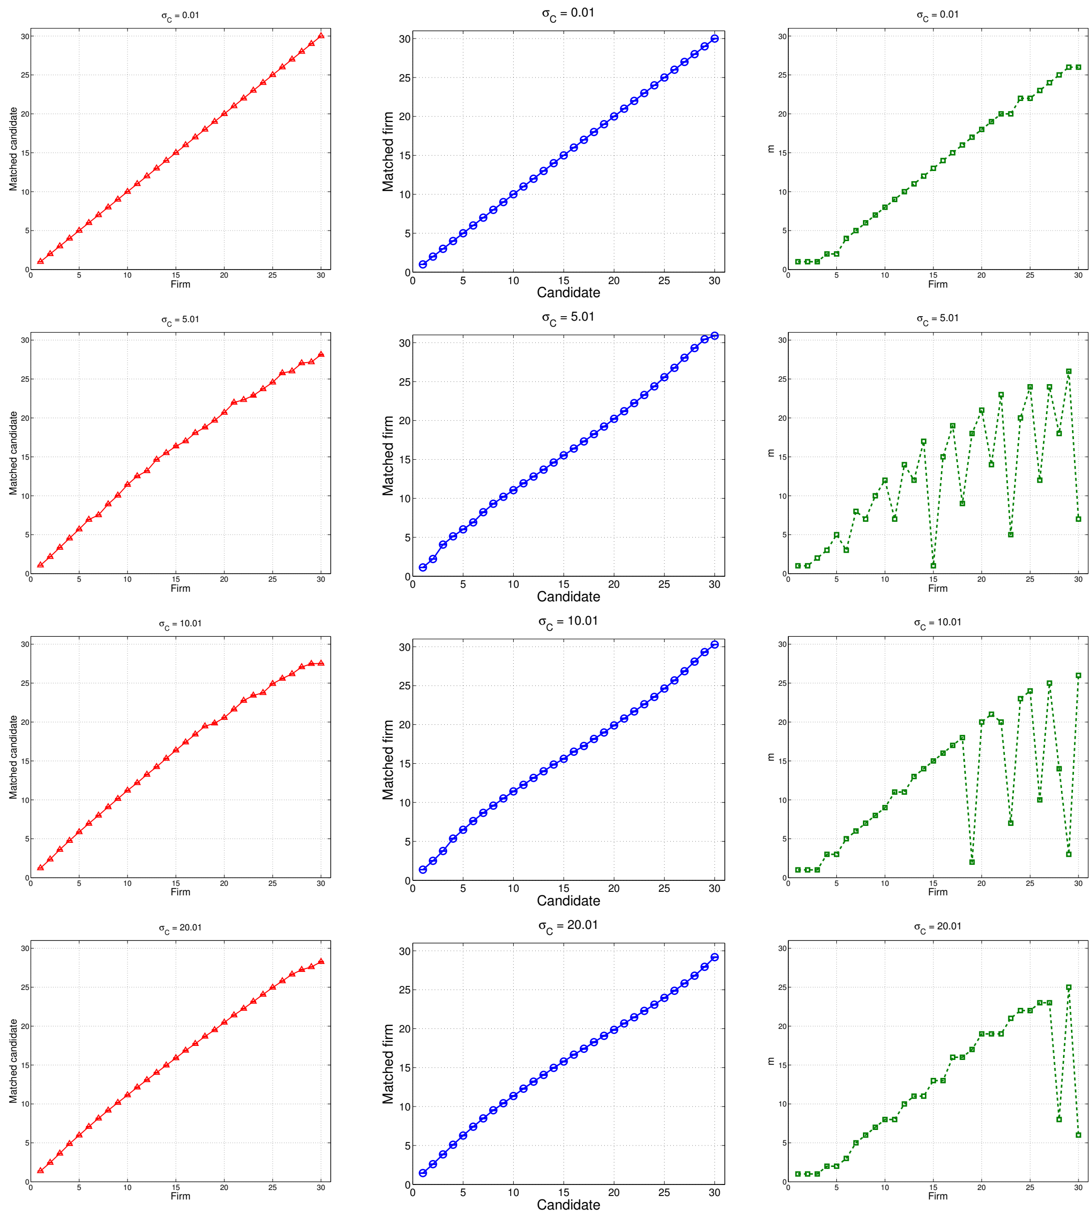

Figure 2.7: Strategies when firms have access to both a common signal and a private signal. Noise in the common signal is increasing as we move down the rows ($\sigma_C = 0.01, 5.01, 10.01, 20.01$ respectively). As in Figure 2.7, the left column shows the average rank of the candidate each firm is matched to, the center column shows the average rank of the firm each candidate is matched to, and the right column shows the strategy employed by each firm.

much societal benefit when private signals are available, but could be beneficial when there are only common signals available.

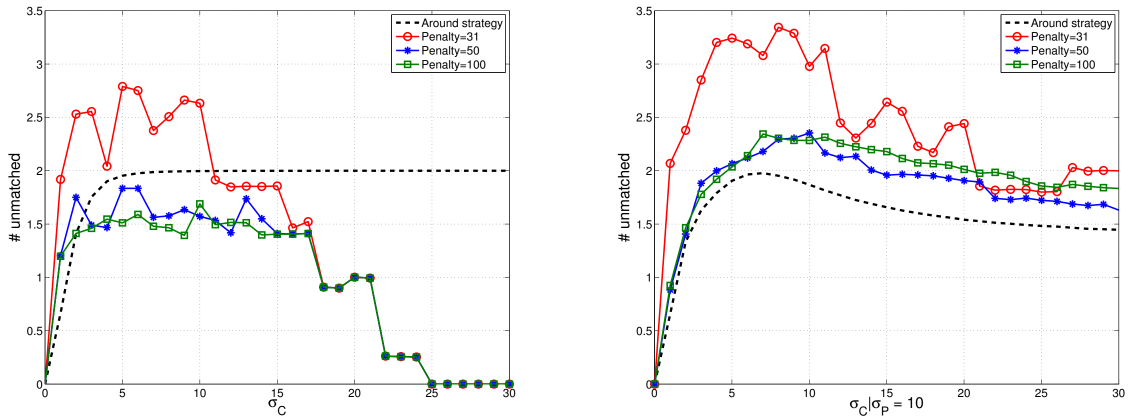

Figure 2.8: Number of agents left unmatched when firms use more sophisticated interviewing strategies with only a common signal (left) and common and private signals (right).

2.5 Discussion

In two-sided matching, firms want to hire the best candidate they can. Complications arise because of the tradeoff between quality and "gettability", and the fact that the game takes place in two stages, interviewing and matching. Our model focuses on capturing the essence of these phenomena. We are particularly interested in the role of information, because mechanisms are harder to change than the information available to participants. In addition to the "common signals" of applicant quality, many job markets allow firms to gather higher quality private signals in advance of making very costly interviewing decisions through phone screens or convention interviews. The main result of this chapter is that inefficiencies that arise through the use of simple, but anecdotally common, interviewing strategies (like interviewing around your own rank), can be alleviated by either the use of more sophisticated interviewing strategies or the use of additional private information. Given the complexities of picking interviewing strategies and the sensitivity to others' choices, institutional encouragement to participants to acquire diverse private signals may be more robust.

Chapter 3

Auction Markets and Information Design

In an auction game, assume a fixed mechanism; can the seller expect to make more revenue if the bidders are more or less informed than the "baseline"? In this chapter, we investigate the revenue-enhancement information design problem in the classic second-price auctions. In a second price, or Vickrey, auction, bidders are asked to submit sealed bids. The bidder who submits the highest bid is awarded the object, and pays the amount of the second highest bid. In particular, we explore the following two auction games: (1) a basic common-value auction model, where the value of the item is determined either by a single attribute or by two independent attributes when each bidder can receive information from exactly one of the attributes; (2) an interdependent-value auction, where the valuation for each bidder is decided by a common value attribute and a private attribute.

We show that in the common-value auction settings, there is no benefit to the auctioneer in terms of expected revenue from sharing information with the bidders, although there are effects on the distribution of revenues. In an interdependent-value model with mixed private-and common-value components, however, we show that asymmetric, information-revealing signals can increase revenue.

Our model is applicable to complex situations where the sender of the signal has the opportunity to communicate in different ways to different receivers. This can happen in situations like corporate mergers [19, 123], where targets (sellers or signal senders in our case) have to communicate with potential acquirers (the signal receivers). It is known that targets often inflate their output [64] or themselves may not be aware of their value to an acquirer due to the complexity and intangible characteristics which cannot be easily observed [81].

We position this work in the persuasion literature [77, 122], where a sender strategically reveals information through signals. Much of this literature focuses on the design of the optimal signaling scheme [63,113]. While this is tractable in some cases, for example with costly signals and a single receiver [62], or when a single buyer is signaling to a single monopolist seller [128], the problem of optimal signal design is not always even computationally, leave alone analytically, tractable [50, 149]. Therefore, the demonstration of a revenue-enhancing signal structure in the game with multiple receivers that we demonstrate here is significant, even if the particular structure we find is not the optimal one.

3.1 Common Value Auctions

We begin by considering a single-item auction with two risk-neutral bidders (agents) $i \in \{1, 2\}$ and a seller. Both bidders value the object identically: the item has a common value of $v \in R^+$ to the two bidders. The realization of v is not observed by either the seller or the bidders. v depends on an underlying state of the world $w \in \Omega$. Without loss of generality, we assume that the item's value is 0 when w's quality is *Bad (B)* and 1 when w's quality is *Good (G)*, and the common prior is represented by $P(G) = x, x \in [0, 1]$. Before bidding, each bidder receives a conditionally independent low (L), or high (H) signal from seller without cost, $s_i \in \{H, L\}$.

$$P[s_1 = H|G] = p_1 \quad P[s_1 = L|B] = q_1$$

$$P[s_2 = H|G] = p_2 \quad P[s_2 = L|B] = q_2$$

where s_i is agent i's signal and all signals have accuracy of $p_i, q_i \in [1/2, 1]$. Thus, a high (low) signal suggests a good (bad) value of the item.

Following prior literature, we make some assumptions.

Assumption 1 *Seller cannot distort or conceal information once the signal realization is known. [77].*

Assumption 2 *Bidders play only weakly undominated strategies. [24]*

The first assumption allows us to abstract from the incentive compatibility issues, while the second helps rule out implausible or uninteresting equilibria.

In the game, the seller decides the signal structure S with the goal of maximizing her expected revenue R and the bidders submit their bids based on their private signals s_i. The seller runs a two-player second-price sealed-bid (SPSB) auction. Define $\text{bid}_{s_{-i}}(s_i)$ as the bid of bidder i given she receives signal s_i and the other bidder receives signal s_{-i}. The seller can either reveal the realization of the signal privately to the corresponding bidder, or reveal it publicly. Here we show the analysis of private revelation, as public revelation follows similarly.

Proposition 1 *If the seller reveals the realization of the signal privately to the corresponding bidder, a unique symmetric equilibrium exists. Each agent bids her expected value conditioned on her opponent's signal being equal to her own,*

$$\begin{aligned} bid_L(L) &= \mathbb{E}[v|s_1 = L, s_2 = L] \\ &= P(G|s_1 = L, s_2 = L) \\ &= \frac{(1-p_1)(1-p_2)x}{(1-p_1)(1-p_2)x + q_1q_2(1-x)}, \end{aligned}$$

$$\begin{aligned} bid_H(H) &= \mathbb{E}[v|s_1 = H, s_2 = H] \\ &= P(G|s_1 = H, s_2 = H) \\ &= \frac{p_1p_2x}{p_1p_2x + (1-q_1)(1-q_2)(1-x)}. \end{aligned}$$

Proof The proof of this proposition is similar to prior work of Hausch [73] and of Brinkman, Wellman, and Page [24]. Assumption 2 (that bidders play only weakly undomainated strategies) restricts an agent with a *Low* signal to bid between $E[v|s_i = L, s_{-i} = L]$ and $E[v|s_i = L, s_{-i} = H]$, and one with a *High* signal to bid between $E[v|s_i = H, s_{-i} = L]$ and $E[v|s_i = H, s_{-i} = H]$. To see that the proposed strategy in proposition 1 is the only symmetric equilibrium, we begin by assuming that there exists a symmetric strategy that, when receiving signal L, the Bidder 1 bids x_1 and the Bidder 2 bids x_2, and when receiving signal H, Bidder 1 bids y_1 and Bidder 2 bids y_2. Suppose $x_1 \geq x_2$, then Bidder 1 will be strictly better off by deviating to $E[v|s_i = L, s_{-i} = L]$ when receiving an L signal, since bidding x_1 could result in negative utility ($E[v|s_i = L, s_{-i} = L] - x_2$) if Bidder 2 also receives an L signal. Similarly, if $y_1 \geq y_2$, Bidder 2 has incentive to switch to $E[v|s_i = H, s_{-i} = H]$ when receiving an H signal to achieve higher expected utility. Thus, the equilibrium bids above constitute the only symmetric equilibrium. □

Equilibrium selection It is well known that the second-price common-value auction generally has many equilibria [3, 73, 84, and so on]. Assumption 2 helps us to rule out all dominated bids. In this game, suppose that Bidder 1 obeys the strategy in Proposition 1. Bidder 2, conditional on receiving signal L bids $b \in (\text{bid}_L(L), \mathbb{E}[v|s_1 = H, s_2 = L]]$ and, conditional on receiving signal H, bids $\text{bid}_H(H)$. These strategies are still Nash equilibria. Thus, Nash equilibrium provides no prediction about revenue beyond an upper bound on the full surplus. For this chapter's purpose, therefore, we only focus on symmetric equilibrium bidding strategies.

In a common value auction, the seller's expected revenue R is the expected value $\mathbb{E}[v]$ of the item, minus the sum of the two bidders' utilities. When each bidder observes a private signal only, we can treat each bidder independently and minimize the utility of each bidder.

Theorem 1 *If each bidder observes her own private signal, the optimal signal structure for the seller in terms of revenue is* $p_1 = p_2 = 1, q_1, q_2 \in [1/2, 1]$*, or* $p_1 = p_2 = q_1 = q_2 = 1/2$*, where* $\max R = \mathbb{E}[v]$.

Proof For revenue maximization, we can treat the two-bidder second-price sealed-bid auction as a three-player, constant-sum game. The revenue

$$R = \mathbb{E}[v] - \mathbb{E}[u_1] - \mathbb{E}[u_2]. \tag{3.1}$$

$\mathbb{E}[u_i]$ is Bidder i's expected utility and R is maximized when $\mathbb{E}[u_1] = \mathbb{E}[u_2] = 0$, where

$$\begin{aligned}
\mathbb{E}[u_1] &= p(s_1 = H, s_2 = L)(\mathbb{E}[v|s_1 = H, s_2 = L] - \text{bid}_L(L)) \\
&= p_1(1-p_2)x - p(s_1 = H, s_2 = L)\text{bid}_L(L), \\
\mathbb{E}[u_2] &= p(s_1 = L, s_2 = H)(\mathbb{E}[v|s_1 = L, s_2 = H] - \text{bid}_L(L)) \\
&= (1-p_1)p_2 x - p(s_1 = L, s_2 = H)\text{bid}_L(L),
\end{aligned}$$

which gives us

$$p_1 = p_2 = 1, q_1, q_2 \in [\frac{1}{2}, 1], \forall x \in [0,1],$$

or

$$p_1 = p_2 = q_1 = q_2 = \frac{1}{2}, \forall x \in [0,1],$$

or

$$p_1, p_2, q_1, q_2 \in [\frac{1}{2}, 1], \text{when } x = 1.$$

When $p_1 = p_2 = q_1 = q_2 = 1$, the seller always reveals complete information, thus the expected revenue R is also $\mathbb{E}[v]$. $\square$

We can see that there is a wide range of signal structures that achieve the maximum revenue in equilibrium, and none of these is better than a policy of revealing no information at all. Another natural question to ask concerns the distribution of revenues to the seller under different signal structures. It is relatively easy to compute the variance of the revenue

$$var(R) = (\text{bid}_L(L) - \text{bid}_H(H))^2(1 - P(HH))P(HH) \tag{3.2}$$

Clearly $var(R)$ is minimized at $q_1 = q_2 = 0.5$. Figure 3.1 shows some illustrative examples of the standard deviation of revenue.

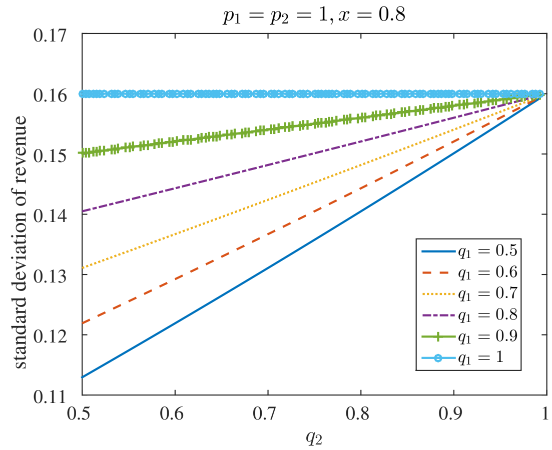

Figure 3.1: Standard deviations of revenue for different revenue-maximizing signal structures in the simple common-value model. While each of these signal structures achieves the same revenue, the risk profiles are substantially different.

3.1.1 Adding an Intermediate Value

Brinkman et al. [24] study a common-value auction setting with intermediate values, which serves as a model for studying signal acquisition by bidders. They motivate this setting with an example of the auction of extraction rights for some resources (say oil and gas) on a specified plot of land. The value to energy companies of these rights depends on the unknown amounts of extractable resources. The question of optimal signaling is motivated in this example by the fact that the government can reveal information about one or both of the specific resources to each energy company. Now the item can take on three possible values, $\{0, g, 1\}$ with $g \in [0, 1]$. The underlying state w which decides the value of the item now has two attributes, $w = (w_1, w_2)$. Each attribute is associated with signals potentially observed by the respective agents. Each bidder can request one signal with no cost. Here we study a variant where the seller can decide which attribute to signal to each bidder and what the corresponding signal structure should be.

Each attribute is still either *Good (G)* or *Bad (B)*, where $P(w_j = G) = x \in [0,1], j \in \{1,2\}$. The realization of each signal is also *High (H)* or *Low (L)*. The signal structure can be represented as ($s_i^j \in \{H, L\}$):

$$
\begin{aligned}
P[s_1^j = H | w_j = G] = p_1 \quad P[s_1^j = L | w_j = B] = q_1 \\
P[s_2^j = H | w_j = G] = p_2 \quad P[s_2^j = L | w_j = B] = q_2
\end{aligned}
$$

where $j \in \{1,2\}$ and s_i^j is Bidder i's signal from attribute j. All signals have accuracy of $p_i, q_i \in [1/2, 1]$.

The value of the good is 0 if neither attribute is G, 1 if both are G, and $g \in [0,1]$ if only one is G.

$$
v = \begin{cases} 0, \text{if } \sum_j \mathbb{I}\{w_j = G\} = 0 \\ g, \text{if } \sum_j \mathbb{I}\{w_j = G\} = 1 \\ 1, \text{if } \sum_j \mathbb{I}\{w_j = G\} = 2 \end{cases}
$$

Figure 3.2 shows the decision flow in this game. The seller's goal is to maximize her expected revenue R. The signal structure and the seller's choice of which attribute to signal to each bidder are both common knowledge.

First, we observe that it must again be the case that the seller's revenue is maximized when revealing no information even in this intermediate value setting, since it can still be modeled as a three-player, constant-sum game, and Equation (3.1) holds. What can we say about signal structures that achieve this revenue? Again, we analyze private revelation.

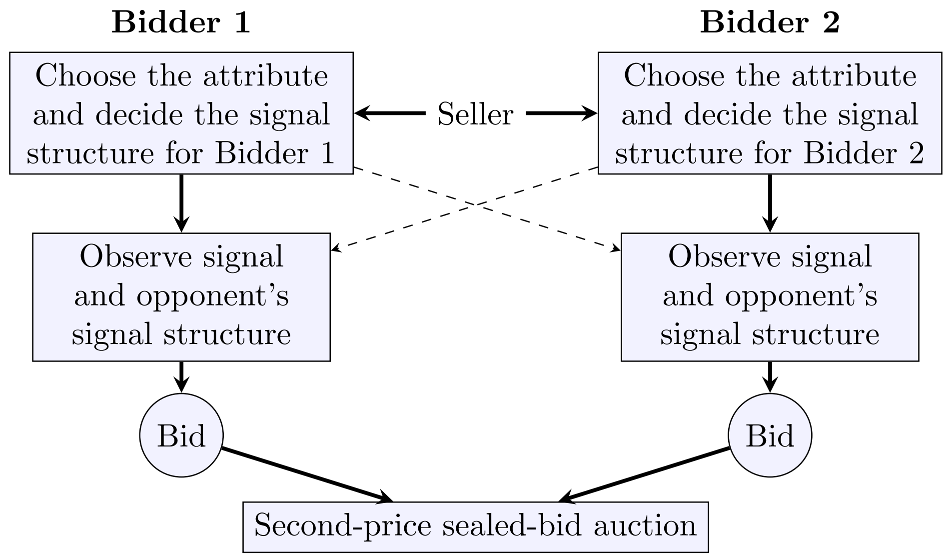

Figure 3.2: The intermediate-value model. Dashed lines mean that the bidder knows the structure of the signal that the other bidder receives, but not the specific realization.

Theorem 2 *In the intermediate value model, (1) if the seller sends signals of different attributes to the two buyers, there is only one signal structure,* $\forall g, x \in [0,1]$, $p_1 = p_2 = q_1 = q_2 = 1/2$ *(equivalent to sending no information) that achieves the maximum possible revenue; (2) if the seller sends signals of the same attribute to both buyers, for* $\forall g, x \in [0,1]$, *there are a number of signal structures that achieve the maximum possible revenue:* $p_1 = p_2 = 1, q_1, q_2 \in [1/2, 1]$ *or* $p_1 = p_2 = q_1 = q_2 = 1/2$.

Proof The seller's revenue still follows Equation (3.1). To maximize R, $\mathbb{E}[u_1] = \mathbb{E}[u_2] = 0$.

- *Sending signals of the same attribute:*

The unique symmetric equilibrium bidding strategy is that each bidder bids her expected value conditioned on her opponent's signal being equal to her own,

$$\mathrm{bid}_L(L) = \mathbb{E}(v|s_i^j = L, s_{-i}^j = L),$$
$$\mathrm{bid}_H(H) = \mathbb{E}(v|s_i^j = H, s_{-i}^j = H).$$

We denote $P(s_i^j = H, s_{-i}^j = L)$ by $P(HL)$,

$$\mathbb{E}[u_i] = P(HL)(\mathbb{E}(v|s_i^j = H, s_{-i}^j = L) - \mathrm{bid}_L(L)).$$

Thus, to maximize R

$$\mathbb{E}(v|s_i^j = H, s_{-i}^j = L) = \mathrm{bid}_L(L). \tag{3.3}$$

The solution of Equation (3.3) is

$$p_1 = p_2 = q_1 = q_2 = \frac{1}{2}, \forall g, x \in [0, 1],$$

or

$$p_1 = p_2 = 1, q_1, q_2 \in [\frac{1}{2}, 1], \forall g, x \in [0, 1],$$

or

$$p_1, p_2, q_1, q_2 \in [\frac{1}{2}, 1], \text{when } x = 1, \forall g \in [0, 1].$$

When $p_1 = p_2 = q_1 = q_2 = 1$, the seller reveals perfect information, thus the expected revenue R is also $\mathbb{E}[v]$.

- *Sending signals of different attributes:*

As the signal accuracy between different attribute is identical, the equilibrium biding strategy is same as above, that is to bid the expected valuation conditioned on the opponent observing the same signal value. Denote $\text{bid}_{-s_{-i}}(s_i)$ as the bid of Bidder i given she receives s_i and the other bidder observes the signal of the other attribute and receives signal s_{-i} ,

$$\text{bid}_{-L}(L) = \mathbb{E}(v|s_i^j = L, s_{-i}^{-j} = L),$$

$$\text{bid}_{-H}(H) = \mathbb{E}(v|s_i^j = H, s_{-i}^{-j} = H).$$

We simplify $P(s_i^j = H, s_{-i}^{-j} = L)$ by $P(H, L)$,

$$\mathbb{E} = [u_i] = P(H, L)(\mathbb{E}(v|s_i^j = H, s_{-i}^{-j} = L) - \text{bid}_{-L}(L)).$$

Thus, to maximize R,

$$\mathbb{E}(v|s_i^j = H, s_{-i}^{-j} = L) = \text{bid}_{-L}(L). \tag{3.4}$$

Solving Equation (3.4) we get,

$$p_1 = p_2 = q_1 = q_2 = \frac{1}{2}, \forall g, x \in [0, 1].$$

□

It is again easy to show that $var(R)$ is minimized at $q_1 = q_2 = 0.5$.

Discussion Brinkman et al. [24] analyze this problem from the perspective of the bidders. In their model, the signal structure is fixed and restricted to the symmetric information case ($p_1 = p_2 = q_1 = q_2$). They show that when the two attributes are sufficiently complementary, that is $g \to 0$, and the signals are noisy, the agents choose to observe the same attribute. When the signal accuracy is high, or the two signals are substitutable $g \to 1$, the agents choose to observe different attributes. Our result above demonstrates that, from the seller's perspective, sending no information can always maximize seller's expected revenue. The seller can also achieve the maximum possible revenue by sending information on the same attribute to both bidders. The corresponding signal structure shows that the bidders always know the item is bad if they see a *low* signal, but they have uncertainty when they see a *high* signal.

3.2 An Interdependent Value Auction

We now move to a setting with an unambiguously positive result for the seller. We consider a classic situation in corporate mergers. A firm (target) can generate synergies if acquired by another firm (bidder) [19]. The source of this synergy may include management, economies of scale, technological matches, tax savings, etc. A sketch of the game is shown in Figure 3.3. The target's quality can be either good or bad, which is unknown to the market and the bidders at the time of bidding. The bidders' types can be high or low tech, privately known to each bidder. The ability of a bidder to generate synergies can be either high or low, which is unknown to the market and to the bidders, but may be discovered by the target (since the target is willing to invest in discovering this prior to making it known that it is open to acquisition). If the type of a bidder is high tech, as long as the ability of the bidder to generate synergies is high, it can get high value ($\alpha > 1$) no matter the target's quality. However, if the type of a bidder is low tech, only when both the ability of the bidder to generate synergies is high and the quality of the target is good, can it get medium value (1).

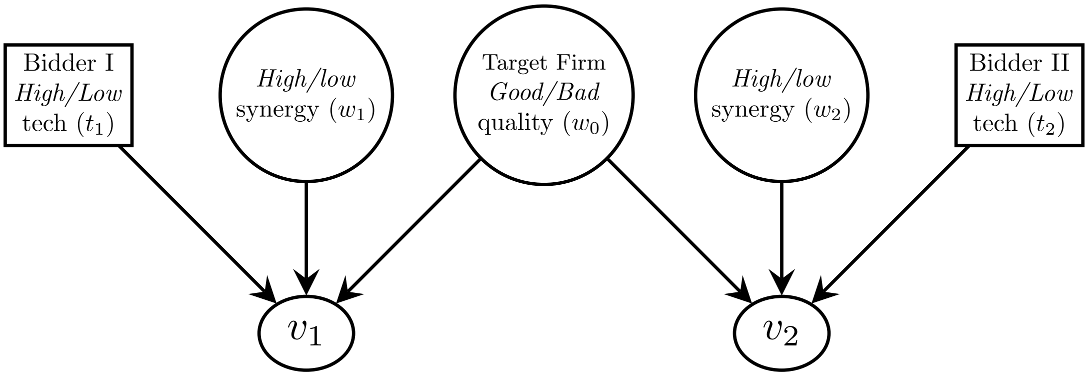

Figure 3.3: A sketch of the interdependent value setting.

3.2.1 Model

We first extend the common-value model of Brinkman *et al* to this situation. The item's value still depends on an underlying state w, which now has three attributes $w = (w_0, w_1, w_2)$. The common attribute w_0 can affect the valuation of both bidders (quality of the target firm), and the private attributes w_1 and w_2 only affect each bidder's own valuation respectively (idiosyncratic synergies). Each attribute takes quality *Good (G)* or *Bad (B)* as above. For simplicity, we assume $P(w_j = G) = x \in [0,1], j \in \{0,1,2\}$ (this assumption can be easily removed and all results hold). The seller sends a signal of the quality of either common or private attribute w_j to each bidder. The realization of each signal is also *High (H)* or *Low (L)*. The signal structure is ($s_i^j \in \{H, L\}$)

$$P[s_1^j = H | w_j = G] = p_1, \quad P[s_1^j = L | w_j = B] = q_1,$$
$$P[s_2^j = H | w_j = G] = p_2, \quad P[s_2^j = L | w_j = B] = q_2.$$

All signals have accuracy of $p_i, q_i \in [1/2, 1]$. Once the signal structure is decided, it becomes common knowledge. The seller can choose to either reveal realizations publicly or privately.

The bidders can be of two types, $t_i \in \{t_l, t_h\}$. The bidders will be of either type with probability $P(t_i = t_l) = P(t_i = t_h) = \frac{1}{2}$. If the bidder is type t_h (high tech firm), then her valuation is only dependent on her private attribute, that is $w_i = G$ with value $\alpha > 1$ (pure strategy Nash equilibrium is not guaranteed if $\alpha = 1$) and $w_i = B$ with value 0. If the bidder is type t_l (low tech firm), her valuation is dependent on both common and private attributes: the bidder's value is 0 if both the common and her private attribute are B, and 1 if both are G. Formally,

$$
\begin{aligned}
& i \in \{1,2\} \\
v_i(w_0, w_i, t_i = t_l) &= \begin{cases} 1, & \text{if } w_0 = G, w_i = G, \\ 0, & \text{else}, \end{cases} \\
v_i(w_0, w_i, t_i = t_h) &= \begin{cases} \alpha, & \text{if } w_i = G, \\ 0, & \text{else}, \end{cases}
\end{aligned}
$$

where $P(t_i = t_l) = \frac{1}{2}$.

3.2.2 Analysis

Before the game, the seller needs to decide which attribute she wants to signal to each bidder and whether the realization of the signal is public or private. The seller still provides one signal to each bidder, but the realization of that signal can be public. The complete results characterizing the best possible revenue impact and the corresponding signal structure based on seller's strategy is shown in Figure 3.1. The main results to note are that there are two signal structures that are revenue enhancing.

When we allow one bidder (w.lo.g. Bidder 2) to observe a signal of her private attribute while the other bidder receives a private signal of the common attribute (case 9), there exists a revenue-enhancing signal structure. In equilibrium, a bidder of type t_h always bids her expected value given the signal realization of private attribute if she receives one. If Bidder 1 is type t_l she bids her expected value given the signal realization she observes. If Bidder 2 is type t_l, if she observes a *low* signal, her bid falls in the range $[\mathbb{E}[v|s_1^0 = L, s_2^2 = L], \mathbb{E}[v|s_1^0 = H, s_2^2 = L]]$ under Assumption 2 and also needs to be smaller than Bidder 1's expected value given Bidder 1 observes a *low* signal $\mathbb{E}[v|s_1^0 = L]$; if she observes a *high* signal, from Assumption 2 her bid falls in the range $[\mathbb{E}[v|s_1^0 = L, s_2^2 = H], \mathbb{E}[v|s_1^0 = H, s_2^2 = H]]$, and also needs to be greater than bidder 1's expected value given Bidder 1 observes a *high* signal $\mathbb{E}[v|s_1^0 = H]$.

Now, suppose the seller chooses signal structure $p_1 \in [0.5, 1], p_2 = 1, q_1 = 1, q_2 = 0.5$. If Bidder 1 observes a high signal, she knows with certainty that the common attribute is good, and is uncertain otherwise. Bidder 2 knows that her private attribute is bad if she observes a low signal, and is uncertain otherwise. Combined with the observation about bid ranges above, it now becomes a simple matter of algebra to show that the expected revenue is greater than that which is achieved when the seller reveals no information or full information, yielding the following theorem:

Theorem 3 *Privately revealing the realization of the common attribute signal to one bidder and privately revealing the realization of the private attribute signal to the other bidder, the* *seller's expected revenue at* $p_1 \in [0.5, 1], p_2 = 1, q_1 = 1, q_2 = 0.5$ *is always better than that she can achieve when revealing no information or full information.*

Proof Without loss of generality, we assume Bidder 1 receives a private signal of the common attribute while Bidder 2 observes a signal of her private attribute.

If Bidder 1 is type t_h, she bids $\text{bid}^1_{t_h} = \alpha x$ since she only receives a signal of the common value. If Bidder 1 is type t_l, sice the private attribute of Bidder 2 does not influence the value of Bidder 1, if she observes a *low* signal, she bids $\text{bid}^1_{t_l}(L) = \mathbb{E}[v|s^0_1 = L]$; if she observes a *high* signal, she bids $\text{bid}^1_{t_h}(L) = \mathbb{E}[v|s^0_1 = H]$.

If Bidder 2 is type t_h, she bids her expected value given the signal realization of the private attribute $\text{bid}^2_{t_h} = \mathbb{E}(v|s^2_2)$. If Bidder 2 is type t_l, if she observes a *low* signal, her equilibrium bid falls in the range

$$\begin{aligned}&\text{bid}^2_{t_l}(L) \in [\mathbb{E}[v|s^0_1 = L, s^2_2 = L], \mathbb{E}[v|s^0_1 = H, s^2_2 = L]],\\ &\text{and } \text{bid}^2_{t_l}(L) \leq \mathbb{E}[v|s^0_1 = L]\end{aligned}$$

under Assumption 2; if she observes a *high* signal, from Assumption 2, her equilibrium bid falls in the range

$$\begin{aligned}&\text{bid}^2_{t_l}(H) \in [\mathbb{E}[v|s^0_1 = L, s^2_2 = H], \mathbb{E}[v|s^0_1 = H, s^2_2 = H]],\\ &\text{and } \text{bid}^2_{t_l}(H) \geq \mathbb{E}[v|s^0_1 = H].\end{aligned}$$

Considering the lower bound of the revenue, we choose the minimum bid under all cases. We separate the revenue into the following four parts,

1. $t_1 = t_l$ and $t_2 = t_l$, where we have $P(t_1 = t_l, t_2 = t_l) = \frac{1}{4}$:

$$\begin{aligned} R_1 = &P(s_1^0 = H)P(s_2^2 = H)\mathbb{E}[v|s_1^0 = H] \\ &+ P(s_1^0 = H)P(s_2^2 = L)\mathbb{E}[v|s_1^0 = L, s_2^2 = L] \\ &+ P(s_1^0 = L)P(s_2^2 = H)\mathbb{E}[v|s_1^0 = L] \\ &+ P(s_1^0 = L)P(s_2^2 = L)\mathbb{E}[v|s_1^0 = L, s_2^2 = L]; \end{aligned}$$

2. $t_1 = t_l$ and $t_2 = t_h$, where we have $P(t_1 = t_l, t_2 = t_h) = \frac{1}{4}$:

$$\begin{aligned} R_2 = &P(s_1^0 = H)P(s_2^2 = H) \\ &\min(\mathbb{E}[v|s_1^0 = H], \mathbb{E}(v|s_2^2 = H)) \\ &+ P(s_1^0 = H)P(s_2^2 = L) \\ &\min(\mathbb{E}[v|s_1^0 = H], \mathbb{E}(v|s_2^2 = L)) \\ &+ P(s_1^0 = L)P(s_2^2 = H) \\ &\min(\mathbb{E}[v|s_1^0 = L], \mathbb{E}(v|s_2^2 = H)) \\ &+ P(s_1^0 = L)P(s_2^2 = L) \\ &\min(\mathbb{E}[v|s_1^0 = L], \mathbb{E}(v|s_2^2 = L)); \end{aligned}$$

3. $t_1 = t_h$ and $t_2 = t_l$, where we have $P(t_1 = t_h, t_2 = t_l) = \frac{1}{4}$:

$$\begin{aligned} R_3 = &P(s_2^2 = H)\max(\mathbb{E}[v|s_1^0 = L, s_2^2 = H], \mathbb{E}[v|s_1^0 = H]) \\ &+ P(s_2^2 = L)\mathbb{E}[v|s_1^0 = L, s_2^2 = L]; \end{aligned}$$

Case	c_1	pr_1	c_2	pr_2	Revenue impact	Maximizing structure	Remarks
1	no	no	no	no	–	–	
2	no	yes	no	no	$\downarrow$	no information	unique eq
3	no	yes	no	yes	$\downarrow$	no information	unique eq
4	publicly	no	no	no	–	any	unique eq
5	publicly	no	publicly	no	–	any	unique eq
6	publicly	no	no	yes	$\downarrow$	private signal no information	unique eq
7	publicly	no	privately	no	$\downarrow$	lower bound maximized at no information	multiple eqs
8	privately	no	no	no	$\downarrow$	lower bound maximized at no information	multiple eqs
9	privately	no	no	yes	$\uparrow$	lower bound better than no information	multiple eqs
10	privately	no	privately	no	$\uparrow$	$p_1 = 1, p_2 = 1,$ $q_1 = 1, q_2 = 0.5$	unique symmetric eq

Table 3.1: Best possible revenue impacts and corresponding signal structures in the interdependent value setting. c_i indicates signaling the common attribute to Bidder i and pr_i indicates signaling the private attribute to Bidder i. For the common attribute, "publicly" means the realization of the signal can be observed by all bidders and "privately" means the realization of the signal can only be observed by the corresponding bidder. Since private values are independent, whether that signal is revealed publicly or privately makes no difference. Note that the order of the two bidders is arbitrary, but the existence of the asymmetry is not.

4. $t_1 = t_h$ and $t_2 = t_h$, $P(t_1 = t_h, t_2 = t_h) = \frac{1}{4}$:

$$\begin{aligned} R_4 =& P(s_2^2 = H) \min(\alpha x, \mathbb{E}(v|s_2^2 = H)) \\ &+ P(s_2^2 = L) \min(\alpha x, \mathbb{E}(v|s_2^2 = L)). \end{aligned}$$

The expected revenue is $R = 0.25(R_1 + R_2 + R_3 + R_4)$. By simple algebra, we find that the proposed signal structure achieves better revenue than either no information or full information. $\square$

An interesting observation about this signal structure is that, while the signal structure conveys more information to Bidder 1, her utility is actually lower compared with when there is no information. Bidder 2's utility improves.

Finally, we see what happens if the seller signals the common attribute to each bidder privately (case 10 in Figure 3.1). In this situation, the equilibrium bidding strategy for t_h type bidder is to bid her expected value regardless of the signal she receives and for t_l type bidder is to bid her expected value conditioned on the other bidder observing same signal. It is easy to show that the signal structure $p_1 = p_2 = 1, q_1 = 1, q_2 = 0.5$, results in higher expected revenue than when the seller conveys no information or full information.

Theorem 4 *When revealing the signal realization of the common attribute privately to each bidder, the seller's revenue is higher at signal structures* $p_1 = p_2 = 1, q_1 = 1, q_2 = 0.5$, *or* $p_1 = p_2 = 1, q_1 = 0.5, q_2 = 1$, *than when revealing no information or full information.*

Proof Since the private attribute of each bidder does not affect the value to the other bidder, if $t_i = t_l$, the equilibrium strategy of bidder i is the same as in Proposition 1.

$$
\begin{aligned}
\mathrm{bid}_L^i(L) &= \mathbb{E}(v|s_i^0 = L, s_{-i}^0 = L) \\
&= v(w_0 = G, w_i = G) \\
&\quad P(w_0 = G|s_i^0 = L, s_{-i}^0 = L)P(w_i = G) \\
&\quad + v(w_0 = G, w_i = B) \\
&\quad P(w_0 = G|s_i^0 = L, s_{-i}^0 = L)P(w_i = B) \\
&\quad + v(w_0 = B, w_i = G) \\
&\quad P(w_0 = B|s_i^0 = L, s_{-i}^0 = L)P(w_i = G) \\
&\quad + v(w_0 = B, w_i = B) \\
&\quad P(w_0 = B|s_i^0 = L, s_{-i}^0 = L)P(w_i = B) \\
&= \frac{(1-p_1)(1-p_2)x^2}{(1-p_1)(1-p_2)x + q_1 q_2 (1-x)}, \\
\mathrm{bid}_H^i(H) &= \mathbb{E}(v|s_i^0 = H, s_{-i}^0 = H) \\
&= \frac{p_1 p_2 x^2}{p_1 p_2 x + (1-q_1)(1-q_2)(1-x)}.
\end{aligned}
$$

If $t_i = t_h$, the bidder's payoff is only related to the relative private attribute, so the equilibrium bidding strategy is always

$$\mathrm{bid}_{t_h}^i = \alpha x.$$

We break the revenue up into the following four parts,

1. $t_1 = t_l$ and $t_2 = t_l$, where we have $P(t_1 = t_l, t_2 = t_l) = \frac{1}{4}$:

$$R_1 = P(s_1^0 = H, s_2^0 = H)\text{bid}_H^i(H) + (1 - P(s_1^0 = H, s_2^0 = H))\text{bid}_L^i(L);$$

2. $t_1 = t_l$ and $t_2 = t_h$, where we have $P(t_1 = t_l, t_2 = t_h) = \frac{1}{4}$:

$$\text{As } \alpha > 1, \alpha x > \text{bid}_H(H);$$
$$R_2 = P(s_1^0 = H)\text{bid}_H^1(H) + P(s_1^0 = L)\text{bid}_L^1(L);$$

3. $t_1 = t_h$ and $t_2 = t_l$, where we have $P(t_1 = t_h, t_2 = t_l) = \frac{1}{4}$:

$$R_3 = P(s_2^0 = H)\text{bid}_H^1(H) + P(s_2^0 = L)\text{bid}_L^2(L);$$

4. $t_1 = t_h$ and $t_2 = t_h$, $P(t_1 = t_h, t_2 = t_h) = \frac{1}{4}$:

$$R_4 = \alpha x.$$

Thus, the expected revenue is $R = \frac{1}{4}(R_1 + R_2 + R_3 + R_4)$. By simple algebra, we can find that the signal structure $p_1 = p_2 = 1, q_1 = 1, q_2 = 0.5$ yields higher revenue than no information or full information, and the revenue increase is from R_3. □

Consider signal structure $p_1 = p_2 = 1, q_1 = 1, q_2 = 0.5$ (the other one is symmetric). Bidder 1 always has perfect information. If Bidder 2 receives a low signal, she is certain w_0 is bad; however, she is uncertain when she gets a high signal. Surprisingly, although Bidder 1 has perfect information, her expected utility is actually lower than that of Bidder 2. It is easy to see that if both bidders are t_l types or t_h types, then the expected utility of each bidder is zero. The interesting case is when one bidder is a t_h type, and the other one is a t_l type. In this situation, the bidder with imperfect information is more likely to receive a high signal than the bidder with perfect information; therefore, in expectation, the perfect information bidder will pay more (since it is a second price auction), hurting her utility.

3.3 Conclusion

The key point in the emerging signaling literature in information economics and computer science is to study what can be achieved through information design, or persuasion, when the mechanism is already fixed. We demonstrate the range of possible outcomes that can be achieved through different signaling schemes in common value auction, and show that the uninformative scheme has the lowest risk among those that extract full surplus. While different signal structures may not help improve revenue in second-price sealed bid common value auctions, there are natural auction models, like the interdependent value model for corporate takeovers we present, in which the optimal design of signal structures can be revenue enhancing.

3.4 Related Work

This part of chapter is related to several literatures. Broadly, this chapter fits into a growing line of literature in AI on how the information environment available to agents influences market outcomes as we address in this thesis. In Chapter 2, we model the effects of common and private signals about quality in matching with interviews. Hajaj and Sarne [70] examine how e-commerce platforms can gain from information withholding policies. Chhabra et al. [31] study the welfare effects of competition between information providers with different levels of information quality. Rabinovich et al [118] present an efficient model for security asset assignment which combines both Stackelberg security games and the Bayesian Persuasion model.

The literature on auctions with signaling, as mentioned earlier, typically analyzes symmetric information structures, where there are few positive results in terms of revenue enhancements. In addition to the literature from economics cited above, recent work in algorithmic economics that assumes symmetric information disclosure includes that of Emek et al [51] as well as Bro Miltersen and Sheffet [25], both of which study second-price auctions of multiple indivisible goods and consider hiding information by clustering. Guo and Deligkas [68] single-item second-price auctions where the item is characterized by a set of attributes and the auctioneer decides whether to hide a subset of attributes.

When we move to asymmetric information, most early work considers the case in which one bidder is perfectly informed about the value of the item, while the other bidders are entirely uninformed [108,148]. Milgrom and Weber [108] show that reducing information asymmetries can increase the seller's expected revenue in a two-bidder first-price common value auction where one bidder is perfectly information and the other bidder is entirely uninformed. Goeree and Offerman [67] also consider public information disclosure in common value auctions, in which the common value is an average of i.i.d. private values (signals) of all bidders. They also conclude that seller's public information disclosure can raises efficiency and seller's revenues. Hausch [73], however, through a simple example in a first price common value auction, shows that reducing information asymmetry may decrease the seller's expected revenue when the better-informed bidder is neither strictly better-informed nor perfectly informed.

Syrgkanis *et al* [134] consider common value hybrid auctions where the payment is a weighted average of the highest and second-highest bids. They show that public revelation of an additional signal to both bidders may decrease the auctioneer's revenue, different from [107]. Parreiras [116] consider continuous signal spaces and also show that second price auction revenue-dominates first price auction. In both of these papers, the seller does not control the information structure for both bidders.

There are also several recent papers considering this question from the optimal mechanism design perspective [18, 42, 129], rather than assuming a fixed structure for the mechanism and analyzing the question of optimal signaling given the mechanism. Very recent work of Alkoby et al [5] analyzes signaling by a third party information provider under a fixed mechanism.

Also related is the literature on deliberative auctions. Deliberation covers any actions that update an agent's belief. In the study of deliberative auctions, research has thus far focused on either the perspective of bidders (receivers) or on optimal mechanism design. Larson and Sandholm [86, 87] provide a very general model for costly information gathering in auctions. They show that under costly deliberation, bidders perform strategic deliberation in equilibrium in most standard auction settings (Vickrey, English, Dutch, first price and VCG). Thompson and Leyton-Brown [135] investigate deliberation strategies for second price auctions where agents have independent private values (IPV) and the impact of agents' strategies on seller's revenue. They perform equilibrium analysis for (1) deliberation with costs, (2) free, but time-limited deliberation. They further show that, in the IPV deliberative-agent setting, the only dominant-strategy mechanism is a sequential posted price auction, in which bidders are sequentially given a posted-price, take-it-or-leave-it offer until the good is sold [136]. Celis *et al* [29] provide an efficient mechanism in IPV deliberative-agent setting to obtain revenue within a small constant factor of the maximum possible revenue. Brinkman *et al* [24] show that the dependence structures among agents' signals of the value of the item they are bidding on can produce qualitatively different equilibrium outcomes of the auction. This literature also typically does not focus on the optimal design of the signal structure from the perspective of the seller.

Chapter 4

Competing Dynamic Matching Markets

From this chapter, we explore dynamic markets. We focus on the dynamics of multiple platform interactions. Specifically, we formalize a two-market model where agents enter one market or both markets; they can then be matched to other agents who have joined the same market or both markets. The markets adhere to different matching policies, with one matching greedily and the other building market thickness through a policy of patience (patient market). We provide an analytic lower bound on the *loss*, or the expected fraction of vertices who enter and leave the pool without finding a match, of the two-market model and show that it is higher than running a single "patient" market. We also provide a quantitative method for determining the loss of the two-market model.

Our work draws motivation from kidney exchange, an instantiation of barter exchange where patients paired with willing but medically incompatible donors swap those donors with other patients. In the United States, multiple fielded kidney exchanges exist, and patient-donor pairs are entered simultaneously into one or more of these markets, based on geographical location, travel preferences, home transplant center preferences, or other logistical reasons.

Individual kidney exchange clearinghouses have incentive to compete on number of matches performed within their specific pools; yet, fragmenting the market across multiple exchanges operating under different matching policies may lower global welfare. In this chapter, we provide the first theoretical and experimental evidence on dynamic kidney exchange graphs showing that this may indeed be the case.

4.1 Kidney Exchange Model

Here, we start from describing the basics of kidney exchange. A kidney exchange can be represented as a directed compatibility graph $G = (V, E)$. Each vertex in the graph is a patient-donor pair in the pool. A directed edge e is constructed from vertex v_i to vertex v_j if the patient v_j is compatible with the donor kidney of v_i. Edges exist or do not exist due to medical characteristics (most importantly blood type, tissue antibodies and antigens) of the patient and the donor. There may also be other logistical constraints, but those are not relevant for our work here. In this pool, the donor of vertex v_i is willing to give her kidney if and only if the patient of v_i receives a kidney. A weight w_e can be assigned to an edge e. The weight typically has been used in the literature to represent the priority of a transplantation (and therefore the utility to the *system* in some senses). We also use it to represent the match quality when recipient v_j receives v_i's donor kidney (this part will be discussed in Chapter 6). In this graph, a sequence of transplants occurs when several vertices form a cycle c. A k-cycle refers to a cycle with exactly k pairs. In this thesis, we only consider 2-cycles and 3-cycles, as is typical in fielded kidney exchange (incorporating cycles longer than 3 offers limited benefit given logistical constraints). Fielded exchanges also gain from chains, where an altruist donor without a paired patient enters the pool and start a directed path of transplants. We do not include chains here.

A *matching* M is therefore a set of disjoint cycles in the compatibility graph G. The cycles must be disjoint because no donor can give more than one of her kidneys (some recent work explores multi-donor donation [54,55] but we do not consider this here). Given a pre-defined utility function $u : \mathcal{M} \to \mathbb{R}$ and the set of all legal matchings $\mathcal{M}$, we are trying to find a matching which maximizes u,

$$M^* \in \arg\max_{M \in \mathcal{M}} u(M).$$

Kidney exchanges typically find the maximum weighted cycle cover, formally,

$$u(M) = \sum_{c \in M} \sum_{e \in c} w_e.$$

In this thesis, we consider two objectives, the number of matches (effectively $w_e = 1, \forall e$), and expected total graft survival (where w_e is defined as the expected graft survival for the recipient in edge e).

An integer programming (IP) solver is usually used to find the optimal solution [2,12,33,45]. We use the position-indexed chain-edge formulation (PICEF) [45] method to find the optimal solution when doing two-&three-cycle swap if necessary.

4.2 Greedy and Patient Exchanges

So far we have described a static matching market, in which all patient-donor pairs are presented, and the market only needs to make a one-time matching decision. However, in practice, patient-donor pairs dynamically join and leave the market. In this thesis, we consider a stylized model of dynamic kidney exchange following the description of Akbarpour et al. [4]. More specifically, an exchange is running in the continuous-time interval $[0, T]$, with agents arriving according to a Poisson process with rate parameter $m \geq 1$. The exchange determines whether potential bilateral transactions between agents are either acceptable or unacceptable. The probability of an *acceptable* transaction existing between any pair of distinct agents is defined as d/m, $0 \leq d \leq m$, and is independent of any other pair of agents in the market. Each agent a remains in the market for a *sojourn* $s(a)$ drawn independently from an exponential distribution with rate parameter $\lambda = 1$; the agent becomes *critical* immediately before her sojourn ends, and this criticality is known to the exchange. An agent leaves either upon being matched successfully by the exchange or upon becoming critical and remaining unmatched, at which point she perishes.

At any time $t \geq 0$, the network of acceptable transactions among agents forms a random graph $G_t = (A_t, E_t)$, where the agents in the exchange at time t form the vertex set A_t, and the acceptable transactions between agents forms the edge set E_t. We assume $A_0 = \emptyset$. Let A_t^n denote the set of agents who enter the exchange at time t, such that with probability 1, $|A_t^n| \leq 1$ for any $t \geq 0$. Finally, let $A = \cup_{t \leq T} A_t^n$.

Akbarpour et al. [4] present a parameterized space of online *matching policies*, with a focus specifically on two: Patient and Greedy. (In the next section, we will present a novel model of two overlapping exchanges, one running the Patient policy and the other running the Greedy policy.) As described above, vertex arrivals are treated as a continuous-time stochastic process. These policies behave as follows.

Greedy. The Greedy matching algorithm attempts to match each entering agent immediately by selecting one of its neighbors (if a neighbor exists at the time of entry) uniformly at random. One obvious consequence of this is that the remaining graph of unmatched agents at any instant is always empty. We refer to a market running this policy as the *Greedy market* or simply *Greedy* for the rest of the thesis.

Patient. The Patient matching algorithm attempts to match each agent only at the instant she becomes critical. As with Greedy, if a critical agent has multiple neighbors, only one is selected uniformly at random. We refer to a market running the Patient policy as a *Patient market* or simply *Patient* when appropriate.

If the random graph model is Erdős-Rényi [53] when not considering arrivals, departures, and matching, then the remaining graph at any instant is also Erdős-Rényi with parameter d/m; furthermore, d is the average degree of the agents. Both the Patient and Greedy policies maintain this observation.

The main result of Akbarpour et al. [4] is that waiting to thicken the market can be substantially more important than increasing the speed of transactions. Formally, the Patient exchange dramatically reduces the number of agents who perish (and thus leave the exchange without finding a match) compared to the Greedy exchange.

In the Akbarpour et al. [4] paper, an agent a receives zero utility if she perishes, or $u(a) = 0$. If she is matched, she receives a utility of 1 discounted at rate δ, or $u(a) = e^{-\delta s(a)}$. In this work, we focus on the special case of $\delta = 0$ in this paper (i.e., we only consider whether or not an agent is matched), and leave the $\delta \neq 0$ case for future research. Let $\text{ALG}(T) := \{a \in A : a \text{ is matched by ALG by time } T\}$. Then, in this model, the loss of an algorithm ALG is defined as the ratio of the expected number of perished agents to the expected size of A, as shown in Equation 4.1.

$$\mathbf{L}(\text{ALG}) = \frac{E[|A - \text{ALG}(T) - A_T|]}{E(|A|)} = \frac{E[|A - \text{ALG}(T) - A_T|]}{mT} \tag{4.1}$$

At any time $t \in [0, T]$, let $Z_{g,t}, Z_{p,t}$ represent the size of the pools under the Greedy and Patient matcing policies, respectively. Then, Akbarpour et al. [4] proved that the Markov chain on $Z_{\cdot,t}$ has a unique stationary distribution under either of those policies. Furthermore, let $\pi_g, \pi_p : \mathbb{N} \to \mathbb{R}^+$ be the unique stationary distribution of the Markov chain on $Z_{g,t}, Z_{p,t}$, respectively, and let $\xi_g := \mathbb{E}_{Z_g \sim \pi_g}[Z_g], \xi_p := \mathbb{E}_{Z_p \sim \pi_p}[Z_p]$ be the expected size of the pool under the stationary distribution under Greedy and Patient. Then, the following observations can be made.

Loss of Greedy. If a Greedy exchange is run for a sufficiently long time, then $\mathbf{L}(\text{Greedy}) \approx \frac{\xi_g}{m}$. The intuition here is that the Greedy pool is (almost) always an empty graph. Equation (4.2) formalizes the loss.

$$\mathbf{L}(\text{Greedy}) = \frac{1}{mT}\mathbb{E}\left[\int_0^T Z_{g,t} dt\right] = \frac{1}{mT}\int_0^T \mathbb{E}\left[Z_{g,t}\right] dt \tag{4.2}$$

Loss of Patient. If a Patient exchange is run for a sufficiently long time, at any point in time it is an Erdős-Rényi random graph. So once an agent becomes critical, she has no acceptable transaction with probability $(1 - d/m)^{Z_{p,t}-1}$. Thus, $\mathbf{L}(\text{Patient}) \approx \frac{\xi_p(1-d/m)^{\xi_p - 1}}{m}$. Equation (4.3) formalizes the loss of a Patient market.

$$\begin{aligned}\mathbf{L}(\text{Patient}) &= \frac{1}{mT}\mathbb{E}\left[\int_0^T Z_{p,t}(1 - d/m)^{Z_{p,t}-1} dt\right] \\ &= \frac{1}{mT}\int_0^T \mathbb{E}\left[Z_{p,t}(1 - d/m)^{Z_{p,t}-1}\right] dt\end{aligned} \tag{4.3}$$

4.3 Overlapping exchanges

The key result of Akbarpour et al. [4] is that a greedy dynamic matching market leads to significantly lower global social welfare than a patient matching market with full knowledge of criticality. The central question of this chapter is what happens in a situation where a greedy exchange and a patient exchange exist *simultaneously* and compete with each other to match some shared portion of the population. Agents in this overlapping subset of the population join both exchanges simultaneously and accept the first match offer from either of the constituent exchanges.

Drawing on Section 4.2, we model this in a similar stochastic, continuous-time framework as follows. Agents arrive at the Competing market (a model for the whole system, incorporating both the Greedy and Patient exchanges) at some rate m according to a Poisson process. For each agent, the probability of entering *both* the Greedy exchange and the Patient exchange is γ, the probability of entering the Greedy exchange alone is $(1-\gamma)\alpha$, and the probability of entering the Patient exchange alone is $(1-\gamma)(1-\alpha)$, where $\gamma, \alpha \in [0, 1]$. The probability that a bilateral transaction between each pair of agents is acceptable remains d/m, *conditioned on both agents being mutually "visible" to an exchange.* The agents' rates of perishing, received utility for being (un)matched, and other settings are otherwise the same as in Section 4.2.

We analyze the Competing market as three separate evolving pools:

Greedy$_c$ is the pool consisting of agents who enter the Greedy exchange only (with probability $\alpha(1-\gamma)$).

Patient$_c$ is the pool consisting of agents who enter the Patient exchange only (with probability $(1-\alpha)(1-\gamma)$).

Both$_c$ is the pool consisting of agents who enter both exchanges (with probability γ).

We use $\hat{Z}_{g,t}, \hat{Z}_{p,t}$ and $\hat{Z}_{b,t}$ to denote the size of Greedy$_c$, Patient$_c$ and Both$_c$, respectively, at any time t. Similar to an exchange running a single Greedy or Patient matching policy, the Markov chain on $\hat{Z}_{\cdot,t}$ also has a unique stationary distribution. Let $\hat{\pi}_\cdot : \mathbb{N} \to \mathbb{R}^+$ be the unique stationary distribution of the Markov chain on $\hat{Z}_{\cdot,t}$, and let $\hat{\xi}_\cdot := \mathbb{E}_{\hat{Z}_\cdot \sim \hat{\pi}_\cdot}[\hat{Z}_\cdot]$ be the expected size of the pool under the stationary distribution. Using this, we will define the loss of Greedy$_c$, $\hat{\mathbf{L}}(\text{Greedy}_c)$, the loss of Patient$_c$, $\hat{\mathbf{L}}(\text{Patient}_c)$, and the loss of Both$_c$, $\hat{\mathbf{L}}(\text{Both}_c)$.

First, note that the graph formed by the agents in Greedy$_c$ is empty, so the loss—as in Equation (4.2)—can be approximated by $\hat{\mathbf{L}}(\text{Greedy}_c) \approx \frac{\hat{\xi}_g}{m}$.

Next, we consider the agents in Both$_c$. If an edge exists between an agent in Both$_c$ and an existing agent in Greedy$_c$ or another agent in Both$_c$, she will be matched immediately by the Greedy exchange (and thus does not contribute to the loss). Similar to the Greedy$_c$ case, at any point in time t, the Both$_c$ pool is an empty graph; thus, any unmatched agents who become critical in Both$_c$ will only be matched to agents in Patient$_c$. Thus, these leftover agents in Both$_c$ have no acceptable transactions with probability $(1 - d/m)^{\hat{Z}_{p,t}}$. Since each agent becomes critical with rate 1, letting Competing market run for a sufficiently long time results in $\hat{\mathbf{L}}(\text{Both}_c) \approx \frac{\hat{\xi}_b(1-d/m)^{\hat{\xi}_p}}{m}$, where $\hat{\xi}_b, \hat{\xi}_p$ are the previously defined expected sizes of Both$_c$ and Patient$_c$.

Finally, we consider the Patient$_c$ pool. At any time t, the agents who remain in Patient$_c$ potentially have acceptable transactions with only the agents in Both$_c$ and the agents in Patient$_c$. Hence, in $\hat{Z}_{p,t}$, once an agent is critical, she has no acceptable transactions with probability $(1 - d/m)^{\hat{Z}_{p,t}+\hat{Z}_{b,t}-1}$. Similarly, each agent becomes critical with rate 1; thus, if we allow the Competing market a sufficiently long execution window, $\hat{\mathbf{L}}(\text{Patient}_c) \approx \frac{\hat{\xi}_p(1-d/m)^{\hat{\xi}_p+\hat{\xi}_b-1}}{m}$.

Because the three pools of agents—Greedy_c, Patient_c, and Both_c—are disjoint (although they may be connected via possible transactions in the ways listed above), we can define the total loss of the Competing market as follows.

$$\mathbf{L}(\text{Competing}) \approx \frac{\hat{\xi}_g + \hat{\xi}_p(1-d/m)^{\hat{\xi}_p+\hat{\xi}_b-1} + \hat{\xi}_b(1-d/m)^{\hat{\xi}_p}}{m}. \tag{4.4}$$

A more precise version of Equation (4.4) follow as Equation (4.5); we will make use of this form in Section 4.5.

$$\begin{aligned}\mathbf{L}(\text{Competing}) &= \frac{1}{mT}\mathbb{E}\Big[\int_0^T \hat{Z}_{p,t}(1-d/m)^{\hat{Z}_{p,t}+\hat{Z}_{b,t}-1} \\ &+ \hat{Z}_{b,t}(1-d/m)^{\hat{Z}_{p,t}} + \hat{Z}_{g,t}dt\Big] \\ &= \frac{1}{mT}\int_0^T \mathbb{E}\big[\hat{Z}_{p,t}(1-d/m)^{\hat{Z}_{p,t}+\hat{Z}_{b,t}-1} \\ &+ \hat{Z}_{b,t}(1-d/m)^{\hat{Z}_{p,t}} + \hat{Z}_{g,t}\big]dt\end{aligned} \tag{4.5}$$

Unfortunately, we do not have a closed form expression for the stationary distribution or the expected size of the pool under the stationary distribution. We note that each of $\hat{\xi}_g, \hat{\xi}_p$, and $\hat{\xi}_b$ can be approximated well using Monte Carlo simulations—thus, Equation (4.4) can be solved numerically. We do this in Section 4.5.1 for two parameterizations of the rival market setting.

4.4 A bound on total loss

While we do not have a closed form for the exact expected loss of the Competing market as described by Equation (4.4), we can provide bounds on the overall loss. In this section, we give one such bound for the global loss under the constraint that Greedy_c is more likely to receive agents than the overlapping Both_c exchange. Formally, this occurs when $\gamma \leq 0.5$ and $\alpha \geq \frac{\gamma}{1-\gamma}$. We also impose some loose requirements on the arrival rate of vertices to the exchange and the probability of an acceptable transaction existing between two agents; intuitively, the exchange cannot be "too small" or "too sparse," which we formalize below. Under these assumptions, we use the bound to prove Theorem 5, which states that a single Patient market outperforms the Competing market.

Theorem 5 *Assume $\gamma \leq 0.5$, $m > 10d$, and $\alpha(1-\gamma) \geq \max\left\{\gamma, \frac{1}{2}e^{-d/2}(1+3d)\right\}$. Then, as $m \to \infty$ and $T \to \infty$, almost surely*

$$\mathbf{L}(\mathit{Competing}) > \mathbf{L}(\mathit{Patient}).$$

Proof We prove the theorem by giving a lower bound on $\hat{\mathbf{L}}(\text{Greedy}_c)$, the loss of only the greedy portion of the Competing market. In our model, the fraction of agents entering only the Greedy_c side of the market is $\alpha(1-\gamma)$; for notational simplicity, we use $x := \alpha(1-\gamma)$ in this proof. Similarly, the fraction of agents entering Both_c is γ; again, for notational simplicity, we use $y := \gamma$ throughout this proof.

As before, let $\hat{Z}_{g,t}$ be the size of Greedy_c at any $t \in [0,T]$, and $\hat{\tau}$ the expected size of the Greedy_c pool. Similarly, let $\hat{Z}_{b,t}$ be the size of Both_c at any $t \in [0,T]$, and $\hat{\eta}$ the expected size of the Both_c pool. That is,

$$\hat{\tau} := \mathop{\mathbb{E}}_{t\sim\text{unif}[0,T]}\left[\hat{Z}_{g,t}\right] \quad \text{and} \quad \hat{\eta} := \mathop{\mathbb{E}}_{t\sim\text{unif}[0,T]}\left[\hat{Z}_{b,t}\right].$$

By assumption, $\alpha(1-\gamma) \geq \gamma$; that is, the arrival rate of Greedy_c is greater than or equal to the arrival rate of Both_c. In this case, $\hat{\tau} \geq \hat{\eta}$; the Greedy matching policy removes verties from both Both_c and Greedy_c, while the Patient matching policy removes vertices from only Both_c, which means the matching rate for Both_c is greater than the matching rate for Greedy_c.

From Akbarpour et al. [4], we know the expected rate of perishing of the individual Greedy exchange is equal to the pool size because the Greedy matching policy does not react to the criticality of an agent at any time t in its pool and each critical agent will perish with probability 1. Therefore, we can draw directly on Equation (4.2) to write

$$\hat{\mathbf{L}}(\text{Greedy}_c) = \frac{1}{xmT}\mathbb{E}[\int_{t=0}^{T} dt\ \hat{Z}_{g,t}] = \frac{\hat{\tau}}{xm}. \tag{4.6}$$

We know x and m, so lower bounding $\hat{\tau}$ will result in an analytic lower bound on $\hat{\mathbf{L}}(\text{Greedy}_c)$. Following the ideas of Akbarpour et al. [4], we do this by lower bounding the probability that an agent a does not *ever* have an acceptable transaction for the duration of her sojourn $s(a)$. Because these agents cannot be matched by any matching policy, this directly gives a lower bound on $\hat{\mathbf{L}}(\text{Greedy}_c)$. Toward this end, fix an agent $a \in A$ who enters Greedy_c at time $t_0 \in \text{unif}[0,T]$ and draws a sojourn $s(a) = t$. Let $f_{s_a}(t)$ be the probability density function at t of $s(a)$. Then we can write the probability that a will never have a neighbor (i.e., possible match) as

$$\mathbb{P}\left[N(a)=\emptyset\right] = \int_{t=0}^{\infty} f_{s_a}(t)\mathbb{E}\left[(1-d/m)^{\hat{Z}_{g,t_0}+\hat{Z}_{b,t_0}}\right]$$
$$\mathbb{E}\left[(1-d/m)^{|AG^n_{t_0,t_0+t}+AB^n_{t_0,t_0+t}|}\right]dt,$$

where $AG^n_{t_0,t_0+t}$ (resp. $AB^n_{t_0,t_0+t}$) denotes the set of agents who enter Greedy_c (resp. Both_c) in time interval $[t_0, t_0+t]$. The first expectation captures the probability that agent a has no matching at the moment of entry and the second expectation considers the probability that no new agents that can match with a arrive during her sojourn.

Using Jensen's inequality, we have

$$\begin{aligned}\mathbb{P}\left[N(a)=\emptyset\right] \geq& \int_{t=0}^{\infty} e^{-t}(1-d/m)^{\mathbb{E}[\hat{Z}_{g,t_0}+\hat{Z}_{b,t_0}]}\\ &(1-d/m)^{\mathbb{E}[|AG^n_{t_0,t+t_0}+AB^n_{t_0,t+t_0}|]}dt\\ =& \int_{t=0}^{\infty} e^{-t}(1-d/m)^{\hat{\tau}+\hat{\eta}}(1-d/m)^{(x+y)mt}dt.\end{aligned}$$

From the assumptions in the theorem statement, $\frac{d}{m} < \frac{1}{10}$, so $1-d/m \geq e^{-d/m-d^2/m^2}$. Also, as described earlier, $\hat{\tau} \geq \hat{\eta}$ (when $\gamma \leq 0.5$ and $\alpha \geq \frac{\gamma}{1-\gamma}$, as assumed). Therefore,

$$\begin{aligned}\hat{\mathbf{L}}(\text{Greedy}_c) &\geq \mathbb{P}\left[N(a)=\emptyset\right]\\ &\geq e^{-(\hat{\tau}+\hat{\eta})(d/m+d^2/m^2)}\\ &\quad\times \int_{t=0}^{\infty} e^{-t-(x+y)td-(x+y)td^2/m}dt\\ &\geq \frac{1-(\hat{\tau}+\hat{\eta})(1+d/m)d/m}{1+(x+y)d+(x+y)d^2/m}\\ &\geq \frac{1-2\hat{\tau}(1+d/m)d/m}{1+(x+y)d+(x+y)d^2/m},\end{aligned} \tag{4.7}$$

where the third inequality is obtained from the fact that $e^{-z} \geq 1 - z$ when $z \geq 0$, here $z = (\hat{\tau} + \hat{\eta})(d/m + d^2/m^2)$.

Combining Equation (4.6) and Equation (4.7),

$$\hat{\mathbf{L}}(\mathrm{Greedy}_c) = \frac{\hat{\tau}}{xm} \geq \frac{1 - 2\hat{\tau}(1 + d/m)d/m}{1 + (x+y)d + (x+y)d^2/m},$$

which gives us a lower bound for $\hat{\tau}$,

$$\hat{\tau} \geq \frac{xm}{1 + (3x+y)d + (3x+y)d^2/m}.$$

Thus, as $m \to \infty$, we get,

$$\begin{aligned} \hat{\mathbf{L}}(\mathrm{Greedy}_c) &\geq \frac{1}{1 + (3x+y)d + (3x+y)d^2/m} \\ &\geq \frac{1}{1+3d}. \end{aligned}$$

We are interested in bounding the total loss of the Competing market, which is $\mathbf{L}(\mathrm{Competing}) = x\hat{\mathbf{L}}(\mathrm{Greedy}_c) + (1-\alpha)(1-\gamma)\hat{\mathbf{L}}(\mathrm{Patient}_c) + y\hat{\mathbf{L}}(\mathrm{Both}_c)$. By definition, both $\hat{\mathbf{L}}(\mathrm{Patient}_c) \geq 0$ and $\hat{\mathbf{L}}(\mathrm{Both}_c) \geq 0$, and by Equation (), $\hat{\mathbf{L}}(\mathrm{Greedy}_c) \geq \frac{1}{1+3d}$. Thus,

$$\mathbf{L}(\mathrm{Competing}) \geq \frac{x}{1+3d}.$$

Akbarpour et al. [4] showed that running an *individual* Patient market results in exponentially small loss $\mathbf{L}(\mathrm{Patient}) < \frac{1}{2}e^{-d/2}$. Thus, as $T, m \to \infty$, we can get,

$$\mathbf{L}(\mathrm{Competing}) > \mathbf{L}(\mathrm{Patient}). \tag{4.8}$$

$\square$

We note that the result of Theorem 5 holds for only a section of the possible parameterizations of a Competing market—specifically, when $\gamma \leq 0.5$ and $\alpha \geq \frac{\gamma}{1-\gamma}$. In the next section, we will give numerical results showing that this result—that the loss of the Competing market is greater than the loss of an individual Patient exchange—appears to hold for a vastly larger space of values of γ and α. Indeed, experimentally, we will see that the loss of the Competing market is sometimes greater than the loss of an individual *Greedy* exchange, which itself is substantially greater than the loss of an individual Patient exchange.

4.5 Experimental validation

In this section, we provide experimental validation of the theoretical results presented in Sections 4.3 and 4.4. Section 4.5.1 quantifies the loss due to competing markets as described by Equation (4.4), while Section 4.5.2 expands the model to *kidney exchange* and draws from realistic data to quantify the loss of competing kidney exchange clearinghouses.

4.5.1 Dynamic matching

In Section 4.3, we gave a method for computing the expected loss due to competing markets as Equation (4.4); however, we were unable to derive closed forms for the expected size of the competing, patient, and greedy pools ($\hat{\xi}_b$, $\hat{\xi}_p$, and $\hat{\xi}_g$, respectively) under the stationary distribution. These quantities can be estimated using Monte Carlo simulation for different entrance rates m. We do that now.

Figures 4.1 and 4.2 simulate agents entering the Greedy_c, Both_c, and Patient_c markets according to a Poisson process with rate parameter $m = 1000$ and remaining for a sojourn drawn from an exponential distribution with rate parameter $\lambda = 1$. An agent chooses to enter Both_c with probability γ, only Greedy_c with probability $\alpha(1-\gamma)$, and only Patient_c with probability $(1-\alpha)(1-\gamma)$, as in the theory above. We vary $\alpha \in \{0, 0.1, \ldots, 1\}$ and $\gamma \in \{0, 0.1, \ldots, 1\}$, and plot the global loss realized for each of these parameter settings.

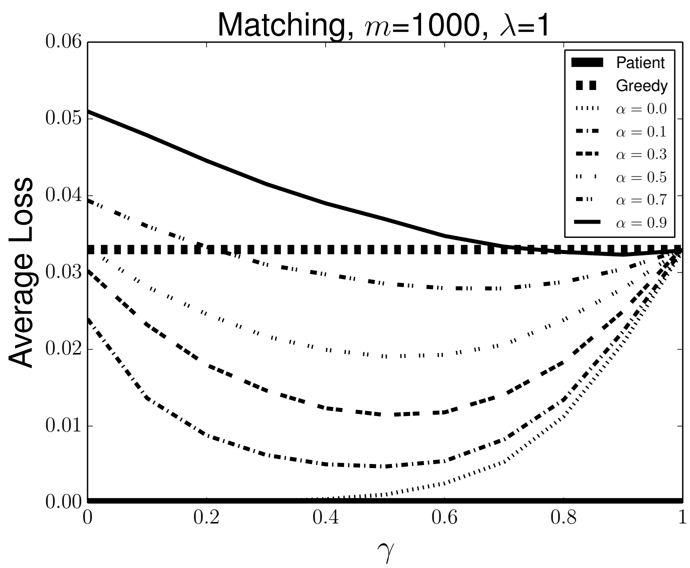

Figure 4.1: Average loss (y-axis) as the overlap between markets γ increases (x-axis), with entrance rate parameter $m = 1000$ and $d = 20$, for different values of α. The loss of individual Patient and Greedy markets are shown as thick black and thick dashed bars, respectively.

Immediately obvious is that running a single Patient market results in dramatically less loss than competing markets, for all different values of α and γ. Furthermore, we see that the loss of a single Greedy market is also dramatically higher than the loss of a single Patient market, as predicted by Akbarpour et al. [4]. Indeed, from Equation (4.3) we would expect the single Patient market to have essentially zero loss, so these experiments show that adding in a rival Greedy_c market increases loss. In fact, as the left side of Figure 4.1 and the right side of Figure 4.2 show, it is the case that if the markets do not overlap substantially (i.e., γ is low)

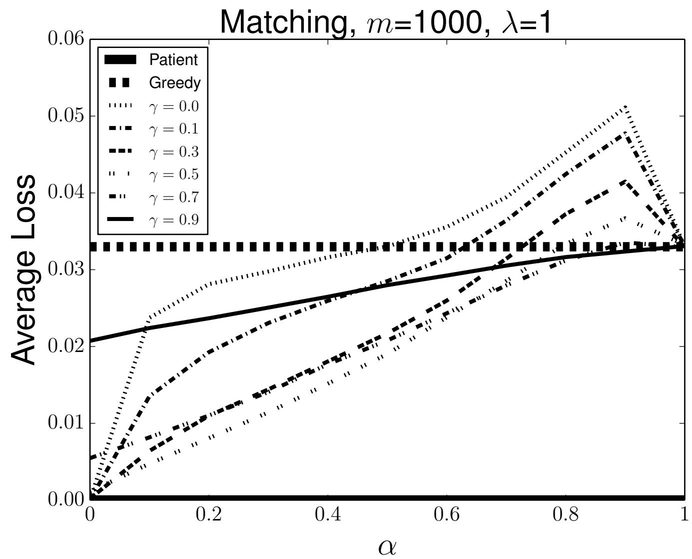

Figure 4.2: Average loss (y-axis) as the probability α of entering Patient$_c$ or Greedy$_c$ changes (x-axis), with entrance rate parameter $m = 1000$ and $d = 20$, for different values of the market overlap γ. The loss of individual Patient and Greedy markets are shown as thick black and thick dashed bars, respectively.

and agents are more likely to enter the greedy side of the market (i.e., α is near 1), then the loss of the competing market is worse than running a single Greedy market! This is due in part to the decrease in market thickness on the Patient$_c$ side of the market—a behavior we will see exacerbated below and in the kidney exchange experiments of Section 4.5.2.

Figure 4.3 decreases the rate parameter of the entrance Poisson process to $m = 100$, while holding the probability of an acceptable transaction between two agents at that of Figures 4.1 and 4.2 (so $d = 2$, leading to $2/100 = 2\%$). With fewer participants in the market overall, all the qualitative results of the $m = 1000$ markets above are amplified. The individual Greedy market's loss is now 5.9% worse than the individual Patient market (as opposed to 3.3% in the $m = 1000$ case); both individual markets' losses are substantially higher as well. Similarly, the parameter settings for which the competing market scenario has higher loss than either individual market are much broader than the $m = 1000$ case, which is a product of market thinness.

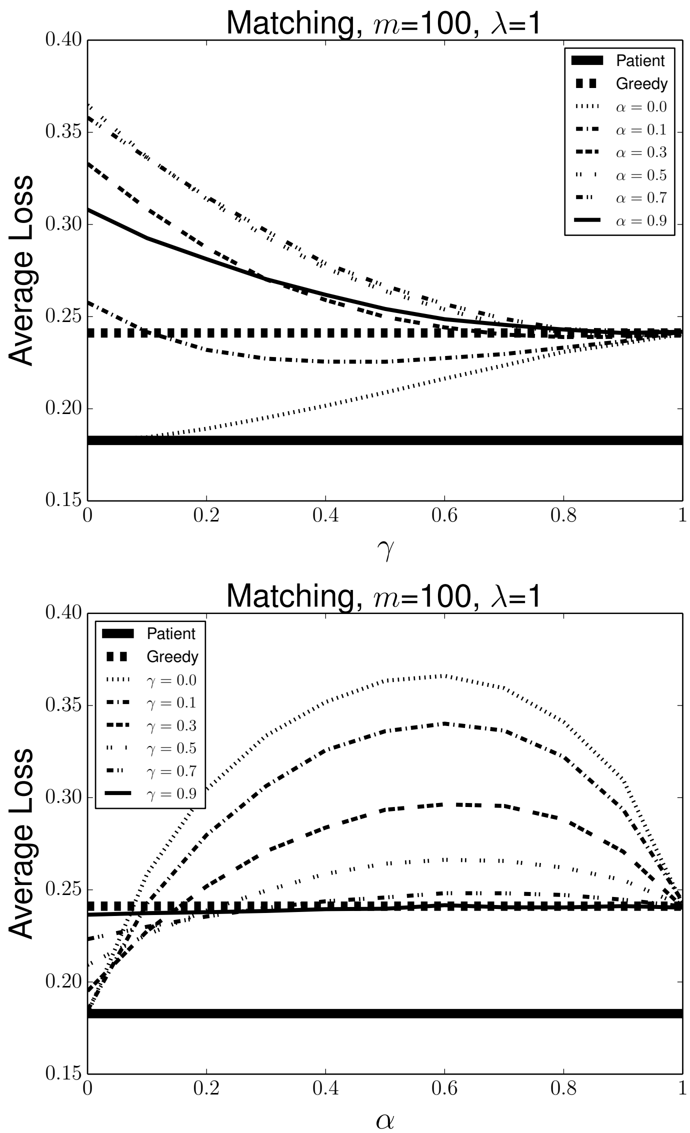

Figure 4.3: Average loss as the probability α of entering Patient$_c$ or Greedy$_c$ (top) or the overlap between the two markets γ (bottom) changes, with entrance rate parameter $m = 100$ and $d = 2$. The loss of individual Patient and Greedy markets are shown as thick black and thick dashed bars, respectively.

4.5.2 Dynamic kidney exchange

In this section, we expand our matching model to one of *barter exchange*, where agents endowed with items participate in directed, cyclic swaps of size greater than or equal to two. One recently-fielded barter application is *kidney exchange*, where patients with kidney failure swap their willing but incompatible organ donors with other patients. We focus on that application here. Dynamic barter exchange generalizes the matching model presented above, so we would not expect the earlier theoretical results to adhere exactly. Interestingly, as we show in Sections 4.5.2 and 4.5.2, the qualitative ranking of matching policy loss (with a patient market outperforming a greedy market, both of which outperform two rival markets) remains.

This section's experiments draw from two kidney exchange compatibility graph distributions. One distribution, which we call SAIDMAN, was designed to mimic the characteristics of a nationwide exchange in the United States in steady state [127]. Yet, kidney exchange is still a nascent concept in the US, so fielded exchange pools do not adhere to this model. With this in mind, we also include results performed on a dynamic pool generator that mimics the United Network for Organ Sharing (UNOS) nationwide exchange, drawing data from the first 193 match runs of that exchange. We label the distribution derived from this as UNOS.

Formally, we represent a kidney exchange pool with n patient-donor pairs as a directed *compatibility graph* $G = (V, E)$, such that a directed edge exists from patient-donor pair $v_i \in V$ to patient-donor pair $v_j \in V$ if the donor at v_i can give a kidney to the patient at v_j. Edges exist or do not exist due to the medical characteristics (blood type, tissue type, relation, and many others) of the patient and potential donor, as well as a variety of logistical constraints. Our generators take care of these details; for more information on how edge existence checking is done in the SAIDMAN and UNOS distributions, see Saidman et al. [127] or Dickerson and Sandholm [48], respectively. Importantly, under either distribution, there is no longer a costant probability "d/m" of an acceptable transaction existing between any two agents.

Vertices arrive via a Poisson process with rate parameter $m = 100$ and depart according to an exponential clock with rate parameter $\lambda = 1$ as before, and choose to enter either exchange or both with the previously-defined probabilities γ and α. However, a "match" now only occurs when a vertex forms either a 2-*cycle* or 3-*cycle* with one or two other vertices, respectively.[7] Section 4.5.2 performs experiments on 2-cycles alone, which adheres more closely to the theoretical setting above (2-cycles can be viewed as a single undirected edge between two vertices), while Section 4.5.2 expands this to both 2- and 3-cycles.

Kidney exchange with 2-cycles only

We now present results for dynamic matching under competing Patient_c and Greedy_c kidney exchanges, both of which use only 2-cycles. Figure 4.4 and Figure 4.5 show losses incurred in our parameterized market when run on SAIDMAN-generated and UNOS-generated pools, respectively.

While the barter exchange environment under either the SAIDMAN or UNOS distributions clearly breaks the structural properties of the stationary distribution of the underlying

[7] In fielded kidney exchange, cycles longer than some short cap L (e.g., $L = 3$ at the UNOS exchange and many others) are typically infeasible to perform due to logistical constraints, and thus are not allowed. We adhere to that constraint here. Fielded exchanges also realize gains from *chains*, where a donor without a paired patient enters the pool and triggers a directed path of transplants through the compatibility pool. We do not include chains in this work.

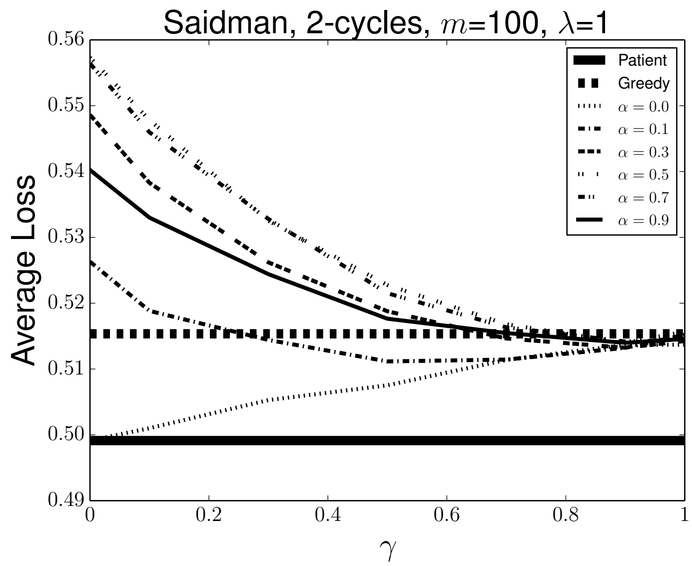

Figure 4.4: Average loss under various values of γ and α for the SAIDMAN distribution with 2-cycles only.

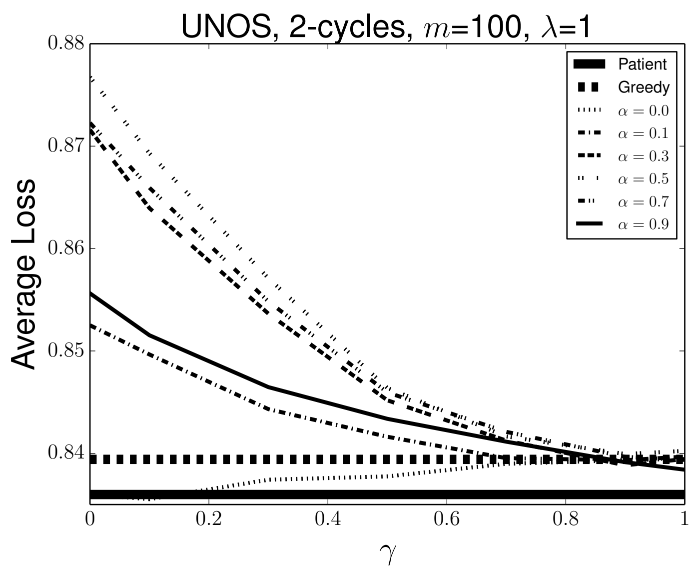

Figure 4.5: Average loss under various values of γ and α for the UNOS distribution with 2-cycles only.

Markov process used in our theoretical results, the qualitative results of these experiments align with the traditional dynamic matching results of Section 4.5.1. The overall loss realized by UNOS is substantially higher than that realized by SAIDMAN because, in general, UNOS-generated graphs are more sparse than those from the Saidman family. Similarly, in either distribution there exist "highly-sensitized" vertex types that are extremely unlikely to find a match with another randomly selected vertex, and thus almost certainly create loss. Indeed, both Figure 4.4 and 4.5 exhibit higher loss than the similarly-parameterized Figure 4.3 of Section 4.5.1.

Kidney exchange with both 2- and 3-cycles

We now extend our experiments to allow for "matches" that include both 2- and 3-cycles. Unlike Section 4.5.1 or 4.5.2, where a matched edge was chosen uniformly at random from the set of all acceptable transactions between a distinguished vertex and its neighbors, in these results we may wish to distinguish a potential match from others (for example, by choosing a 3-cycle before a 2-cycle, as the former results in a larger myopic decrease in the market's loss). Thus, given a set of possible 2- and 3-cycle matches, we consider two matching policies: Uniform selects a cycle at random from the set of possible matches, regardless of cycle cardinality, while Uniform3 selects a 3-cycle randomly (if one exists), otherwise a random 2-cycle.

Figures 4.6 and 4.7 show results for the Saidman and UNOS distributions, respectively, under the Uniform match selection policy. Intuitively, one might expect the loss of a matching policy run in the 2- and 3-cycle case to be less than the same policy run in the 2-cycle case alone, as the set of possible matches weakly increases in the former case. We see this behavior when comparing the Saidman results of Figure 4.6 to the earlier 2-cycle-only Saidman results of Figure 4.4, witnessing a drop in global loss of around 4% for any parameter setting. We see a similar decrease in loss when comparing the new UNOS results of Figure 4.7 to those in the 2-cycle case shown in Figure 4.5.

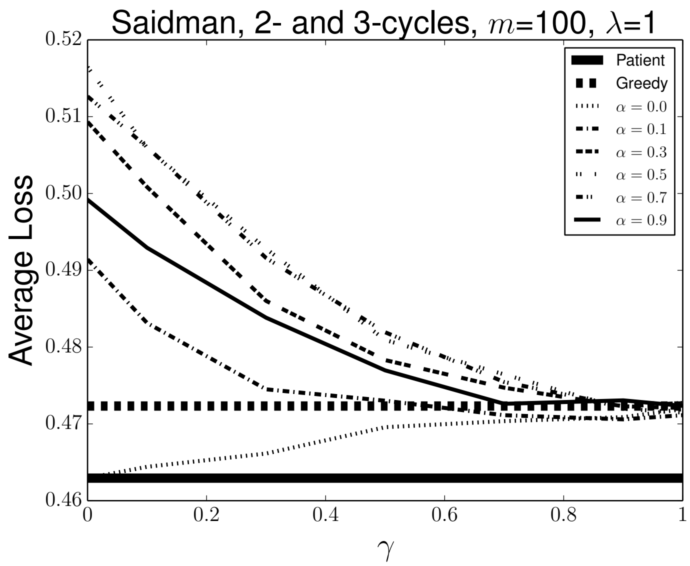

Figure 4.6: Average loss under various values of γ and α for the SAIDMAN distribution with both 2- and 3-cycles, under the UNIFORM matching policy.

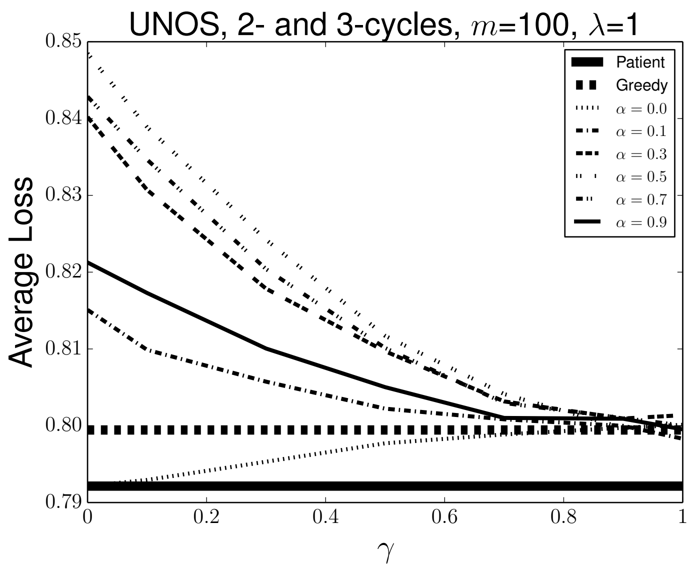

Figure 4.7: Average loss under various values of γ and α for the UNOS distribution with both 2- and 3-cycles, under the UNIFORM matching policy.

We now consider the UNIFORM3 matching policy, which would likely be closer to how a fielded exchange would act. Figures 4.8 and 4.9 show results for the SAIDMAN and UNOS families of compatibility graphs, respectively. The loss of the individual Patient market does not change in either distribution, which is likely a byproduct of the thicker markets induced by its match cadence. Curiously, the loss of the individual Greedy market drops dramatically—to around the Patient loss in the UNOS case, and *below* Patient in the SAIDMAN case. This large drop in Greedy loss is likely due in part to Greedy now "poaching" larger 3-cycles from the leftover market from which the Patient policy draws. The other qualitative results of earlier sections are repeated, with rival markets hurting global loss relative to either individual market for nearly all settings of γ and α.

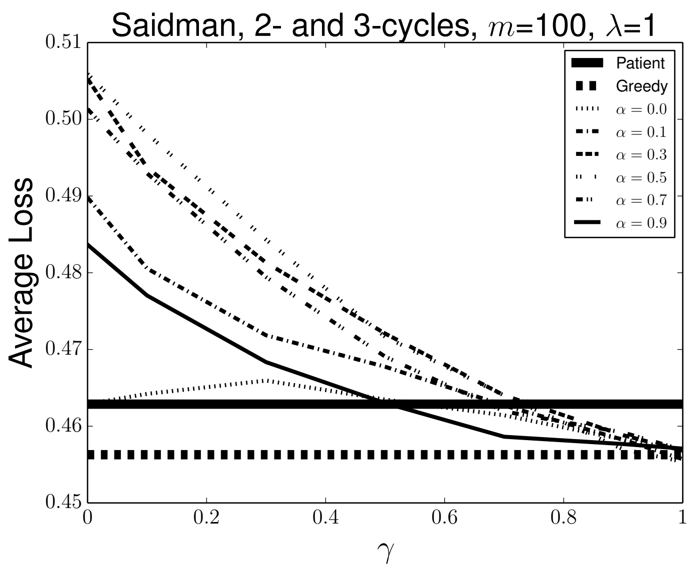

Figure 4.8: Average loss under various values of γ and α for the SAIDMAN distribution with both 2- and 3-cycles, under the UNIFORM3 matching policy.

4.6 Discussion

Our main goal is to study the impact of competition between exchanges in a dynamic matching setting. In this chapter, we extended the recent dynamic matching model of Akbarpour et al. [4] to two rival matching markets with overlapping pools. Specifically, we formalized a two-market model where agents enter one market or both markets; they can then potentially

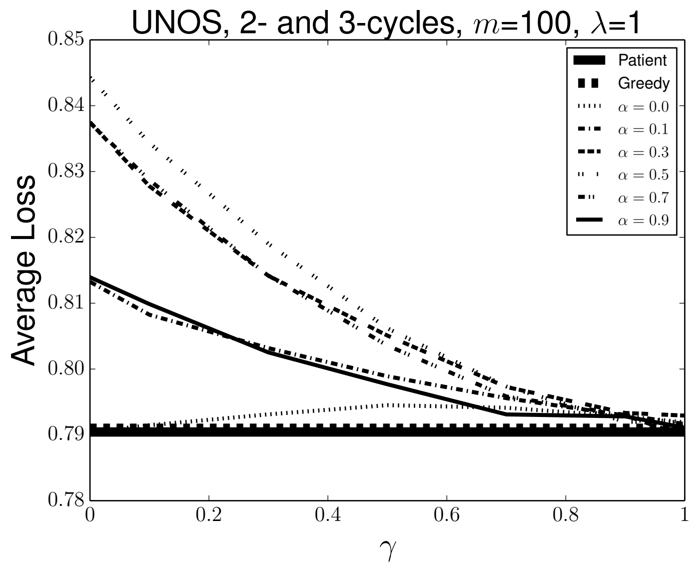

Figure 4.9: Average loss under various values of γ and α for the UNOS distribution with both 2- and 3-cycles, under the UNIFORM3 matching policy.

be matched to other agents who have joined the same market or both markets. The markets, called Greedy and Patient, adhere to different matching policies. We provided an analytic lower bound on the *loss* of the two-market model and showed that it is higher than running a single Patient market. We also provided a quantitative method for determining the loss of the two-market model. We supported these theoretical results with extensive simulation. We also looked at competing *kidney exchanges*, and provided (to our knowledge) the first experimental quantification of the loss in global welfare in a setting with two clearinghouses using realistic kidney exchange data drawn from a generator due to Saidman et al. [127] and another based on the United Network for Organ Sharing (UNOS) program.

We see competing dynamic matching markets as fertile ground for future research, with a trove of both theoretical and practical questions to answer. First, the model of Akbarpour et al. [4] discounts the utility of a match by the time the matching agent has already waited in the pool; this is well motivated in a variety of settings, including kidney exchange. Our results in this chapter assume a discount factor of zero, so it would be valuable to consider the impact on discounted loss for non-zero cases. Second, in our model the choice of market to enter is exogenously determined for each agent. In reality, agents with different levels of knowledge, wealth, etc. may make strategic decisions on which markets to enter. Thus, one could approach this dynamic matching problem from a game-theoretic point of view. Similarly, taking network effects (where more popular exchanges have an easier time attracting agents, lower operating costs, higher probabilities of two agents forming an acceptable transaction, and other advantages) into account would make these models more applicable to many real-world settings. Finally, we only looked at two overlapping markets; generalizing this to any number of overlapping markets would also be of interest.

In terms of barter exchange and, specifically, kidney exchange, the question of how clearinghouses interact is a timely one. In the United States and, eventually, elsewhere, multi-center and single-center exchange clearinghouses are already competing, each drawing from some (often overlapping) subset of the full set of patient-donor pairs available. Indeed, the dynamic barter exchange problem *in a single market* is still not fully understood (barring very promising recent work due to Anderson et al. [6]). We saw in Section 4.5.2 that including 3-cycles in the matching process results in lower loss, even when two markets overlap, compared to including only 2-cycles (a result that has been shown repeatedly in the static [124] and dynamic [6] single clearinghouse setting), so extending the theoretical underpinnings of our framework to a more general setting would be of great value. Finally, it is curious that the UNIFORM3 policy had such a large effect on the loss of the individual Patient and Greedy exchanges compared to the UNIFORM policy; further exploration of different matching policies (including those that use a strong prior to consider possible future states of the pool when matching now) would be helpful in making policy recommendations to fielded exchanges.

Chapter 5

Equilibrium Behavior in Competing Dynamic Matching Markets

Built on the two-market model of Chapter 4, this chapter addresses questions about the equilibrium behavior of competing dynamic matching market systems. Instead of using stochastic participants and fixed policies, we now consider the strategic behaviors: first from the points of view of individual participants when market policies are fixed, and then from the points of view of markets when agents are stochastic.

First, we analyze models where individual market participants have agency. These participants can be of different types (short-lived or long-lived) and may choose entrance into the market system such that their individual utility is maximized. Different types of agents may have different preferences, and we analyze equilibrium behavior in both continuous (Section 5.2) and discrete (Section 5.3) time settings. We show that even with just two types of agents, strategic market choice can induce market fragmentation—while there are some pooling equilibria where all strategic agents choose the same market (which is socially preferable), separating equilibria become significantly more likely, as the proportion of agents who are assigned to a particular market increases (these agents may be constrained by geography or cost, for example), and with short-lived agents choosing the patient market and long-lived agents the greedy market. This is because the patient market is typically thicker, giving a higher probability of matching during an agent's sojourn, and short-lived agents suffer less penalty because the market attempts to match them sooner relative to arrival. Unfortunately, the fragmentation comes at significant social cost in reduced thickness.

Second, in Section 5.4, we prescribe agency to the markets themselves, allowing them to choose overall matching policies (defined by the frequency at which they decide to match) strategically to maximize their overall utility. In this case, the agents are *stochastic* in their choice to join one or the other market, or to enter both markets. We quantify via best response dynamics the social welfare loss of this competitive marketplace under a variety of initial conditions, and compare that loss to the lower bound provided by a single market running an optimal matching policy.

Overall, our results demonstrate the serious concern of a "race to the bottom" when multiple matching markets compete. This is due to both fragmentation and the choice of socially suboptimal matching policies by individual markets. When agents choose markets strategically, differences in their types and utilities can lead to preferences for one or the other markets and induce separating equilibria and fragmented markets. Even when agents do not have market choice, if markets can choose their matching policies, individual markets may be incentivized to match as early as possible an inefficient fraction of the time in the race to match more agents. The intersection of differential impact on different types and competing matching platforms raises important ethical issues in allocation and regulation. Such discussions can be informed by our models. Further, our models can also provide the foundation for future models that consider situations where *both* agents and markets can be strategic. Section 5.5 concludes with some recommendations for policymakers derived from our results.

5.1 Preliminaries

We still follows Greedy polices and Patient polices as we discussed in Chapter 4. Besides, the market can also choose a clearing rule that interpolates between the Patient and Greedy clearing rules (the so-called Patient(α) clearing rule), which allows tuning of the matching rate. Specifically, a market matching with the Patient(α) strategy draws an exponential random variable C_v with rate parameter $1/\alpha$ for each vertex v. If vertex v entering at time t becomes critical at time $t_c < t + C_v$, she matches at t_c, as in the Patient matching algorithm. Otherwise the vertex matches at time $t + C_v$. Note that when $\alpha \to 0$ we will have $C_v \xrightarrow{p} 0$, which corresponds to a Greedy matching algorithm.

5.1.1 Model I: Strategic Agents

Our first model considers two types of agents in terms of length of life, *short-lived* and *long-lived*. Short-lived agents come into the markets with a length of life T_s and long-lived agents have a length of life T_l, where $T_s < T_l$. Each agent (who is aware of her own type) decides which market to enter upon arrival. A fraction θ of agents are short-lived and the remaining $1 - \theta$ fraction are long-lived. We allow a ϕ fraction of random-choice agents (*random* agents) to choose either market with 0.5 probability. The remaining $1 - \phi$ fraction of agents are *strategic*. For these models, we restrict attention to models in which one Greedy and one Patient market compete. The action space for agents is the market choice, $B = \{Greedy, Patient\}$. We want to analyze the equilibrium strategies of strategic agents given the setting of θ and ϕ. Here, the market choice becomes a tradeoff between matching probability and utility. That is, entering a Patient market may give an agent a higher matching probability but lower utility as the agent has a higher expected sojourn time; in contrast, immediate matching from a Greedy market provides a higher utility but may lower the probability of matching since the market is not thick enough.

We investigate the behavior of strategic agents in the two-market Model I in both continuous time (Section 5.2) and discrete time (Section 5.3) models.

5.1.2 Model II: Strategic Markets

Our second, complementary, direction is to model the situation where agent behaviors are stochastic, but markets themselves make strategic decisions. We define each market's utility as the aggregate utility of the (non-strategic) agents it matches (it is reasonable to assume that the market can capture some fraction of this utility). We follow the model in Section 4.3 for assigning agents to one or both of the two competing markets. A γ_1 fraction of agents are assigned to both markets; the market which successfully matches the agent first will receive utility from the match. The remaining agents are only assigned to one market: a γ_2 fraction enter the first market, while a $1 - \gamma_2$ fraction enter the second market.

The action chosen by a market is its choice of market-clearing rule, parameterized by the matching rate α described above. The market-clearing rule choice involves a tradeoff: if Market 1 chooses a fast matching rate, it will match more agents assigned to both markets, but will match fewer agents which are only assigned to Market 1. The relative market sizes,

	Shared Notation		
p	Probability of potential transaction existing between two agents		
T	Number of time periods the markets operate		
k	Exponential rate of Poisson process entering market		
	Model I Notation		Model II Notation
δ	Discount rate used by short- & long-lived agents	$M_{\{1,2\}}$	Strategic market 1 (2)
a	A short- or long-lived agent	γ_1	Probability a vertex enters both markets
$s(a)$	The sojourn time of agent a	γ_2	Probability a vertex enters M_1 only
u	The utility function of a discounting agent	$\alpha_t^{\{1,2\}}$	Exponential rate parameter for matching
$T_{\{s,l\}}$	The length of life of a short-(long-)lived agent		rate in market M_1 (M_2) at time t
θ	Fraction of short-lived agents	d	Exponential rate parameter for criticality
ϕ	Fraction of random-choice agents	T_R	Time between updates of clearing rates
$m_{g,e}$	Pr. acceptable transaction when entering Greedy (Patient)	$u_{\{M_1,M_2\}}$	Utility function of strategic market $M_{\{1,2\}}$
$m_{\{g,p\},s}$	Pr. acceptable transaction staying in Greedy (Patient)		
$m_{p,c}$	Pr. of acceptable transaction when critical in Patient		
$\Lambda_{p,t}$	Number of agents critical at time t in Patient market in the discrete model, or rate of perishing in the continuous model		
$Z_{\{g,p\},t}$	Size of Greedy (Patient) market at time t		
$U_{\{s,l\},\{g,p\}}$	Expected utility of short- (long-)lived agents choosing Greedy (Patient) market		

Table 5.1: Variable definitions in Chapter 5.

parameterized by γ_1 and γ_2 are factors in the optimal choice. We investigate equilibrium behavior via simulation of two markets in Model II in Section 5.4.

5.2 Strategic Agents in Continuous Time

We consider two markets operating simultaneously, one Greedy and one Patient. For simplicity, we assume that lengths of life T_s and T_l for short-lived and long-lived agents are constants that are fixed across the same type of agents.[8]

The markets run in the continuous-time interval $[0, T]$. Agents arrive according to a Poisson process with rate parameter $k \geq 1$ ($k = 100$ in our simulations). The type of each arriving agent is stochastic; with probability θ, the arriving agent is a short-lived type; and with

[8]We also ran experiments where T_s and T_l are sampled from two exponential distributions with different rate parameters λ_s and λ_l, truncated so that $T_s < 1$ and $T_l \geq 1$. The results were qualitatively very similar to the case where T_s and T_l are constants.

probability $1-\theta$, she is a long-lived type. Parameter ϕ controls whether an agent is random or strategic, that is, with probability ϕ, she is a random agent and w.p. $1-\phi$ she is a strategic agent. Upon arrival, the agent needs to decide which market to enter. Random agents choose a market uniformly at random and strategic agents choose a market based on comparing the expected utilities of entering each market.

We first consider agents entering the Greedy market. As the Greedy market matches agents immediately upon entry, the probability of an agent having acceptable transactions immediately after entering at any time t is $m_{g,e}(t) = (1-(1-p)^{Z_{g,t}-1})$, where $Z_{g,t}$ represents the size of the pool under the Greedy matching policy at time t. To be noticed, t here is an infinitesimal time. Since entry occurs stochastically in continuous time, only one agent enters exactly at time t. Therefore, as long as there exist any acceptable transactions, the entering agent will be matched immediately. Once the moment of entry has passed, an agent can only be matched at the point in time when some other agent enters the market. The probability of an agent who was not matched at entry having an acceptable transaction at the time of entry of some other agent is $m_{g,s}(t) = \frac{(1-(1-p)^{Z_{g,t}-1})}{Z_{g,t}-1}$. Denote the probability of an agent entering the Greedy market at any point in time t as $P[\text{Entry}_g^t]$. Thus, the expected utility of an agent for choosing the Greedy market $U_{type,g}(t)$ at time t given she knows her type is

$$\begin{aligned} U_{type,g}(t) = & m_{g,e}(t) + \textstyle\int_0^{T_{type}} P[\text{Entry}_g^t] \\ & m_{g,s}(t+s(a))e^{-\delta s(a)} ds(a), \end{aligned} \tag{5.1}$$

where $type \in \{short, long\}$.

Now consider agents entering the Patient market. The Patient market attempts to match agents at the instant they become critical. The probability of an agent having acceptable transactions during their stay (before perishing) at any time t is $m_{p,s}(t) = \Lambda_{p,t}\frac{(1-(1-p)^{Z_{p,t}-1})}{Z_{p,t}-1}$, where $\Lambda_{p,t}$ is the rate of perishing in the Patient market and $Z_{p,t}$ is the size of the pool under the Patient matching policy at time t. The probability of an agent having acceptable transactions at the instant she becomes critical is $m_{p,c}(t) = (1 - (1-p)^{Z_{p,t}-1})$. Denote the probability that some agent in the Patient market becomes critical at any time t as $P[\text{Exit}_p^t]$. The expected utility of an agent for choosing patient market $U_{type,p}(t)$ at time t given her type is

$$U_{type,p}(t) = \textstyle\int_0^{T_{type}-\epsilon} P[\text{Exit}_p^t] m_{p,s}(t+s(a)) \\ e^{-\delta s(a)} ds(a) + m_{p,c}(t+T_{type}) e^{-\delta T_{type}}, \tag{5.2}$$

where ϵ is an infinitesimal amount of time right before an agent perishes and $type \in \{short, long\}$.

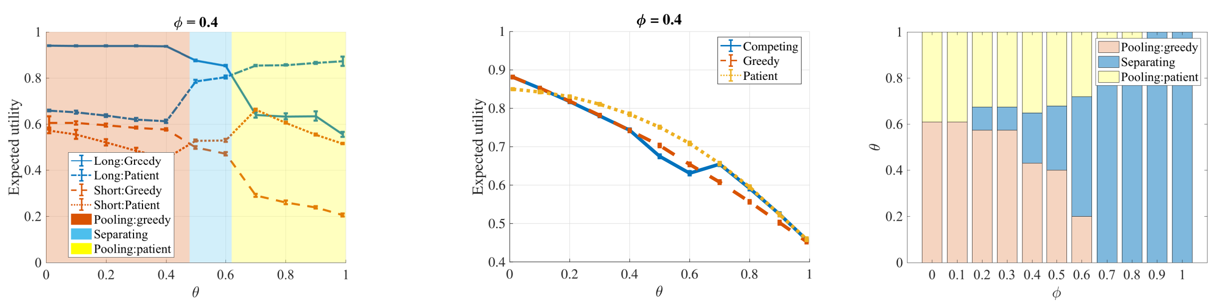

(a) (b) (c)

Figure 5.1: Results of continuous market for $p = 0.02, \delta = 0.05, T_s = 2, T_l = 3$. *Long:Greedy* (*Short:Greedy*) and *Long:Patient* (*Short:Patient*) show the expected utility of a strategic long-type (short-type) agent if she chooses Greedy and Patient respectively.

Equations (1) and (2) clarify the tradeoffs agents face. In general, while the patient market may give a higher probability of finding a match, the fact that an agent typically has to wait longer diminishes her expected utility. Since agents start with the same utility and it diminishes at the same rate, this means that short-lived agents will have a relatively higher preference for the Patient market compared with long-lived agents (who have to wait longer until the point in time when they are most likely to get matched, the time of criticality, in the patient market). Since there are positive externalities to entering a market and making it thicker, we may expect that the market-choice game may have both pooling and separating equilibria, where either both types of agents enter one market or short-lived agents enter the Patient market while long-lived agents enter the Greedy market.

Since the equations above do not admit closed-form solutions, we use empirical game-theoretic analysis to find equilibria in the game with strategic market-choice for each type. The strategy space is $B = \{Greedy, Patient\}$. For different values of θ and ϕ, we compute the utilities of strategic short-lived and long-lived agents if they choose the Greedy market or the Patient market respectively using Monte Carlo simulations holding the strategies of the other agents fixed, and ascertain whether or not pooling or separating equilibria exist in different regions of the θ, ϕ space. As conjectured, we do see an overall pattern of pooling and separating equilibria in different regions. Figure 5.1a shows an example of the results when the fraction of random agents is $\phi = 0.4$. These results can be broken up into three regions: The red region represents pooling equilibria where both long-type and short-type strategic agents choose the Greedy market; long-lived and random agents are the majority in this region, thus the Greedy market can be thick enough.[9] The yellow region represents pooling equilibria where both types choose the Patient market; in this range of settings, we have more short-lived agents and the Greedy market is not thick enough as the short-lived agents perish too soon.[10] In the blue region, we find a separating equilibrium exists: strategic short-lived types choose the Patient market ($U_{s,p} > U_{s,g}$) and long-lived types choose the Greedy market ($U_{l,g} > U_{l,p}$).

[9]Note that, in this region, both types of agents choosing the Patient market is also an equilibrium, albeit one with overall lower social welfare.

[10]Similarly, in this region, both types of agents choosing the Greedy market is also an equilibrium, lower in social welfare than the Patient pooling equilibrium.

Figure 5.1b shows overall social welfare in the Competing system with a single Greedy market and a single Patient market under different settings of θ when the fraction of random agents $\phi = 0.4$. We can see the market fragmentation caused by competition, separating equilibria ($\theta \in [0.4, 0.6]$), lowers the social welfare when compared to a single market. This pattern holds across the whole range of ϕ.

Finally Figure 5.1c shows the range of separating and pooling equilibria as a function of ϕ, the proportion of random agents. As the proportion of random agents increases, the portion of the θ domain covered by separating equilibria increases, since the thicknesses of the two markets are determined almost entirely exogenously, and the main consideration is an optimization of utility rather than equilibrium considerations of what other strategic agents are doing.

5.3 Strategic Agents in Discrete Time

While the model of Section 5.2 uses essentially the same models of utility as prior work, we are restricted by the lack of analytical tractability. We now consider a discrete time version of Model I that captures the same basic intuitions and can be used more directly in modeling strategic market choice. We believe this model is more amenable for further work on these questions. Now, agents enter the market in "batches", that is, $k \geq 1$ agents enter the Competing market at each time step t. Short-lived agents live for T_s time steps and long-lived agents live for T_l time steps ($T_l > T_s$), where T_s and T_l are fixed constants for each type. At each time step t, each market operates as follows: o_1: agents enter $\rightarrow$ o_2: market clears $\rightarrow$ o_3: agents perish. Random agents choose a market to enter uniformly at random and strategic agents choose a market based on comparing the expected utilities of entering each market.

We first analyze the utility of agents choosing the Greedy market. As the Greedy policy will match agents immediately after they enter the market, the probability of an agent having acceptable transactions immediately after entering is $m_{g,e}(t) = (1-(1-p)^{Z_{g,t}-1})$, where $Z_{g,t}$ represents the size of the pool under the Greedy matching policy at time t. The market will run a maximum matching algorithm at each time step during o_2 as $k \geq 1$ agents enter the market at the same time. This means that agents may be unmatched even if they have potential acceptable transactions. We define the probability of being matched in the maximum matching given the agent has acceptable transactions as $\chi_g(t)$ in the Greedy market at time t. The probability of an agent having acceptable transactions when they stay in the market (that is, not at their time-step of entry) at time t is $m_{g,s}(t) = (1-(1-p)^{k_{g,t}})$, where $k_{g,t}$ is the number of agents entering to the Greedy pool at time t. Thus, the expected utility of an agent for choosing Greedy market $U_{type,g}$ at time t given she knows her type is

$$\begin{aligned} U_{type,g}(t) = m_{g,e}(t)\chi_g(t) + (1 - m_{g,e}(t)\chi_g(t)) \\ \Big[e^{-\delta} m_{g,s}(t+1)\chi_g(t+1) + \\ \textstyle\sum_{s(a)=2}^{T_{type}-1} e^{-\delta s(a)} m_{g,s}(t+s(a))\chi_g(t+s(a)) \\ \textstyle\prod_{j=1}^{s(a)-1}(1 - m_{g,s}(t+j)\chi_g(t+j)\Big], \end{aligned} \tag{5.3}$$

where $type \in \{short, long\}$ and $T_{type} \geq 3$. We have two special cases, where

$$U_{type,g}(t) = m_{g,e}(t)\chi_g(t) \text{ when } T_{type} = 1;$$

and

$$U_{type,g}(t) = m_{g,e}(t)\chi_g(t) + (1 - m_{g,e}(t)\chi_g(t))m_{g,s}(t+1)\chi_g(t+1)e^{-\delta} \text{ when } T_{type} = 2.$$

We next consider the expected utility of agents choosing the Patient market. The Patient market will match agents only at the instant they become critical. The probability of an agent having acceptable transactions when they stay in the Patient market at each time step t is $m_{p,s}(t) = (1-(1-p)^{\Lambda_{p,t}})$, where $\Lambda_{p,t}$ is the number of agents becoming critical in the Patient market at time t. As there may be more than one agent becoming critical at each time t, the Patient market will also run a maximum matching at o_2. We define the probability of being matched in the maximum matching given the agent has acceptable transactions as $\chi_p(t)$. The probability of an agent having acceptable transactions when she is critical is $m_{p,c}(t) = (1-(1-p)^{Z_{p,t}-1})$, where $Z_{p,t}$ is the size of the pool under the Patient matching policy at time t. Thus, the expected utility of an agent for choosing Patient market $U_{type,p}$ at time t given she knows her type is

$$\begin{aligned} U_{type,p}(t) = m_{p,s}(t)\chi_p(t) + \\ &\Big[\textstyle\sum_{s(a)=1}^{T_{type}-2} e^{-\delta s(a)} m_{p,s}(t+s(a))\chi_p(t+s(a)) \\ &\textstyle\prod_{j=0}^{s(a)-1}(1-m_{p,s}(t+j)\chi_p(t+j))\Big] + \\ &\textstyle\prod_{j=0}^{T_{type}-2}(1-m_{p,s}(t+j)\chi_p(t+j)) \\ &m_{p,c}(t+T_{type}-1)\chi_p(t+T_{type}-1)e^{-\delta(T_{type}-1)}, \end{aligned} \tag{5.4}$$

where $type \in \{short, long\}$ and $T_{type} \geq 3$. We also have two special cases, where

$$U_{type,p}(t) = m_{p,c}(t)x_p(t) \text{ when } T_{type} = 1;$$

and

$$U_{type,p}(t) = m_{p,s}(t)\chi_p(t) + (1-m_{p,s}(t)\chi_p(t))m_{p,c}(t+1)\chi_p(t+1)e^{-\delta} \text{ when } T_{type} = 2.$$

At any time $t \in [0, T]$, $Z_{g,t}, Z_{p,t}$ represent the sizes of the pools under the Greedy and Patient matching policies, respectively. The Markov chain on $Z_{\cdot,t}$ has a unique stationary distribution under either of those policies. Let $\pi_g, \pi_p : \mathbb{N} \to \mathbb{R}^+$ be the unique stationary distributions of the Markov chain on $Z_{g,t}, Z_{p,t}$, respectively, and let $\xi_g := \mathbb{E}_{Z_g \sim \pi_g}[Z_g], \xi_p := \mathbb{E}_{Z_p \sim \pi_p}[Z_p]$ be the expected sizes of the pool under the stationary distribution under Greedy and Patient. After mixing, we represent the expected sizes of the pools at any time as ξ_g, ξ_p respectively. Similarly, $k_{g,t}, \Lambda_{p,t}, \chi_g(t)$ and $\chi_p(t)$ also can be represented by expected values k_g, Λ_p, χ_g and χ_p. We use Monte Carlo simulations to estimate $\xi_g, \xi_p, k_g, \Lambda_p, \chi_g$ and χ_p respectively. This then allows us to numerically compute the expected utilities in Equation (5.3) and (5.4) and derive the equilibria for different parameter settings.

Figure 5.2 shows an example of the results when the fraction of short-lived agents is $\phi = 0.4$. The results are qualitatively very similar to those from the continuous-time model, but the additional analytical tractability of the model presented here makes it promising for future development of models of competing markets.

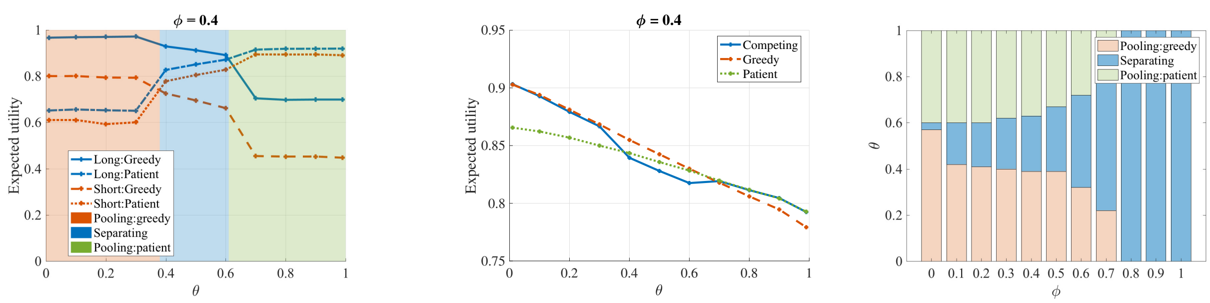

(a) (b) (c)

Figure 5.2: Results of discrete market for $p = 0.02, \delta = 0.05, T_s = 2, T_l = 3$. *Long:Greedy* (*Short:Greedy*) and *Long:Patient* (*Short:Patient*) show the expected utility of a strategic long-type (short-type) agent if she chooses Greedy and Patient respectively.

5.4 Strategic Markets

In the previous two sections, we assumed all markets operated with *fixed* matching policies, and strategic agents entered that system in a way that maximized their individual expected utility. Here, under Model II, we prescribe agency onto the markets themselves, allowing them to strategically *adjust* their matching policies under best response dynamics to maximize their expected aggregate utility. We investigate equilibrium behavior in this model, and measure overall social welfare loss relative to a single-market baseline.

5.4.1 Experimental Setup

We are interested in modeling the behavior of a two-market system where the markets respond to each other under best response dynamics. Formally, at any time period, one market observes the matching rate of its competitor and then chooses, *for the next time period*, its own matching rate that will yield maximum payoff for perpetuity,[11] even though the market will change its best response within a short span of time T_R. Best response dynamics have been shown to mimic many settings where agents operate reactively or with bounded expertise [146], and can be used in some cases to find equilibria [109].

We simulated the long-term utilities for two markets M_1 and M_2 with Patient(α^1) and Patient(α^2) matching policies, respectively, for $(\alpha^1, \alpha^2) \in \mathbb{R}^{\geq 0} \times \mathbb{R}^{\geq 0}$, for $T = 250$ periods, and 100 trials. We estimated the best response functions BR_1 and BR_2 for markets M_1 and M_2, respectively by simulating two markets with overlaps $\gamma_1 \in \{0.1, 0.2, \ldots, 0.9\}$ over a grid of patience parameters α^1, α^2. As a reminder, higher values of α correspond to more

[11] For a formal overview of best response dynamics, see, for example, the book by [110].

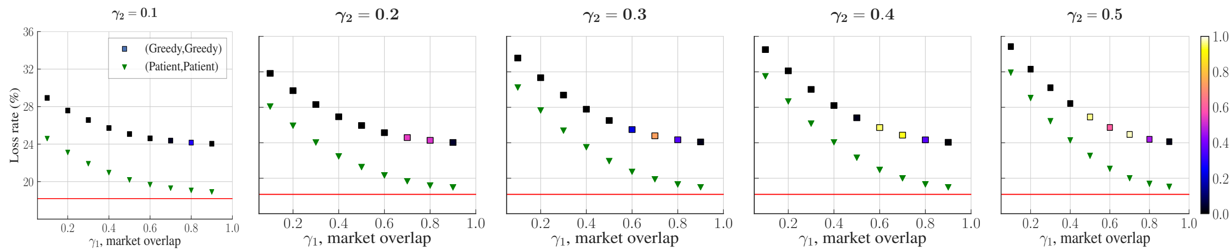

Figure 5.3: Simulation results for $d = 1, k = 100, p = 0.02$. γ_2 denotes the fraction of vertices, not in both markets, that enter only market M_1. The red line denotes the loss rate with a single Patient market. The green triangles denote the loss rate of a (Patient,Patient) equilibrium (an equilibrium in essentially all bootstrap samples). The squares denote the loss rate of a (Greedy,Greedy) equilibrium and the color of squares denotes the proportion of bootstrap samples which reach a fast matching outcome (defined in §5.4.2) from initial conditions (Greedy,Greedy).

patience—i.e., matching less frequently—and higher values of γ_1 indicate higher overlap—i.e., more agents entering both markets.

We assumed the markets have bounded rationality in their computations of best response functions. From the set S of all Monte Carlo simulations, we took $X = 2500$ bootstrap samples of size $n = 50$, $\{S'_i\}_{i=1}^{X}$ where $S'_i \subset S$. Each bootstrap sample represents simulations that a boundedly rational market would run. Thus, given a single bootstrap sample $S'_i \subset S$:

$$\mathrm{BR}_i(\alpha^1) = \arg\max_{\alpha^2} \mathbb{E}_{s \in S'_i}[u_{M_2}(\alpha^1, \alpha^2)]$$

$$\mathrm{BR}_i(\alpha^2) = \arg\max_{\alpha^1} \mathbb{E}_{s \in S'_i}[u_{M_1}(\alpha^1, \alpha^2)]$$

Under best response dynamics, the matching rate will now change over time, so we let α_t^1 and α_t^2 denote the matching rates at time t of market M_1 and M_2, respectively. We iterated best responses until convergence or cycles occurred over initial conditions of α_0^1, α_0^2 values.

5.4.2 Experimental Results

In general, we observe two main phenomena for the best response dynamics. First, we observe convergence to the Patient strategy under appropriate initial conditions (α_0^1, α_0^2) for any constituent in the competing market system. Second, for markets with sufficient overlap, and sufficiently low initial values of (α^1, α^2), we observe convergence to a (Greedy,Greedy) equilibrium or (α^1, α^2) parameters very close to (Greedy,Greedy). No other phenomena occur in more than 5% of bootstrap samples.

To simplify the description of results, we refer to convergence to (Greedy,Greedy), or cycles or equilibria involving solely $0 \leq \alpha^{\{1,2\}} \leq 1/100$, as fast matching. For the parameter range chosen for the simulations (specifically $d = 1$), $0 < \alpha^{\{1,2\}} \leq 1/100$ rarely impacts the matching choice. Furthermore, the social welfare for these outcomes only differ by at most 0.3%. We describe the notable effects of the parameter choices on best response dynamics below.

Market overlap. The impact of the market overlap γ_1 on the best response dynamics can be characterized by the effect on the range of initial matching rates (α_0^1, α_0^2) which converged to a fast matching outcome in "many" bootstrap samples—here, we use a cutoff of 25%. Figure 5.3 visualizes this behavior for increasing values of market overlap γ_1.

When market overlap $\gamma_1 \leq 0.4$, less than 0.1% of bootstrap samples converge to a fast matching outcome for any chosen initial matching rates. This is expected, as with $\delta = 0$, a faster matching rate increases utility u_{M_1} of M_1 primarily when M_1 successfully matches an agent that enters both markets before M_2 can. When $\gamma_1 \in [0.4, 0.8]$, the range of initial conditions that converged to a fast matching outcome rose to a peak at or before $\gamma_1 = 0.7$, then fell off.

Surprisingly, when $\gamma_1 = 0.9$, no initial conditions converged to fast matching in more than 4.4% of bootstrap samples.

Market asymmetry. We investigate the impact of γ_2, which controls the balance of agents entering only market M_1 (which occurs as $\gamma_2 \to 1$) or market M_2 ($\gamma_2 \to 0$). When $\gamma_2 \leq 0.1$, no conditions outside of fast matching converged to fast matching in more than 25% of samples.

5.4.3 Welfare Loss

We now measure the impact of competition on global social welfare. As with MODEL I, we define social welfare as the discounted total number of matches; here, however, we set $\delta = 0$. As before, we compute the distribution of social welfare for a range of γ_1 and γ_2 with respect to our bootstrap samples of Monte Carlo simulations. Figure 5.3 shows social welfare for outcomes (Patient,Patient) and (Greedy,Greedy)—which approximates the loss rate of fast matching—as well as the proportion of bootstrap samples that converge to fast matching from initial conditions (Greedy, Greedy). Note that, just as in MODEL I, all experimental outcomes are strictly worse (i.e., result in lower social welfare) than that of a single Patient market.

As expected, when the overlap γ_1 increases, the expected loss rate decreases due to a larger network of potential matches. However, as γ_1 increases, some initial conditions also become more likely to result in a fast matching outcome. For example, under two equally-sized markets ($\gamma_2 = 0.5$), the Greedy loss rate of 26.9% is higher than the loss rate of 24.3% for $\gamma_1 = 0.4$, where the only outcome that occurs with meaningful probability is (Patient,Patient). This additional equilibrium occurs in 96.4% of bootstrap samples. The additional overall welfare loss of 4.4%-5.0% incurred by a fast matching outcome for the same initial parameters is shown in Figure 5.3.

We also observe the effect on social welfare of a thicker market as market asymmetry increases; at its most extreme ($\gamma_2 = 0$ and $\gamma_2 = 1$), all vertices are effectively in a single market. As such, as γ_2 moves toward its bounds, again there are stronger network effects on social welfare.

Welfare losses arise both from matching speed and market fragmentation. As a baseline, the loss rate that occurs from a *single* Patient market—one with no competition—under the same model parameters is 18.2%. As shown in Figure 5.3, all other market conditions result in greater overall loss. In the succeeding section we explore policy options that could help a central planner mitigate this loss due to competition.

5.5 Policy Implications & Future Directions

Our results indicate that, left to themselves, matching markets that compete with each other can cause significant social welfare losses through fragmentation (§5.2 and §5.3) and suboptimal matching policies (§5.4). Our results are a proof-of-concept support of the "race to the bottom" seen in many real competing matching market systems. For example, in the US, multiple kidney exchanges compete over patient-donor pairs and/or hospitals. Two of the largest US exchanges are the National Kidney Registry (NKR) and the United Network for Organ Sharing (UNOS). NKR matches in an essentially greedy fashion. UNOS started by matching once per month, then moved to twice per month, weekly, and now 2+ times per week in part to reduce the "failure rate" caused by competition with the fast-matching

NKR. We see this behavior replicated in our model, and can quantify social welfare loss as well. Combinations of analytic and simulation results of this nature have set policy in kidney exchanges before (e.g., Dickerson *et al* [47] and Dickerson and Sandholm [48] have set parts of UNOS policy), and our model could help inform this debate.

While our research can inform policy discussions, it is important to have a separate conversation about the ethics of different regulatory and policy changes and how these can impact different populations (for example, by better-serving short-lived patients at the expense of long-lived ones), and our research is not intended to be prescriptive on those issues. That said, since we cannot use money directly to match supply and demand in a matching market, the market/policy designer's toolkit must consider other options. For example, in kidney exchange, one could offer increased priority in the future on the deceased-donor waitlist (living donor kidney grafts typically survive 10-15 years before another transplant is needed) if they were to go to a patient market rather than a greedy [14, 69, 131, 138].

Chapter 6

Modeling Quality in Matching Markets: The Case of Kidney Exchange

To improve the performance of kidney exchange for both number of matches and individual quality, one proposal, which does not rely on the ability of solving market fragmentation or setting up a national center, has been to incorporate compatible pairs into exchanges.

There is the potential for significant benefit from including directed donation pairs in kidney exchanges that also include incompatible pairs. The benefit can arise from two fronts: (1) a significant increase in the number of incompatible donors who find matches; (2) an increase in the quality of matches, since factors like HLA match [104, 127] etc. play a role in expected graft survival. The main goal of this chapter is to estimate the potential benefits along both these fronts in a realistic manner. In doing so, we will also contribute to the literature on matching with cardinal utilities by providing a realistic data-generation mechanism for cardinal utilities.

Overall, in this chapter, we model quality in matching markets and investigate the benefits of incorporating compatible pairs in kidney exchange. We estimate expected survival of a graft from the recently proposed *Living Donor Kidney Profile Index (LKDPI)* [104], and use this as our measure of quality. [12] We impose the basic incentive compatibility constraint that, for compatible pairs to be transplanted through exchange instead of directly, each recipient must receive a graft with lower LKDPI, or increased expected survival time, compared with that of her original donor.

During the years from 2014 to 2016, we were able to obtain data on 184 living donor kidney transplantations that took place at the transplant center of Barnes-Jewish Hospital (henceforth "Center"). Of these 184, 171 were directed donations from a compatible donor to his/her paired recipient. We obtained complete information that enabled computation of the LKDPI on 166 of these pairs, which we use to estimate distributions of LKDPI scores (and hence expected graft survival) within compatible pairs and across pairs. We were able to obtain complete antibody and antigen data on 121 of these pairs, which enables donor-recipient compatibility checking.[13]

The first question we can ask is about the heterogeneity of match qualities across pairs and the effects of this heterogeneity on the quality of the final matching. At one extreme, LKDPIs across pairs could be completely independent of the original LKDPIs within the pairs. This would correspond to maximally heterogeneous match qualities and offer the highest possible benefits to recipients in compatible pairs of participating in the exchange.

[12]LKDPI itself is a somewhat complex number to interpret. It is intended to be on the same scale as the KDPI for cadaveric kidneys, which is a percentile measure. Thus an LKDPI of 10 indicates that the kidney is comparable to the 10th percentile of cadaveric kidneys in terms of quality (with lower numbers being better). However, since some living donor kidneys can be better than *any* cadaveric kidney, LKDPI values can also be negative.

[13]We do not have data for the remaining 45 pairs because of a change in the software system, so there is no selection bias.

At the other extreme, LKDPIs could be completely determined by the characteristics of the donor or the recipient in a pair, in which case there would be no social gains from trade [10]. In reality, LKDPI does take into account various match characteristics (for example, HLA mismatches and body weight ratios), but where the gains from trade may fall in the spectrum is an empirical question. Our experiments confirm that the distribution of match quality (LKDPIs) from "external" donors is far from independent of the match quality within a compatible pair, and this has significant implications on the possible gains from trade to the compatible pairs. As a benchmark, we conduct counterfactual tests that assume no incompatibilities among any of the pairs, and that all 166 pairs participate in a pareto-improving kidney exchange with 2 and 3 cycle swaps. This improves the average LKDPI of transplanted kidneys from 37.15 to 25.5, corresponding to about 1.5 years of expected graft survival. We can estimate the hypothetical benefit if all donor-recipient pair LKDPIs were independent draws from the same distribution, and we find that the new average LKDPI achieved would be 2.67, corresponding to more than a 5 year benefit in terms of expected graft survival. Interestingly, we provide evidence that the variability is largely driven by characteristics of the donor rather than the recipient, so there could be benefits from increasing the pool of possible donors, as has recently been suggested [55].

Based on this observation, we argue for the importance of constructing a minimal simulator that produces realistic LKDPI / match quality values, and describe the construction of such a simulator, which closely matches the characteristics we observe. Having established the potential for gains from matching simply among already compatible pairs, we then turn to estimating impacts in differently-sized populations when both compatible and incompatible pairs are present using this simulator, paired with the standard *Saidman simulator* [127] for generating recipient-donor pairs and compatibilities. These two simulation mechanisms together enable us to simulate realistic living donor kidney scenarios of any size with compatible and incompatible pairs. We use the simulator to estimate both the increase in the number of recipients in incompatible pairs who would be matched if compatible pairs participated in the exchange, as well as the increase in the expected graft survival for recipients in compatible pairs that participate in the exchange.

We find that with compatible pairs joining the kidney exchange, the percentage of matched incompatible pairs almost doubles. For example, with a small pool of 50 donor-recipient pairs, 74% of incompatible pairs are matched, compared with 39% when the two-&three-cycle swap is only run within the incompatible pairs. With a large pool of size 600, the percentage of matched incompatible pairs reaches 91%, compared with 54% if we only run two-&three-cycle swap within incompatible pairs. These results are similar to those of Gentry [61], who also estimate that the proportion of incompatible pairs matched can be doubled by participation of compatible pairs. They focus only on compatible recipients gaining a donor age benefit. Our methods, combined with the LKDPI, also allow us to estimate the benefits to recipients in compatible pairs. If the optimizer maximizes expected survival of grafts over the entire population, there is an increase of 1.4-2.5 years in expected graft survival among recipients from compatible pairs. If the optimizer instead maximizes number of transplants, this number is between 0.9 and 1.23 years.

An important practical consideration is likely to be that of waiting time. Compatible pairs may not be willing to wait even in order to find a potentially better match. Therefore, we consider a dynamic matching model where the incompatible pool matches either in a greedy or patient fashion (a la [4]), but the compatible pairs match greedily (from the incompatible pool if it improves the match for the compatible-pair recipient and directly from donor to recipient otherwise). Even with this pessimistic restriction, we estimate substantial benefits, going from matching 35% of incompatible pairs to 55% for the arrival and departure rates we examine.

In this dynamic setting, we also look at the effects on two hard-to-match subpopulations, namely blood group O recipients and highly sensitized patients. We estimate that the positive impacts on blood group O recipients are more substantial than in the general population (an increase from 18% to 46%), while those on the highly sensitized population are similar to the general population (an increase from 24% to 37%).

By bringing quantitative estimates of these benefits into the light, we can inform policy debates. For example, how much expected benefit would be needed to convince compatible pairs to enter an exchange? How long would they be willing to wait in a dynamic setting? These are all questions that can begin to be addressed from the foundation of the models and simulator we develop in this work.

6.1 Modeling Match Quality

Historically, much work on matching (and welfare economics broadly) has focused on ordinal preferences rather than cardinal utility [9]. This sidesteps the problem of having to make interpersonal comparisons of utility, and research has focused on outcomes in terms of objectives like stability and Pareto optimality [110]. However, with the increasingly important social roles played by matching mechanisms [43, 44, 112, 126], it is imperative to understand the outcomes of mechanisms in terms of overall social welfare (however this is defined for a given application) as well as distributional effects. Doing so necessitates considering specific models of utility [14, 69, 94, 105].

There is value in traditional parametric models that are used for utility, and these have been central to model development. Examples of such models include utilities that decay exponentially in waiting time [4, 7], and random utility models for specific match pairs [38]. However, a common criticism of such models is that it is unclear how general or valuable results are when the utility model itself is not grounded in reality. In our case, we are explicitly looking for a realistic model that can be used for decision-making. Further, in order to convince compatible pairs to enter kidney exchanges, we must be able to quantify the expected benefit to them in some meaningful manner, therefore, we need an individual model of match quality that can be reasoned about from the perspectives of agents in the market. One important consideration that we defer to future work is the waiting cost to agents in terms of cost and quality of life. For compatible pairs, this is a complex modeling problem from a practical standpoint, because the baseline waiting time is itself highly variable. The time from initial workup to transplantation for a compatible pair is at least several months long because of the barrage of necessary testing, and for part of this time the pair is not even sure that they will be judged compatible. Therefore, in this chapter we focus on match quality, and subject our analyses to pessimistic assumptions (greedy dynamic matching), and various robustness checks (varying pool sizes can proxy for match frequency, for example).[14]

Quantifying match quality. Transplant surgeons often have to make decisions on whether a proposed transplant is worthwhile to proceed with. The Kidney Donor Profile Index (KDPI) was developed as a means of assessing the quality of a cadaveric (deceased donor) kidney [119]. Recently, the Living Kidney Donor Profile Index (LKDPI) has been proposed

[14] The question of how to analyze waiting cost from the perspective of the matching market is also complex; however, one reasonable way to think about it is as costs to the healthcare *system*. For example, dialysis costs $70,000-$100,000 per year [74, 83, 100], and this is a cost that must be borne by some agent (individuals, private insurance, hospitals, or the government). Incorporating this can be useful when the modeling task is to assess matching policies and how they change costs over the entire system, rather than from the perspective of individual agents, hospitals, and so on.

as an analog for living donations [104]. LKDPI takes into account characteristics of both the donor and the recipient.

KDPI is a percentile score. For example, a score of 4 implies that the kidney is in the "top 4%" of cadaveric kidneys. LKDPI is intentionally designed to be on the same scale (as mentioned in the Introduction, since living donor kidneys can be better than any cadaveric kidney, the LKDPI often takes negative values as well). Therefore, optimizing for LKDPI, while a useful proxy, is semantically ill-founded. However, since LKDPI is computed based on a survival model (Cox regression [34]), one can translate the model to a model of expected graft survival (or graft half-life), the survival time of the transplanted organ in the donor [72].[15] We have found that an exponential curve fits the graft half-life as a function of LKDPI almost perfectly (see Figure 6.1), and can thus estimate expected graft survival as $14.78e^{-0.01239x}$ where x is the LKDPI. We can use this measure in place of LKDPI where it is more appropriate.[16] Thus the edge weight w_e in each cycle is defined as the estimated expected graft survival of the recipient.

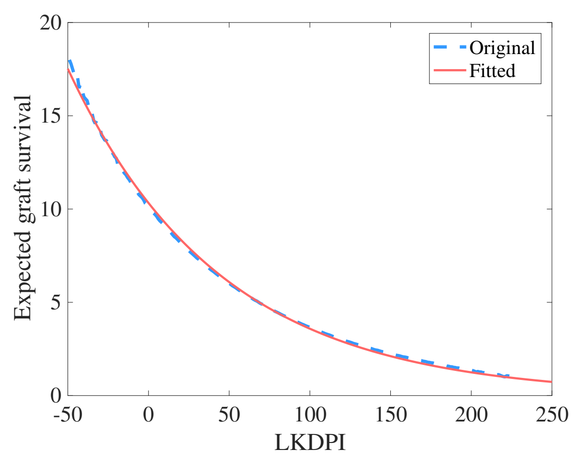

Figure 6.1: An exponential curve fits the graft half-life as a function of LKDPI.

[15] After graft failure, the donor typically needs another transplant.

[16] One could also use expected graft survival as input to an expected "Quality Adjusted Life Year" (QALY) [23, 137, 140] computation over the lifetime of the recipient.

6.2 Exchanges between compatible pairs: A single center analysis

6.2.1 Data description

Massie *et al* [104] come up with the LKDPI measure based on several important characteristics for determining graft survival. We gathered de-identified data on all donor and recipient characteristics that are used in computing LKDPI from all directed living-donor transplants performed at the center in a three year period (2014-2016). There were 166 such transplants with complete characteristics for calculating LKDPI and graft survival; 121 of them also include complete HLA antibody and antigen information. The distribution of each characteristic is shown in Table 6.1. We also analyze the correlation of every pair of characteristics, shown in Figure 6.2, which serves as a fundamental building block for designing the simulator in Section 6.3.

6.2.2 Counterfactual analysis within the center

Typically, if a donor and recipient are deemed medically compatible, a directed transplant is performed, with the donor's kidney going to the recipient. However, there may be cases where the match quality is low even if they are compatible, and perhaps the recipient could receive a better kidney through an exchange; for example, they may be able to receive a kidney from a younger donor, or avoid an immunologically risky donor/recipient combination, like child to mother or husband to wife [61]. Such scenarios are hypothetical, and may seem unlikely at first glance. To validate our conjecture, for these donor-recipient pairs, we computed the

	Mean	s.d.
Donor Age	48.22	12.68
Donor eGFR	98.11	15.08
Donor Systolic BP	124.14	13.11
Donor BMI	27.78	4.46
Recipient Weight (Female)	180.7	42.26
Recipient Weight (Male)	190.34	39.9
Donor Weight (Female)	160.75	30.06
Donor Weight (Male)	200.8	32.8
Donor Sex	F: 0.7	M: 0.3
Rec Sex	F: 0.35	M: 0.65
Donor African-American	Y: 0.05	N: 0.95
Donor Cigarette Use	Y: 0.32	N: 0.68
Donor/Rec Related	Y: 0.50	N: 0.50
Donor Blood Type	O: 0.6, A: 0.3, B: 0.07, AB: 0.03	
Rec Blood Type	O: 0.46 A: 0.39 B: 0.12 AB: 0.03	
Donor/Rec ABO compatible	Y: 0.88	N: 0.12
	Donor/Rec related	Donor/Rec unrelated
Donor/Rec HLA-B Mismatches	0: 0.18, 1: 0.32, 2: 0.5	0: 0.01, 1: 0.1 , 2: 0.89
Donor/Rec HLA-DR Mismatches	0: 0.13, 1: 0.06, 2: 0.81	0: 0.01, 1: 0.06, 2: 0.93
	Counterfactual Matrix: all unrelated	
Donor/Rec HLA-B Mismatches	0: 0.009, 1: 0.091, 2: 0.9	
Donor/Rec HLA-DR Mismatches	0: 0.02, 1: 0.04, 2: 0.94	

Table 6.1: Distribution of each characteristic of the center's data. F/M means Female/Male, Y/N represents Yes/No, and Rec is a shortening of Recipient.

expected graft survival (EGS) of each pair, and then performed counterfactual simulations to assess the potential to improve outcomes. The two counterfactual simulations share a Pareto improvement restriction—no recipient may receive a kidney with a shorter (i.e., worse) EGS for them than the EGS for them of the kidney from their original paired donor.

Optimal. In the first simulation, we find the best matching, allowing arbitrary length cycles; this can also be treated as a bipartite matching problem (with the restriction that the matching must be perfect) between donors on one side and recipients on the other.

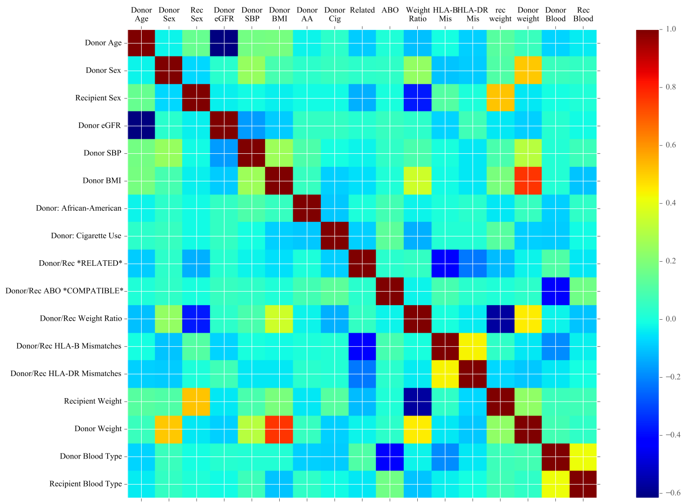

Figure 6.2: Correlation matrix of each pair of characteristics.

Two and Three-cycle swap. In the second simulation, we only allow either a direct donation from the donor to the recipient or through a two- and three -cycle kidney exchange, to more closely approximate realistic logistical constraints.

We consider two subsets of the data. The "complete" 166-pair subset, assuming no HLA incompatibilities, and the "restricted" subset of 121 pairs for which we have complete antibody/antigen information and can determine all incompatibilities and rule out such transplants. The distribution of EGS and corresponding LKDPI among the real pairs and in the results of our counterfactual simulations are shown in Figure 6.3. The mean and median EGS and LKDPIs are given in Tables 6.2 and 6.3 below.

We can see there is a median improvement of 1.93 years of expected graft survival for the Optimal and 1.38 years for the two-&three-cycle swap (over a median half-life of 10.84 years). We also see that including compatibility constraints itself does not have a huge effect on the

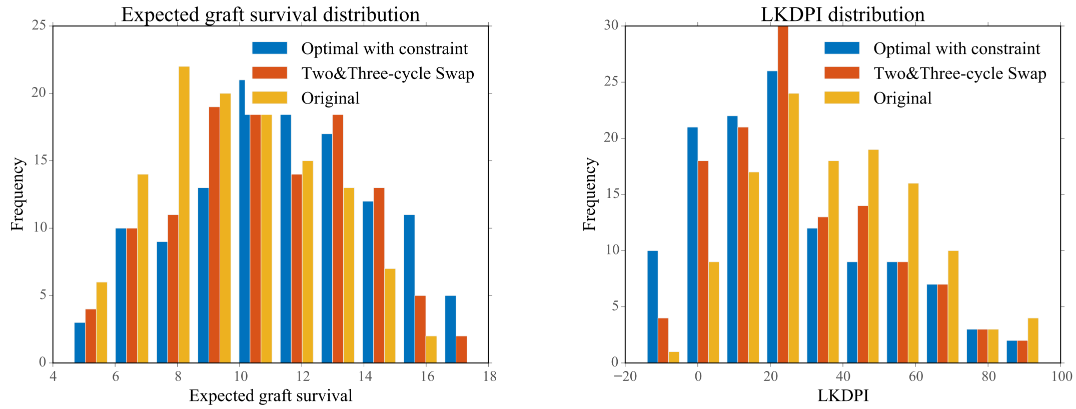

Figure 6.3: Distribution of the expected graft survival (left) and LKDPI (right) of the original matched pairs and matched pairs in the two counterfactual simulations, using 121 subset of real data with HLA antigens and antibodies from the center over the last three years.

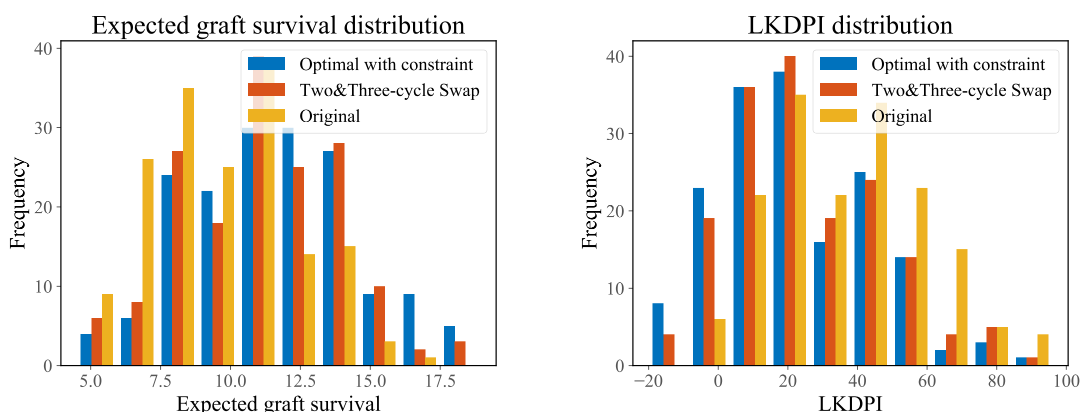

Figure 6.4: Distribution of the expected graft survival (left) and LKDPI (right) of the original matched pairs and matched pairs in the two counterfactual simulations, using the 166 full dataset of real data from the center over the last three years. We can see the distribution is similar to Figure 6.3.

	Optimal	Two&Three -cycle Swap	Original
Mean	11.44	10.88	9.80
Median	11.39	10.84	9.46

(a) 121 subset with HLA antigens and antibodies

	Optimal	Two&Three -cycle Swap	Original
Mean	11.58	11.14	9.67
Median	11.67	11.18	9.34

(b) 166 full data without HLA antigens and antibodies

Table 6.2: Mean and median EGS for two counterfactual simulations, compared to reality over the last three years at the center. Figure (a) shows the 121-subset of data with HLA antigens and antibodies, and Figure (b) shows the 166-subset assuming no incompatibilities.

	Optimal	Two&Three -cycle Swap	Original
Mean	23.99	27.69	36.10
Median	21	25	36

(a) 121 subset with HLA antigens and antibodies

	Optimal	Two&Three -cycle Swap	Original
Mean	23.46	25.50	37.15
Median	19	22.5	37

(b) 166 full data without HLA antigens and antibodies

Table 6.3: Mean and median LKDPI for two counterfactual simulations, compared to the reality over the last three years at the center. Figure (a) shows the 121-subset of data with HLA antigens and antibodies, and Figure (b) shows the 166-subset assuming no incompatibilities.

results (some of the improvement in the larger set is simply due to having a thicker market). Beyond the specific results, it is surprising to see the high number of transplants that were performed with LKDPIs above 50, since these indicate that the average *cadaveric* kidney would have been better for the recipient, in contrast to the conventional wisdom that living donor kidneys are always better. The optimized matches from the counterfactual “exchange” are much better, with many fewer “bad” matches and many more with LKDPI of 20 or lower, predictive of excellent outcomes.

6.2.3 Discussion

This is a proof-of-concept for the potential of improving quality of matching. One immediate question arises from the fact that we are using three years worth of data on recipients and donors in a static setting; this is obviously unrealistic. However, the main point is to estimate realistic distributions from data; we can use projections to then analyze differently-sized static markets (from smaller ones to larger ones that could be realized through regional pooling or already-functioning national exchanges). We turn to these questions and beyond in the next section.

6.3 Including compatible pairs in kidney exchanges

In addition to improving match quality, we may also be able to improve the *number* of matches by including compatible pairs to thicken the exchange with incompatible pairs. This could also lower costs for transplant centers by allowing for more internal matches where the transplant center does not need to go to a regional or national exchange to find a match for an incompatible pair.[17] In order to estimate the possible benefits more systematically over different possible population sizes, we need to efficiently and correctly simulate LKDPIs over donor and recipient populations.

[17] This could have positive and negative effects overall, by perhaps increasing fragmentation, but lowering costs. However, many centers choose not to participate in broader exchanges much of the time in practice, for a variety of reasons.

6.3.1 Determinants of match quality, and design of a Compatibility+LKDPI simulator

This would be simple if LKDPIs were distributed in a manner that was easy to correctly estimate, for example, independently, or independently conditional on the LKDPI of the original compatible pair. Unfortunately, this turns out not to be the case. To get a simple benchmark of how much this may affect the results, we can simulate different distributions based on data from the center.

We first build a counterfactual matrix of estimated graft survival based on the original (166-pair) data by calculating LKDPI values for each of these 166 pairs. We then investigate the expected graft survival of donor-patient pairs under the Optimal and two-&-three cycle swap matching algorithms when resampling the matrix in different ways. To simulate independent LKDPIs, we resample individual LKDPIs from the whole matrix. To simulate donor-dependent LKDPIs, we shuffle all donors for a given recipient, and to simulate recipient-dependent LKDPIs, we shuffle all recipients for a given donor. The results are shown in Table 6.4. The first row shows statistics from the original compatible matching. As we see, most of the methods for generating LKDPIs vastly overestimate the possible gains, and the evidence is consistent with the observation that the determination of LKDPI/expected graft survival is largely based on the donor's characteristics [104]. These results demonstrate the need for a good simulator.

The central empirical facts that allow us to construct an efficient simulator are analyses of the joint distributions of variables involved in determining compatibility (PRA and ABO compatibility, based on the simulator of [127]) and computing LKDPI (See Table 6.1), and analysis of the possible underlying mechanisms of dependence. In particular, compatibility is solely a function of blood type and antibodies, while LKDPI considers many other factors, most of which have limited relationship to those (Figure 6.2). Since the state of practice in kidney transplantation has been to always assume that any living donor is excellent (a practice called into question by our results above), it is unlikely that there is any selection bias in the characteristics we sample for typical compatible pair arrivals. We first generate a donor-recipient pair, with all LKDPI-related characteristics generated sequentially in a manner that respects the data distributions in Table 6.1 and the correlation structure shown in Figure 6.2. We then generate the PRA (percentage reactive antibodies) and compatibility based on the Saidman model. Details of our simulator are in Appendix B, Algorithm 3. The last line of Table 6.4 shows that the simulator produces results very close to the real data.

	EGS original	EGS 2&3 swap	EGS Optimal	LKDPI original	LKDPI 2&3 swap	LKDPI Optimal
Original 166 dataset	9.67	11.14	11.58	37.15	25.50	22.46
Sample from the whole matrix	9.23	14.40	15.30	40.51	2.67	-2.5
Shuffle all donors per recipient	9.19	14.16	14.94	40.92	4.11	-0.47
Shuffle all recipients per donor	9.21	11.74	12.50	40.70	20.6	15.49
Sample from the simulator	9.38	11.40	11.80	39.21	24.50	20.09

Table 6.4: The EGS and LKDPI comparison of different sampling methods and different market clearing algorithm.

6.3.2 Experimental results using the LKDPI simulator

We can now use the LKDPI simulator to estimate the benefits in terms of both quality and quantity of transplants. We study the impact of different optimization objectives (survival and number of matches) on outcomes for both compatible and incompatible pairs. We

C-Or	Compatible pairs original donation		
ENM	Expected number of matched pairs		
P-I	Pool with only incompatible pairs		
P-CI	Pool with both incompatible and compatible pairs		
Maximize EGS		Maximize ENM	
I-maxSur	P-I $2-\&3-cycleswap$	*I-MaxNum*	P-I $2-\&3-cycleswap$
I-O-MaxSur	P-I Optimal	*I-O-MaxNum*	P-I Optimal
CI-MaxSur	P-CI $2-\&3-cycleswap$	*CI-MaxNum*	P-CI $2-\&3-cycleswap$
CI-O-MaxSur	P-CI Optimal	*CI-O-MaxNum*	P-CI Optimal

Table 6.5: Table indexing abbreviations we use corresponding to different optimization objectives, matching methods, and different subpopulation measurements.

are most interested in the improvement of (1) expected graft survival of compatible pairs compared with their original donation, since the incentive for compatible pairs to enter is to seek a better organ for the recipient; (2) the number of matched incompatible pairs compared with the number when running two-&three-cycle swap only on incompatible pairs. We find the maximum weighted cycle cover, where the weight can be (1) $w_e =$ expected graft survival of recipient, (2) $w_e = 1$, (maximizing the number of matched pairs). Table 6.5 summarizes the possible objectives and the metrics that we measure.

In our experiments, we fix the size of the pool and generate donor-recipient pairs using the simulator. We find that the sizes of the compatible and incompatible pool are roughly even. This matches the statistics of the center we have data from. In 2017, 217 compatible pairs and 181 incompatible pairs registered for initial transplant workups (though only 1/3 of them ended up having a transplantation procedure in the center).

We then run two-&three-cycle swap under the Pareto improvement restriction, where compatible pairs only swap if their expected graft survival increases. Our results for the proportion of incompatible pairs matched for different pool sizes can be summarized as follows (Figure 6.5 shows more detail):

	Without compatible	With compatible
Size of pool: 50 (25+25)	39%	74%
Size of pool: 100 (50+50)	48%	83%
Size of pool: 600 (300+300)	54%	91%

These results are similar to the results of Gentry [61], who also estimate that the proportion of incompatible pairs matched could be doubled by participation of compatible pairs that would gain a donor age benefit. From the perspective of compatible pairs, there is a 2.15-2.61 improvement in expected years of graft survival improvement (for those whose donor changes) when we maximize expected survival of the whole population, and 1.36-1.63 years when we maximize the number of matched incompatible pairs.

While the rate of entry of compatible and incompatible pairs may be similar, it is possible that one or the other population is less likely to go through with a transplant. In order to study how our results would vary with different assumptions about this, we hold the number of compatible pairs fixed and vary the number of incompatible pairs. These results can be seen in Appendix C, and qualitatively still suggest substantial benefits from incorporating compatible pairs.

6.4 Modeling Dynamic Markets

While different sizes of pools can proxy for different match frequencies, and patience in matching can improve outcomes [4], there are also good arguments and practical concerns that favor greedy or frequent matching [7,37]. In particular, it is a reasonable, if pessimistic,

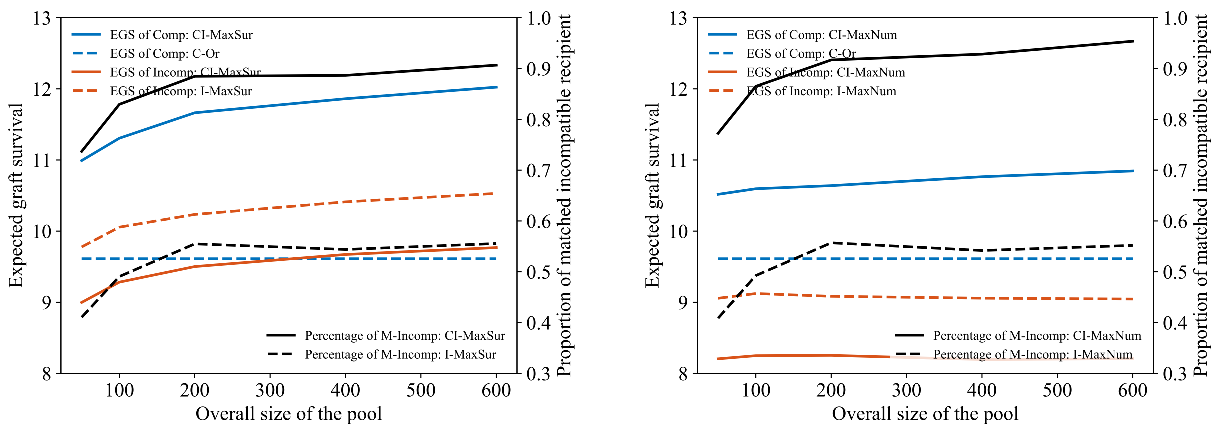

(a) Maximizing EGS.

(b) Maximizing number of matched pairs.

Figure 6.5: The comparison between expected graft survival of compatible pairs by participating two-&three-cycle swap (blue solid line) and their original matching (blue dash line), expected graft survival of incompatible pairs when compatible pairs participate two-&three-cycle swap (red solid line) and only within incompatible pairs (red dashed line), and proportion of matched incompatible pairs when compatible pairs participate two-&three-cycle swap (black solid line) and only within incompatible pairs (black dashed line), where Figure (a) shows the results of maximizing the expected graft survival across the whole graph G, and Figure (b) shows the results of maximizing the number of matched pairs.

assumption, that compatible pairs would be completely unwilling to wait, and would therefore insist on an immediate exchange, or else they would want to go ahead with the direct donation.

In this section, we build a dynamic model where patient-donor pairs arrive gradually over time. Incompatible pairs stay in the market until they find an acceptable swap or they perish (they may leave the market if the patient's condition deteriorates to the point where kidney transplants become infeasible, for example). Compatible pairs must either be matched with an incompatible pair at the moment of arrival, or else the donor gives directly to the recipient immediately.

Arriving patient-donor pairs are still generated from our simulation model described above. Pairs arrive at the market according to a Poisson process, with rate parameter $m \geq 1$. The sojourn of an agent is draw from an exponential distribution, with rate parameter $\lambda = 1$. If a compatible pair arrives, all feasible swaps for that pair are considered (where feasibility means both that compatibility requirements are satisfied and the recipient in the compatible pair receives a higher-quality kidney match). If there is more than one acceptable swap the newly entered pair chooses the one with the longest expected graft survival for its own recipient; ties are broken uniformly at random. Incompatible pairs can be matched either greedily, in the same manner as above, or using a *patient* algorithm [4] which waits to match until the moment an agent is about to perish (with the caveat that incompatible pairs with possible matches in the incompatible pool are not considered as matches for compatible pairs).

We again find a substantial benefit in terms of the number of incompatible pairs matched under either mechanism (from approximately 35% to approximately 55%). Figure 6.6 also shows that compatible pairs for whom the new mechanism changes the match improve their expected graft survival by almost two and a half years. There is some effect of "competition" – since match quality is largely a function of the donor, and compatible pairs need to receive good donors in order to participate, the average expected graft survival of those who are transplanted in the incompatible pool actually goes down; however, the huge increase in the number of matches more than compensates in terms of the sum total of years of graft life (where pairs that don't receive a transplant are assigned an EGS of 0). Therefore, our results demonstrate that the potential value of incorporating compatible pairs is high even under pessimistic assumptions about what wait times they would be willing to tolerate.

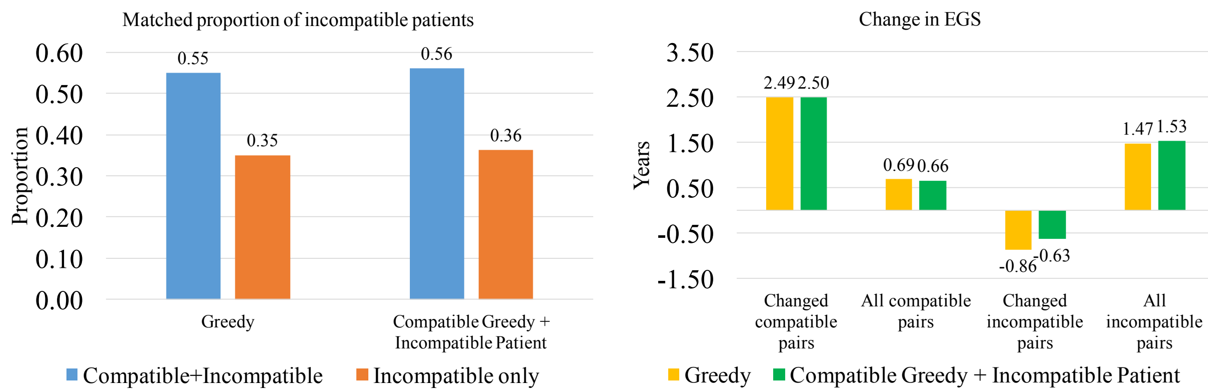

Figure 6.6: Comparison of matched proportion of incompatible patients (left) and change in EGS (right) when running different matching algorithms for the incompatible pool in the dynamic setting.

6.4.1 Fairness considerations: Hard to match types

An important consideration in kidney exchanges is how they may differentially affect different populations. The populations one often worries about are those who are harder to match. Therefore, we consider the effects on two groups of hard-to-match patients, those with blood group O (patients with blood group O have fewer ABO-compatible living donors [65]) and highly sensitized patients, who are likely to have antibodies to a significant fraction of the population. We define highly sensitized patients as those whose PRA is greater than 80%, constituting approximately 30% of the patient population.

Figure 6.7-left shows that there is a significant improvement for the matched proportion of incompatible blood type O patients (from 0.18 to 0.46) when incorporating compatible pairs. Therefore the relative benefit to this group is actually higher than to the rest of the population. The matched proportion of incompatible highly sensitized patients improves to 37% from 24% when compatible pairs are included, a rate of increase roughly similar to that in the overall population.

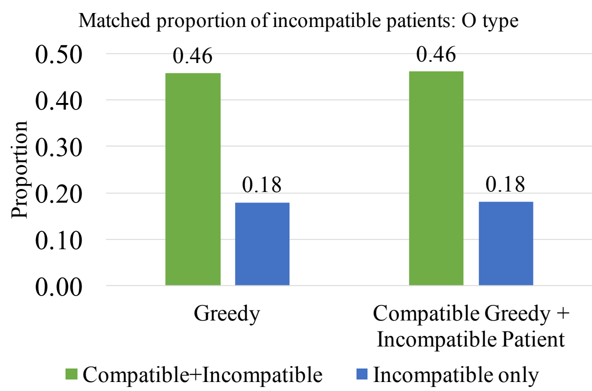

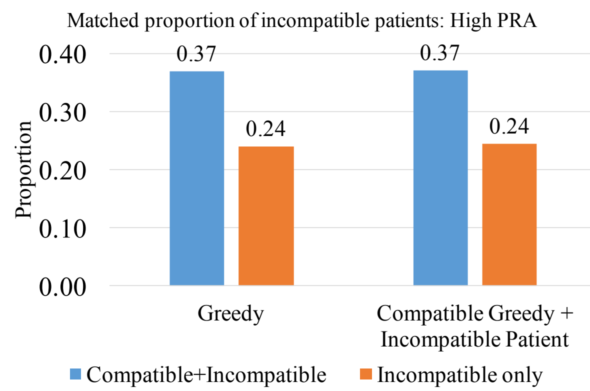

Figure 6.7: Comparison of matched proportion of incompatible patients of Blood Type O (left) and high PRA (right) when running different matching algorithms for the incompatible pool in the dynamic setting.

6.5 Conclusion and Future research

Living donor kidney transplantation has proven to be an important domain for the development of matching theory and algorithms. It is becoming increasingly important to study cardinal utilities in kidney exchange, and we believe this could open up more fertile avenues for research. Our main goal in this chapter is to develop the framework and a robust framework for analyzing match quality in models of kidney exchange. Our framework is based on real donor and recipient data from a major transplant center. We have also used the model to estimate the benefits, in terms of both quantity and quality of transplants, of including compatible pairs in kidney exchange. We find that if we were able to induce compatible pairs to join kidney exchanges, the percentage of matched incompatible pairs would increase dramatically, and there would also be a substantial increase in expected graft survival for recipients in compatible pairs. Quantifying the potential quantitative benefits of participating through LKDPI may also make compatible pairs more likely to join.

While our work here is largely in a static setting and a simple dynamic setting, the development of our realistic LKDPI simulator allows for the investigation of many different matching models. Of particular interest will be questions related to matching policy in the dynamic setting with both compatible and incompatible pairs, incorporating wait times, and possible systemic effects of including compatible pairs in exchanges, for example, changes in incentives for centers.

Chapter 7

Competition Between Financial Exchanges

In this chapter, we turn our eyes to financial markets and investigate the market competition in this domain. As we discussed in Section 1.2.2, most modern financial exchanges operate using the continuous double auction (CDA) mechanism, which is a greedy fashion mechanism. The existence of this continuous time feature has led to the development of the phenomenon called High frequency trading in the last two decades. With companies keeping investing in faster infrastructure for trading and events like the "flash crash" of May 2010, high frequency trading (HFT) has become an increasingly debated topic in both the media and policy spheres [93]. There is increasing evidence that at least one form of high frequency trading, namely latency arbitrage, has reached a point of socially diminishing returns. Frequent batch auctions is recommended as a market structure that could replace CDAs, since the minimum time period between trades is specified, and there is no benefit to being faster than that.

In this chapter, we first provide a baseline of social welfare in a single market with high-frequency traders. We then consider the competition between platforms that employ different microstructures: one, a continuous double auction, and the other, frequent batch auction.

We ask: (1) Which of these markets would traders choose? (2) Could the frequent batch auction market replace the CDA market simply by entering the marketplace of exchanges, or would it require a regulatory push?

Our key measure of welfare is the price of immediacy – the expected loss suffered by background traders. This measures the cost that the "average trader" pays in order to execute transactions. This is a different measure than that of Wah and Wellman [144] or Wah et al [142], who use surplus. These are both reasonable measures, but surplus is most meaningful in private value models, where some meaning can be attributed to different agents having different valuations for an asset. Our model follows in a tradition of common value models, where the asset has a true underlying value, and different traders may have different estimates of that true value. The existence of background traders in our model provides a useful proxy for estimating the cost of trading. It is worth noting that this doesn't mean that background traders are necessarily losing money – typically such traders would stay in the market for much longer, and under reasonable models of price appreciation, these "losses" can be thought of as transaction costs for buy-and-hold type investors.

First, we look at simple models of individual markets and confirm that our model satisfies the basic intuitions one would expect. Namely, informed traders (in particular, low latency traders) make more profit (and background traders are consequently made worse off) in CDA markets than in frequent call markets. A zero-profit market maker (with no specialized information) can greatly improve the position of background traders, taking away most of the profit opportunities from informed traders in CDA markets.

Next, we model competition between a CDA market and a frequent call market when informed traders pick which market to place their orders in based simply on which market is more mispriced with respect to their current belief. We show that the informed traders do better overall when they choose to place orders in the market that is more mispriced from their perspective. We show that, when informed traders are all using this strategy, a majority of orders flows to the call market, and background traders are better off than in a single CDA market.

Note that all of the above analysis is not in an equilibrium setting – we assume that all informed traders use the same strategy. We can use the insights developed in these models to begin analyzing strategic market choice. We do so by introducing a learning framework, where informed traders learn a parametric form for the expected profit of choosing to place an order in a market (and a non-parametric probability of order execution) given the distance of that market's "current price" from the trader's estimate of the true value of a stock. We show that, when all agents use this learning approach, they converge to an approximate equilibrium where a majority of trades again flow to the frequent call market.

7.1 Market Model

7.1.1 The CDA and Call Markets

Our model of competing markets consists of two markets, one employing a *continuous double auction* (CDA) mechanism and the other one employing a *frequent call* (CALL) mechanism. We begin by describing the details of each individual market, which will serve as the foundation for our model of competing markets.

Each market is running in the continuous-time interval [0,T]. A single security is traded in the market. There is an underlying "true value" process. The initial true value of the security v_0 is drawn from a Gaussian distribution with mean v_{initial} and standard deviation σ_{initial}. Then, the true value jumps according to a Poisson process with rate parameter λ_{jump}. If the true value jumps, the new true value v_t is generated from $v_t \sim \mathcal{N}(v_{t-dt}, \sigma_j)$, where v_{t-dt} is the price instantaneously before the jump (we restrict $v_t \geq 0$, so all values are truncated at 0).

In the CDA market, outstanding orders are maintained in two priority queues: one for bids (the buy orderbook) and one for asks (the sell orderbook). Bids and asks are prioritized by price first and time second. When a new order comes in, it is added to the corresponding order book. A trade is executed if the highest bid exceeds or is equal to the lowest ask. The execution involves the orders at the top of the bid and ask queues, at the price of the older of the two orders involved.

The CALL market is similar to that described by Budish et al [27]. It clears in fixed intervals of time τ (the *call interval*). At each clearing time, the market collates all of the orders and computes the aggregate demand and supply functions of all bids and asks, respectively. The market clears where supply equals demand, with all executions occurring at the same price, called the market-clearing price. None of the orders are visible to any traders during the call interval. The market announces the market-clearing price after each clearing (*market announcement*). When no order was executed at the last clearing time, if both the buy and sell orderbooks are not empty, the market announcement will be the mid-point of the highest bid and lowest ask, otherwise it will be the most recent available market-clearing price. All untraded orders roll into the next call.

7.1.2 Valuation model

Each trader has a private valuation for the security (or equivalently for our purposes, a private signal of the true value). We have two types of traders, informed (IF) traders and background (BG) traders. Each informed trader IF_i receives a private signal of the security value, $w_{i,t} \sim \mathcal{N}(v_{\hat{t}}, \sigma_{\text{trader}})$, where $v_{\hat{t}}$ is the underlying true value of the security at some time $\hat{t}$, where $\hat{t} \leq t$, and σ_{trader} is a noise parameter. We define two types of informed traders, namely, low latency (LL) traders and high latency (HL) traders. High latency traders have staler information, i.e., $w_{i,t} \sim \mathcal{N}(v_{t-\delta}, \sigma_{\text{trader}})$, and low latency traders observe information with no delay, thus $w_{i,t} \sim \mathcal{N}(v_t, \sigma_{\text{trader}})$. Background traders do not have any private information; each arriving background trader wishes to either buy or sell one unit (with equal probability). They demand immediacy, that is, they want to get their orders executed as soon as possible, so they are willing to take any market price.

7.1.3 Agent arrival process

There are a fixed number of traders of each type. Informed traders and background traders both arrive at the market according to separate Poisson processes, with informed traders arriving with rate λ_{IF} and background traders' arriving with rate λ_{BG}. In the event that an informed trader arrival occurs, a specific IF trader is selected uniformly at random from all the IF traders to place/replace an order; similarly, if a background trader arrival occurs, a specific BG trader is selected uniformly at random from all the BG traders to place/replace an order.

7.1.4 Agent strategies in individual markets

Each of the informed traders and background traders is only allowed to maintain a single unit order in the market. When informed and background traders reenter the market, they can replace existing orders that have not yet been executed. We model informed traders as using limit orders and background traders as using market orders exclusively.

Informed traders' strategy: $S_{\textbf{individual}}$ We consider trading strategies in the *Zero Intelligence* (ZI) family for informed traders. There is a large literature involving ZI strategies, including some controversy, which we will not rehash here [32, 115]. While ZI strategies are clearly not the "best" trading strategies in isolation, it is also generally believed that they model order arrival processes well, and they are a standard method for choosing prices in complex agent-based market simulations [57, 114]. We first define, in the CDA market,

$$p^*_{t,\text{CDA}} = \begin{cases} \dfrac{(BID_t.p) + (ASK_t.p)}{2}, \text{if } BID_t \text{ and } ASK_t \text{ exist,} \\ \text{the most recent execution price, if any order} \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{book is empty,} \end{cases}$$

where $(BID_t.p)$ and $(ASK_t.p)$ refer to the price of BID_t and ASK_t respectively. And in the CALL market,

$$p^*_{t,\text{CALL}} = \text{the most recent } \textit{market announcement}.$$

When an informed trader IF_i places an order, a limit price is generated from

$$p_{i,t} \sim \mathcal{N}(p^*_{t,\text{market}}, \sigma_{\text{price}}),$$

where market $\in \{\text{CDA}, \text{CALL}\}$. Based on $p_{i,t}$ and $w_{i,t}$, the informed trader IF_i's strategy at time t is as follows,

$$p_{i,t} \begin{cases} >w_{i,t}, \text{places a unit sell order,} \\ <w_{i,t}, \text{places a unit buy order,} \\ =w_{i,t}, \text{uniformly at random places a unit buy or sell order.} \end{cases}$$

Note that $p_{i,t} = w_{i,t}$ is a zero probability event. If $p_{i,t} > (ASK_t.p)$ and the order is a buy order, then it executes immediately and therefore effectively functions as a market order. Similarly if $p_{i,t} < (BID_t.p)$ and the order is a sell order. We call the strategy above $S_{\text{individual}}$.

Background traders' strategy The background traders choose whether they want to buy or sell a unit uniformly at random. Once the direction is decided, the order is routed to the market and handled in a special manner as a market order through a "waiting" mechanism. The market is aware of the direction of a market order and the fact that this indicates the trader would like to execute the order at any available market price. However, market orders are not visible to any other traders in both the CDA and CALL markets, since they may need to wait for execution if there is no corresponding limit order on the other side in the CDA market, and at least until the next call in the CALL market.

Market maker's strategy In the CDA market, we also incorporate a market maker in some of our experiments. To increase the liquidity of the market, the market maker maintains a unit buy order and a unit sell order at all times. This market maker is implemented using the Bayesian market making algorithm (BMM) of Brahma et al [22], with parameters tuned to maintain near zero-profit. BMM is a learning algorithm that learns from the current bid and ask prices and the direction of incoming trades, augmented with jump prediction and a technique to widen its spread in times of uncertainty. BMM updates its own belief whenever there is an execution, and it immediately replaces its orders. Our implementation closely follows that of Brahma et al, except that we only need to use it for unit orders in our model.

7.1.5 CDA and CALL market operation

CDA market operation In the CDA market, BID_t and ASK_t are based only on orders from informed traders (and also possibly the market maker). If the order book only has market orders that are waiting from background traders, the market shows the order book as empty. The scenario that we want to simulate is that background traders are waiting in the market to buy or sell; as soon as an unfilled corresponding order becomes available, they will immediately take the other side of that order. We need to specify the execution priority in the situation where one side of the market has both market orders and limit orders from informed traders. In this case it must be that the other side of the market is empty (note that this never happens with a market maker present), otherwise the market orders on the first side would have executed. In this situation, we prioritize by time. Procedure 2 illustrates the operation of a CDA market when a new buy order arrives (a sell order arrival is similar).

CALL market operation The main difference from a standard aggregation mechanism in our implementation involves the background traders. Background traders would like to buy or sell at any price, so all the market orders are always at the top of both the sell orderbook and the buy orderbook in the CALL market. All the market orders in each orderbook are prioritized by submission time, with earlier submissions having higher priority. At each clearing time, the market collates all of the orders and computes the aggregate demand

ALGORITHM 2: CDA market operation when a new buy order arrives

Input: buy orderbook, sell orderbook

A new buy order OD_1 arrives, $(\mathrm{OD}_1.p)$ is the price of OD_1
if not empty(sell orderbook) **then**
  if OD_1 is a limit order **then**
    if $(\mathrm{OD}_1.p) \geq (ASK_t.p)$ and ASK_t comes earlier than any market order **then**
      Execution(ASK_t, OD_1) at $(ASK_t.p)$
      (*in which case, OD_1 is the highest bid and sell orderbook has limit orders from IF traders or BMM*)
    else
      if sell orderbook contains market orders **then**
        Execution(the oldest market ask, OD_1) at $(BID_t.\mathrm{p})$
        (*in which case, OD_1 is the highest bid*)
      end if
    end if
  else {OD_1 is a market order}
    if ASK_t is available and comes earlier than any market order **then**
      Execution(ASK_t, OD_1) at $(ASK_t.\mathrm{p})$
    else
      Execution(the oldest market ask, OD_1) at *the most recent execution price*
    end if
  end if
end if

and supply functions of all bids and asks, respectively. The market clears where supply equals demand, with all executions occurring at the same price, the market-clearing price. If the market only clears market orders, the market-clearing price is the most recent market announcement. If only market orders clear on one side of the market, while some limit orders clear on the other side, the market-clearing price is determined by the side that has limit orders being cleared. More specifically, if cleared buy orders consist of only market orders and cleared sell orders include limit orders, the clearing price will be the highest ask of all the cleared limit orders; if cleared sell orders consist of only market orders and cleared buy orders include limit orders, the clearing price will be the lowest bid of all the cleared limit orders. When some limit orders clear on both sides of the market, the clearing price is the midpoint of the highest ask and lowest bid of all cleared limit orders.

7.2 Competing markets

In the competing markets model, we assume that one CDA market and one CALL market run simultaneously. A single security is traded in both markets. Thus, there is only one underlying "true value" process, but the CDA and CALL markets can price the security differently. The traders choose to place orders in only one market at a time, although they can switch markets each time they re-enter. Each market is running in the same manner as when there is an individual market and each trader can maintain only one unit order in the whole system.

7.2.1 Agent strategies in competing markets

Informed traders' strategy On what basis should a trader choose which market to place an order in? One important factor is the distance between the trader's belief and $p^*_{t,\text{market}}$,

$$d^{\text{market}}_{i,t} = |w_{i,t} - p^*_{t,\text{market}}|, \tag{7.1}$$

where market $\in \{\text{CDA}, \text{CALL}\}$. We call this the *belief distance*. Comparing $d^{\text{CDA}}_{i,t}$ with $d^{\text{CALL}}_{i,t}$, the informed trader IF_i can choose one of two strategies. One is to place the order in the market that has larger $d^{\text{market}}_{i,t}$, S_{LARGE}, and the other is to place the order in the market that has smaller $d^{\text{market}}_{i,t}$, S_{SMALL}. The tradeoff here is that IF_i gets lower probability of execution but higher profit if she places the order in the market that has larger $d^{\text{market}}_{i,t}$. We will discuss the effects of these two different strategies in Section 7.3. Strategy 1 shows a summary of IF_i's strategy in the competing markets. After IF_i decides in which market to place the limit order, she follows $S_{\text{individual}}$ to decide the direction and price of the order in the selected market.

ALGORITHM 2: Strategy 1 IF_i's strategy in the competing markets at time t

1: **if** following S_{LARGE} **then**

2: $$d^{\text{CDA}}_{i,t} \begin{cases} > d^{\text{CALL}}_{i,t}, & \text{places an order at CDA market following } S_{\text{individual}} \\ < d^{\text{CALL}}_{i,t}, & \text{places an order at CALL market following } S_{\text{individual}} \end{cases}$$

3: **else** {following S_{SMALL}}

4: $$d^{\text{CDA}}_{i,t} \begin{cases} < d^{\text{CALL}}_{i,t}, & \text{places an order at CDA market following } S_{\text{individual}} \\ > d^{\text{CALL}}_{i,t}, & \text{places an order at CALL market following } S_{\text{individual}} \end{cases}$$

5: **end if**

Background traders' strategy When a background trader enters the market to place or replace a new order, she first compares the price in both markets. For instance, if she wants to buy, she will compare ASK_t and $p^*_{t,\text{CALL}}$ (if they are available), and select the market which has the lower price to place a market order there. If ASK_t is not available, that is the sell orderbook in the CDA market shows as empty, she will place the order in the CALL market. She follows a similar process for sell orders.

After every market clearing in the CALL market, the background traders check whether their orders have been executed. If not, and the corresponding order books in the CDA market are not empty, that is the sell orderbook is not empty if a BG trader wants to buy and the buy orderbook is not empty if a BG trader wants to sell (here empty means the order books do not have orders from informed traders – market orders from background traders are not visible to any trader), the background traders move their existing orders from the CALL market to the CDA market. In the implementation, the background traders with orders that did not execute are randomly permuted. Each of them moves their order to the CDA market in this random order, until there are no corresponding limit orders on the other side of the market in the CDA. This process is atomic in time.

Strategy 2 shows the overall framework for implementing buy orders for background traders in the competing markets model. The sell strategies are similar.

ALGORITHM 2: Strategy 2 The overall framework for implementing the buy orders for background (BG) traders in the competing markets model

$t = 0$
while $t \leq T$ **do**
  An event happens after Δt
  $t = t + \Delta t$
  if A BG trader arrives in the market **then**
    BG_i is selected from all BG traders uniformly at random
    if ASK_t is available **then**
      if ASK_t ¡= $p^*_{t,\mathrm{CALL}}$ **then**
        BG_i places a market order in CDA market
      else
        BG_i places a market order in CALL market
      end if
    else {sell orderbook in CDA market shows as empty}
      BG_i places a market order in CALL market
    end if
  end if
  if CALL market clear **then**
    After each market clearing
    A = {All the BG traders who have buy orders in CALL market}
    while A is not empty and ASK_t is available **do**
      Select BG_i uniformly at random from A
      $A = A - \{\mathrm{BG}_i\}$
      BG_i moves her order to the buy orderbook in CDA market, the new order age is re-generated from the age counter.
      CDA market clears
    end while
  end if
end while

7.3 Simulation results

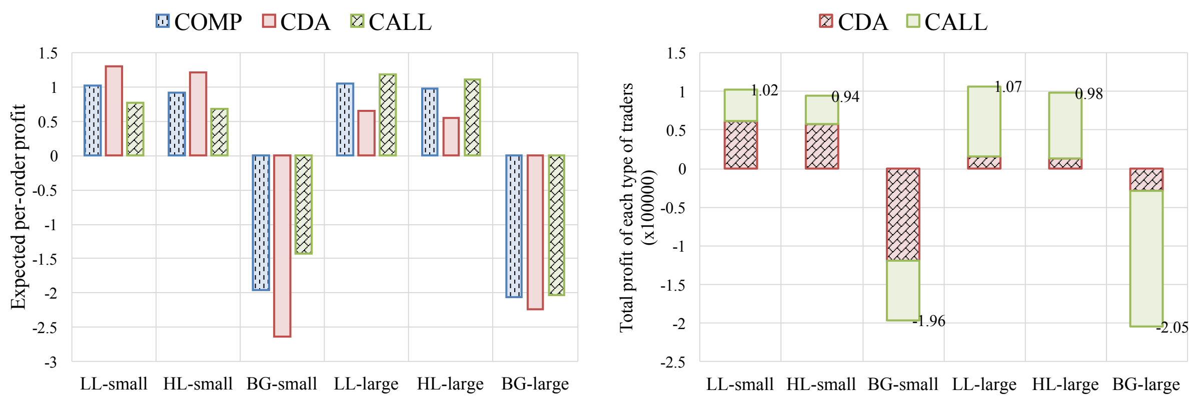

Figure 7.1: Comparison of expected per-order (left) and total (right) profit in the competing markets under the small distance (S_{SMALL}) and large distance strategies (S_{LARGE}). LL-*, HL-* and BG-* represent low latency, high latency, and background traders respectively. On the left, the first bar in each group shows the expected per-order profit in the whole competing system, while the second and third show the contributions of the CDA and CALL markets to that total. On the right, the stacked bars show total profit, with contributions from each of CDA and CALL shown within in different shades.

In this section, we simulate four different environments, namely CDA vs CALL competing markets (competing markets), an individual CDA market (i-CDA market), an individual CALL market (i-CALL market) and an individual CDA market with BMM (i-CDA-BMM market). The parameters are set as follows. Each simulation run lasts $T = 100,000$ units of time. The initial true value of the security v_0 is drawn from $\mathcal{N}(v_{\text{initial}} = 50, \sigma_{\text{initial}} = 4)$. The true value jump parameter $\sigma_j = 4.0$ and the rate parameter for the jump is $\lambda_{\text{jump}} = 0.0001$, which means there is a jump every 10000 units of time on average. We have 20 informed traders, 10 high latency traders and 10 low latency traders, and 20 background traders. Reentry rates are fixed across the environments, with informed traders arriving in the market at rate $\lambda_{\text{IF}} = 2$, and background traders entering at rate $\lambda_{\text{BG}} = 1$. In all settings, CALL markets clear every 1 unit of time, $\tau = 1$. The standard deviation of informed traders' belief is $\sigma_{\text{trader}} = 2.0$, and the high latency traders' information delay is $\delta = 1000$ units of time. The time between jumps is 10000 units of time on average, so high latency traders have information with no delay a significant fraction of the time, as $v_{t-\delta} = v_t$. Following a jump, high latency traders receive staler information for the next 1000 units of time. The standard deviation of the distribution from which informed traders draw ZI prices is $\sigma_{\text{price}} = 4.0$. We simulate both S_{LARGE} and S_{SMALL} strategies for informed traders.

Across our experiments, we are interested in the total profit, expected per-order profit and order execution percentage for each trader type. At time T, all shares held by traders are liquidated at price v_T. Unfilled orders are abandoned. The expected per-order profit of each trader type is total profit divided by the total number of executed and replaced un-traded orders. In the environment with competing markets, we also calculate the total and expected per-order profit in CDA market and CALL market separately.

Figure 7.1 shows that the informed traders make higher profit in both per-order and in total when using S_{LARGE}. The difference are small but statistically significant. Therefore, we would expect the informed traders to choose S_{LARGE} if given these two options (if they had to choose one as a group), confirming our intuition that traders gravitate to markets in which they perceive more mispricing. Because of this, for the rest of our analysis, we use S_{LARGE} as the strategy for informed traders in the competing markets.

As mentioned in the introduction, one measure of social welfare is the "price of immediacy" which is the loss suffered by background traders. Figure 7.2 shows that, for the non-competing settings, background traders perform better in the i-CALL market than the i-CDA market in terms of both the expected per-order (left figure) and total (right figure) profit (consequently, informed traders have lower profit in the i-CALL market than the i-CDA market). The i-CDA-BMM market has much higher social welfare, as measured by

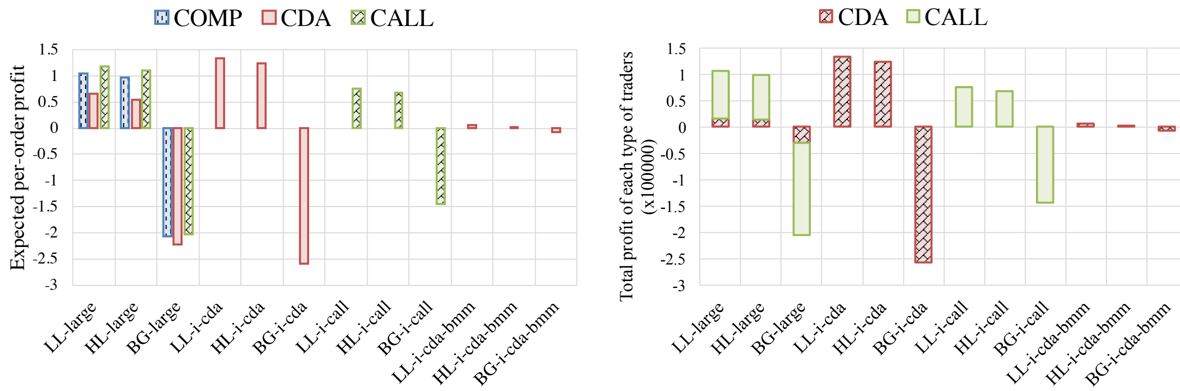

Figure 7.2: The expected per-order (left) and total profit (right) of different types of traders in the competing markets (*-large), the i-CDA market (*-i-cda), the i-CALL market (*-i-call) and the i-CDA-BMM market (*-i-cda-bmm).

background trader losses, than both i-CALL and i-CDA markets. This confirms some of the results of Wah and Wellman [144] in a completely different model and setting. One possible solution to the problems resulting from HFT may then be to have market-making agents who are regulated and deployed to perform this specific role in CDAs. They could be compensated separately for this role. However, (1) there are additional risks associated with this role [36] and (2) markets have been moving away from having designated specialists and allowing HFTs and others to fulfil the role of market-makers. Given these practical realities, it is important to understand how markets can function without them, so we focus on comparing situations with no market maker.

In the competing markets (first 3 panels of Figure 7.2), we analyze the expected per-order and total profit in the whole system, and also in the CDA market and CALL market separately. Similar to running an individual market, the order execution percentage is close to 100 for background traders in the competing markets (shown in Figure 7.3), and so background traders are not losing out in terms of order execution. Considering expected per-order profit, background traders do better in the CALL market than the CDA market in the competing markets (see left figure of Figure 7.2, *-large). Overall, they are doing worse in the CALL

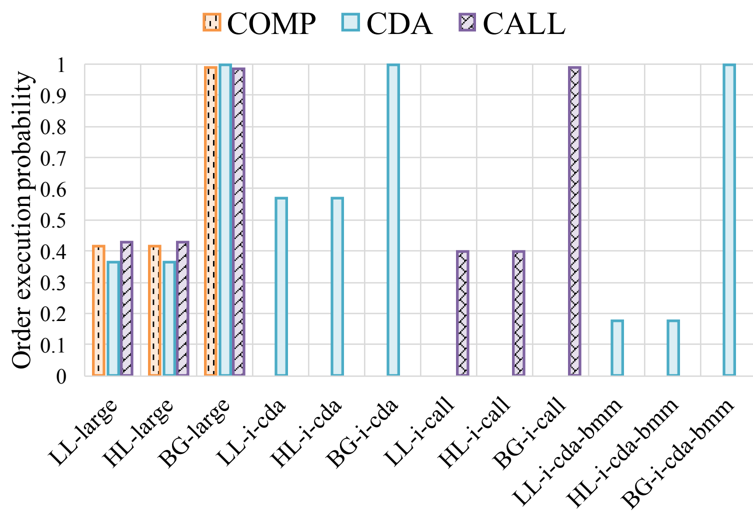

Figure 7.3: Order execution percentage in the competing markets (*-large), the i-CDA market (*-i-cda), the i-CALL market (*-i-call) and the i-CDA-BMM (*-i-wmm).

market than the CDA market in the competing system (shown on the right of Figure 7.2, *-large), but this is because a vast majority of orders are going to the CALL market (see Figure 7.4), and they lose more money there. In sum, they are doing slightly better in the competing markets than they do in the i-CDA market in terms of both the expected per-order and total profit. The better news is that the CALL is absorbing a large fraction of the orders, driving trade away from the CDA (see Figure 7.4). This is promising, because if the CALL could absorb all the trades, the BG traders would be better off, as the system would reduce to the i-CALL market. We note that these results are robust for a wide range of strategy parameters, information delay and arrival rates.

7.4 Learning traders

The analysis above shows that frequent call markets absorb a large fraction of trade when we assume that all informed traders use the same strategy. We now use the insights developed

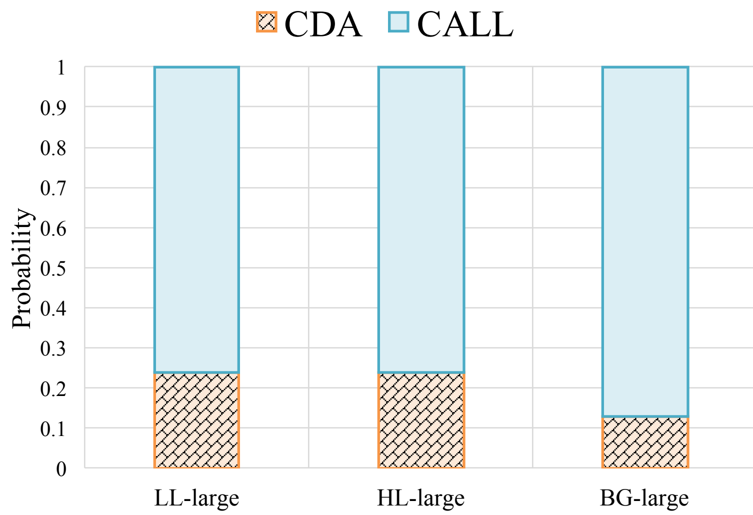

Figure 7.4: Illustration of proportion of orders entering CDA market vs. CALL market in the competing markets under $\mathrm{S}_{\mathrm{LARGE}}$.

in these models to analyze strategic market choice. In this section, we introduce a learning framework where informed traders learn a parametric form for the expected profit of choosing to place an order in a market, and also a non-parametric probability of order execution given the distance between that market's price and the trader's belief (the *belief distance* defined previously). We then analyze the behavior of the system with these learning traders.

7.4.1 Learning algorithm

The expected profit π_t of an order placed at time t, contingent on its execution, according to a trader's belief, is

$$\pi_t = \begin{cases} w_t - \text{execution price, if buy} \\ \text{execution price } - w_t, \text{if sell} \end{cases} \tag{7.2}$$

where w_t is the trader's belief about the true value at time t. The trader computes expected profit assuming her belief is correct. The main idea here is to predict the expected profit if the order is placed in a particular market. Traders must learn an estimate of this as a function of the belief distance. In this chapter, we allow all the traders to learn an expected profit function that is quadratic in the belief distance. Each learning trader LIF_i uses an online regression algorithm for reinforcement learning based on one developed by Walsh et al [145]. The form of the learning model is:

$$y_{i,t}^{\text{market}} = \mathbf{D}_{i,t}^{\text{market}} \mathbf{Q}_i^{\text{market}}, \tag{7.3}$$

where $y_{i,t}^{\text{market}}$ predicts the expected per-order profit contingent on execution, market $= \{\text{CDA}, \text{CALL}\}, \mathbf{D}_{i,t}^{\text{market}} = [(d_{i,t}^{\text{market}})^2, d_{i,t}^{\text{market}}, 1]$ and $\mathbf{Q}_i^{\text{market}}$ contains the weight parameters of the model for market $= \{\text{CDA}, \text{CALL}\}$. Thus, the predicted expected per-order profit not contingent on execution is given by

$$\mathbb{E}(y_{i,t}^{\text{market}}) = \Pr(\text{exe—}d_{i,t}^{\text{market}}) \mathbf{D}_{i,t}^{\text{market}} \mathbf{Q}_i^{\text{market}} \tag{7.4}$$

The probability of execution $\Pr(\text{exe—}d_{i,t}^{\text{market}})$ is learned non-parametrically by counting successful and unsuccessful executions in bins of the belief distance.

The trader uses an ϵ-greedy algorithm to select a market to trade in along the learning path. Whenever the trader makes a decision, with probability $1 - \epsilon$, she places an order in the market with higher predicted expected profit, and with probability ϵ, she randomly picks one market to place the order ($\epsilon = 0.1$ in our case). After market selection, the trader chooses a price based on the ZI strategy $S_{\text{individual}}$.

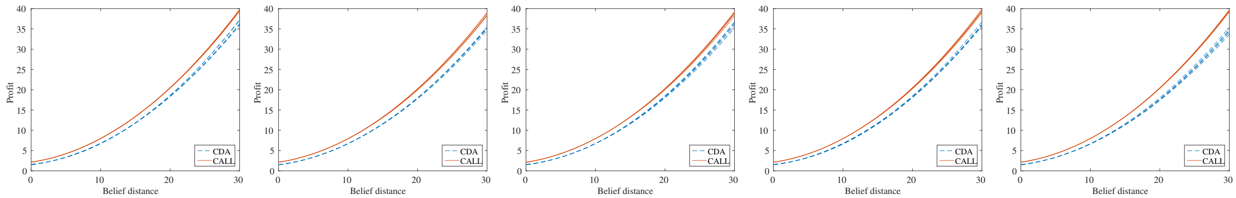

Figure 7.5: Five example plots of the curves traders learn for profits in the CDA and CALL markets. Each graph shows the learned curves of 5 traders in one instantiated learning simulation.

7.4.2 Results

Our goal is to use this model of learning traders to investigate two questions: (1) If all the informed traders use the same learning algorithm, do they converge to (approximate) equilibrium strategies? (2) Can we characterize any equilibria of the competing markets system?[18] We use an experimental framework similar to Section 7.3. A CALL market and a CDA market run simultaneously from $[0, T]$, $T = 100,000$. The CALL interval $\tau = 1$. There are 20 learning informed traders (LIF). These are all low latency traders who observe information with no delay; therefore $w_{i,t} \sim \mathcal{N}(v_t, \sigma_{\text{trader}})$, where v_t is the underlying true value of the security at time t, and σ_{trader} is the noise parameter. The reentry rate of learning informed traders is $\lambda_{\text{LIF}} = 2$. There are 20 background traders (with reentry rate $\lambda_{\text{BG}} = 1$) following Strategy 2. In addition, we also simulate the existence of a pool of *fixed* informed traders (FIF) who are committed to a particular market, either CDA or CALL. This is to ensure that there is some flow of trade in each market – otherwise there are degenerate equilibrium paths where all traders start off by going to one of the markets, and there is never incentive to deviate to the other. There are 5 fixed (low-latency) informed traders in

[18]Note that traders only choose which market to place an order in (albeit as a function of the belief distances to both markets, so this can be a complex decision space). Once the market is determined, the choice of price is according to the ZI strategy. The problem becomes exponentially more complex if traders can strategize over both market choice and price. Over time, traders should learn the expected profit of the ZI strategy in a market, a useful proxy for the profit potential of that market. It could be interesting to interact this learning problem with different pricing strategies, but some restriction will always be necessary to gain any traction.

each market who place orders following $S_{\text{individual}}$. The reentry rate of fixed informed traders is defined as λ_{FIF}. We vary the reentry rate of fixed informed traders $\lambda_{\text{FIF}} = \{1, 0.1, 0.0005\}$.

Outcomes of the learning process The first question is whether the learning process followed by the informed traders converges, and, if so, whether the learned representations are a good approximation to the true profit function. Empirically, we find that the estimates of $\mathbf{Q}_i^{\text{CDA}}, \mathbf{Q}_i^{\text{CALL}}$ under different λ_{FIF} settings all do converge. Further, each trader's parameters converge to very similar ranges (see Figure 7.5). We check whether these parameters are a good approximation by fixing the parameters of all the traders except one and having them play strategies using those parameters. For the remaining trader, we flip a coin to determine the choice of market, and test whether the profit achieved is well-fit by the curve given by the learned parameters. Figure 7.6 shows that the learning curves of the trader for each market are very close to the polynomial curves that are best-fit to the profits achieved using the randomized strategy, confirming both that the quadratic space is a good fit and that the learned parameters are correct for the environment.

Equilibrium As mentioned above, there are two main questions we would like to engage. First, since all the traders are converging to a particular set of learned parameters, do these parameters constitute an equilibrium or an approximate equilibrium (under the specified space of strategies – i.e., where the strategy is a mapping from d^{CALL} and d^{CDA} to one of the two markets, which can be specified by the quadratic form of the expected profit function and the nonparametric probability of execution model)? Is there a profitable deviation (some other set of parameters that one of the traders could use and increase her profits)?

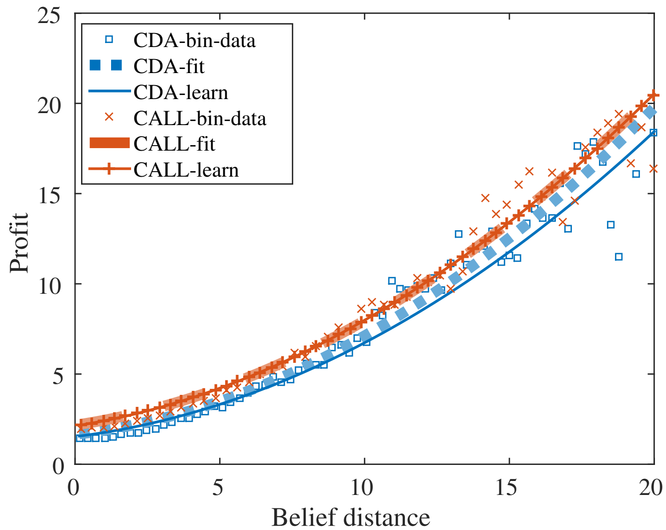

Figure 7.6: The learned curve vs. the best fit to realized after-the-fact data for profit as a function of belief distance in CDA and CALL markets.

To find deviations, we search the parameter space (holding the execution probability model constant) for this trader using Bayesian Optimization (BO), a powerful framework for optimization of a black-box function or expensive objective function that uses very few function evaluations [130]. Here, the objective function is the expected profit of a trading strategy that uses the parameters $\mathbf{Q}_i^{\mathrm{CDA}}, \mathbf{Q}_i^{\mathrm{CALL}}$ when the other traders are using their learned strategies. Market selection is determined by $y_{i,t}^{\mathrm{CDA}}$ and $y_{i,t}^{\mathrm{CALL}}$ based on Equation (7.4), so the actual values of $\mathbf{Q}_i^{\mathrm{CDA}}$ and $\mathbf{Q}_i^{\mathrm{CALL}}$ are not important in themselves. The important thing is how they decide the relation between $y_{i,t}^{\mathrm{CDA}}$ and $y_{i,t}^{\mathrm{CALL}}$ at each prediction. Therefore, based on the value of parameters of the learned curves from Figure 7.5, we constrain our search space from $[-\mathbf{10}, \mathbf{10}]$, as this is enough to represent the relation between predicted profit in these two markets. We utilize an existing code base for BO [60] to search the space.

Our results show that the learned parameters yield an approximate equilibrium, achieving between 90-95% (0.91, 0.95, 0.95 for $\lambda_{\mathrm{FIF}} = .0005, 0.1, 1$ respectively) of the profit of the best response strategy found by BO. In the learned approximate equilibrium, typically above 90%

of orders are placed in the CALL market. Interestingly, the best response strategy found by BO always resulted in the deviating trader placing *every single order* in the call market. So we then asked whether all informed traders placing all their orders in the call market is an equilibrium, and found that, except under exceptional conditions, it is (that is, BO returned a set of parameters for the remaining trader that resulted in that trader placing all its orders in the CALL market as well). The only condition which we found under which it is not an equilibrium is when there is very little liquidity from FIF traders in the CDA market, but still some background traders – in this case the deviating informed trader can essentially become a price setter and trade with the background traders at whatever prices it chooses.

7.5 Conclusion

We have developed an agent-based model in the tradition of classic microstructure models to engage the question of whether frequent call markets can drive liquidity away from CDA markets. If they could do so, this would have the potential to increase welfare both by reducing transaction costs for average market participants and by reducing the incentive for firms to engage in the latency "arms race." Our results are promising. Even in the presence of impatient background traders who primarily demand immediacy and are willing to pay for it, we show in both a simple zero-intelligence model, and more sophisticated learning and equilibrium settings, that call markets have the potential to attract a large fraction of the order flow. In addition to the policy implications, we believe the modeling approach taken in this chapter constitutes a useful bridge between classic financial microstructure models and more complex agent-based models, preserving intuition from the former, while allowing us to examine richer environments and questions.

Chapter 8

Conclusion and Future Work

8.1 Conclusions

This thesis addresses the design, analysis, and modeling of mechanisms and information structures for real-world applications. We provide new and richer models for understanding static and dynamic markets across different domains, including kidney exchange and financial markets. We address the outcome of single platforms, as well as fill in gaps in the study of the dynamics of multiple platform interactions. Within static markets, we show how the information structure and environment influence the outcomes of matching markets and auction markets. Because mechanisms are harder to change than the information available to participants, understanding the effect of information structure becomes valuable. For dynamic markets, we first provide baselines of different market-clearing rules on market outcomes. Our main contribution in this part is modeling the dynamics of multiple platform competition, which is understudied in most of the domains. We focus on social welfare, as well as the equilibrium behaviors of both individuals and markets in competing-market systems. Through these efforts, we can make better-grounded policy recommendations for both kidney exchange and financial markets, and help better inform the debates in both areas.

8.2 Future Work

While we provide suggestions for future work within its domain in each chapter, here I provide a broader picture of the overall thesis, especially with regard to dynamic matching markets.

8.2.1 Static Markets

For information design in static markets, we investigated the role of information on aggregate and distributional outcomes, and also designed revenue enhancement signal structures. In this thesis, we look at stylized models for both matching markets and auction markets. One direction is to consider a richer and more complex model. For example, in matching markets, we study the situation where there is a universally shared, common knowledge ranking of all firms, and there is a "true" universally shared ranking of all candidates. A natural extension is to examine situations with more diverse preferences, and where interviews are costly but not necessarily budgeted (allowing employers to decide strategically how many candidates to interview). In auction markets, we consider the second-price auction with one seller and two bidders. This model can be extended to the auction with more bidders, though the complexity for determining the equilibrium strategies and revenue-enhancing signal structure will be dramatically increased. Another direction is to investigate what happens when the market outcome is not decided in a centralized manner (using the Gale-Shapley algorithm as above, or in a second-price auction, for example), but must instead take place through explicit offers and acceptances.

8.2.2 Dynamic Markets

In dynamic markets, we mainly focus on the dynamics of multiple markets' interaction. On this front, there are several more issues can be further studied. One is to consider more complex models which would make these models more applicable to many real-world settings. For example, in dynamic kidney exchange problem, the future work can take network effects (where more popular exchanges have an easier time attracting agents, lower operating costs, higher probabilities of two agents forming an acceptable transaction, and other advantages) into account. Besides, we look at only two overlapping markets for both kidney exchange and financial markets; generalizing this to any number of overlapping markets would also be of interest.

Chapter 6 provides a cardinal utility model. Another important line to study is how to use this cardinal utility model is to further improve outcomes in kidney exchange in the dynamic setting. For example, we have done related work which develops a matching algorithm for kidney exchange with compatible pairs. The algorithm utilizes the estimation of "shadow survival" (as opposed to the more common shadow prices) by the cardinal utility model and is mainly based on the online primal-dual technique. A future interesting direction is to incorporate the waiting time in this kind of utility model. Here, future work can analyze the relationship between match quality and waiting time to assess whether there is a significant change in the expected quality of a match as a function of waiting time; and can also predict the relationship between waiting time and match quality. The transplant center's utility can also be taken into account: the costs of performing local vs. non-local transplants, entering pairs into kidney exchanges, actual procedure costs (which can vary based on donor and recipient characteristics), and insurance reimbursements.

A richer strategy space can also be considered. For example, in competing financial markets, we consider the learning strategies only in quadratic space. One way to release this constraint is to consider non-parametric strategies using Gaussian process. In dynamic matching markets, we simplify the strategy by types (short-lived and long-lived), where it is interesting to study the relation between agent strategies and their criticality.

Appendix A

Additional Experiments For Competing Dynamic Matching Markets

In this part, we provide additional results supporting the dynamic kidney exchange experiments of Section 4.5.2. Figure A.1 corresponds to the 2-cycle-only experiments of Figures 4.4 and 4.5 in the body of the Chapter 4; instead of varying the market overlap parameter γ on the x-axis, they vary the probability α of entering either the Greedy$_c$ or Patient$_c$ market, while holding γ constant for a variety of values. Similarly, Figure A.2 corresponds to the 2- and 3-cycle UNIFORM matching policy experiments of Figures 4.6 and 4.7. Finally, Figure A.3 corresponds to the UNIFORM3 matching policy results shown in Figures 4.8 and 4.9.

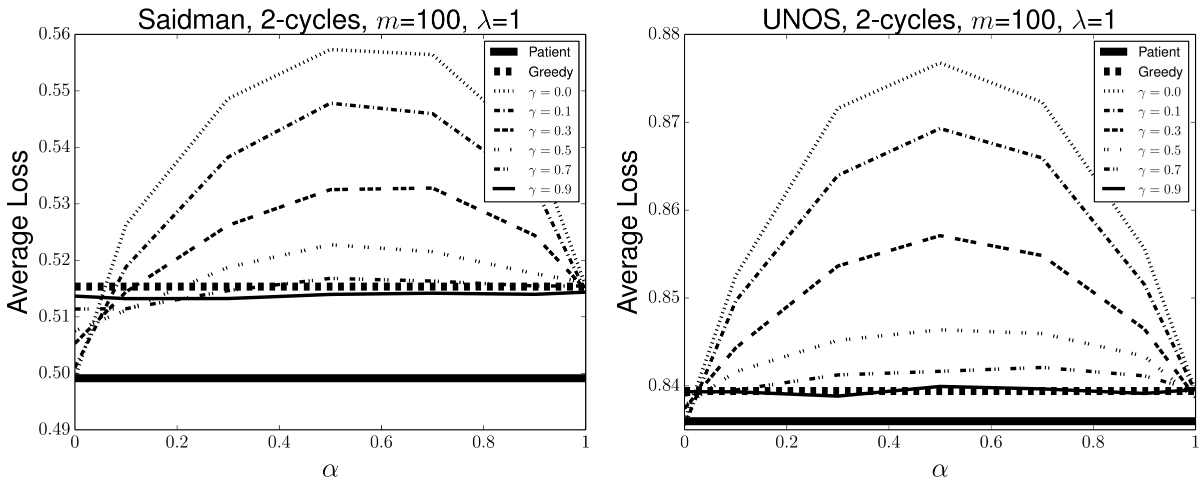

Figure A.1: 2-cycles-only experiments, paired with Figure 4.4 (left) and Figure 4.5 (right).

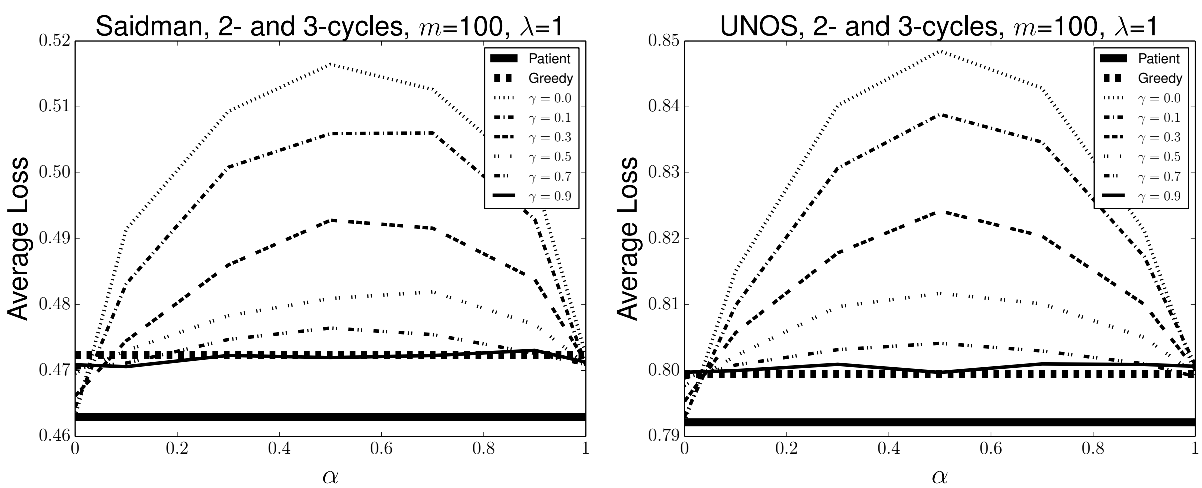

Figure A.2: 2- and 3-cycle Uniform experiments, paired with Figure 4.6 (left) and Figure 4.7 (right).

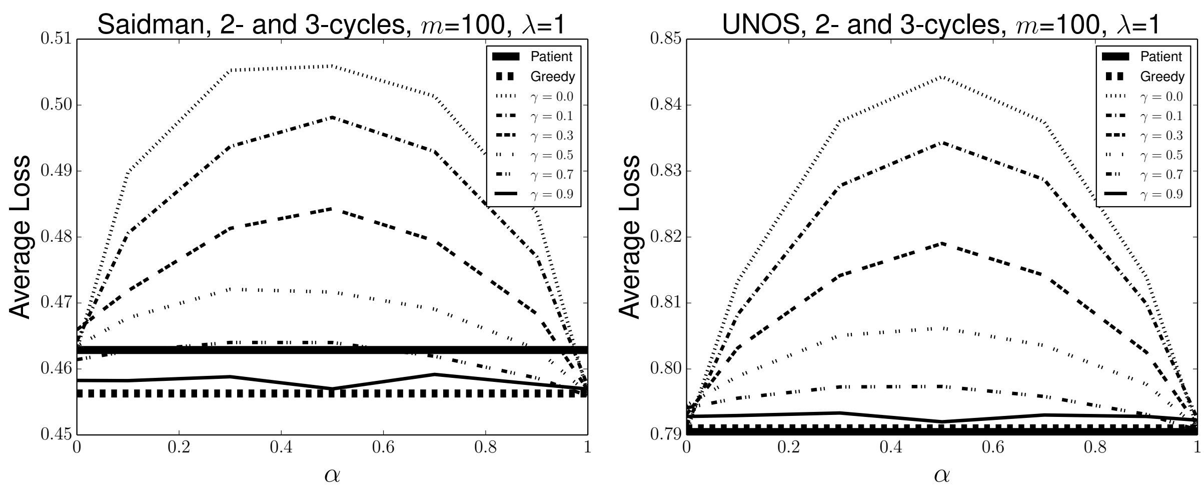

Figure A.3: 2- and 3-cycle UNIFORM3 experiments, paired with Figure 4.8 (left) and Figure 4.9 (right).

Appendix B

Simulator Details

This part provides the details of the simulator in Section 6.3. Our basic simulation model is based on the distribution of all relevant recipient and donor characteristics from the data of the center. The characteristics of each donor-recipient pair are generated from the distribution of the center's data (See Table 6.1). We determine compatibility based on the simulator from Saidman et al [127], which utilizes PRA and ABO compatibility. More specifically, we first generate a donor-recipient pair, with all LKDPI-related characteristics generated sequentially in a manner that respects the data distributions in Table 6.1 and the correlation structure shown in Figure 6.2. We then generate the PRA (percentage reactive antibodies) and compatibility based on the Saidman model. The exact details of how we generate the characteristics can be found in Algorithm 3.

Age(Years)	Average Measured GFR (ML/min/1.73m$\hat{2}$)
20-29	116
30-39	107
40-49	99
50-59	93
60-69	85
70+	75

Table B.1: Average measured GFR by age in people.

ALGORITHM 3: Details of Generating Living Donor Pair

Input: Pair ID
Output: A living donor pair
Sample following characteristics based on the distribution of Table 6.1:
Donor Age $\sim \mathcal{N}(48.22, 12.68)$;
Donor Sex: $P(F) = 0.7, P(M) = 0.3$;
Rec Sex: $P(F) = 0.35, P(M) = 0.65$;
Donor eGFR: Table-B.1 based on Donor Age;
Donor SBP Table-B.1 based on Donor eGFR;
Donor Weight: Sample based on Donor Sex;
Rec Weight: Sample based on Recipient Sex;
Donor BMI: 0.0948 (Donor Weight) + 11.387;
Donor/Rec Weight Ratio: Donor Weight/Rec Weight;
Donor Blood Type: Based on Saidman's simulator ;
Recipient Blood Type: Based on Saidman's simulator ;
Donor is African American: Based on Donor Blood Type ;
Donor cigarette use: $P(Y) = 0.32, P(N) = 0.68$;
Donor&Rec Related: Based on Table 6.1;
Check Donor&Rec ABO compatibility;
Donor&Rec HLA-B mismatches and Donor&Rec HLA-DR mismatches: Jointly sample from Table 6.1 based on whether the pair is related or not;
Donor&Rec isWifePatient: Based on Saidman's simulator; if the recipient is female and the donor-recipient pair is unrelated, the probability that the donor is the recipient's spouse is 0.4897;
Recipient PRA: Based on Saidman's simulator;
Generate crossmatch incompatibility: Based on PRA and isWifePatient;
Determine compatibility: The pair is compatible if and only if both ABO compatible and a negative crossmatch.

To note, in this simulator, (1) the estimated GFR (line 5) is generated from Table B.1 which depends on age instead of using the distribution from Table 6.1[19]; (2) The BMI (line 9) is generated based on a regression on data from the transplant center; (3) When we consider a counterfactual pair, we always assume they are unrelated. (4) HLA-B and HLA-DR mismatches of a donor-recipient pair are generated based on whether the donor and recipient are related or not. When we need to decide the HLA-B and HLA-DR mismatches of a counterfactual pair, we use the distribution from the counterfactual matrix instead of the distribution from the original dataset.

[19]See https://www.kidney.org/sites/default/files/docs/12-10-4004_abe_faqs_aboutgfrrev1b_singleb.pdf.

Appendix C

Varying the Size of the Incompatible Pair Pool

This appendix provides more experimental results corresponding to Section 6.3.2. While the rate of entry of compatible and incompatible pairs may be similar, it is possible that one or the other population is less likely to go through with a transplant. This could result in different ratios between the sizes of the two pools. In order to study how our results would vary with different assumptions about this, we hold the number of compatible pairs fixed and vary the number of incompatible pairs. Both compatible and incompatible pairs are randomly generated using the population characteristics from Table 6.1 and following Algorithm 3, where the compatibility is decided by Saidman's simulator. The number of compatible pairs we consider are 50, 100, and 200, while the number of incompatible pairs ranges from 10 to 200.

The performance of incompatible pairs – *The number of matched pairs.* For incompatible pairs, we are primarily interested in the increase in the number of matches when

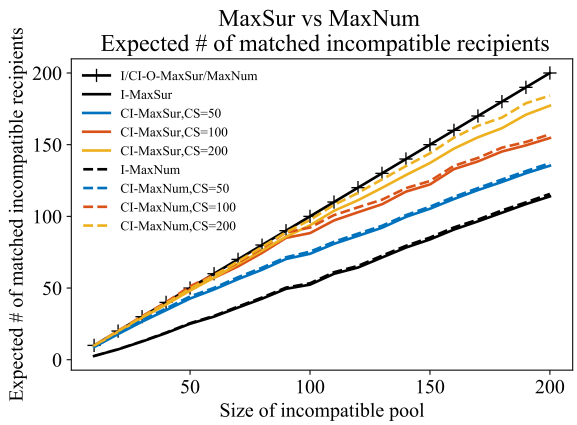

(a)

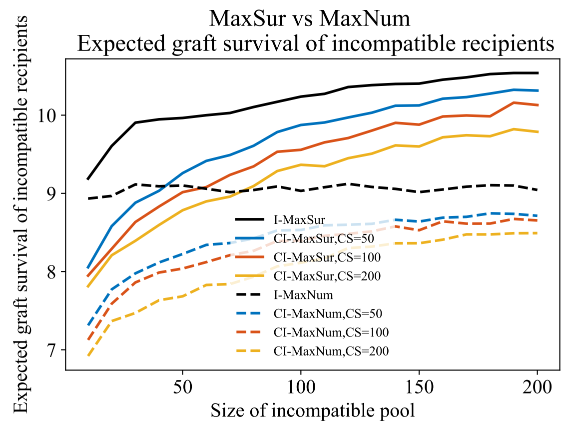

(b)

Figure C.1: (a) Expected number of matched **incompatible** pairs under maximizing expected graft survival (solid lines) and expected number of matched recipients (dash lines) when holding the number of compatible pairs (CS) as 50, 100, 200; (b) Expected graft survival of **incompatible** recipients under maximizing expected graft survival (solid lines) and expected number of matched recipients (dash lines) when holding the number of compatible pairs (CS) as 50, 100, 200 and varying the size of incompatible pairs from 10 to 210. Each point in the graph is an average of 500 simulations.

compatible pairs join the pool. Figure C.1a shows the expected number of matched incompatible pairs/recipients when maximizing expected graft survival of all cycles (*-MaxSur) and maximizing the number of matched pairs ($w_e = 1$, *-MaxNum). For both objective functions, the optimal matching will match all the pairs (I/CI-O-MaxSur/MaxNum). In two-&three-cycle swap, both objective functions achieve similar performance (though *-MaxNum are slightly better then *-MaxSur). When the market is thick enough (compatible size is 200, CS=200), the number of matched incompatible pairs is very close to the optimal solution. In general, for two-&three-cycle swap, the pool with compatible pairs (CI-*) matches far more incompatible recipients then only running two-&three-cycle swap within the incompatible pairs (I-*).

–***Expected graft survival.*** We now investigate how expected graft survival of incompatible pairs changes when compatible pairs join the pool. The results of comparing *-MaxSur and

*-MaxNum can be found in Figure C.1b. Overall, *-MaxSur (solid lines) has longer expected graft survival than *-MaxNum (dash lines) as we expect. When compatible pairs participate, expected graft survival of incompatible pairs is lower than when running two-&three-cycle swap within incompatible pairs (I-*). Another interesting observation is that the expected graft survival of incompatible recipients decreases as the number of compatible pairs increases for both *-MaxSur and *-MaxNum.

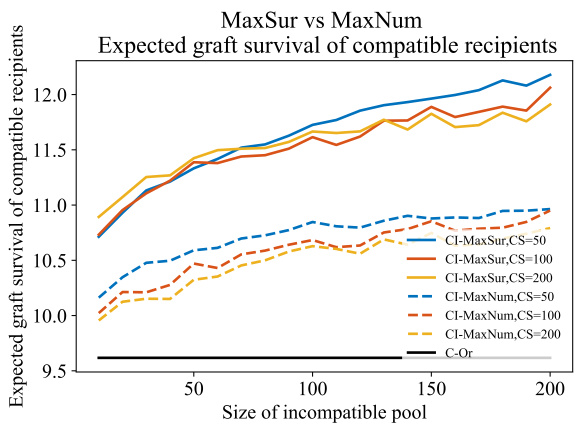

(a)

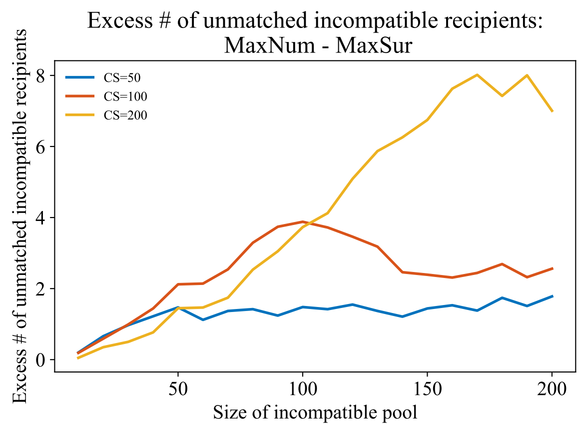

(b)

Figure C.2: two-&three-cycle swap : Expected graft survival of **compatible** recipients under maximizing expected graft survival (solid lines) and expected number of matched recipients (dash lines) when holding the number of compatible pairs (CS) as 50, 100, 200 and varying the size of incompatible pairs from 10 to 210.

The performance of compatible pairs – *Expected graft survival.* Under the Pareto improvement restriction, the compatible pairs are guaranteed to match with their original donor at least and they only swap if they can find a better organ for both the Optimal and two-&three-cycle swap. From Figure C.2a we can see that for both objective functions (MaxSur and MaxNum), compatible pairs have a substantially longer graft survival for participating two-&three-cycle swap (CI-*) than if matched with their original donor (C-or). The size of the compatible pool does not have major influence on the performance. It is also obvious that the compatible pairs benefit more when the market clearing algorithm maximizes the expected graft survival rather than the number of matched pairs. The number of incompatible pairs who are not matched when maximizing graft survival, but who would have been matched when maximizing the number of matches, is shown in Figure C.2b.

Bibliography

[1] Atila Abdulkadiroğlu and Simon Loertscher. Dynamic house allocations, 2007. Working paper.

[2] David J. Abraham, Avrim Blum, and Tuomas Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pages 295–304, New York, NY, USA, 2007. ACM.

[3] Ittai Abraham, Susan Athey, Moshe Babaioff, and Michael Grubb. Peaches, lemons, and cookies: Designing auction markets with dispersed information. Technical report, January 2011.

[4] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Thickness and information in dynamic matching markets. 2017.

[5] Shani Alkoby, David Sarne, and Igal Milchtaich. Strategic signaling and free information disclosure in auctions. In *Proc. AAAI*, 2017.

[6] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. A dynamic model of barter exchange. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1925–1933, 2015.

[7] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017.

[8] Elliot Anshelevich, Meenal Chhabra, Sanmay Das, and Matthew Gerrior. On the social welfare of mechanisms for repeated batch matching. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 60–66, 2013.

[9] Elliot Anshelevich and Sanmay Das. Matching, cardinal utility, and social welfare. *ACM SIGECom Exchanges*, 9(1):4.1–4.7, 2010.

[10] Elliot Anshelevich, Sanmay Das, and Yonatan Naamad. Anarchy, stability, and utopia: Creating better matchings. *Autonomous Agents and Multi-Agent Systems*, 26(1):120–140, 2013.

[11] Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. On matching and thickness in heterogeneous dynamic markets. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 765–765, New York, NY, USA, 2016. ACM.

[12] Itai Ashlagi, Felix Fischer, Ian A Kash, and Ariel D Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, 91:284–296, 2015.

[13] Itai Ashlagi, Patrick Jaillet, and Vahideh H. Manshadi. Kidney exchange in dynamic sparse heterogenous pools. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 25–26, New York, NY, USA, 2013. ACM.

[14] Itai Ashlagi and Alvin E Roth. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*, 9(3):817–863, 2014.

[15] Lawrence M. Ausubel. "auctions: Theory" for the new palgrave, 2 nd edition. 2006.

[16] Pranjal Awasthi and Tuomas Sandholm. Online stochastic optimization in the large: Application to kidney exchange. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 405–411, 2009.

[17] Mariagiovanna Baccara, SangMok Lee, and Leeat Yariv. Optimal dynamic matching. 2015. Working paper.

[18] Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. *Journal of Economic Theory*, 137(1):580 – 609, 2007.

[19] Elazar Berkovitch and Naveen Khanna. A theory of acquisition markets: Mergers versus tender offers, and golden parachutes. *Review of Financial Studies*, 4(1):149–174, 1991.

[20] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.

[21] Joseph H Bragdon, Ali Hortaprivatelycsu, and Dan Ariely. Matching and sorting in online dating. *American Economic Review*, 100(1):130–163, 2010.

[22] Aseem Brahma, Mithun Chakraborty, Sanmay Das, Allen Lavoie, and Malik Magdon-Ismail. A Bayesian market maker. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 215–232, New York, NY, USA, 2012. ACM.

[23] R Scott Braithwaite, David O Meltzer, Joseph T King Jr, Douglas Leslie, and Mark S Roberts. What does the value of modern medicine say about the $50,000 per quality-adjusted life-year decision rule? *Medical Care*, 46(4):349–356, 2008.

[24] Erik Brinkman, Michael P. Wellman, and Scott E. Page. Signal structure and strategic information acquisition: Deliberative auctions with interdependent values. In *Proc. AAMAS*, pages 229–236, 2014.

[25] Peter Bro Miltersen and Or Sheffet. Send mixed signals: Earn more, work less. In *Proc. ACM EC*, pages 234–247, 2012.

[26] Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. High-frequency trading and price discovery. *Review of Financial Studies*, 27(8):2267–2306, 2014.

[27] Eric B Budish, Peter Cramton, and John J Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, pages 1547–1621, 2015.

[28] K. Burdett and R. Wright. Two-sided search with nontransferable utility. *Review of Economic Dynamics*, 1:220–245, 1998.

[29] L. Elisa Celis, Anna R. Karlin, Kevin Leyton-Brown, C. Thach Nguyen, and David R. M. Thompson. Approximately revenue-maximizing auctions for deliberative agents. In *Proc. AAAI*, pages 1313–1318, 2012.

[30] Hector Chade. Matching with noise and the acceptance curse. *Journal of Economic Theory*, 129(1):81–113, 2006.

[31] Meenal Chhabra, Sanmay Das, and David Sarne. Competitive information provision in sequential search markets. In *Proc. AAMAS*, pages 565–572, 2014.

[32] Dave Cliff, Janet Bruten, et al. Zero not enough: On the lower limit of agent intelligence for continuous double auction markets. *HP Laboratories Technical Report HPL*, 1997.

[33] Miguel Constantino, Xenia Klimentova, Ana Viana, and Abdur Rais. New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research*, 231:57–68, 2013.

[34] David Roxbee Cox and David Oakes. *Analysis of survival data*, volume 21. CRC Press, 1984.

[35] S. Das and E. Kamenica. Representations of information structures. In *Allerton Conf. on Communication, Control, and Computing*, pages 737–743, Sept 2015.

[36] Sanmay Das. The effects of market-making on price dynamics. In *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 887–894, Estoril, Portugal, May 2008.

[37] Sanmay Das, John P Dickerson, Zhuoshu Li, and Tuomas Sandholm. Competing dynamic matching markets. In *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*, volume 8, page 19, 2015.

[38] Sanmay Das and Emir Kamenica. Two-sided bandits and the dating market. In *Proceedings of IJCAI*, pages 947–952, 2005.

[39] Sanmay Das, Emir Kamenica, and Renee Mirka. Reducing congestion through information design. In *Proceedings of the 55th Allerton Conference on Communication, Control, and Computing*, pages 1279–1284, 2017.

[40] Sanmay Das and Zhuoshu Li. The role of common and private signals in two-sided matching with interviews. In *Proc. WINE*, pages 492–497, 2014.

[41] Sanmay Das and John N. Tsitsiklis. When is it important to know you've been rejected? A search problem with probabilistic appearance of offers. *Journal of Economic Behavior and Organization*, 74:104–122, 2010.

[42] Constantinos Daskalakis, Christos Papadimitriou, and Christos Tzamos. Does information revelation improve revenue? In *Proc. ACM EC*, pages 233–250, 2016.

[43] Nicholas Diakopoulos. Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*, 3(3):398–415, 2015.

[44] Nicholas Diakopoulos. Accountability in algorithmic decision making. *Communications of the ACM*, 59(2):56–62, 2016.

[45] John P. Dickerson, David Manlove, Benjamin Plaut, Tuomas Sandholm, and James Trimble. Position-indexed formulations for kidney exchange. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 2016.

[46] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Dynamic matching via weighted myopia with application to kidney exchange. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1340–1346, 2012.

[47] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Optimizing kidney exchange with transplant chains: Theory and reality. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 711–718, 2012.

[48] John P. Dickerson and Tuomas Sandholm. FutureMatch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[49] Joanna Drummond and Craig Boutilier. Elicitation and approximately stable matching with partial preferences. In *Proceedings of IJCAI*, 2013.

[50] Shaddin Dughmi. On the hardness of designing public signals. *Games and Economic Behavior*, 2018.

[51] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. In *Proc. ACM EC*, pages 514–531, 2012.

[52] Leah Epstein, Asaf Levin, Julián Mestre, and Danny Segev. Improved approximation guarantees for weighted matching in the semi-streaming model. *SIAM Journal on Discrete Mathematics*, 25(3):1251–1265, 2011.

[53] Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.

[54] Haluk Ergin, Tayfun Sönmez, and M. Utku Ünver. Multi-donor organ exchange, 2017. Working paper.

[55] Gabriele Farina, John P. Dickerson, and Tuomas Sandholm. Operation frames and clubs in kidney exchange. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI'17, pages 199–205. AAAI Press, 2017.

[56] J Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460(7256):685–686, 2009.

[57] J Doyne Farmer, Paolo Patelli, and Ilija I Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences of the United States of America*, 102(6):2254–2259, 2005.

[58] Daniel Friedman. The double auction market institution: A survey. 01 1993.

[59] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[60] Jacob Gardner, Matt Kusner, Kilian Q. Weinberger, John Cunningham, and Zhixiang Xu. Bayesian optimization with inequality constraints. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 937–945, 2014.

[61] SE Gentry, DL Segev, Mary Simmerling, and RA Montgomery. Expanding kidney paired donation through participation by compatible pairs. *American Journal of Transplantation*, 7(10):2361–2370, 2007.

[62] Matthew Gentzkow and Emir Kamenica. Costly persuasion. *American Economic Review*, 104(5):457–62, May 2014.

[63] Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Working paper*, 2016.

[64] Ronald J. Gilson and Alan Schwartz. Understanding macs: Moral hazard in acquisitions. *The Journal of Law, Economics, and Organization*, 21(2):330–358, 2005.

[65] Petra Glander, Klemens Budde, Danilo Schmidt, T. Florian Fuller, Markus Giessing, Hans-Hellmut Neumayer, and Lutz Liefeldt. The ?blood group o problem? in kidney transplantation?time to change? *Nephrology Dialysis Transplantation*, 25(6):1998–2004, 2010.

[66] Lawrence R Glosten and Paul R Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

[67] Jacob K. Goeree and Theo Offerman. Competitive bidding in auctions with private and common values. *The Economic Journal*, 113(489):598–613, 2003.

[68] Mingyu Guo and Argyrios Deligkas. Revenue maximization via hiding item attributes. In *Proc. IJCAI*, pages 157–163, 2013.

[69] Chen Hajaj, John P. Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 921–928, 2015.

[70] Chen Hajaj and David Sarne. Strategic information platforms: Selective disclosure and the price of "free". In *Proc. ACM EC*, pages 839–856, 2014.

[71] Jonathan V. Hall, John J. Horton, and Daniel T. Knoepfle. Labor market equilibriation: Evidence from Uber, 2017. Working paper Revision requested from American Economic Review.

[72] Sundaram Hariharan, Christopher P Johnson, Barbara A Bresnahan, Sarah E Taranto, Matthew J McIntosh, and Donald Stablein. Improved graft survival after renal transplantation in the united states, 1988 to 1996. *New England Journal of Medicine*, 342(9):605–612, 2000.

[73] Donald B. Hausch. An asymmetric common-value auction model. *The RAND Journal of Economics*, 18(4):611–621, 1987.

[74] Philip J Held, F McCormick, A Ojo, and John P Roberts. A cost-benefit analysis of government compensation of kidney donors. *American Journal of Transplantation*, 16(3):877–885, 2016.

[75] Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. In *Proceedings of ACM EC*, pages 359–376, 2014.

[76] Patrick R Jordan, Yevgeniy Vorobeychik, and Michael P Wellman. Searching for approximate equilibria in empirical games. In *Proceedings of AAMAS*, pages 1063–1070, 2008.

[77] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615, 2011.

[78] Edward H Kaplan. *Managing the demand for public housing.* PhD thesis, Massachusetts Institute of Technology, 1984.

[79] Edward H Kaplan. Tenant assignment models. *Operations Research*, 34(6):832–843, 1986.

[80] Andrei A Kirilenko, Albert S Kyle, Mehrdad Samadi, and Tugkan Tuzun. The flash crash: The impact of high frequency trading on an electronic market. *working paper*, 2014.

[81] Franz Knecht and Vera Calenbuhr. Using capital transaction due diligence to demonstrate csr assessment in practice. *Corporate Governance: The international journal of business in society*, 7(4):423–433, 2007.

[82] Scott Duke Kominers, Alexander Teytelboym, and Vincent P Crawford. An invitation to market design. *Oxford Review of Economic Policy*, 33(4):541–571, 2017.

[83] Kimberly D Krawiec and Michael A Rees. Reverse transplant tourism. *Law & Contemporary Problems*, pages 145–173, 2014.

[84] Vijay Krishna. *Auction Theory.* Academic Press, 2002.

[85] Albert S Kyle. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335, 1985.

[86] Kate Larson and Tuomas Sandholm. Bargaining with limited computation: Deliberation equilibrium. *Artificial Intelligence*, 132(2):183 – 217, 2001.

[87] Kate Larson and Tuomas Sandholm. Costly valuation computation in auctions. In *Proc. TARK*, pages 169–182, 2001.

[88] Kate Larson and Tuomas Sandholm. Mechanism design and deliberative agents. In *Proc. AAMAS*, pages 650–656, 2005.

[89] Blake LeBaron, W Brian Arthur, and Richard Palmer. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23(9):1487–1516, 1999.

[90] Robin S Lee and Michael Schwarz. Interviewing in two-sided matching markets. Technical report, National Bureau of Economic Research, 2009.

[91] Robin S. Lee and Michael Schwarz. Interviewing in two-sided matching markets. *The RAND Journal of Economics*, 48(3):835–855, 2017.

[92] Jacob D. Leshno. Dynamic matching in overloaded waiting lists, 2015. Working paper.

[93] Michael Lewis. *Flash Boys: A Wall Street Revolt*. WW Norton & Company, 2014.

[94] Jian Li, Yicheng Liu, Lingxiao Huang, and Pingzhong Tang. Egalitarian pairwise kidney exchange: Fast algorithms via linear programming and parametric flow. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 445–452, 2014.

[95] Zhuoshu Li and Sanmay Das. An agent-based model of competition between financial exchanges: Can frequent call mechanisms drive trade away from CDAs? In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 50–58, 2016.

[96] Zhuoshu Li and Sanmay Das. On the design of revenue-enhancing signal structures in common- and interdependent-value second price auctions. In *IJCAI Workshop on Algorithmic Game Theory*, 2017.

[97] Zhuoshu Li, Sanmay Das, Sofia Carrillo, and Jason Wellen. Estimating the benefits of incorporating compatible pairs in kidney exchange. Working paper. Presented at the AAMAS-IJCAI Workshop on Agents and Incentives in Artificial Intelligence, 2018.

[98] Zhuoshu Li, Neal Gupta, Sanmay Das, and John P. Dickerson. Equilibrium behavior in competing dynamic matching markets. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 389–395. International Joint Conferences on Artificial Intelligence Organization, 7 2018.

[99] Steven A Lippman and John W Mamer. Exploding offers. *Decision Analysis*, 9(1):6–21, 2012.

[100] Thaminda Liyanage, Toshiharu Ninomiya, Vivekanand Jha, Bruce Neal, Halle Marie Patrice, Ikechi Okpechi, Ming-hui Zhao, Jicheng Lv, Amit X Garg, John Knight, et al. Worldwide access to treatment for end-stage kidney disease: A systematic review. *The Lancet*, 385(9981):1975–1982, 2015.

[101] Tyler Lu and Craig Boutilier. Learning Mallows models with pairwise preferences. In *Proceedings of ICML*, pages 145–152, 2011.

[102] James B MacQueen. Optimal policies for a class of search and evaluation problems. *Management Science*, 10(4):746–759, 1964.

[103] Colin L Mallows. Non-null ranking models. I. *Biometrika*, 44(1/2):114–130, 1957.

[104] Allan B Massie, Joseph Leanza, LM Fahmy, EKH Chow, Niraj M Desai, X Luo, EA King, MG Bowring, and DL Segev. A risk index for living donor kidney transplantation. *American Journal of Transplantation*, 16(7):2077–2084, 2016.

[105] Nicholas Mattei, Abdallah Saffidine, and Toby Walsh. Mechanisms for online organ matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI'17, pages 345–351, 2017.

[106] Albert J Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740, 2013.

[107] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, pages 1089–1122, 1982.

[108] Paul R Milgrom and Robert J Weber. The value of information in a sealed-bid auction. *Journal of Mathematical Economics*, 10(1):105–114, 1982.

[109] Dov Monderer and Lloyd S Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[110] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic Game Theory*. Cambridge University Press Cambridge, 2007.

[111] Jinzhong Niu, Kai Cai, Simon Parsons, Peter McBurney, and Enrico H. Gerding. What the 2007 TAC market design game tells us about effective auction mechanisms. *Autonomous Agents and Multi-Agent Systems*, 21(2):172–203, 2010.

[112] Cathy O'Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books, 2017.

[113] Michael Ostrovsky and Michael Schwarz. Information Disclosure and Unraveling in Matching Markets. *American Economic Journal: Microeconomics*, 2(2):34–63, May 2010.

[114] Abraham Othman. Zero-intelligence agents in prediction markets. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 879–886, 2008.

[115] Imon Palit, Steve Phelps, and Wing Lon Ng. Can a zero-intelligence plus model explain the stylized facts of financial time series data? In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 653–660. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

[116] Sérgio O Parreiras. Affiliated common value auctions with differential information: the two bidder case. *Contributions in Theoretical Economics*, 6(1):1–19, 2006.

[117] Tao Qin, Xiubo Geng, and Tie-Yan Liu. A new probabilistic model for rank aggregation. In *NIPS*, pages 1948–1956, 2010.

[118] Zinovi Rabinovich, Albert Xin Jiang, Manish Jain, and Haifeng Xu. Information disclosure as a means to security. In *Proc. AAMAS*, pages 645–653, 2015.

[119] Panduranga S Rao, Douglas E Schaubel, Mary K Guidinger, Kenneth A Andreoni, Robert A Wolfe, Robert M Merion, Friedrich K Port, and Randall S Sung. A comprehensive risk quantification score for deceased donor kidneys: The kidney donor risk index. *Transplantation*, 88(2):231–236, 2009.

[120] F. T. Rapaport. The case for a living emotionally related international kidney donor exchange registry. *Transplantation Proceedings*, 18:5–9, 1986.

[121] Baharak Rastegari, Anne Condon, Nicole Immorlica, and Kevin Leyton-Brown. Two-sided matching with partial information. In *Proceedings of ACM EC*, pages 733–750, 2013.

[122] Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, 2010.

[123] Jeffrey J Reuer, Tony W Tong, and Cheng-Wei Wu. A signaling theory of acquisition premiums: Evidence from ipo targets. *Academy of Management Journal*, 55(3):667–683, 2012.

[124] Alvin Roth, Tayfun Sönmez, and Utku Ünver. Efficient kidney exchange: Coincidence of wants in a market with compatibility-based preferences. *American Economic Review*, 97:828–851, 2007.

[125] Alvin E Roth. Two-sided matching with incomplete information about others' preferences. *Games and Economic Behavior*, 1(2):191–209, 1989.

[126] Alvin E Roth. *Who Gets What—and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt, 2015.

[127] Susan L Saidman, Alvin E Roth, Tayfun Sönmez, M Utku Ünver, and Francis L Delmonico. Increasing the opportunity of live kidney donation by matching for two- and three-way exchanges. *Transplantation*, 81(5):773–782, 2006.

[128] Weiran Shen, Pingzhong Tang, and Yulong Zeng. A closed-form characterization of buyer signaling schemes in monopoly pricing. In *AAMAS*, 2018.

[129] Vasiliki Skreta. On the informed seller problem: optimal information disclosure. *Review of Economic Design*, 15(1):1–36, 2011.

[130] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2951–2959, 2012.

[131] Tayfun Sönmez and Utku Unver. Market design for living-donor organ exchanges: An economic policy perspective. Technical report, Boston College Department of Economics, 2017.

[132] Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973.

[133] Xuanming Su and Stefanos A. Zenios. Patient choice in kidney allocation: A sequential stochastic assignment model. *Operations Research*, 53:443–455, 2005.

[134] Vasilis Syrgkanis, David Kempe, and Eva Tardos. Information asymmetries in common-value auctions with discrete signals. In *Proc. ACM EC*, pages 303–303, 2015.

[135] David Robert Martin Thompson and Kevin Leyton-Brown. Valuation uncertainty and imperfect introspection in second-price auctions. In *Proc. AAAI*, pages 148–153, 2007.

[136] David Robert Martin Thompson and Kevin Leyton-Brown. Dominant-strategy auction design for agents with uncertain, private values. In *Proc. AAAI*, pages 745–750, 2011.

[137] George W Torrance and David Feeny. Utilities and quality-adjusted life years. *International Journal of Technology Assessment in Healthcare*, 5(4):559–575, 1989.

[138] Panos Toulis and David C. Parkes. Design and analysis of multi-hospital kidney exchange mechanisms using random graphs. *Games and Economic Behavior*, 91(0):360–382, 2015.

[139] Utku Ünver. Dynamic kidney exchange. *Review of Economic Studies*, 77(1):372–414, 2010.

[140] Yolanda Bravo Vergel and Mark Sculpher. Quality-adjusted life years. *Practical Neurology*, 8(3):175–182, 2008.

[141] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[142] Elaine Wah, Dylan R Hurd, and Michael P Wellman. Strategic market choice: Frequent call markets vs. continuous double auctions for fast and slow traders. In *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*, 2015. To appear.

[143] Elaine Wah and Michael P Wellman. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pages 855–872. ACM, 2013.

[144] Elaine Wah and Michael P Wellman. Welfare effects of market making in continuous double auctions. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 57–66. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[145] Thomas J Walsh, István Szita, Carlos Diuk, and Michael L Littman. Exploring compact reinforcement-learning representations with linear regression. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 591–598. AUAI Press, 2009.

[146] Jörgen Weibull. *Evolutionary Game Theory*. MIT Press, 1995.

[147] Michael P Wellman. Methods for empirical game-theoretic analysis. In *Proceedings of AAAI*, volume 21, pages 1552–1555, 2006.

[148] Robert B. Wilson. Competitive bidding with asymmetric information. *Management Science*, 13(11):816–820, 1967.

[149] Haifeng Xu, Rupert Freeman, Vincent Conitzer, Shaddin Dughmi, and Milind Tambe. Signaling in bayesian stackelberg games. In *AAMAS*, 2016.

[150] S. A. Zenios. Optimal control of a paired-kidney exchange program. *Management Science*, 48(3):328–342, 2002.

Vita

Zhuoshu Li

Degrees	B.S., Automation and Information Science, Beihang University, July 2013 Ph.D., Computer Science, Washington University in St. Louis, December 2018

Publications

Zhuoshu Li, Sanmay Das, Sofia Carrillo and Jason R. Wellen (2018). Estimating the Benefits of Incorporating Compatible Pairs in Kidney Exchange. Working paper. Presented at the AAMAS-IJCAI Workshop on Agents and Incentives in Artificial Intelligence.

Zhuoshu Li, Neal Gupta, Sanmay Das and John P. Dickerson (2018). Equilibrium Behavior in Competing Dynamic Matching Markets. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, pp. 389–395.

Zhuoshu Li, Zhitang Chen, Pascal Poupart, Sanmay Das, and Yanhui Geng (2018). Faster Policy Adaptation in Environments with Exogeneity: A State Augmentation Approach. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-18)*, pp. 1035–1043.

Zhuoshu Li and Sanmay Das (2017). On the Design of Revenue-Enhancing Signal Structures in Common- and Interdependent-Value Second Price Auctions. Presented at IJCAI Algorithmic Game Theory Workshop.

Zhuoshu Li and Sanmay Das (2016). An Agent-Based Model of Competition Between Financial Exchanges: Can Frequent Call Mechanisms Drive Trade Away from CDAs? In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-16)*, pp. 50–58.

Sanmay Das, John P. Dickerson, Zhuoshu Li, and Tuomas Sandholm (2015). Competing Dynamic Matching Markets. In *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA-15)*, pp. 19.

Sanmay Das and Zhuoshu Li (2014). The Role of Common and Private Signals in Two-Sided Matching with Interviews. In *Web and Internet Economics – 10th International Conference (WINE-14)*, pp. 492–497.

December 2018