

Nikhil Agarwal^{a,b} and Eric Budish^{c,b,*}^aMIT Department of Economics, Cambridge, MA, United States^bNBER, Cambridge, MA, United States^cUniversity of Chicago Booth School of Business, Chicago, IL, United States^{*}Corresponding author: e-mail address: eric.budish@chicagobooth.edu**CONTENTS**

1	Introduction	2
2	Theoretical framework	6
2.1	Taxonomy of market design problems	8
2.1.1	Matching or allocation?	8
2.1.2	Transferable utility or non-transferable utility?	9
2.1.3	Single-unit vs. multi-unit demand?	9
2.1.4	Endowments?	10
2.1.5	Clarification: are schools in school-choice agents or objects?	10
2.2	Canonical market design problems	10
2.3	Canonical market-design mechanisms	14
2.3.1	Gale-Shapley deferred acceptance	14
2.3.2	Immediate acceptance (“Boston mechanism”)	18
2.3.3	Random serial dictatorship	20
2.3.4	Top trading cycles	21
2.3.5	Hylland and Zeckhauser pseudomarket	22
2.3.6	Draft approaches to multi-unit assignment	23
2.3.7	Competitive equilibrium approaches to multi-unit assignment	24
3	Empirical frameworks and applications	25
3.1	Non-transferable utility models.....	26
3.1.1	Random utility model	27
3.1.2	Analysis with data from assignment mechanisms	28
3.1.3	Analysis with data on final outcomes	36
	Continuum many-to-one matching	36
	One-to-one or few-to-one matching	41
3.2	Transferable utility models	46

☆ We thank Eduardo Azevedo, Jacob Leshno, Jeremy Fox, Parag Pathak, Al Roth, Dan Waldinger, Alex Wolitzky, and two anonymous reviewers for helpful comments. Agarwal is grateful to Paulo Somaini for collaboration on surveys of related topics that this article draws on. With permission of the respective publishers and authors, this chapter reuses portions from Agarwal and Somaini (2020; 2021b) and adapts certain figures from Agarwal and Somaini (2018).

3.2.1	<i>Models with separable unobserved heterogeneity</i>	46
3.2.2	<i>Semi-parametric approaches</i>	48
3.3	<i>Extensions</i>	50
4	Designing markets	52
4.1	<i>Diagnosing market failures</i>	52
4.2	<i>Evaluating and comparing designs</i>	58
4.3	<i>Proposing new market designs</i>	62
5	Conclusion	68
	References	70

Abstract

This Handbook chapter seeks to introduce students and researchers of industrial organization (IO) to the field of market design. We emphasize two important points of connection between the IO and market design fields: a focus on market failures — both understanding sources of market failure and analyzing how to fix them — and an appreciation of institutional detail.

Section 2 reviews theory, focusing on introducing the theory of matching and assignment mechanisms to a broad audience. It introduces a novel “taxonomy” of market design problems, covers the key mechanisms and their properties, and emphasizes several points of connection to traditional economic theory involving prices and competitive equilibrium.

Section 3 reviews structural empirical methods that build on this theory. We describe how to estimate a workhorse random utility model under various data environments, ranging from data on reported preference data such as rank-order lists to data only on observed matches. These methods enable a quantification of trade-offs in designing markets and the effects of new market designs.

Section 4 discusses a wide variety of applications. We organize this discussion into three broad aims of market design research: (i) diagnosing market failures; (ii) evaluating and comparing various market designs; (iii) proposing new, improved designs. A point of emphasis is that theoretical and empirical analysis have been highly complementary in this research.

Keywords

Market power, Microeconomic theory, Industrial organization, Competitive equilibrium approach, Market design

1 Introduction

Textbook models envision markets as abstract institutions that clear supply and demand. Real markets have specific designs and market clearing rules. These features affect market participants and their allocations in various ways — they determine the actions an agent can take, the incentives for taking those actions, the information

environment, the interactions between agents' actions, and, ultimately, the final allocation. Well-designed markets have rules that coordinate and incentivize behavior in ways that lead to desirable outcomes. But it is not a given that all markets have good design. The Market Design field studies these rules in order to understand their implications, to identify potential market failures, and to remedy them by designing better institutions.

This focus, on identifying and remedying market failures – with careful attention to specific market institutions – is closely shared with Industrial Organization (IO). Like market designers, IO economists are also interested in analyzing the structure and behavior of agents in a market, and their impact on market performance. Researchers in both fields are directly interested in shaping regulation and policy. Moreover, both fields combine theoretical and empirical approaches in their analysis. In fact, there has been recent convergence in the empirical approaches, with market design drawing from methods used in empirical IO.

Rather than a difference in objectives or methods, the fields mostly differ in their relative emphasis on sources of market failures and the reforms that are recommended. Whereas imperfect competition is the most common cause of market failure that is studied in IO, a market designer more commonly focuses on the market mechanism itself. A poorly designed market mechanism can directly result in inefficient allocations or provide incentives that undermine efficiency, irrespective of the presence of market power. The remedies recommended by the two fields correspondingly address either the source of market power or the rules of the market that are the cause of failure.

A simple but obvious point is that a holistic approach to studying markets would consider both types of problems. Diagnosing the cause of a market failure requires investigating both the design and the nature of competition. Remedying the market failure may also require considering the two and their interaction. Moreover, as we will see below, a market's design can influence the exercise of market power and vice-versa. Thus, studying only the rules of the market or only market power issues may miss the complete picture.

The first of at least three reasons for considering both types of problems is to correctly trace the cause of a potentially undesirable outcome. Consider the well-known example of the medical residency match, which assigns newly minted medical school graduates to hospital training positions. This centralized labor market uses rank-ordered preference lists submitted by both sides as inputs into an algorithm that determines final matches (Roth and Peranson, 1999). A lawsuit argued that the coordination afforded by this market clearing mechanism allowed residency programs to collude and suppress salaries to about \$40,000 a year (Jung et al. versus Association of American Medical Colleges et al., 2002), much lower than salaries of other medical professionals performing similar work. The lawsuit was dismissed after Congress enacted an exception to anti-trust law for the medical match, but did not resolve the cause of low salaries.

While the rules of the medical residency market are one potential reason for wage suppression, so are other traditional sources of market power. Bulow and Levin

(2006) indicted the market's rules – specifically an implicit prohibition on personalized wage bargaining because salaries are set before the match – as the cause of low salaries. The market design solution to this problem, suggested in Crawford (2008), is to modify the algorithm and to allow residents to submit salary contingent preferences. The alternative explanation for low salaries is that residency programs are exercising monopsony power. Agarwal (2015) argues that accreditation requirements limit the number and size of residency programs and, combined with heterogeneity in program quality, can depress salaries. Estimates of residents' willingness to pay for training at high-quality programs indicate substantial wage markdowns in competitive equilibrium with restricted entry, suggesting that the market's design may not be the problem.

The second reason to carefully analyze both market design and market power is that they interact. In fact, the exercise of market power is mediated through the market's design. A salient example is supply reduction in the context of the most recently concluded FCC spectrum auctions. The purpose of the auction was to repurpose spectrum allocated for television broadcasting to broadband internet. This auction required the FCC to simultaneously buy and sell spectrum while respecting a complex set of engineering constraints on the feasible transactions. As is now well documented, small private equity firms purchased multiple small, low-value television stations and withheld some from the auction (Doraszelski et al., 2019; Ausubel et al., 2017). It has been argued that the motivation and effect of this strategy is to increase the selling prices for the other stations. That is, private equity firms acquired and exploited market power in a manner that was sensitive to the details of the market's design.

Finally, market power itself can affect a market's design because influential market participants may not have the incentive to adopt a good design. An example is the design of financial markets. Budish et al. (2015) showed that the predominant market design used by financial exchanges around the world has a design flaw that gives rise to an arms race for trading speed and harms market liquidity. The effects are quantitatively important: in one setting, over 20% of trading volume takes place in trading races, and trading in races constitutes 33% of the literature's standard measure of the market's cost of liquidity (Aquilina et al., 2021). Yet, to date, the market design reform suggested in Budish et al. (2015) remains essentially unadopted. Budish et al. (2020) suggest that the reason why may be that exchanges profit from the source of inefficiency: exchanges earn significant revenue from selling speed (e.g., fast connections to their venues), and this source of revenue would dry up under a reform that addressed the arms race. Since it may not be in the interest of influential market participants (the exchanges) to adopt an efficient design, the adoption of a better design may require intervention from a government regulator.

Another point of similarity between market design and IO is that research in both fields is often motivated not only by understanding the sources of market failures, but by understanding how to *fix* them. Within IO, this objective is central in the study of competition policy (Whinston, 2006) and the regulation of natural monopolies (Laffont and Tirole, 1993), to give two prominent examples. These areas consider the

circumstances in which particular policy tools should be deployed in order to mitigate the harms caused by market power. Similarly, research in market design often seeks to understand when and how a market's design fails, and the solutions to those design flaws.

One common type of failure that needs fixing occurs when markets simply fail to aggregate information about preferences when matching demand and supply. For example, Abdulkadiroglu et al. (2005) describe a 2003 reform of New York City's high school match, which moved from a decentralized waitlist-offer process that left many students unmatched to an economist-advised design. The reform resulted in many more students being placed to more preferred schools, thereby reducing exits from the school system and increase enrollment (Abdulkadiroglu et al., 2017b). Thus, a well-designed system can improve outcomes by improving co-ordination. Many additional examples along these lines – where simply coordinating demand and supply on a centralized platform creates substantial value – come from internet marketplaces (see Levin, 2013; Einav et al., 2016).

Markets can also fail for a variety of more subtle reasons. Poorly designed markets may have rules that result in suboptimal participation incentives; rent-seeking behavior; exploitable frictions that distort allocations; or incentives for strategic preference reporting that results in avoidable inefficiency. These four failures have been documented in the contexts of, respectively, kidney exchange markets (Agarwal et al., 2019); high-frequency trading of financial assets (Budish et al., 2015); the market for clinical psychologists (Roth and Xing, 1997); and course allocation problems (Budish and Cantillon, 2012). Each of these papers also suggest market design solutions.

The goal of this chapter is to introduce ideas from the market design field to a graduate student of IO. One challenge in a comprehensive survey of the field is that markets vary in a number of ways. They differ in the types of agents that participate in the market, the types of transactions facilitated by the market (buying/selling or forming partnerships), the constraints on these transactions (e.g. timing, transfers), and the relevant frictions. In the interest of maintaining a well-defined focus, our methodological discussion will emphasize models of matching and assignment (hereafter, matching for short): students to schools, workers to firms, kidney patients to organs, courses to students, food to food banks.

This chapter begins by reviewing the core tools in the analysis of matching markets before describing various ways and domains in which the principles of market design have been applied. Sections 2 and 3 review the economic theory and empirical methods respectively. The theory underlying the study of matching market institutions is based on the language of mechanism design. In these models, agents in the market are endowed with preferences over their match partners. The rules of the market determine how these preferences can be expressed through messages and how these messages translate to allocations. Our approach to this body of theory introduces a novel “taxonomy” of market design problems, and emphasizes several points of connection to traditional economic theory involving prices and competitive equi-

librium. The hope is that this approach can make this body of theory easily accessible to researchers in IO.

Since preferences are a core primitive in this theory, the empirical methods draw on a large literature estimating random utility models of consumer preferences (Block and Marshak, 1960; Berry et al., 1995). The tools in this literature use revealed preference arguments and data from a matching market on the realized matches or reported preferences. Much like modern empirical IO, they rely on a market's specific features and an appropriately chosen theoretical framework to formulate an empirical approach. The typical goal is to use these estimates to analyze the design of alternative market institutions or the incentives to exercise market power and strategically distort decisions. One of our pedagogical aims in this methods section is to emphasize how tools familiar to IO economists have been adapted and applied in market design settings, and to give a sense of some uncharted terrain.

Although our methodological discussion is purposefully narrow, many market design principles are more general. In fact, some of the most important market design applications involve auctions, such as the FCC spectrum auctions, double auctions for financial assets, and procurement auctions. These institutions also constitute a specific set of rules for allocating goods and services to market participants. We refer the interested reader to Chapter 11 in this handbook, which includes a deeper discussion of auction markets.

Section 4 discusses applications of the market design toolkit. Here, we expand our focus from narrowly discussing matching markets to other examples highlighting the main types of challenges that market design research confronts. We classify the goals of research in market design into three types: diagnosing market failures, evaluating and comparing various market designs, and proposing new, improved designs. These goals are interrelated and can be achieved via either theoretical or empirical analysis. Indeed, theory and empirical analysis are often complementary in this research. While these goals are probably not exhaustive, our hope is that the classification is useful for understanding the contributions of papers in the area.

2 Theoretical framework

The goal of this section is to introduce the reader to the theoretical ideas they are most likely to encounter in the literature at the intersection of market design and IO. As noted, it focuses on the part of the market design literature that studies matching markets, as auction markets and platform markets are covered elsewhere in this Handbook.¹

We organize the section along three dimensions. In Section 2.1, we provide some notation and a taxonomy of types of market design environments. In Section 2.2, we then formally describe many of the canonical market design environments. In Sec-

¹ We encourage readers to consult Chapter 11 (this volume) and Chapter 7 (Volume 4).

tion 2.3, we provide an overview of some of the canonical market design mechanisms. This section also defines some of the key design objectives and constraints along the way, some of which are meaningful across many different kinds of problems and some of which are tailored to specific problem environments.

Throughout, our goal is to provide a kind of theoretical “orientation” to the reader at a level appropriate for a graduate student in IO.² When encountering a specific paper on a specific market design problem — and in market design, the specificity of the environment is often what necessitates exciting new research — it can sometimes be difficult to have a sense of the general economics principles involved. For example, we guess that many readers who first encounter the school choice literature will not understand that “no justified envy” — which is a property specific to school-choice models — is conceptually just a special case of Gale-Shapley stability, which in turn is closely related to Pareto efficiency. Nor that Pareto efficiency, which is, justifiably, a central design objective in many allocation environments, can in fact be quite a poor proxy for social welfare in multi-unit allocation environments without transfers.

In addition to providing such an orientation, we try to emphasize two methodological points throughout this section. First, while matching theory often looks quite different from mechanism design theory in the Myerson-ian tradition, at a high level the goals are very similar — produce desirable allocations given the constraints of the environment, including both technological constraints and incentive constraints. One reason for the difference in appearance is that market design problems often involve multiple goals that cannot obviously be collapsed into a single objective function, such as considerations of both efficiency and fairness. A second reason is that the tractability of the Myerson-ian approach often relies on a numeraire good, and many matching problems lack one. Thus, rather than collapse multiple distinct objectives into a single objective function, which is intractable to maximize anyways, the researcher instead tries to identify a mechanism which performs well along each of the dimensions of interest, e.g., by satisfying attractive properties of both efficiency and fairness.³ Again, this might look different from mechanism design, but it is ultimately an effort in maximizing objectives subject to constraints.

Second, and relatedly, while much of matching theory at first glance appears to have little to do with the traditional microeconomic theory that involves prices, many matching mechanisms in fact have a price-like structure. That is, there are prices or price-like statistics, that can be used to understand participants’ choice sets, and ultimately their allocations. Often, the prices or price-like statistics are personalized — for example, in the matching of students to colleges, the “price” of a particular college for a particular student might depend on how highly the college ranks the student (e.g., based on test scores) (Azevedo and Leshno, 2016). In some mechanisms,

² Roth and Sotomayor (1990) is the classic textbook treatment of matching theory and remains an essential reference, even though it predates and hence does not discuss many of the subsequent applications of this theory.

³ See Budish (2012), from which this section draws, for discussion of the relationship between the matching theory approach and Myerson-ian mechanism design.

the prices or price-like statistics are explicit in the design, even though there is not money per se in the allocation environment (Hylland and Zeckhauser, 1979; Budish, 2011; Prendergast, 2017). More recently, empirical researchers have started to utilize the price-like structure underlying some matching markets to import tools from other parts of economics that study discrete choice from choice sets (see Section 3).

We will try to bring out these connections to mechanism design and prices throughout this section, as we suspect it will increase the accessibility of this exciting area of theory for IO readers.

2.1 Taxonomy of market design problems

Throughout this section we will work with the following unified notational conventions:

- Sets of agents or objects. There are two sets, \mathcal{I} and \mathcal{J} , with generic elements i and j . The set \mathcal{I} is always a set of agents (workers, students, consumers, etc.), whereas the set \mathcal{J} will sometimes be another set of agents and will sometimes be a set of objects.
- Quantities. We will let q_i denote the quantity of agents of type i , and similarly q_j denote the quantity of agents/objects of type j . If an agent/object has multiple units, all units are identical.
- Numeraire good. In most of the problems we will discuss in this section there are no monetary transfers. When there are monetary transfers, we will use the notation t_i to denote the transfer to agent i (with negative transfers $t_i < 0$ denoting payments made by the agent).
- Preferences. Agents have preferences over outcomes, which could include who or what they are matched with on the other side of the market, and/or monetary transfers. We will mostly use the notation u_i to describe agent i 's cardinal utility from a given outcome, with \succsim_i the associated ordinal preference relation. We will try to be clear throughout whether a particular mechanism relies on cardinal preference information (and if so of what form) or just ordinal preference information.

We find it helpful to categorize market design settings into an informal taxonomy. The dimensions of the taxonomy are as follows:

2.1.1 Matching or allocation?

In a matching problem, \mathcal{I} and \mathcal{J} are two distinct sets of agents: firms and workers, customers and suppliers, etc. Agents have preferences that depend on who they match with on the other side of the market.

In an allocation problem \mathcal{I} is a set of agents but now \mathcal{J} is a set of objects. Agents have preferences over what objects they receive, but objects do not have preferences over who they are assigned to.

2.1.2 Transferable utility or non-transferable utility?

Transferable utility settings have a numeraire good, whereas non-transferable utility settings do not.

We emphasize that “non-transferable utility” includes but is not limited to environments “without money”. For example, in the National Residency Matching Program, the match of doctors and hospitals is at fixed wages, i.e., wages that are exogenous to the matching process. This means that the market design does not allow for transfers of utility between the agents using a numeraire good, but there is of course “money” paid by hospitals to the doctors they match with.

We will adopt the terminology that an allocation problem with non-transferable utility is an “assignment problem”,⁴ whereas an allocation problem with transferable utility is an “auction”.

2.1.3 Single-unit vs. multi-unit demand?

The next dimension in the problem taxonomy is whether agents demand just a single unit or demand multiple units.

Cases in which at least some of the agents have single-unit demand include: school choice (students can only attend a single school); workers in labor markets (doctors match to a single residency); kidney exchange (patients require a single kidney); public housing allocation (households require a single apartment).

In matching settings, it is common for one side of the market to have single-unit demand while the other side requires multiple units. For example, while students require a single school, schools admit many students; workers seek a single job, but firms hire multiple workers. This is known as the “many-to-one” matching problem.

Single unit demand occupies a special place in the literature because many of the canonical mechanisms, described below in Section 2.2, are developed for either (i) single-unit demand, or (ii) multi-unit demands that can be treated as, in effect, multiple separable single-unit demands. For example, the many-to-one matching model of Gale and Shapley (1962) treats each “position” at a residency program or school as, in effect, its own unit-demand entity. This approach works under the assumption that the residency program or school’s preferences are, roughly, additive separable (Roth, 1984 calls these “responsive” preferences, in that preferences over bundles “respond” to preferences over individual objects). The insight is that each position at a firm can be treated as its own party to the match, with all positions at the firm having the same preferences over workers, allowing the machinery of one-to-one matching to carry through.

Analysis of market design problems in which agents have multi-unit demands that cannot be treated as multiple single-unit demands is much more common in the

⁴ This usage of the term “assignment” to mean allocation without monetary transfers is common but by no means universal. Shapley and Shubik (1971) describe versions of what they refer to as the “assignment problem” both with and without monetary transfers. Demange et al. (1986) use the term “assignment problem” to describe what most would now call a multi-object auction.

auctions part of market design, where combinatorial auctions have been a central focus over the past several decades, than in the matching or assignment parts of market design. Exceptions include Budish (2011) and related works on course allocation and combinatorial assignment, and Prendergast (2017)'s work on the allocation of food to food banks. Not coincidentally, these papers develop mechanisms in which prices play a central role; once agents are making tradeoffs across multiple types of heterogeneous objects, prices are helpful for setting marginal rates of substitution and determining choice sets.

2.1.4 Endowments?

A last dimension to include in the taxonomy is whether or not agents have any endowed property rights. Many of the canonical market design problems start from a blank slate, but in some models and some applications, agents begin with endowments of various sorts. In the house allocation model of Shapley and Scarf (1974), the model begins with each agent owning their current house. In kidney exchange, each patient-donor pair begins endowed with its own kidney to donate (i.e., the donor's).

Additionally, in some allocation problems, agents' outside options can vary in important ways that affect market design.

2.1.5 Clarification: are schools in school-choice agents or objects?

In the school-choice literature, schools are sometimes modeled as agents with preferences over the other side of the market, but are sometimes modeled as a sort of hybrid between objects and agents. In this latter case, schools have what looks like a preference-ordering over students, but this ordering is interpreted as an administrative "priority" that does not count towards evaluation of economic efficiency or incentive compatibility (Abdulkadiroglu and Sonmez, 2003). The interpretation is that students' priorities at different schools are somewhat like property rights to attend them — for instance, a student who lives near to a particular school or whose sibling attends the school may have high priority to attend it — and the school administrator views it as desirable for a mechanism to respect those property rights, while not having preferences per se about what student goes to what school. Put differently, the mechanism design objective is student welfare, while school priorities place additional constraints on the problem beyond traditional technology and incentive constraints.

We will come back to this issue below, and discuss what are really three distinct versions of the school-choice problem, depending on the interpretation of schools' orderings over students.

2.2 Canonical market design problems

One-to-one matching without transfers

In the one-to-one matching problem without transfers, introduced by Gale and Shapley (1962): \mathcal{I} and \mathcal{J} are two distinct sets of agents; each agent has unit demand for a match on the other side of the market; and utility is non-transferable.

Formally, each agent $i \in \mathcal{I}$ has complete, transitive preferences over all potential matching partners in \mathcal{J} and the possibility of being unmatched, while each agent $j \in \mathcal{J}$ has complete, transitive preferences over all potential matching partners in \mathcal{I} and the possibility of being unmatched. That is, the domain of agent i 's utility function, $u_i(\cdot)$, is the set $\mathcal{J} \cup \emptyset$, while for j 's the domain is $\mathcal{I} \cup \emptyset$. In the Gale and Shapley algorithm (described below) agents report only their ordinal preferences over match partners. It is common in this literature to denote these ordinal preferences by a rank-ordered list ("ROL"). For example, the notation $\succsim_i: j_1, j_2, j_3, \emptyset, \dots$ would denote that agent i 's first choice match is j_1 , second choice is j_2 , third choice is j_3 , and then that they prefer going unmatched to all other potential matching partners.

A matching is a function $\mu: \mathcal{I} \cup \mathcal{J} \rightarrow \mathcal{I} \cup \mathcal{J}$ with the following properties. First, each agent is either matched to an agent on the other side or is unmatched: $\mu(i) \in \mathcal{J} \cup \emptyset$ for all $i \in \mathcal{I}$, $\mu(j) \in \mathcal{I} \cup \emptyset$ for all $j \in \mathcal{J}$. Second, no agent is matched to multiple partners: $\mu^{-1}(i) \in \mathcal{J} \cup \emptyset$ for all $i \in \mathcal{I}$, $\mu^{-1}(j) \in \mathcal{I} \cup \emptyset$ for all $j \in \mathcal{J}$.

Many-to-one matching without transfers

In the many-to-one matching problem without transfers, also introduced by Gale and Shapley, agents in \mathcal{I} are just like in the one-to-one case but the agents in \mathcal{J} now have capacity for, and demand for, multiple match partners. Formally, each $j \in \mathcal{J}$ has $q_j \in \mathbb{Z}^+$ positions and j 's preferences are defined over sets of up to j match partners.

In general, this could be captured with a utility function $u_j(\cdot)$ defined on subsets of \mathcal{I} with cardinality of up to q_j . However, the literature mostly works with preferences that are essentially additive-separable over the agents in \mathcal{I} . Formally, ordinal preferences over sets of match partners are said to be *responsive* to ordinal preferences over individual match partners if: for any set $S \subseteq \mathcal{I}$ with $|S| < q_j$, and any pair of agents i', i'' that are each not in set S , if $i' \succsim_j i''$, then $\{i'\} \cup S \succsim_j \{i''\} \cup S$ (Roth, 1985). In words, responsiveness imposes that the agent's preferences over bundles are consistent with preferences over individuals. Responsiveness is not quite as restrictive as assuming additive-separable utility (the latter implies the former), but it rules out many forms of substitutes and complements.

Matching with transfers

Becker (1974) introduced the problem of matching with transferable utility. We focus on describing the one-to-one version of the model for simplicity.

As in Gale and Shapley, \mathcal{I} and \mathcal{J} are two distinct sets of agents, and each has capacity for at most one match partner on the other side of the market. Utility can be perfectly transferred among partners to a match. For example, in worker-firm matching, utility can be transferred with money. Thus, worker i 's utility from matching with firm j and being paid a transfer t_{ij} is $u_{ij} + t_{ij}$, and if firm j values the worker's services at v_{ji} , then the firm's payoff is $v_{ji} - t_{ij}$.

Once utility is perfectly transferable within a match, the economically central description of utilities is the matching surplus created by a particular match. We will denote this by $\phi_{ij} = u_{ij} + v_{ji}$. This surplus ϕ_{ij} , which describes the total utility created if i and j match, can then be divided up between the two match partners in

any arbitrary fashion. We will denote by u_i the utility i creates if unmatched, and similarly v_j the utility j creates if unmatched. These match utilities, ϕ_{ij} , u_i , and v_j , each defined over the relevant sets, are primitives of the problem.

A matching is now a function μ like above that describes who matches with whom, and a transfer function $t : \mathcal{I} \cup \mathcal{J} \rightarrow \mathcal{R}$ that describes how the surplus in the match gets split.

Single-unit assignment

In the single-unit assignment problem, \mathcal{I} is a set of agents, \mathcal{J} is a set of objects, and agents have unit demand preferences over the objects. Typically, each object is modeled as being available in unit supply, and the number of objects is either equal to or greater than the number of agents, i.e., $|\mathcal{J}| \geq |\mathcal{I}|$.

Depending on the nature of the exercise, agents' preferences are sometimes modeled as ordinal and sometimes as cardinal. With unit demand, ordinal preferences are a rank-ordered list over the objects and possibly the empty allocation. Cardinal preferences can be represented by a utility function $u_i : \mathcal{J} \rightarrow \mathcal{R}_+$; since there are no transfers, it usually is appropriate to use a utility function that takes the von Neumann Morgenstern form.

House allocation

The house allocation problem, introduced by Shapley and Scarf (1974), is like single-unit assignment except that each individual is initially endowed with one of the objects. Our notational preference is to number the elements of the object set \mathcal{J} based on the corresponding agent who owns the house — agent i_1 initially is endowed with object j_1 , agent i_2 with object j_2 , etc. This then allows preferences to be defined over the objects, just as in single-unit assignment.

Another notational convention sometimes used in the literature, which is mathematically equivalent but potentially confusing, is to work with just a single set, \mathcal{I} , and understand that each element in the set describes both an agent and his endowed object. So, agent i_1 might have ordinal preferences $\succsim_{i_1} : i_2, i_1, i_3, \dots$ which means that i_1 's first choice object is the one that is endowed to i_2 , his second choice object is his current endowment, his third choice object is that endowed to i_3 , etc.

Kidney exchange

Kidney exchange is an important application of the house-allocation model. In this case, the agent-object pair corresponds to a patient and their associated donor. The patient, who is in need of a kidney, has preferences over potential donor kidneys (his own donor and other pairs' donors). The key theory references for kidney exchange are Roth et al. (2004), Roth et al. (2005), and Roth et al. (2007). We will return to kidney exchange as an application throughout Section 4.

School choice

In the school choice problem, introduced by Abdulkadiroglu and Sonmez (2003), \mathcal{I} is a set of students and \mathcal{J} is a set of schools. Students have ordinal preferences

over schools and schools have ordinal rankings over students. As discussed above in Section 2.1.5, the schools' ordinal rankings over students are sometimes interpreted as preferences and are sometimes interpreted as administrative priorities that are not welfare relevant. Sometimes this ranking is strict (e.g., based on test scores), but often it is coarse, such as ranking all students who live in the immediate vicinity of the school above students who do not, or ranking all students with a sibling at the school above students without.

If school priorities are indeed student property rights, this raises the further question of whether or not these property rights are tradable. Thus, depending on the interpretation of schools' rankings over students, there are really three distinct versions of the school-choice problem. First, in which agents are matching to agents, as in Gale and Shapley (1962). Second, in which agents are matching to objects, as in single-unit assignment, and school priorities place additional constraints on the allocation as non-tradable property rights. As discussed below, this case also leads to Gale and Shapley (1962)'s deferred acceptance algorithm being an attractive solution. Last, in which agents are matching to objects, and school priorities give agents tradable property rights — in this case, school choice is more closely related to house allocation (Shapley and Scarf, 1974), in that priorities function as tradable endowments.

Multi-unit and combinatorial assignment

The multi-unit and combinatorial assignment problems, studied in Budish and Cantillon (2012) and Budish (2011), are generalizations of the single-unit assignment problem in which agents have preferences over bundles of objects and objects are in multi-unit supply. The terminology “multi-unit assignment” is typically used to describe cases where agents' preferences over objects are additive separable or responsive, as in many-to-one matching as described above. The terminology “combinatorial” is typically used to describe cases where agents' preferences over objects are more general.

Formally, there is a capacity vector $q_1, \dots, q_{|\mathcal{J}|}$ that describes each objects' supply, each agent i has a set of feasible consumption bundles, denoted X_i , and each agents' preferences are represented by a utility function $u_i : X_i \rightarrow \mathbb{R}_+$. For example, in the course allocation setting, the set of feasible consumption bundles can encode constraints such as (i) each student takes at most one seat in each class; (ii) each student takes at most a certain number of classes overall; (iii) students cannot take courses that meet at the same time (scheduling constraints) or that violate curricular requirements (curricular constraints). In the problem COVAX faces for global allocation of Covid-19 vaccines across countries, the set of feasible consumption bundles might encode which vaccines have regulatory approval in which countries, or a countries' capacity to distribute a given quantity of vaccines before expiry (Castillo et al., 2021).

In the multi-unit assignment problem, each u_i is assumed to be additive-separable: for each agent i there exist item values $\alpha_{i1}, \dots, \alpha_{i|\mathcal{J}|}$ such that $u_i(x_i) = \sum_{j \in x_i} \alpha_{ij}$ for all $x_i \in X_i$. In the combinatorial assignment problem, no such assumption is imposed,

though additive-separability can be a useful starting point for describing preferences (Budish and Kessler, 2021).

An allocation is feasible if (i) each agent obtains a consumption bundle $x_i = (x_{i1}, \dots, x_{i|\mathcal{J}|})$, that is feasible for them, denoted $x_i \in X_i$, and (ii) the allocation satisfies the capacity constraints for each object, which we can write as $\sum_i x_{ij} \leq q_j$ for all $j \in \mathcal{J}$.

2.3 Canonical market-design mechanisms

2.3.1 Gale-Shapley deferred acceptance

We define the Gale-Shapley Deferred Acceptance algorithm for the case of one-to-one matching. We will let the set \mathcal{I} describe a set of workers, the set \mathcal{J} firms, and describe the worker-proposing version of the algorithm.

What preference data are reported to the algorithm: Each worker i reports a rank-ordered list of firms. Leaving a firm off the list denotes preferring to go unmatched to being matched with that firm.

Similarly, each firm j reports a rank-ordered list of workers. Leaving a worker off the list denotes preferring to go unmatched to being matched with that worker.

Round 1: In the first round of the algorithm, each worker i “proposes to” the highest-ranked firm on their preference list. If a firm receives one or more proposals in this round, the firm says “maybe” to the single proposal they prefer the most (assuming at least one is acceptable) and “reject” all others. Note that when we say “proposes to”, “maybe”, and “reject” we are trying to give anthropomorphic description to computer code. We trust the formal mathematical meaning is sufficiently clear.

Round k: In each subsequent round of the algorithm, any worker i who has been rejected in the previous round proposes to the highest-ranked firm on their preference list that they have not yet proposed to. If the worker has no more firms to propose to, the worker will go unmatched.

Firms that have one or more active proposals after this (including if there is a proposal held over from prior rounds) say maybe to the single proposal they prefer the most and reject all others. Note that a maybe from a previous round can become a rejection if the firm receives a proposal they like better.

Ending condition: The algorithm ends when either (i) there is a round with no rejections, or (ii) there is a round with no new proposals. If either of these conditions is satisfied, then all “maybes” become matches.

Key property: stability

A matching μ is said to be *stable* if no pair of agents who are not matched to each other in μ prefer to be matched to each other over their match in μ . Moreover, stability requires that no individual prefers to be unmatched to their match in μ . In either event, the pair or individual is said to “block” the matching μ .

Gale and Shapley's (1962) paper both introduced the concept of stable matching, and proved that their deferred acceptance algorithm produces a stable match (with respect to the reported preferences, incentives for which we turn to next).

Key property: strategy-proof or approximately strategy-proof

A mechanism is said to be *strategy-proof* (SP) if reporting truthfully is a dominant strategy.

The deferred acceptance algorithm is SP for the side of the market that makes proposals (Roth, 1982b; Dubins and Freedman, 1981). The rough intuition is seen by considering an agent who is on the proposing side who reports truthfully and is matched with her third most-preferred firm. Could she have done better, perhaps matched to her second-most preferred alternative, by misreporting this second choice as her first choice? No. The reason why is that, to have even reached the part of the algorithm where she proposed to her third choice, she must have earlier been rejected by her second choice firm. For her second choice firm to have rejected her, that firm in turn must have had a proposal from some other worker it preferred to her. If our protagonist ranked this firm higher, it's possible she would have initially been told maybe when otherwise she would have initially been told no, but either way, once the more-preferred worker comes along and proposes, our protagonist will get told no.

The deferred acceptance algorithm is in fact *not* strategy-proof on the receiving side of the market. There is a specific type of potentially profitable manipulation, known in the literature as a “truncation strategy” (Roth and Rothblum, 1999) in which a participant on the receiving side of the market reports their preferences in the honest order, but truncates the list, i.e., reporting some match partners as unacceptable who in fact are acceptable, just less preferred. This truncation can set off a chain reaction, whereby in rejecting one potential match partner strategically reported as unacceptable, say Alice, that match partner applies to another firm, who in turn likes Alice better than their current match Bob. This then causes Bob to be rejected (by the firm who prefers Alice), which in turn leads Bob to apply to the firm that did the manipulation in the first place, who like Bob very much. Whereas, before the manipulation, this firm never would have gotten an application from Bob in the first place, and instead would match to Alice.

There is a rough analogy between a truncation strategy in matching and demand reduction in uniform-price auctions — by saying the marginal units are unacceptable, one gets better pricing on the inframarginal units. Roth and Peranson (1999) showed in computational simulations that opportunities to successfully manipulate the market are rare in realistic size markets. Immorlica and Mahdian (2015) and Kojima and Pathak (2009) then showed theoretically that such manipulations become rare in a theory model as the market grows large. Azevedo and Budish (2019) define a notion of approximate strategy-proofness called *strategy-proof in the large* (SP-L), and show that deferred acceptance is SP-L for the receiving side of the market.

Key property: no justified envy

As described above in Section 2.1.5, in school choice the schools are sometimes treated as agents with preferences for match partners on the other side, and sometimes are treated as objects to be allocated to agents on the other side, but with “priority rankings” over the agents they match with. In this latter case, the phrase “stability” is inappropriate, because a student and a school are not both agents, so cannot form a blocking pair in the traditional sense. Abdulkadiroglu and Sonmez (2003) therefore introduce the term *no justified envy* — the idea being that if student a likes school s better than his assigned school, and school s also ranks student a higher than one of their assigned students, say student b , then student a is said to have justified envy of student b . The terminology captures that a likes b 's outcome better than their own (known as “envy”), and is more deserving of the slot at school s than is this student, because of the priority ranking.

We emphasize that justified envy is just a different term from blocking pair for describing the same mathematical property: if a prefers s to her current match, s is matched to b , and s ranks a higher than b . Therefore, just as the Gale-Shapley algorithm yields an allocation that is stable in the case of agent-to-agent matching, the Gale-Shapley algorithm yields an allocation with no justified envy under this interpretation of schools in the school-choice problem.

In the case where schools' preferences over students are strict, the Gale-Shapley algorithm yields an allocation that maximizes student welfare (in the sense of Pareto efficiency) subject to no justified envy as a constraint. In fact, Gale and Shapley (1962) themselves noted that their algorithm can be interpreted as the solution to a constrained optimization problem (see Budish, 2012). In the case where schools' preferences over students have indifferences, it is possible to improve student welfare relative to the Gale-Shapley allocation, while preserving no justified envy, but at the cost of strategy-proofness. See Erdil and Ergin (2008) and Abdulkadiroglu et al. (2009).

Many-to-one variant

The many-to-one variant of the Gale-Shapley algorithm is nearly identical to the one-to-one version. Firm j still submits a single rank-order list over individual workers. The difference is that if the firm has q_j positions, it can hold as “maybe” up to q_j workers, only issuing rejections if it has more than q_j proposals in a given round.

Under the assumption that firms' preferences are responsive and the workers are the side of the market doing the proposing, the many-to-one Gale-Shapley algorithm inherits the stability and incentives properties from the one-to-one case. The insight that allows the results to translate is that each “position” at a firm can be treated as its own party to the match — with, under responsiveness, all positions at the firm having the same rank-order preference list over workers — allowing the machinery of one-to-one matching to carry through. If the firms are the ones doing the proposing, there are some important differences, and in particular the algorithm is no longer strategy-proof for the proposing side (Roth, 1985).

Roth and Peranson (1999) describe a modification to the algorithm that allows for pairs of workers to identify as couples who seek two positions in the same city. The couple submits a rank-ordered list over pairs of positions, while firms rank workers individually as before. The modification is highly non-trivial, and in particular has to deal with the possibility, first noted in Roth (1984), that no stable matching exists.

Azevedo-Leshno price-theoretic interpretation of Gale-Shapley (“cutoff structure”)

Azevedo and Leshno (2016) provide a price-theoretic interpretation of the Gale-Shapley algorithm. Consider a many-to-one matching market in which students are matching to colleges. Each student i has a utility for each college j , denoted by u_{ij} , and each college j assigns a score to each student i , denoted v_{ji} . The score represents the college’s preference for the student. Azevedo and Leshno (2016) show that the Gale-Shapley algorithm produces a set of price-like statistics, called “cutoffs”, where p_j denotes the minimum score threshold necessary to be admitted to college j . Each student i , in turn, matches to their most preferred college out of the set of colleges for which the student’s score is above the threshold. That is, student i matches to their most preferred element in the set

$$S(\mathbf{v}_i; \mathbf{p}) = \{j : v_{ji} \geq p_j\}.$$

The vector of cutoffs \mathbf{p} is such that the total number of students i with school j as their most-preferred option within the set $S(\mathbf{v}_i; \mathbf{p})$ is equal to the capacity of school j if $p_j > 0$, or weakly less than the capacity of school j if $p_j = 0$ — analogously to competitive equilibrium prices in general equilibrium theory. Fig. 1 illustrates the cutoff idea visually in an example with two colleges. The horizontal axis represents scores at college 1, the vertical axis scores at college 2. The lines represent the cutoffs for each school — roughly 0.6 for school 1 and 0.4 for school 2. The top-right quadrant contains students whose scores are sufficiently high at both colleges that they can choose either one (the green ones choose school 1, the blue ones choose school 2). The top-left quadrant contains students whose scores are high enough to get into college 2 but not college 1; the bottom-right quadrant the students whose scores are high enough to get into college 1 but not college 2; and the bottom-left quadrant the students whose scores are too low at both colleges.

Intuitively, it is helpful to normalize scores to live on $[0,1]$, in which case the cutoff for a school represents the percentile of the school’s preferences that one must be above, in order to gain admission. These cutoffs are sort of like the price of a given school, the difference versus traditional prices being that each student has a different “budget” at each school, depending on the school’s evaluation of them v_{ji} . But, as with traditional prices, a low-dimensional set of statistics, the p_j ’s, determines agents’ choice sets, i.e., the sets $\{j : v_{ji} \geq p_j\}$.

This insight about the underlying price-theoretic structure of Gale-Shapley has already proved useful in some empirical work, as described in Section 3. We suspect other researchers at the intersection of IO and market design will find additional applications in the future.

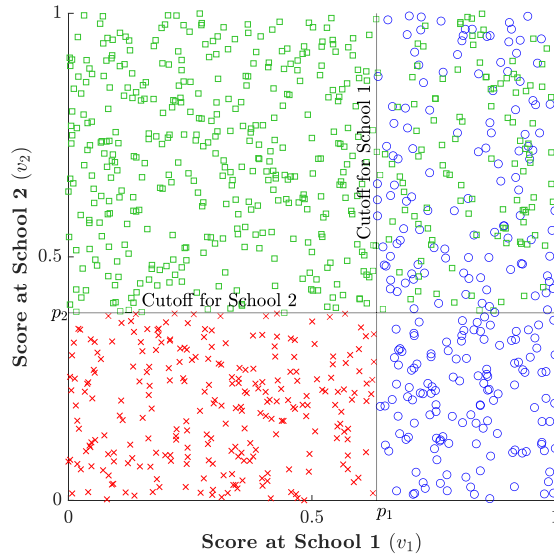


FIGURE 1 Price-Theoretic Interpretation of Gale-Shapley, following Azevedo and Leshno (2016).

Notes: This figure follows Azevedo and Leshno (2016), Figure 3. See the text for description. Students' scores are uniformly distributed on the unit square. Students are equally likely to prefer School 1 to School 2 as vice versa. School 1 has fewer seats and therefore a higher cutoff. The blue circles represent students who match to School 1, the green squares students who match to School 2, and the red crosses students who do not match to either.

Hatfield-Milgrom connection between Gale-Shapley and simultaneous ascending auctions

Azevedo and Leshno (2016)'s insight about Gale-Shapley in turn builds on foundational work of Hatfield and Milgrom (2005), who show that there is a deep connection between the Gale-Shapley algorithm and monotonic combinatorial auction designs such as the simultaneous ascending auction. In particular, the Gale-Shapley algorithm can be understood as a monotonic process in which, in the student-college example, students' choice sets start as the set of all colleges and then monotonically shrink (as they accumulate rejections in the deferred acceptance algorithm), whereas colleges' choice sets start as empty and then monotonically grow (as they accumulate applications in the deferred acceptance algorithm).

2.3.2 Immediate acceptance (“Boston mechanism”)

A common mechanism used for school choice is called either the “Immediate Acceptance” algorithm, based on its conceptual relationship to deferred acceptance, or the “Boston Mechanism”, after the school district where its usage was first noted in the academic literature (by Abdulkadiroglu and Sonmez, 2003).

What preference data is reported to the algorithm: We will refer to \mathcal{I} as students and \mathcal{J} as schools. As in deferred acceptance, the students in \mathcal{I} report preferences over schools in \mathcal{J} in the form of a rank-order list. The schools in \mathcal{J} either report their rank-order preferences over the students in \mathcal{I} in the form of a rank-order list, or these preferences are randomly generated using some kind of lottery (see Pathak and Sethuraman, 2011). We will describe the many-to-one version of the algorithm, in which school j has q_j slots.

Round 1: In the first round of the algorithm, each student i proposes to the highest-ranked school on their preference list. If a school receives proposals in this round, the school says “yes” (not “*maybe*”) to the up-to q_j proposals they prefer the most, and reject all others.

Round k : In each subsequent round of the algorithm, any student i who was rejected in the previous round proposes to the highest-rank school on their preference list that they have not yet proposed to. If a student has no more schools to propose to, they go unmatched.

If school j has already said yes to q_j students in previous rounds, the school rejects all applications in this round. If the school has said yes to fewer than q_j students, and has say q'_j slots remaining, the school says “yes” to the up to q'_j proposals they prefer the most in this round, rejecting the rest.

Ending condition: The algorithm ends when either (i) there is a round with no rejections, or (ii) there is a round with no new proposals. All of the “yes’s” along the way become the matches.

Key difference versus deferred acceptance

The key difference versus the deferred acceptance algorithm is that the most-preferred proposals in a given round, for a school that still has remaining capacity, are responded to with “yes” rather than “maybe”. This means that if a student proposes to a school in round $k > 1$ the school might reject the student even if it previously said yes to a student the school likes less.

The cutoff structure of immediate acceptance

Building off of Azevedo and Leshno (2016), Agarwal and Somaini (2018) show that the Immediate Acceptance algorithm also has a price-like cutoff structure, which they then use in estimation. This structure is based on a re-interpretation of immediate acceptance as a mechanism that gives students priority based on the position where the school is ranked. Specifically, students that rank a school first receive the highest priority, followed by students that rank a school second, and so on. Other school priorities and tie-breakers can be used to order students within each group. A cutoff can then be constructed analogously as the Deferred Acceptance case. In fact Agarwal and Somaini (2018) show that most school choice mechanisms used in practice can be represented in a similar manner.

Discussion: strategic manipulability and efficiency

Immediate Acceptance is not strategy-proof. This is easy to see: a student whose favorite school, say a , is very highly demanded, and whose second-favorite school, say b , is less so, might prefer to rank b first, getting it with high probability, than to take a gamble on getting the highly-popular a . The reason is that, if the student reports a first, and does *not* get it, then by the time they ask for b in the second round, it might already be full — and because the algorithm is one of *immediate* as opposed to deferred acceptance, the student will be locked out of b , even if they have relatively high priority there.

This concern about the manipulability of the mechanism, in conjunction with anecdotal evidence that less-sophisticated families had difficulty with strategizing, and the availability of the Gale-Shapley mechanism as a celebrated and sensible alternative, led market design researchers to initially conclude that the Boston mechanism is a flawed mechanism, and Gale-Shapley should be used instead (Abdulkadiroglu and Sonmez, 2003; Pathak and Sonmez, 2008).

That said, it may be possible to salvage a case for Immediate Acceptance by studying its Bayes-Nash equilibria. As noted by Miralles (2008) and Abdulkadiroglu et al. (2011), Immediate Acceptance has Bayes-Nash equilibria, under stylized circumstances, in which students strategically misreport their preferences, optimally, and these equilibria are actually more efficient than the dominant-strategy equilibria of the deferred acceptance algorithm. The rough intuition is that these equilibria get some information about cardinal preference intensity into the allocation (by *how much* does the student like a better than b), whereas deferred acceptance only gets ordinal preference information into the allocation. Probabilities of getting into a particular school in a particular round then play a price-like role in equilibrium; see Section 2.3.5 for further discussion. In principle, methods developed in Azevedo and Budish (2019) could implement these Bayes-Nash equilibria of Immediate Acceptance in a manner that is SP-L.

However, the empirical evidence to date (discussed in Section 4) suggests that the potential welfare gains, even if students play the Bayes-Nash equilibria perfectly, are relatively small, whereas various other evidence suggests that strategic mistakes are likely to be a prominent concern in this setting.

2.3.3 Random serial dictatorship

Random serial dictatorship is a commonly-studied mechanism in single-unit assignment problems.

What preference data is reported to the algorithm: We will refer to \mathcal{I} as agents and \mathcal{J} as objects. Assume each object is in unit supply. Agents in \mathcal{I} report their preferences over objects in \mathcal{J} in the form of a rank-order list.

Random serial ordering: The algorithm begins by choosing a random ordering over the agents. With $|\mathcal{I}|$ agents there are $|\mathcal{I}|!$ possible random orders.

Round 1: Whichever agent is 1st in the serial order chooses their most-preferred object.

Round k: The agent who is k th in the serial order chooses their most-preferred object, from whatever is still remaining.

Connection to deferred acceptance

Random serial dictatorship is equivalent to the following version of deferred acceptance: generate a random serial order over the agents, and then run deferred acceptance using this serial order as the preference ordering for every object in the set \mathcal{J} . Either way, the 1st agent gets their most preferred object, the 2nd agent gets their most preferred object other than the one taken by agent 1, the 3rd agents gets their most preferred object other than the ones taken by agents 1 and 2, etc.

Key properties: strategy-proof and ex-post Pareto efficient

The random serial dictatorship mechanism is strategy-proof and ex-post Pareto efficient. In this context, ex-post Pareto efficiency means that, for any realization of the random serial order over agents, the resulting allocation is such that there is no other allocation that all agents weakly prefer, with at least some strict.

2.3.4 Top trading cycles

Top trading cycles was invented by David Gale as a solution to the house-allocation problem (as reported in Shapley and Scarf, 1974). The algorithm works as follows:

What preference data is reported to the algorithm: We will refer to $\mathcal{I} = i_1, \dots, i_n$ as agents, $\mathcal{J} = j_1, \dots, j_n$ as objects, and assume i_1 is endowed with j_1 , i_2 is endowed with j_2 , etc. Assume each object is in unit supply. Each agent reports a rank-ordered list over all objects they prefer to their own endowment; any object not included in the rank-order list is understood to be less preferred to the agent's own endowment.

Round 1: Create a directed graph in which each node is an agent-object pair, and out of each node is exactly one directed edge to the remaining object the agent likes best. If the agent likes their own object best, they point to themselves.

Observe that this graph has at least one cycle: there are n nodes and each node has a directed edge emanating from it. (This potentially includes a cycle in which an agent points to themselves).

For every cycle: execute the trades described by that cycle. For example, if there is a cycle $\{i_1, j_1\} \rightarrow \{i_2, j_2\} \rightarrow \{i_3, j_3\} \rightarrow \{i_1, j_1\}$, then i_1 gets object j_2 , i_2 gets object j_3 , and i_3 gets object j_1 . Then remove these nodes from the graph.

Round k: Create a directed graph in which each node is an agent-object pair that has not yet traded, and out of each node is exactly one directed edge to the remaining object the agent likes best.

Again, this graph has at least one cycle. Execute the trades described by that cycle. Remove these nodes from the graph.

Ending condition: The algorithm ends when there are no remaining nodes in the graph. This means that every agent has either traded or reached a stage in the algorithm where their most-preferred alternative was to keep their endowment.

Connection to competitive equilibrium

While Shapley and Scarf (1974) did joke that their model bore little resemblance to real-world housing markets, they noted that the solution did have a connection to competitive equilibrium. Formally, assign price p_1 to all objects that trade in round 1, price $p_2 < p_1$ to all objects that trade in round 2, etc. At these prices, (i) each agent sells and buys a house at the same price (i.e., the trade is in their budget set), and (ii) more subtly, each agent gets the most-preferred house in their budget set. This can be seen by noting that an agent who trades in round k got a house that they like better than any other house that traded in round k or later; that is, they got the best house they can afford, given their budget p_k .

Leshno and Lo (2021) provide a characterization of top trading cycles in terms of cutoffs. The TTC cutoffs are higher-dimensional than the cutoffs for deferred acceptance found in Azevedo and Leshno (2016), reflecting that an agent's threshold for obtaining a given object depends not only on how much that object likes the agent, but how much the agent's endowment is liked by those who can facilitate a trade for the object.

Key properties: strategy-proof and ex-post Pareto efficient

The key properties of TTC are that it is strategy-proof and ex-post Pareto efficient (Roth, 1982a). These are the same properties as random serial dictatorship, and in fact these two mechanisms are in some environments two different ways of implementing the exact same distribution over allocations (Abdulkadiroglu and Sonmez, 1998).

2.3.5 Hylland and Zeckhauser pseudomarket

Hylland and Zeckhauser (1979) proposed a competitive equilibrium approach to the single-unit assignment problem.

What preference data is reported to the algorithm: Agents in \mathcal{I} report their preferences over objects in \mathcal{J} in the form of a von-Neumann Morgenstern utility function. Formally, each agent i reports a vector $u_{i1}, \dots, u_{i|\mathcal{J}|}$ where element u_{ij} indicates the agent's vNM utility for object j .

The pseudomarket equilibrium

Conceptualize each object j as perfectly divisible into probability shares. Each agent i will be allocated a vector of probability allocations $x_{i1}, \dots, x_{i|\mathcal{J}|}$ with the properties that: (i) each $x_{ij} \in [0, 1]$, and (ii) $\sum_j x_{ij} \leq 1$. Summed over all agents, an allocation is feasible if, in addition, (iii) $\sum_i x_{ij} \leq 1$ for all j . This notion of feasibility relies on the idea of “implementing” a probabilistic allocation, via appeal to the Birkhoff-von Neumann theorem: that is, finding a convex combination of sure allocations that correspond to the intended probabilistic allocation.

With these concepts in hand, the Hylland Zeckhauser pseudomarket mechanism is simple to describe. First, agents report their vNM preferences. Then, the mechanism finds a vector of competitive equilibrium prices $p_1^*, \dots, p_{|\mathcal{J}|}^*$: prices such that, when each agent i is allocated their most-preferred affordable bundle at these prices given a common budget b (wlog, b can be normalized to 1), the market clears. Formally, agent i is assigned the bundle $x_{i1}^*, \dots, x_{i|\mathcal{J}|}^*$ that maximizes their utility $u_i \cdot x_i$ subject

to the budget constraint $p^* \cdot x_i \leq b$ and the unit-demand constraint $\sum_j x_{ij} \leq 1$; and market clearing means that $\sum_i x_{ij}^* \leq 1$ for all j .

One of the key theoretical contributions of Hylland and Zeckhauser (1979) is an existence theorem for such prices. Computing pseudomarket equilibrium prices remains non-trivial. For recent progress, please see Eraslan et al. (2021).

Connection to the immediate acceptance algorithm

Miralles (2008) discovered that there exists a Bayes-Nash equilibrium of the immediate acceptance algorithm (i.e., Boston mechanism) that coincides with the outcome from truthful play of the Hylland-Zeckhauser mechanism. A rough intuition for how this is possible is that even though the immediate acceptance algorithm mechanism asks agents to report ordinal preference information, while the Hylland-Zeckhauser mechanism requires vNM utilities, how agents strategically choose to report their ordinal preferences depends on their underlying vNM utilities, in just the right way.

Essentially, the *probability* of obtaining a good in equilibrium of the Boston mechanism, serves an analogous role to the *prices* in the HZ mechanism.

Properties: ex-ante efficient, SP-L

As noted above in the discussion of RSD, an allocation is ex-post Pareto efficient if there is no other allocation that all agents weakly prefer, with at least some strict. A probability distribution over allocations is *ex-ante Pareto efficient* if there is no other probability distribution over allocations that all agents weakly prefer, with at least some strict. Ex-ante Pareto efficiency implies ex-post Pareto efficiency, in the sense that for a random allocation to be ex-ante Pareto efficient, it must be the case that any sure allocation that occurs with positive probability is ex-post Pareto efficient. The reverse need not be true, as shown in Bogomolnaia and Moulin (2001).

The Hylland-Zeckhauser pseudomarket mechanism is ex-ante Pareto efficient. It is not strategy-proof, but it is SP-L.

2.3.6 Draft approaches to multi-unit assignment

Draft mechanisms, in which agents take turns choosing objects one-at-a-time over a series of rounds, are a common approach to multi-unit assignment problems in practice. A prominent example is sports teams choosing players, and Harvard Business School has long used a draft for course allocation.

Draft mechanisms are not strategy-proof. Intuitively, if an agent's most-preferred object is not widely liked by others, but their second most-preferred object is widely liked by others, they might do better by trying to choose their second most-preferred earlier than their most-preferred, which they anticipate being able to obtain later. In equilibrium, drafts are not Pareto efficient, a point first made by Brams and Straffin (1979). See also Brams and Taylor (1996) for many variations on the draft idea.

The multi-object version of random serial dictatorship — which is like a draft except that there is one round in which agents choose complete bundles, instead of many rounds where agents choose one-at-a-time — is both strategy-proof and Pareto efficient. Nevertheless, Budish and Cantillon (2012) show that the draft performs

better on both welfare and fairness grounds than does the dictatorship. This example serves as a reminder that the fact that a mechanism satisfies attractive properties is not always a reliable guide to mechanism performance (Budish, 2012).

2.3.7 Competitive equilibrium approaches to multi-unit assignment

A series of recent papers has developed competitive equilibrium approaches to assignment problems in which agents have multi-unit demands. In this subsection we describe the variations on this idea. For formal properties of each variation, please see the underlying papers.

Approximate Competitive Equilibrium from Equal Incomes (A-CEEI)

Budish (2011) considers an environment in which agents have unrestricted preferences over bundles of objects, like in a combinatorial auction, but there are exogenous restrictions against the use of monetary transfers. A motivating problem at the time was course scheduling at universities. More recently, an instance of this problem is vaccine allocation across countries from vaccine doses jointly purchased by the COVAX consortium (Castillo et al., 2021).

The mechanism works as follows.

Step 1: Agents report ordinal preferences over bundles of objects.

Step 2: Agents are assigned approximately equal budgets of an artificial currency; formally, for some $\beta > 0$ but arbitrarily small, all budgets are drawn from the interval $[1, 1 + \beta]$.

Step 3: The mechanism finds prices that approximately clear the market. Formally, this is a set of prices where, when each agent is allocated their most-preferred affordable bundle at these prices (based on the ordinal preferences reported in Step 1 and the random budget drawn in Step 2), market clearing error is smaller than a small bound.

Step 4: These allocations are implemented.

Notice that agents are allocated a sure bundle rather than a probabilistic allocation as in Hylland and Zeckhauser.

The reason for the approximations is that exact competitive equilibrium from equal incomes might not exist. The main theorem in Budish (2011) shows that small amounts of budget inequality and market-clearing error are sufficient to restore existence. See Reny (2017) for a generalization of the existence theorem.

Multi-unit generalization of Hylland-Zeckhauser probability shares market

Budish et al. (2013) provide a generalization of the Hylland and Zeckhauser pseudo-market mechanism to accommodate certain kinds of multi-unit demands and certain kinds of additional constraints on demand and supply. Relative to Budish (2011), the key advantage is that market clearing is exact rather than approximate, whereas the key disadvantage is that preferences are required to be additive-separable over objects.

Feeding America artificial-currency market

Prendergast (2017) reports on an artificial currency market the author and colleagues implemented for the Feeding America system of food banks across the United States. Whereas the competitive equilibrium mechanisms described above are static (or one-shot, i.e., run just a single time) the Feeding America market mechanism is infinitely repeated.

In each allocation round, food banks submit bids, in an artificial currency, for the available truckloads of food (e.g., a truckload of chicken coming from Nebraska next week). The highest bidder wins and pays their bid amount (i.e., it is a first-price auction). The key idea is that the currency paid in the bid is then distributed to all of the other food banks, in proportion to their populations. In this way, the overall amount of currency in the system stays fixed over time, so there is some consistency of prices over time (e.g., the price of a truckload of chicken).

From a mechanism design perspective, the key thing to point out about the Prendergast (2017) artificial-currency market is that, since the market is infinitely repeated, fake money becomes like real money that enters the utility function — the fake money always has a future use. Theoretically, this is important because it makes it reasonable to model market participants as having quasi-linear preferences over the objects and the fake money, and to then utilize the auction theory enabled by quasi-linear utility. Whereas, in artificial-currency markets that are one-shot, the fake money has no future use so it is not appropriate to put money into the utility function. Instead, prices just encode choice sets, given agents' budgets, more like in general equilibrium theory than in auction theory.

3 Empirical frameworks and applications

We now turn our attention to empirical models of matching and assignment markets. A primary goal of these models is to estimate agents' preferences or payoffs. These quantities are the primitives in the theoretical models described above. Their estimates will allow us to compare and evaluate alternative market designs, understand their welfare and distributional properties, and engage in counterfactual analysis.

Specific to market design, there are certain instances in which centralized matching mechanisms directly yield data on ordinal preferences. In particular, if a mechanism is strategy-proof and agents understand this property, then data on the preferences agents report to the mechanism can be used directly to analyze certain questions of interest, without estimating a parametric model of utility. For example, such data enable comparison with other strategy-proof alternatives that also use ordinal preferences.

Although these cases represent an ideal case for empirical work, they are comparatively rare. It is much more common for the analyst to confront data limitations or mechanisms that are not strategy-proof. Additionally, even when reliable ordinal preference information is available, it is sometimes useful to estimate a cardinal representation of the distribution of preferences.

The methods described in this section fill this important gap. We begin by discussing empirical models with non-transferable utility (section 3.1) before turning to models with transferable utility (section 3.2). Section 4 briefly discusses empirical results based on these models, but expands to include other market design applications as well.

3.1 Non-transferable utility models

Recall that non-transferable utility models assume that transfers between agents are either prohibited or are exogenously determined. The empirical leading examples include matching students to schools or colleges (Abdulkadiroglu et al., 2017b; Agarwal and Somaini, 2018), entry-level jobs without salary negotiations (Agarwal, 2015), oil drilling (Vissing, 2018), and marriage markets.

Empirical approaches in these settings build from random utility models, which are commonly used in IO to represent consumer preferences in discrete choice settings (McFadden, 1973; Manski, 1977; Berry et al., 1995). In the matching context, these models parametrize utility of an agent from matching with potential partners as a function of characteristics observed in the data. The key difference from the consumer demand context lies in how the data are used to learn about preferences. Whereas a consumer can pick their most preferred product at the posted prices, in the matching context the assignment that results from a choice is determined by the market's design.

The empirical approach depends both on the rules of the market being analyzed and the type of data that is available. Rules matter because they specify how agents express their preferences, may give agents incentives to manipulate their reported preferences, and determine the final allocation. Therefore, the mechanism's properties shape the assumptions that will be used during estimation. And, given a set of rules, the approach will depend on whether we observe data on rank-order lists submitted to a mechanism or only the final matches.

We begin by describing the random utility model (section 3.1.1) before proceeding to estimation methods. The methods are separated into two cases, based on whether we have access to data on reported rank-order lists from an assignment mechanism (section 3.1.2) or data only from observed matches (section 3.1.3). In both cases, we assume that the researcher has access to rich micro-data on individual characteristics.⁵

⁵ This data requirement differs from methods for demand estimation that use aggregate data on market shares (e.g. Berry et al., 1995). There are two reasons for this difference. First, our goal is to understand partnerships between the two sides of the market. Thus, we require micro-data (Berry et al., 2004; Petrin, 2002). Second, the shares are driven primarily by the number of matches that are feasible for each side of the market.

3.1.1 Random utility model

Denote i 's utility if matched to j with u_{ij} and j 's utility from being matched to i with v_{ji} . Let u_{i0} and v_{j0} denote the utilities of remaining unmatched. This formulation implicitly assumes that an agent's utility from a match does not depend on the other matches in the economy. A typical empirical exercise involves identifying and estimating the joint distributions of the vectors of random utilities $\mathbf{u}_i = (u_{i1}, \dots, u_{iJ})$ and $\mathbf{v}_i = (v_{i1}, \dots, v_{iJ})$ conditional on observable characteristics.⁶

We start by representing the preferences of agents on one side of the market. The most general form for the utility of i being matched with j that we will employ is given by

$$u_{ij} = u(\mathbf{x}_j, \mathbf{z}_i, \xi_j, \epsilon_i) - d_{ij}, \quad (1)$$

where \mathbf{z}_i and \mathbf{x}_j are vectors of observed characteristics for i and j , respectively, and d_{ij} is a scalar observable that potentially varies with both i and j . The term ϵ_i captures unobserved determinants of agent i 's preferences. It may be multi-dimensional and include j -specific taste shocks. The term ξ_j includes unobserved characteristics of j . The term d_{ij} is an observable, match-specific characteristic (e.g. distance between i and j) that will be used as the metric for utility in some applications.

A random utility model requires scale and location normalizations because choices (under uncertainty) are invariant to a positive, affine transformation of utilities. Accordingly, we will normalize the value of the outside option to zero, i.e. $u_{i0} = 0$. Observe that the unit coefficient on d_{ij} represents a scale normalization.⁷

While the analysis of identification will often allow for general functional forms, the empirical methods below typically use additional parametric assumptions to ease the computational burden and to achieve statistically precise estimates with finite sample sizes. The most convenient functional forms depend on available data and the mechanism or setting being analyzed. For example, a commonly used parametric form encompassed by the model above assumes that

$$u_{ij} = \mathbf{x}'_j \boldsymbol{\beta} + \mathbf{x}'_j \bar{\boldsymbol{\gamma}} \mathbf{z}_i + \xi_j + \mathbf{x}'_j \boldsymbol{\gamma}_i + \epsilon_{ij} - d_{ij}, \quad (2)$$

where $\boldsymbol{\gamma}_i$ and ϵ_{ij} are mean-zero, normally distributed random variables with variances to be estimated, and $\bar{\boldsymbol{\gamma}}$ is a matrix conformable with \mathbf{x}'_j and \mathbf{z}_i . The vector $\boldsymbol{\theta}$ denotes the model's unknown parameters, namely $(\boldsymbol{\beta}, \bar{\boldsymbol{\gamma}}, \xi_1, \dots, \xi_J)$ and the parameters governing the distribution of ϵ_{ij} and $\boldsymbol{\gamma}_i$. This formulation is both tractable and flexible. Such specifications are commonly used in empirical models of consumer demand because they capture many preference determinants such as a vertical index of quality that is valued equally by every agent and heterogeneous preferences based on observables as well as unobservables.

⁶ We refer the reader to Matzkin (2007) for the formal definition of identification that we employ in this chapter.

⁷ The specification above also assumes that all agents dislike increases in d_{ij} . This restriction is not essential in many cases discussed below, and the sign of this coefficient can be estimated.

In some applications, the preferences on the other side of the market v_{ji} may be known from administrative data or institutional knowledge. For example, many schools and colleges use exam scores to rank students (e.g. Fack et al., 2019; Akyol and Krishna, 2017), while other school districts use different but still well-defined priorities. In these cases, v_{ji} does not need to be estimated.

When v_{ji} is unknown, one can specify an analogous model for the preferences of agents on the other side of the market. Specifically, the utility of agent $j \in \mathcal{J}$ for matching with agent $i \in \mathcal{I}$ is given by

$$v_{ji} = v(\mathbf{x}_j, \mathbf{z}_i, \eta_i) - w_{ji}, \quad (3)$$

where η_i is unobserved and w_{ji} has an interpretation analogous to d_{ij} . In this case, we would also normalize v_{i0} to zero.

The preference model includes two noteworthy assumptions. First, the baseline model has no externalities. An agent's utility depends on only their own matches.⁸ This rules out preferences for attending school with specific peers or working with specific colleagues. In other applications involving matching between firms, the lack of externalities rules out preferences that depend on a competitor's matches. We will discuss extensions that incorporate some of these features in Section 3.3. Second, the model abstracts away from costs of acquiring information about the other side of the market by assuming that preferences are well formed. An exception is Narita (2018), which considers the possibility that preferences evolve after agents receive an initial assignment.

3.1.2 Analysis with data from assignment mechanisms

Two common sources of data relevant for preference analysis are rank-order lists (e.g. Hastings et al., 2009; Abdulkadiroglu et al., 2017b) and participant surveys (Budish and Cantillon, 2012). Correspondingly, a well-developed literature has taken advantage of the rich information contained in these reports and used them to derive methods to estimate agents' preferences. These methods employ revealed preference implications of assumed participant behavior.

We will use school choice mechanisms as our central example because these mechanisms are widely used around the world. In this section, we refer to agents on side \mathcal{I} as students and agents on side \mathcal{J} as schools. Our goal will be to estimate students' preferences for schools. An analogous exercise analyzes schools' preferences for students when rank-ordered data are available. The discussion below is brief as we point the reader to Agarwal and Somaini (2020) for a more thorough review of methods for estimating preferences in school choice settings.⁹

⁸ This assumption is sometimes referred to as "responsive preferences" (see Roth and Sotomayor, 1990, Chapter 5).

⁹ The methods discussed below are generally applicable to other settings where a researcher can obtain data on preferences. For example, Hitsch et al. (2010) estimate preferences in an online dating context by analyzing the decision to contact a potential date. They interpret the decision to contact a potential date as indicative of high utility. Their approach allows them to estimate flexible preferences for men and women.

We will typically assume the following conditional independence condition of the form:

$$\epsilon_i \perp \mathbf{d}_i | \mathbf{z}_i, \mathbf{x}, (\xi_j)_{j=1}^J, \quad (4)$$

where $\mathbf{d}_i = (d_{i1}, \dots, d_{iJ})$ and $\mathbf{x} = (x_1, \dots, x_J)$. The independence condition (4) assumes that agent i 's unobserved taste shocks are conditionally independent of the vector of numeraire match-specific characteristics \mathbf{d}_i given the other observed characteristics of i , \mathbf{z}_i , and the observed vector of observed characteristics for agents on the other side of the market characteristics of market configuration \mathbf{x} . The assumption must be evaluated within each empirical application, and it is typically reasonable if \mathbf{x} is a sufficiently rich control.¹⁰

In the school choice example, if d_{ij} is the distance from i 's residence to j 's location (as in Abdulkadiroglu et al., 2017b, for example), then students' preferences can be summarized in terms of their "willingness to travel." The conditional independence assumption above requires that distance to school is independent of other unobserved determinants of preferences for schools. This assumption may be a good approximation if \mathbf{z}_i includes sufficiently rich data about a student's achievements, demographics and socio-economic characteristics. Relaxing this assumption would be likely to require a model of residential choice and sorting based on unobserved factors that influence preferences for schools.

The main challenge in analyzing rank-order lists submitted to mechanisms is that student reports may not reflect their true preferences. For example, students may avoid ranking schools to which they are unlikely to be admitted. In other cases, students may have an incentive to manipulate their rankings in order to gain priority at certain schools. Below, we discuss two cases, one in which students report truthfully and the other in which students manipulate their preferences.¹¹ We discuss parametric assumptions most convenient for estimation after laying out the core empirical strategy in these two cases.

Truthful reports: Reports made to a strategy-proof school choice mechanism enable a straightforward empirical strategy if agents understand it and follow this recommendation.¹² In some cases it is possible to answer the relevant questions of interests directly using these reported rank-order lists. For example, Abdulkadiroglu et al. (2009) use data from New York City's implementation of the Deferred Acceptance

¹⁰ Relaxing this assumption is a fruitful avenue for future research. Such work will likely require augmenting the model to incorporate other sources of exogenous variation and specify alternative data-generating processes.

¹¹ We focus on estimating preferences on one side of the market. The approaches discussed below can be applied to both sides of the market (see He and Magnac, 2020; Aue et al., 2020, for example).

¹² Evidence from both experiments and the field suggests that students are more likely to report their preferences truthfully when interacting with a strategy-proof mechanism (Chen and Sonmez, 2006; de Haan et al., 2018). Nonetheless, comprehending that a mechanism is strategy-proof may be complicated (Li, 2017) and some students are liable to mistakenly submit rankings that are not truthful (Rees-Jones, 2018; Shorrer and Sovago, 2019; Hassidim et al., 2020; Artemov et al., 2021; Budish and Kessler, 2021).




FIGURE 2 Revealed Preferences – Truthful Reports.

Algorithm to simulate and compare with alternative ordinal school choice mechanisms while assuming truthful reporting.

In many cases, for either welfare or counterfactual analysis, it becomes necessary to estimate a cardinal representation of preferences from the ordinal information. Specifically, if agent i ranks j above j' , then we can infer that

$$u_{ij} > u_{ij'}.$$

This inequality is similar to the standard discrete choice case in which we infer that $u_{ij} > u_{ij'}$ for all j' if a consumer picks j . But, in the case of a truthfully reported rank-order list, we learn more fine-tuned information about more than just the most-preferred option.

It is less clear how to treat schools that are not ranked on the list. One common approach is to assume that students rank all schools that are acceptable, i.e. preferable to the outside option. Thus, if \underline{j} is the lowest-ranked school, then $u_{i\underline{j}} > u_{i0} > u_{ij'}$ if j' is not ranked. In this model, the various rank-order lists partition the space of utilities, as shown in Fig. 2 for when $J = 2$. The five regions in the figure correspond to the various ways in which two schools can be ranked, including the possibility that only one school is submitted.

An alternative reason why a student may not rank a school is that she is ineligible at that school. In this case, a researcher may want to limit the revealed preference inequalities above to the set of schools where a student is eligible. A related, but conceptually distinct model, is when a student omits a school because she believes that her chances of admission at that school are low but ranks the remaining schools truthfully. This model is a special case of models that consider reports from a manipulable mechanism.

Observe that truthfully reported rank-order lists are similar to, but provide richer information about preferences than, standard discrete choice models in which a consumer picks only their favorite product. Specifically, if a consumer picks option 1 in a standard discrete choice setting, then we can only deduce that the consumer's utilities are in either the region labeled "Rank 1" or "Rank 1>2" in Fig. 2, but we cannot distinguish between these two regions. The richer information in ordered lists can help identify heterogeneity in preferences (Beggs et al., 1981; Berry et al., 2004). In the school choice context, students often rank many more schools, allowing for very rich specifications for the distribution of utilities (see Abdulkadiroglu et al., 2017b, for example). Accurately estimating this heterogeneity is important for analyzing the value of improving assignments because swapping the allocations of any two agents leaves welfare unchanged in models with homogeneous preferences.

Our goal is to identify the cdf F_{U^*} , the joint cdf of the random vector \mathbf{u}_i^* with the j -th element equal to $u(\mathbf{x}_j, \mathbf{z}_i, \xi_j, \epsilon_i)$. We drop the explicit conditioning on $\mathbf{z}_i, \{\mathbf{x}_j, \xi_j\}_{j=1}^J$ for notational simplicity and assume the condition in Eq. (4) holds. Under this assumption, the probability that i submits the rank-order list $R = (j_1, j_2, \dots, j_J)$ can be written as

$$\mathbb{P}(R | \mathbf{d}_i = \mathbf{d}; F_{U^*}) = \int 1 \left\{ \mathbf{u}_{j_k}^* - d_{j_k} \geq \mathbf{u}_{j_{k+1}}^* - d_{j_{k+1}} \text{ for all } k \in \{1, \dots, J-1\} \right\} dF_{U^*}.$$

Convenient functional forms for estimating this model via maximum likelihood are further discussed below.

Manipulable mechanisms: Although strategy-proof mechanisms are desirable on theoretical grounds, many school districts use manipulable mechanisms. The widely criticized but still commonly used Immediate Acceptance mechanism, for example, prioritizes students who rank a school higher, generating strategic incentives. Lab studies (Chen and Sonmez, 2006), survey data (de Haan et al., 2018), and signs of strategic reporting in administrative data (Calsamiglia and Güell, 2018; Agarwal and Somaini, 2018) suggest that students do respond to these incentives.

To empirically analyze reports in manipulable mechanisms, it is useful to think about reports as actions in a game. Each action is associated with an expected payoff. If agents maximize expected utility, the observed report must yield the highest expected payoff. This approach assumes a considerable degree of sophistication as it requires agents to perform two cognitively demanding tasks. First, they must be able to calculate the expected payoff for each possible report. Second, they must maximize over all possible reports. We focus on the case where agents have rational expectations and can optimize before discussing extensions.

Let $\mathbf{L}_R \in \Delta^J$ be a probability vector representing an agent's beliefs about the probabilities with which she will be assigned to each of the J schools if she submits the report $R \in \mathcal{R}_{\mathcal{I}}$. The expected utility of this report is $\mathbf{u}_i \cdot \mathbf{L}_R$. If we observe the report R_i from student i , then optimality implies that $\mathbf{u}_i \cdot \mathbf{L}_{R_i} \geq \mathbf{u}_i \cdot \mathbf{L}_R$ for all $R \in \mathcal{R}_{\mathcal{I}}$. Let C_{R_i} be the set of utilities \mathbf{u}_i such that the report R_i maximizes expected utility. This set is a convex cone in the space of utilities that contains the origin. Moreover,




FIGURE 3 Manipulable Mechanisms.

the collection of sets C_R for $R \in \mathcal{R}_{\mathcal{I}}$ partitions the space.¹³ Fig. 3 illustrates these sets for our simplified case with two schools. In this example, $\mathbf{u}_{R,R'}$ represents utilities for which the student is indifferent between submitting R and R' . Similarly, a student with utilities given by $\mathbf{u}_{R,R''}$ is indifferent between R and R'' . The students with utility vectors in the set C_R (weakly) prefer R to the other reports.

Notice that this representation is general and does not directly assume that the mechanism is manipulable. In fact, in the special case of a strategy-proof mechanism, the assignment probabilities \mathbf{L}_R are such that reporting preferences truthfully are optimal. This is because an implication of the strategy-proof assumption is that $\mathbf{u}_i \cdot \mathbf{L}_{R_i} \geq \mathbf{u}_i \cdot \mathbf{L}_R$ if R_i corresponds to the truthful report. In addition, this framework can also be used to analyze the case in which students rank schools in order of preference, but omit some schools where admission probabilities are low because only a limited number of schools can be ranked.

The discussion above implicitly assumes that the vectors \mathbf{L}_R for $R \in \mathcal{R}_{\mathcal{I}}$ are known to the analyst. In practice, these have to be estimated. Under rational expectations, these beliefs are objective assignment probabilities. Towards constructing an estimator for these probabilities, Agarwal and Somaini (2018) noticed that almost all the mechanisms used in practice can be described using a cutoff structure analogous to one that applies to stable allocations discovered in Azevedo and Leshno (2016). The distribution of these cutoffs in equilibrium determines the objective assignment

¹³ More precisely, every $\mathbf{u} \in \mathbb{R}^J$ belongs to the interior of at most one of the sets in the collection and belongs to at least one set C_R . There is one exception. If two reports R_i and R'_i result in the same vector of probabilities, then the sets C_{R_i} and $C_{R'_i}$ will be identical to each other.

probabilities. Thus, instead of estimating L_R , one can instead estimate the cutoff distribution, which is a lower dimensional object. The cutoff structure is also useful for estimating beliefs under alternative assumptions on the belief formation process.

In this model, the probability that i submits the rank-order list R_i can be written as

$$\mathbb{P}(R_i | \mathbf{d}_i = \mathbf{d}; F_{U^*}) = \int 1 \{(\mathbf{u}^* - \mathbf{d}) \cdot L_{R_i} \geq (\mathbf{u}^* - \mathbf{d}) \cdot L_R \text{ for all } R \in \mathcal{R}_{\mathcal{I}}\} dF_{U^*}.$$

This expression follows because $(\mathbf{u}^* - \mathbf{d}) \cdot L_{R_i}$ is the expected utility from reporting R_i , which must be larger than the expected utility from any alternative report R . This expression also forms the basis of estimation via maximum likelihood. We provide further details below.¹⁴

Several extensions that vary behavioral assumptions have been based on this approach. Kapor et al. (2020a) propose estimating L_R by surveying agents. The survey focused on families participating in the school choice mechanism in New Haven and found significant differences between elicited and objective assignment probabilities. He (2017) and Hwang (2014) do not impose all the conditions imposed by optimality. Instead, they derive a few intuitive necessary conditions that reports have to satisfy and use the implied revealed preference relations to estimate preferences. One benefit of this approach is that it relies less heavily on optimal play. A cost is computational and statistical complexity because incomplete models of behavior do not admit maximum likelihood methods. Agarwal and Somaini (2018) and Calsamiglia et al. (2020) estimate mixture models in which some agents behave optimally while others behave naively; i.e., agents report their true ordinal preferences even if it is in their interest to report something else. For a more detailed survey of methods for incomplete and mixture models, see Agarwal and Somaini (2020).

Parametric assumptions and estimation: A common feature of the models described above is that they result in linear restrictions on the vector of utilities \mathbf{u}_i for each agent i . This simple structure allows for likelihood-based estimation methods. These methods typically employ specific functional form and distributional assumptions on Eq. (1) in order to limit the dimension of parameters to be estimated. The two most commonly used functional forms are based on logit and probit errors.

Logit models

Consider the special case of Eq. (2) in which

$$u_{ij} = \delta_j + \mathbf{x}_j \bar{\gamma} \mathbf{z}_i - d_{ij} + \varepsilon_{ij} \quad (5)$$

¹⁴ Agarwal and Somaini (2018) show conditions under which variation in \mathbf{d} can be used to identify the distribution of utilities F_{U^*} . An alternative approach, developed in Carvalho et al., 2019 for the two-school case and generalized in Agarwal and Somaini (2018), is to use variation in assignment probabilities L_R that is orthogonal to preferences to identify F_{U^*} .

and $u_{i0} = \varepsilon_{i0}$, where ε_{ij} follows an extreme-value type I distribution with location parameter 0 and scale parameter σ . In addition to the distributional assumption on ε_{ij} , this specification excludes the terms γ_i and folds $\mathbf{x}_j\beta + \xi_j$ into the fixed effect δ_j . Fack et al. (2019) used this parametric form to estimate high school preferences in Paris under both stability and truth-telling. For notational convenience, collect the parameters of the model in the vector $\theta = (\delta, \bar{\gamma}, \sigma)$.

This functional form is useful when rank-order lists are assumed to be truthful since the probability that student i submits the rank-order list $R_i = (j_1, j_2, \dots, j_J)$ can be written in closed form. It is given by

$$\mathbb{P}(R_i | \mathbf{x}_j, \mathbf{z}_i; \theta) = \prod_{k=1}^J \frac{\exp\left(\frac{1}{\sigma} (\delta_{j_k} + \mathbf{x}_{j_k} \bar{\gamma} \mathbf{z}_i - d_{ij_k})\right)}{1 + \sum_{j \neq j_{k'} \text{ for } k' < k} \exp\left(\frac{1}{\sigma} (\delta_j + \mathbf{x}_j \bar{\gamma} \mathbf{z}_i - d_{ij})\right)}. \quad (6)$$

The term corresponding to $k = 1$ is the probability that the school ranked first, j_1 , has the highest utility. This term is identical to the standard discrete choice case since j_1 is the most-preferred school. The term corresponding to the general k is the probability that the school ranked in position k has the highest utility amongst the schools not ranked any higher. This multiplicative form is specific to the logit model and its independence of irrelevant alternatives property (Beggs et al., 1981). A benefit of this assumption is that the parameters θ can be estimated by Maximum Likelihood. Eq. (6) reveals that each rank-order list contains strictly more information than the observed assignment. Using data on reports will typically yield more precise estimates and allow for more flexible parameterizations than only using data on allocations will allow.¹⁵

In many contexts, we expect a student who ranks a school with, say, good math outcomes at the highest position will also rank other schools with good math outcomes near the top of their list. Such patterns motivate introducing the random coefficients γ_i in Eq. (5). In these models, students with a high coefficient on a particular school characteristic will tend to rank many schools with high values of that characteristic. To incorporate this heterogeneity, the likelihood functions need to be modified. For example, suppose that γ_i is assumed to be distributed $\gamma_i \sim \mathcal{N}(0, \Sigma_\gamma)$ with density $\phi(\cdot; \Sigma_\gamma)$, as is common practice (see Berry et al., 1995, for example). In the case of truthful preferences, the likelihood is now

$$\mathbb{P}(R_i | \mathbf{x}_j, \mathbf{z}_i; \theta, \Sigma_\gamma) = \int \mathbb{P}(R_i | \mathbf{x}_j, \mathbf{z}_i; \theta) \phi(\gamma; \Sigma_\gamma) d\gamma. \quad (7)$$

An analogous change is required when the analysis is conducted assuming that only final matches are observed and are stable.

¹⁵ Relying on rank-ordered data requires a specific model of student behavior. Artemov et al. (2021) argue that relying on allocation stability yields results that are robust to mis-specifications of the model of behavior that could bias approaches that rely on reports.

A challenge with specifications that include random coefficients is that closed-form expressions for the probabilities are not typically available. Estimation techniques for these models typically require simulation, even in the simpler discrete choice context. Provided that the number of random coefficients is small, this expression can be approximated by numerical integration or simulation methods. However, approximation error in this integral can result in bias in the final estimates if the objective function is non-linear in the approximation error. We refer the reader to Train (2009) for recommendations and results on simulation-based estimators.

While the logit model has closed-form expressions when reports are truthful, it is not tractable for manipulable mechanisms. Next, we discuss an alternative parametrization, based on the probit model, that is useful in this case.

Probit models

Another popular approach is based on the probit model with random coefficients. In this model, we specify

$$u_{ij} = \delta_j + \mathbf{x}_j \bar{\gamma} z_i + \mathbf{x}_j \gamma_i - d_{ij} + \varepsilon_{ij}, \quad (8)$$

where

$$\gamma_i \sim \mathcal{N}(0, \Sigma_\gamma) \quad \text{and} \quad \varepsilon_{ij} \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

The model parameters can be estimated using a Markov Chain Monte Carlo (MCMC) technique called a Gibbs sampler with an appropriate conjugate prior distribution for the parameters θ . This method generates a Markov chain by iterating between drawing the model parameters (including the random coefficients) conditional on simulated utilities u_{ij} and drawing the utilities u_{ij} conditional on the parameters θ and the observed assignments or reports.¹⁶ That is, the sampler iterates through the following three steps:

$$\begin{aligned} & u_{ij} | \mathbf{u}_{i,-j}, \gamma_i, \theta; \text{data for each } j \\ & \gamma_i | \mathbf{u}_i, \theta; \text{data} \\ & \theta | \{\mathbf{u}_i, \gamma_i\}_{i=1}^N; \text{data}, \end{aligned}$$

where $\mathbf{u}_{i,-j} = (u_{i1}, \dots, u_{ij-1}, u_{ij+1}, \dots, u_{iJ})$. The procedure can begin from an arbitrary value of the parameters and an initial value of \mathbf{u}_i that is consistent with the reported preferences. Iterating through this procedure results in a long sequence of simulated draws known as a Markov chain. A portion from the beginning of the chain is discarded and the distribution of the remainder can be used to compute both the point estimates and credible sets (Bayesian variants of confidence intervals) simultaneously. This Bayesian technique yields estimates that are asymptotically equivalent to the maximum likelihood estimator (see the Bernstein-von Mises Theorem. van der

¹⁶ Drawing the utilities in this second step is known as data augmentation.

Vaart, 2000, Theorem 10.1).¹⁷ These features have made the probit model popular for discrete choice models in the marketing literature.

The second and third steps are standard and identical across the models discussed above. Specifically, assuming a normal prior on θ , γ_i is also normally distributed given the conditional distributions above. Within θ , the coefficients $\bar{\gamma}$ and δ_j are normally distributed if their prior distributions are normal, and the covariances Σ and Σ_γ follow inverse-Wishart distributions.

The first step draws the utility of each school conditional on the current draws of the utilities of the other schools and the observed rank-order list or outcome. The model assumptions imply that $u_{ij} | \mathbf{u}_{i,-j}, \gamma_i, \theta; data$ has a truncated normal distribution with truncation points determined by the observed report. The truncation points are determined based on whether preferences are truthfully reported or whether the mechanism is manipulable. We refer the reader to Abdulkadiroglu et al. (2017b) for further details on estimating models with truthful reporting and to Agarwal and Somaini (2018) for the case of manipulable mechanisms.

3.1.3 Analysis with data on final outcomes

This section reviews approaches when we only observe data on the realized matches from a single large matching market. Throughout, we will assume that these matches are stable. The main difficulty in learning about preferences is that the final assignments depend on the preferences of agents on both sides of the market. Accordingly, we must disentangle two sets of preferences from only observing one set of matches.

In this more limited data environment, we will need to strengthen the independence assumption in Eq. (4) to

$$(\eta_i, \epsilon_i) \perp (\mathbf{d}_i, \mathbf{w}_i) | z_i, \mathbf{x}, (\xi_j)_{j=1}^J, \quad (9)$$

where $\mathbf{w}_i = (w_{1i}, \dots, w_{Ji})$. Therefore, the unobservable on both sides of the market are independent of the preference shifters.

We will distinguish between two types of markets. The first type is a continuum many-to-one matching model with a large number of agents on one side a few agents on the other. In this type of market, each agent in \mathcal{J} can match with many agents in \mathcal{I} . We assume that I is large but J is small. The most common example of this type is student assignment to schools or colleges. The second type of market has a large number of agents on each side.

Continuum many-to-one matching

Consider settings in which agents on side \mathcal{J} can match with a large number of agents on the side \mathcal{I} , while the number of agents on \mathcal{J} is small. Because there are many agents on side \mathcal{I} , each individual agent on this side is strategically and outcome irrelevant for others. However, their aggregate preferences and strategies may influence

¹⁷ We refer the reader to Gelman and Rubin (1992) for a textbook treatment of Gibbs sampling and to McCulloch and Rossi (1994) for a discussion more specific to discrete choice models.

outcomes. School and college admissions are key examples of this setting type. We will therefore refer to agents on side \mathcal{I} as students and agents on side \mathcal{J} as schools.

There are two relevant types of data in these settings. The first is known preferences or priorities used by schools to admit students, so that the researcher can directly ascertain how two students will be ranked, possibly up to a random tie-breaker. For example, many school districts prioritize students in their walk-zone and students who have siblings already enrolled, and many college systems prioritize students with high school grades or entrance exam scores. When this type of data is available, the researcher need only estimate the students' preferences for the schools. When such data are inaccessible, the researcher must estimate preferences on both sides of the market, which is more challenging. This case is relevant to college admissions systems and entry-level job settings in which the rules used by agents on side \mathcal{J} are unknown.

In both cases, we consider the problem when only data on final matches are available, assuming that pairwise stability is satisfied. As before, this assumption requires justifications based on theory and institutional background on the process used in the market to assign students to schools. The main implication of the assumption is that the stable matches can be characterized by a cutoff rule. Recall the result from Azevedo and Leshno (2016) that, in a stable match, each student i is assigned to her most preferred school in the set

$$S(\mathbf{v}_i; \mathbf{p}) = \{j : v_{ji} \geq p_j\},$$

where $\mathbf{v}_i = (v_{1i}, \dots, v_{Ji})$ and $\mathbf{p} = (p_1, \dots, p_J)$ is the vector of cutoffs. We will now use this formulation to learn about student preferences in two cases. The first is when v_{ji} and \mathbf{p} are data and the second is when these quantities need to be estimated.

Known priorities (school choice): Suppose the researcher knows each student's eligibility score for each school, denoted v_{ji} , up to a tie-breaker, and the final assignment is stable. That is, $v_{ji} = v_j(\mathbf{z}_i)$ where the function $v_j(\cdot)$ is known and \mathbf{z}_i is observed. The cutoff scores p_j can be computed as the lowest eligibility score v_{ji} of a student who was matched to school j if the school does not have available capacity. Otherwise, the cutoff p_j is equal to 0. The goal is then to estimate and identify the specification of preferences defined in Eq. (1).

This model is used by Fack et al. (2019) to study Parisian high school admissions, which are determined by a deferred acceptance mechanism, and by Akyol and Krishna (2017) to study Turkish high schools that use an entrance exam to make admissions decisions. This assumption can also be used to study higher education settings that use an entrance exam. For example, Bordon and Fu (2015) and Bucarey (2018) uses stability to estimate preferences for colleges in Chile.

To see what can be learned with this information and the final assignments, consider the case with only two schools, 1 and 2, and an outside option, 0. Fig. 4 shows five regions of utilities denoted by Roman numerals. Each region implies different ordinal preferences except for region V, which pools the cases when $u_{i0} > u_{i1} > u_{i2}$ and $u_{i0} > u_{i2} > u_{i1}$. A student who is eligible for both schools will be assigned to

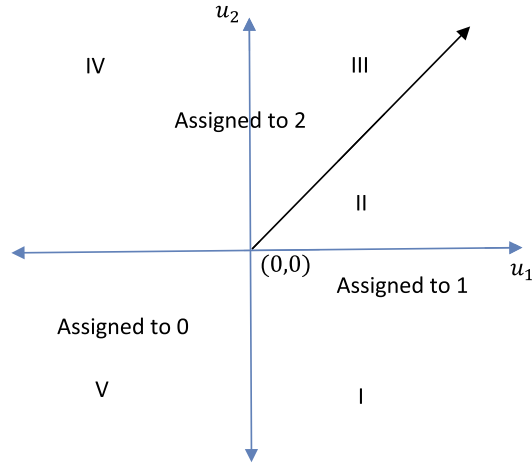


FIGURE 4 Stability – Both schools are feasible.

school 1 if her utilities belong to either region I or II. Therefore, the share of students assigned to school 1 amongst those eligible for both schools is an estimate of the total probability mass of the distribution of utilities in regions I and II. Similarly, the share assigned to school 2 is an estimate of the total probability mass in regions III and IV.

A student eligible only for school 1 can either be assigned to that school or remain unassigned. In the former case, we can infer that $u_{i0} < u_{i1}$ which is the darkly-shaded region in Fig. 5. In the latter case, we infer $u_{i1} < u_{i0}$ which is shaded lightly. The share of students assigned to school 1 amongst these students is an estimate of the total probability in regions I, II, and III of Fig. 4.

These arguments are similar to those for standard consumer choice models but differ crucially in that not all students are assigned to their first-choice school. In this context, a student's choice set is constrained by her eligibility. Thus, observed assignments provide no information about preferences for schools that are not in a student's choice set. Learning about the full distribution of ordinal preferences for students with a vector of eligibility score \mathbf{v}_i will require extrapolation using data from students with larger choice sets. Fack et al. (2019) perform this extrapolation by assuming that the unobserved determinants of preferences in Eq. (1) are conditionally independent of eligibility given the observables included in the model. Formally, they require that

$$\epsilon_i \perp \mathbf{v}_i \mid \mathbf{z}_i, \mathbf{d}_i, \{\mathbf{x}_j, \xi_j\}_{j=1}^J. \quad (10)$$

This assumption may be a reasonable approximation if \mathbf{z}_i contains a rich set of student characteristics but can be violated, for example if eligibility scores are correlated with both unobserved student ability and unobserved preference parameters.




FIGURE 5 Stability – Only one school is feasible.

Under this assumption, the probability of each observed assignment can be used to construct a likelihood function given a parametrization of utilities. Specifically, let F_{U^*} denote the joint cdf of the random vector \mathbf{u}_i^* with the j -th element equal to $u(\mathbf{x}_j, \mathbf{z}_i, \xi_j, \epsilon_i)$. We will drop the conditioning on $\mathbf{z}_i, \{\mathbf{x}_j, \xi_j\}_{j=1}^J$ for notational simplicity. The independence assumptions in Eqs. (4) and (10) obviate the need to condition on \mathbf{d}_i and \mathbf{v}_i . Under this assumption, the probability that i is assigned to j given the parameter F_{U^*} can be written as

$$\begin{aligned} \mathbb{P}(\mu(i) = j | \mathbf{v}_i = \mathbf{v}, \mathbf{p}, \mathbf{d}_i = \mathbf{d}; F_{U^*}) \\ = \int 1 \left\{ u_j^* - d_j \geq u_{j'}^* - d_{j'} \text{ for all } j' \in S(\mathbf{v}; \mathbf{p}) \right\} dF_{U^*}. \end{aligned}$$

This expression enables estimation via maximum likelihood or other likelihood-based methods. Both the logit and the probit models are amenable for estimation. Further details are discussed in Agarwal and Somaini (2020).

Additionally, this expression shows that the preference shifter \mathbf{d} plays a crucial role in identification. Under our assumptions, \mathbf{d} changes the desirability of each school exogenously and consequently alters the schools to which students are assigned. This source of variation provides a wealth of information about agents' preferences. Consider the probability that $\mu(i) = 0$, which is equal to the probability that $\mathbf{u}_i^* - \mathbf{d}_i$ belongs to region V in Fig. 4. This probability is identified in the two-school case if both schools are in the choice set or if the assumption (10) above holds. It is equal to:

$$\mathbb{P}(\mu(i) = 0 | \mathbf{d}_i = \mathbf{d}) = \mathbb{P}(\mathbf{u}_i^* - \mathbf{d} \leq 0) = F_{U^*}(\mathbf{d}).$$

Thus, we identify $F_{U^*}(\mathbf{d})$ by the share of students in region V for $\mathbf{d}_i = \mathbf{d}$. Variation in \mathbf{d} allows us to identify F_{U^*} evaluated at different values. Finally, Eqs. (1) and (4)

imply that the joint cdf of $\mathbf{u}_i = (u_{i1}, \dots, u_{iJ})$ conditional on \mathbf{d} is given by $F_{U|\mathbf{d}}(\mathbf{u}) = F_{U^*}(\mathbf{u} + \mathbf{d})$, implying that the former is nonparametrically identified.¹⁸

Unknown priorities (college admissions): We now consider the implications of stability in many-to-one matching environments where preferences on both sides of the market have to be estimated. We will use college admissions as an example. The empirical challenge is not limited to estimating preferences for colleges. The revealed preference arguments for students that we derived in the school choice context above are not possible here, because college preferences are unknown.

Nonetheless, a considerable amount of information is available in the matches. Consider the simple case of $J = 2$. If student i is observed attending college $j = 1$, then we can deduce the following:

- Student i prefers college 1 to remaining unassigned: $u_{i1} \geq 0$.
- Student i clears the threshold for college 1: $v_{1i} \geq p_1$.
- Either student i prefers college 1 to college 2, or student i does not clear the threshold for college 2: $u_{i1} \geq u_{i2}$ or $v_{2i} < p_2$.

These restrictions define a set in a four-dimensional space that rationalizes the allocation of student i to college 1.

Agarwal and Somaini (2021a) show how to learn about preferences on both sides of the market simultaneously for the model described by Eqs. (1) and (3). In the model discussed in section 3.1.2, variation in \mathbf{d}_i is used to identify the joint distribution of the J -dimensional vector of students' preferences (u_{i1}, \dots, u_{iJ}) . Similarly, exogenous variation in \mathbf{d}_i and \mathbf{w}_i can be used to nonparametrically identify the joint distribution of the $2J$ dimensional vector $(u_{i1}, \dots, u_{iJ}, v_{1i}, \dots, v_{Ji})$ conditional on all observables, up to appropriate scale and location normalizations. A closely related prior argument in He et al. (2021) shows a similar result under more stringent restrictions on Eq. (1).¹⁹

This joint distribution allows for a host of economic phenomena based on unobservable factors. For example, correlation between u_{ij} and $u_{ij'}$ implies that colleges j and j' are closer substitutes, i.e. students who like one tend to also like the other; correlation between v_{ji} and $v_{j'i}$ suggests that colleges j and j' tend to prefer the same set of students; and correlation between u_{ij} and v_{ji} suggests that students tend to like colleges that also like them.

A detail about the location normalization in this model is worth noting. As before, it is possible to normalize u_{i0} to zero for all i and v_{j0} to zero for all j . Moreover,

¹⁸ It is also possible to develop the same identification argument using any other region in Fig. 4. We choose region V because it is the negative orthant, which results in simpler expressions. Therefore, this model is over-identified.

¹⁹ Specifically, He et al. (2021) assumes that $u_{ij} = u(\mathbf{x}_j, \mathbf{z}_i, \xi_j) - d_{ij} + \epsilon_{ij}$ and $v_{ji} = v(\mathbf{x}_j, \mathbf{z}_i, \xi_j) - w_{ji} + \eta_{ij}$, whereas Agarwal and Somaini (2021a) can work with the general case in which $u_{ij} = u(\mathbf{x}_j, \mathbf{z}_i, \xi_j, \epsilon_i) - g(d_{ij})$ and $v_{ji} = v(\mathbf{x}_j, \mathbf{z}_i, \xi_j, \eta_i) - w_{ji}$ for a general function $g(\cdot)$.

if the researcher has information on the capacity of each college, then it is possible in principle to learn the distribution of v_{ji} for a college that does not fill its seats. This inference is based on students who have characteristics \mathbf{d}_i that indicate strong preferences for college j , but did not attend the college and therefore must have been unacceptable to the college. Unfortunately, it is not possible to use a similar reasoning for colleges that do not have spare capacity. For students who strongly prefer college j but were not admitted, we can only deduce that $v_{ji} < p_j$, and we cannot determine the location of v_{ji} because p_j is not observed. One alternative is to set $p_j = 0$ and to treat these colleges in the same way as those with spare capacity, in order to obtain the distribution of v_{ji} . However, the data are also consistent with any $p_j > 0$ and a distribution of v_{ji} that is shifted by p_j . This ambiguity prevents us from identifying the location parameter of v_{ji} when capacity is not known or when we know that capacity limits are binding. Either case identifies the distribution of the difference $v_{ji} - p_j$.

Methods for estimating this model are a subject of ongoing research. He et al. (2021) and Agarwal and Somaini (2021a) propose a method based on Gibbs sampling and a probit model. In principle, a simulated minimum distance estimator similar to the one used in Agarwal (2015) offers another approach. The interested reader should consult these references for further details.

One-to-one or few-to-one matching

We now consider a model in which the number of agents on both sides is large. Agents on side \mathcal{J} may match with more than one agent in some models, but the number that match with the other side is finite. Thus, as opposed to the continuum model, each agent is now strategically relevant for others in the market. This feature of the model complicates the analysis because cutoffs that do not depend on unobservables cannot be used to simplify the problem.²⁰ This version of the model emphasizes the issues that arise from such strategic interactions and was the initial approach taken by the literature.

This model can be analyzed in two ways. The first is based on the canonical single index model (e.g. Becker, 1973) in which each side of the market is differentiated only by a vertical quality index. The second is when preferences are heterogeneous so that two agents may have differing preferences over agents on the other side of the market. We discuss both approaches below.

Double-vertical preferences: In this model, all agents on one side of the market share the same preferences over all agents on the other side. In our notation, the utility of agent i from matching with j is

$$u_{ij} = u_j = u(\mathbf{x}_j) + \xi_j,$$

where we replace the assumption in Eq. (4) with $\xi_j \perp \mathbf{x}_j$. This model omits both observed and unobserved sources of quality heterogeneity, resulting in a desirability

²⁰ Although we did not explicitly write it as such, in the continuum model, equilibrium cutoffs are only a function of the population distribution of preferences and the mass of agents.

index for each agent j denoted by u_j . The term $u(\mathbf{x}_j)$ is the component explained by observables \mathbf{x}_j , and ξ_j the unobserved component. The preferences on the other side of the market are analogous:

$$v_{ji} = v_i = v(\mathbf{z}_i) + \eta_i,$$

where $\eta_i \perp \mathbf{z}_i$. The location for utilities is normalized by either setting the value of the outside option to 0 or picking an arbitrary value \bar{x}_j and setting $u(\cdot)$ to zero at that value. Because the model does not have a quasilinear term d_{ij} for normalizing the scale, we also set the slope of $u(\bar{\mathbf{x}}_j)$ with respect to one of its components to one. The normalization on the other side of the market is analogous.

Chiappori et al. (2012) analyze a one-to-one matching model with these preferences. They assume that the researcher has access to data on the agents' observable characteristics in a matching market. This approach therefore observes the joint distribution $F_{X,Z}$ of matched agents' observable characteristics. Here we follow their convention in referring to side \mathcal{I} as men and to side \mathcal{J} as women.

In this model, a match is stable if and only if it exhibits perfect assortative matching on u_j and v_i . In such a set of matches, the t -th most desirable man matches with the t -th most desirable woman. Therefore, if F_U and F_V are the cumulative distribution functions of u_j and v_i , respectively, then an agent with characteristics (\mathbf{x}_j, ξ_j) is matched with an agent with characteristics (\mathbf{z}_i, η_i) only if:

$$u(\mathbf{x}_j) = F_U^{-1}(F_V(v(\mathbf{z}_i) + \eta_i)) - \xi_j. \quad (11)$$

Now, consider two men i and i' with identical values of the observed index, $v(\mathbf{z}_i) = v(\mathbf{z}_{i'})$. These two men could have different values of η , and therefore their mates may differ. However, if we consider two populations of men, one with observed characteristics \mathbf{z}_i and the other with observed characteristics $\mathbf{z}_{i'}$, then the distribution of their desirability to women including the η terms will be identical. Thus, the two populations of men will have the same marriage prospects and the women they match with will have the same distribution of observed characteristics.

In the terminology employed by Chiappori et al. (2012), this reasoning allows us to identify "iso-attractiveness profiles" for men by looking at which vertical types end up matching with women with the same distributions of observable characteristics. The same reasoning allows us to identify iso-attractiveness profiles for women. Chiappori et al. (2012) posits v_i as depending on observable characteristics. Formally, they show that for any function ϕ_x of observables \mathbf{x}_j there exists a function ϕ_v of the index v so that $\mathbb{E}(\phi_x(\mathbf{x}_j) | \mathbf{z}_i) = \phi_v(v(\mathbf{z}_i))$, where $\mathbb{E}[\cdot]$ is the expectation operator. The left-hand side is observable, and the right-hand side is a composition of two unknown functions. Differentiating both sides with respect to two components of \mathbf{z}_i , we can measure the following marginal rate of substitution

$$\frac{\partial v(\mathbf{z}_i) / \partial z_{i,k}}{\partial v(\mathbf{z}_i) / \partial z_{i,l}} = \frac{\partial \mathbb{E}(\phi_x(\mathbf{x}_j) | \mathbf{z}_i) / \partial z_{i,k}}{\partial \mathbb{E}(\phi_x(\mathbf{x}_j) | \mathbf{z}_i) / \partial z_{i,l}}$$

because the right-hand side is observed.

A natural question to ask is whether it is also possible to sort the level curves according to their desirability level. Since the argument above only provides the ratio of derivatives, there is no a priori way to know if desirability is increasing in any specific component. It is therefore necessary to assume there is a characteristic that is known to be valued monotonically and is desirable.

A limitation of the argument described above is that we are only able to assess the relative importance of two different components of the observables. In other words, the marginal rate of substitution between $x_{j,k}$ and $x_{j,k'}$ can be determined for any k and k' , but we cannot determine the marginal rate of substitution between $x_{j,k}$ and ξ_j . More broadly, it is not possible to determine the contribution of the observables on either side to the overall variation in preferences.²¹

One conjecture is that it is not possible to identify preferences on both sides of the market in a one-to-one matching market. Diamond and Agarwal (2017) prove this conjecture for double-vertical preferences. As argued above, it is possible to learn the functions $u(\cdot)$ and $v(\cdot)$ under mild restrictions. However, if there are unobserved determinants of preferences on either side of the market, then the matching will not be perfectly assortative in these indices. This is because the match is assortative on $u_j = u(x_j) + \xi_j$ and $v_i = v(z_i) + \eta_i$, not only on the components $u(x_j)$ and $v(z_i)$ that can be predicted by observables. However, the data can be rationalized by either setting $\xi_j \equiv 0$ for all j or $\eta_i \equiv 0$ for all i . This result follows because the double-vertical model only places a single restriction expressed in Eq. (11), but there are two unobservables in the model, ξ_j and η_i . In other words, the matches are governed by unobserved determinants of preferences on both sides of the market, making them hard to disentangle.

Diamond and Agarwal (2017) go on to show that this problem can be solved in many-to-one matching markets, since a setting in which each agent j can match with multiple agents i on the other side has significantly more information than a market with one-to-one matching. An iconic example is the National Residency Matching Program, which uses a variant of the deferred acceptance algorithm (Roth and Peranson, 1999). While each resident is assigned to at most one program, each program can match with several residents. The number of residents to which each program j is matched can be as low as two. For this reason, we term such markets few-to-one matching markets.

As before, if preferences on both sides are vertical, matches are stable if and only if they exhibit perfect assortative matching. In other words, in any stable match, the most preferred residents are allocated to the most preferred hospital until its vacancies are filled. The second most preferred hospital takes the most preferred remaining residents and so on.

²¹ Observe, however, that the polar cases when η_i and ξ_j are both identically equal to zero for all i and j can be ruled out. This is because Eq. (11) above reduces to $u(x_j) = F_U^{-1}(F_V(v(z_i)))$. In this case, men with a given set of characteristics z_i match with women whose observables lie exactly on the same iso-attractiveness curve.

More formally, in such a market, consider a pair of residents i and i' matched to the same hospital j . Eq. (11) generalizes to:

$$\begin{aligned} u(z_i) &= F_U^{-1}(F_V(v(x_j) + \xi_j)) - \eta_i \\ u(z_{i'}) &= F_U^{-1}(F_V(v(x_j) + \xi_j)) - \eta_{i'} \end{aligned} \quad (12)$$

Similarly to the marriage market problem, the lack of perfect sorting based on observables indicates the presence of the errors η_i and ξ_j . However, the composition of the incoming cohort in each program provides additional information about each error term's contribution. The expressions in Eq. (12) suggest that dispersion in the η terms, the unobserved shocks affecting residents' desirabilities, will cause a program to admit residents with heterogeneous observable determinants of human capital. Thus, the unobservables η_i contribute to the variance in the observable characteristics of residents within each program.

This model can be estimated using a simulated minimum distance estimator (Agarwal, 2015; Diamond and Agarwal, 2017), which consists of the following steps. First, define a set of moments in the data m to be matched with our model. Second, fix a vector of parameters $\theta = \{\beta, \gamma, \sigma_\eta, \sigma_\xi\}$ for the model and use them to simulate stable matches and obtain a simulated set of moments $m(\theta)$ as a function of the parameters. Third, compute the distance between the simulated moments and the moments observed in the data, e.g. $\|m - m(\theta)\|_W = \sqrt{(m - m(\theta))' W (m - m(\theta))}$. Fourth, search over θ to minimize the distance.²²

Agarwal (2015) uses three sets of moments for estimation. The first set of moments summarizes the general sorting patterns of residents across programs. Recall that x_j and z_i are column vectors; thus, $x_j z_i'$ is a matrix. Averaging this matrix over all matches yields:

$$\frac{1}{I} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} 1\{\mu(i) = j\} x_j z_i'$$

The second set of moments computes the within-program variances of resident observables for each component $z_{i,\ell}$ of z_i :

$$\frac{1}{I} \sum_{i \in \mathcal{I}} (z_{i,\ell} - \bar{z}_{i,\ell})^2,$$

where \bar{z}_i is the vector of average characteristic values of i 's peers, that is, of residents matched with the same program. The third set of moments computes the correlation between residents' characteristics and the average characteristics of the residents' peers for each set of components $z_{i,\ell}$ and $z_{i,k}$ for $k \neq \ell$:

$$\frac{1}{I} \sum_{i \in \mathcal{I}} z_{i,\ell} \hat{z}_{i,k},$$

where \hat{z}_i is the average characteristics of i 's peers excluding i .

²² Train (2009) provides an overview of best practices.

The first set of moments summarizes the same type of information about the allocation as regressions of an individual's characteristic on those of her match partner (see Chiappori et al., 2012, for example).²³ Both summarize the aggregate sorting patterns based on observable characteristics. The second and third sets of moments include additional information that is required to identify each error term's contribution.

An open question is whether it is possible to relax the assumptions of the double vertical model. Agarwal (2015) assumes that programs have vertical preferences over doctors, but doctors' preferences over programs include some horizontal components. For example, doctors are more likely to be allocated to programs in their birth state or medical school state. This feature cannot be rationalized by a double-vertical model. Instead, it indicates a resident's geographical preference for training close to home.

Heterogeneous preferences: To make progress on allowing for heterogeneity in preferences, Menzel (2015) considers a model in which the utilities are parametrized as follows:

$$\begin{aligned} u_{ij} &= u(\mathbf{x}_j, \mathbf{z}_i) + \varepsilon_{ij} \\ v_{ji} &= v(\mathbf{x}_j, \mathbf{z}_i) + \eta_{ji}, \end{aligned}$$

$u_{i0} = 0 + \max_{k=1, \dots, J} \{\varepsilon_{i0,k}\}$, and $v_{j0} = 0 + \max_{k=1, \dots, J} \{\eta_{j0,k}\}$. The error terms ε_{ij} , $\varepsilon_{i0,k}$, η_{ji} , and $\eta_{j0,k}$ are independent and identically distributed with an upper tail that is of type I. The paper considers the limit of a sequence of economies indexed by J with an equal number of agents on each side and J growing large. Notice that the outside option also becomes more attractive as J increases. This choice is made in order to make sure that the remaining unmatched stays attractive even in a large market with many draws of ε_{ij} and η_{ji} .

Under these assumptions, Menzel (2015) shows that the limiting probability density function of the agent types matched with each other, denoted $f(\mathbf{x}, \mathbf{z})$, has a very tractable functional form. Specifically,

$$\log \frac{f(\mathbf{x}, \mathbf{z})}{f(\mathbf{x}, *)f(\mathbf{z}, *)} = \exp(u(\mathbf{x}, \mathbf{z}) + v(\mathbf{x}, \mathbf{z})),$$

where $f(\mathbf{x}, *)$ is the density of agents on side \mathcal{I} remaining unmatched, and analogously, $f(\mathbf{z}, *)$ is the density of agents on side \mathcal{J} remaining unmatched. This convenient functional form is derived from the core insight that if some set \mathcal{J}_i is willing to match with agent i , then the probability i gets matched with j is the probability that j is i 's most preferred option in the set \mathcal{J}_i . Similarly, i must be j 's most preferred option. These probabilities are given by a logit-like formula in a large market

²³ Sorting patterns and simulation-based estimation methods have also been used in Boyd et al. (2013) to estimate teachers' preferences for working at various schools. Although Boyd et al. (2013) have access to data from many-to-one matches, they do not use this information to construct the latter two sets of moments. As a result, their approach may be susceptible to the non-identification issues discussed above.

and are therefore proportional to $\exp(u(\mathbf{x}, \mathbf{z}))$ and $\exp(v(\mathbf{x}, \mathbf{z}))$ for sides \mathcal{I} and \mathcal{J} , respectively. Hence, $f(\mathbf{x}, \mathbf{z})$ is proportional to the product $\exp(u(\mathbf{x}, \mathbf{z}) + v(\mathbf{x}, \mathbf{z}))$. The probability of remaining unmatched provides the right normalizing constant.

Another approach, from Sorensen (2007), is to assume that matches depend only on the joint surplus $S(\mathbf{x}_j, \mathbf{z}_i, \varepsilon_{ij}, \eta_{ji}) = u(\mathbf{x}_j, \mathbf{z}_i) + v(\mathbf{x}_j, \mathbf{z}_i) + \varepsilon_{ij} + \eta_{ji}$, but the partners split this surplus via Nash bargaining after the match is formed. That is, side \mathcal{I} receives a fraction $\lambda S(\mathbf{x}_j, \mathbf{z}_i, \varepsilon_{ij}, \eta_{ji})$ from a realized match and side \mathcal{J} receives $(1 - \lambda) S(\mathbf{x}_j, \mathbf{z}_i, \varepsilon_{ij}, \eta_{ji})$ for some $\lambda \in [0, 1]$. Using the terminology of Niederle and Yariv (2009), this model exhibits aligned preferences, resulting in a unique pairwise stable match. Sorensen (2007) uses a Bayesian approach to estimate the joint surplus in the market for venture capital, targeting the joint surplus function $S(\mathbf{x}_j, \mathbf{z}_i, \varepsilon_{ij}, \eta_{ji})$ directly.

These results suggest a different limitation of data from one-to-one matches, this time in a model with heterogeneous preferences. Namely, only the sum of the surplus on the two sides of the market is identified.²⁴ The difficulty lies in trying to determine whether the preferences on side \mathcal{I} or on side \mathcal{J} are driving the observed matches.

3.2 Transferable utility models

We now turn to models of (perfectly) transferable utility. As a reminder, the equilibrium in these models govern the match between agents as well as the transfer. Such models have been used to analyze the marriage market, partnerships between car manufacturers and their suppliers (Fox, 2018), and spectrum auctions (Fox and Bajari, 2013). Empirical analysis in these settings is typically conducted only with data on final matches, usually with transfers that are not observed. The data include information on agents and their match partners. The goal is typically to estimate the joint surplus as a function of these characteristics. This surplus can be flexibly split between the match partners.

There are two canonical approaches for analyzing models of transferable utility.²⁵ The first approach is based on Choo and Siow (2006) and its generalizations (e.g. Galichon and Salanie, 2020). These models assume parametric forms of unobserved heterogeneity in payoffs that are additively separable. The second approach, based on Fox (2018), is semi-parametric. It considers deviations from the observed match to construct an estimator. The discussion below explores these two approaches.

3.2.1 Models with separable unobserved heterogeneity

Each agent is characterized by one of many discrete types. For simplicity, we will refer to these sides as workers and firms and assume that each firm can hire only

²⁴ In early work, Logan et al. (2008) used a Bayesian approach to estimate both men and women's heterogeneous preferences regarding their partners. However, we are not aware of results that show identification of this model.

²⁵ We refer the reader to Chiappori and Salanié (2016) and Galichon and Salanie (2021) for surveys of this literature, which discusses related approaches and issues specific to these models in greater detail.

one worker. Let $x_i \in \mathcal{X}$ denote the observed type of worker i . The set \mathcal{X} is finite so that each worker belongs to one of many discrete types. Similarly, let $y_j \in \mathcal{Y}$ be the observed type of firm j , with \mathcal{Y} finite. We assume there are infinitely many workers and firms of each type participating in a single market. We observe the match frequencies between various types of firms and workers, either in the population or through a large enough random sample.

The payoff to worker i from matching with firm j is given by

$$u_{ij} = \alpha_{x_i y_j} + t_{x_i y_j} + \varepsilon_{i y_j},$$

where α_{xy} is the systematic surplus a worker of type x receives from matching with a firm of type y , t_{xy} is the equilibrium transfer between a worker of type x and a firm of type y , and ε_{iy} is an idiosyncratic unobserved payoff to worker i from matching with a firm of type y . The value of remaining unmatched is given by

$$u_{i0} = \alpha_{x_i 0} + \varepsilon_{i0}.$$

Similarly, the payoff of firm j from matching with worker i is given by

$$v_{ji} = \gamma_{x_j y_i} - t_{x_j y_i} + \eta_{j x_i},$$

where γ_{xy} is the systematic surplus a firm of type y attains from matching with a worker of type x , and η_{jx_i} is an idiosyncratic unobserved payoff to firm j from matching with a worker of type x . The value of remaining unmatched is given by $v_{j0} = \gamma_{y_j 0} + \eta_{j0}$. Throughout, we assume that ε_{iy} and η_{jx} are independent and identically distributed. Since we are interested in estimating payoffs relative to remaining unmatched, we normalize α_{x0} and γ_{y0} to zero.

Since each worker is indifferent between matching with any two firms of the same type, given a vector of transfers $t = (t_{xy})_{x \in \mathcal{X}, y \in \mathcal{Y}}$, worker i would choose to match with a firm of type y only if

$$y \in \arg \max_{y'} (\alpha_{x_i y'} + t_{x_i y'} + \varepsilon_{i y'}).$$

Analogously, each firm would choose to match with a worker of the type that maximizes its payoff. The vector of transfer t is an equilibrium if the total number of workers of type x that demand a match with firms of type y is equal to the total number of firms of type y that demand a match with workers of type x . The restriction that the transfers depend only on the types of the workers and firms can be supported in equilibrium.

Choo and Siow (2006) studied the special case of this model when the unobserved payoffs ε_{iy} and η_{jx} each follow a type I extreme value distribution. Under this assumption, the demand for $x - y$ jobs from the worker side is given by

$$\ln \mu_{xy}^d = \ln \mu_{x0}^d + \alpha_{xy} + t_{xy},$$

where μ_{xy}^d is the share of $x - y$ jobs demanded in the economy, and μ_{x0}^d is the share of workers that demand unemployment at the transfer vector t . The resulting equation can be derived from the familiar logit model. Likewise, the supply of $x - y$ jobs from the firm side is given by

$$\ln \mu_{xy}^s = \ln \mu_{0y}^s + \gamma_{xy} - t_{xy},$$

where μ_{xy}^s is the share of $x - y$ jobs supplied and μ_{0y}^s is the share of unfilled jobs at firm types y . The market clears if $\mu_{xy} = \mu_{xy}^d = \mu_{xy}^s$. Summing the two equations, we get

$$2 \ln \mu_{xy} - (\ln \mu_{x0} + \ln \mu_{y0}) = \alpha_{xy} + \gamma_{xy} \equiv \Phi_{xy}.$$

Thus, the systematic part of total surplus Φ_{xy} can be identified from observing the shares of matches between each pair of types $(x, y) \in \mathcal{X} \times \mathcal{Y}$. In addition, the systematic gain a worker of type x acquires from matching with a firm of type y , given by $\alpha_{xy} + t_{xy}$, can be identified from the demand equation. Likewise, the gain a firm of type y receives from matching with a worker of type x can be identified from the supply equation.²⁶

Galichon and Salanie (2020) generalize this argument by accommodating other parametric assumptions on ε_{iy} and η_{jx} . The arguments are based on recognizing that a pairwise stable equilibrium of a transferable utility model maximizes the social surplus and then working with the dual version of the problem. The framework maintains the additive separability and independence of ε_{iy} and η_{jx} from the systematic component payoffs. This assumption allows the total social surplus function Φ_{xy} and the distributions of ε_{iy} and η_{jx} to be flexibly parametrized. We refer the reader to Galichon and Salanie (2021) for a more detailed review of related approaches.

This baseline model assumes that utility is perfectly transferable between match partners via t_{xy} . Galichon et al. (2019) extend this framework to include cases in which transfers are imperfect. This approach considers a feasible set of utilities for each match partner as a function of the underlying match values, α_{xy} and γ_{xy} . Indeed, this model incorporates important examples such as taxes or frictions that may result in loss when transfers are necessary.

3.2.2 Semi-parametric approaches

A line of research initiated in Fox (2010; 2018) develops a semi-parametric approach for estimating matching models with transferable utility that relaxes the distributional assumptions discussed above. We start with a one-to-one matching market for simplicity, say between upstream and downstream firms. Consider the problem faced by an analyst who observes the matches but not the transfers or underlying contract terms. The goal is to estimate the underlying surplus generated by the matches as a function of a set of observed characteristics.

²⁶ These quantities can be estimated either via maximum likelihood or by simply inverting for the quantities of interest using the equations above.

Assume the payoff that downstream firm i receives from matching with firm j is given by $\pi_{ij}^d - t_{ij}$, where t_{ij} is the transfer from the downstream firm to the upstream firm. Similarly, let $\pi_{ij}^u + t_{ij}$ be the profits the upstream firm j accrues from matching with the downstream firm i and receiving the transfer t_{ij} . Thus, the total surplus from a match between firms i and j is equal to $f_{ij} = \pi_{ij}^u + \pi_{ij}^d$. Normalize the profits of remaining unmatched to zero on both sides of the market for each firm.

As discussed in Section 2, a pairwise stable match maximizes the total surplus in the economy subject to feasibility constraints. Let μ be such a match, where $\mu_{ij} = 1$ if i is matched with j , and zero otherwise. Feasibility requires that $\sum_i \mu_{ij} \leq 1$ and $\sum_j \mu_{ij} \leq 1$ since each agent can match with at most one agent on the other side of the market. Since μ maximizes the total surplus, under any alternative feasible match μ' , it must be that $\sum_{ij} \mu'_{ij} f_{ij} \leq \sum_{ij} \mu_{ij} f_{ij}$. Suppose we observe that i is matched with j and i' is matched with j' . Because swapping the partners of i and i' cannot increase the total surplus, it must be that

$$f_{ij} + f_{i'j'} \geq f_{ij'} + f_{i'j}.$$

Another way to derive this inequality is to note that if i matches with j , then it must be that $\pi_{ij}^d - t_{ij} \geq \pi_{ij'}^d - t_{ij'}$ and $\pi_{ij}^u + t_{ij} \geq \pi_{i'j}^u + t_{i'j}$. An analogous pair of inequalities holds because i' matches with j' . Summing these four inequalities yields the version of interest since the transfers cancel.

It is worth noting that the inequality above depends only on the joint surplus, not the underlying transfers. This suggests we may be able to circumvent the data limitation of unobserved transfers. Fox (2018) uses this insight to make progress using a maximum score inequality akin to those developed in Manski (1975) for binary choice models. Specifically, let $x_{ij}\theta$ be an approximation for f_{ij} , where x_{ij} denotes a vector of observable characteristics that vary across partnerships. If there are no unobservables, then we would get that

$$x_{ij}\theta + x_{i'j'}\theta \geq x_{ij'}\theta + x_{i'j}\theta.$$

Based on this observation, Fox (2018) proposes maximizing the objective function

$$S(\theta) = \sum_{i=1}^{N-1} \sum_{i'>i}^N 1 \{x_{ij}\theta + x_{i'j'}\theta \geq x_{ij'}\theta + x_{i'j}\theta\}$$

in order to estimate θ . This objective function counts the total number of pairwise inequalities correctly predicted by a value of θ , a measure of statistical fit. The maximizer yields the least number of violations.²⁷

²⁷ We refer the reader to Fox (2010) for a more detailed discussion of the asymptotic properties of this estimator and an extension to multiple matching markets. The maximum score estimator has convergence rates that are slower than root- N .

One consideration in interpreting this approach is that no value of θ is likely to correctly predict all the inequalities. A common reason is that there are unobserved characteristics that influence payoffs. Thus, the question is whether a value of θ that maximizes $S(\theta)$ also maximizes a version with unobservable terms in the payoff. Graham (2011; 2014) shows that this property, known as the rank-order property, holds when payoffs are of the form $f_{ij} = x_{ij}\theta + \varepsilon_{ij}$ with independent and identically distributed ε_{ij} . Fox et al. (2018) extends this result to more general forms of unobserved heterogeneity, relaxing the additive separability of error terms.

As is clear from the maximum score inequality, the approach allows us to estimate the effects of interactions between the characteristics of a pair of agents. In particular, coefficients on any component of x_{ij} that only varies across i or only across j cannot be identified.

This basic approach can allow for many-to-many matching that involves trading networks in Fox (2018), which builds on equilibrium models in a continuum economy (Azevedo and Hatfield, 2018). Variants of the approach have been analyzed to study the efficiency of FCC spectrum auctions (Fox and Bajari, 2013) as well as mergers (Akkus et al., 2016).

3.3 Extensions

Dynamics

The models described so far in this chapter are static, in that all agents arrive to the market simultaneously and match once and for all. There are many cases in which matching occurs over time. For example, in the markets for child care, public housing, and organ transplants agents or units on one side of the market arrive over time, while agents on the other side can wait. These allocation systems often use waitlists to prioritize agents on the waiting side. Agents on the waiting side have to decide whether to match with a unit that has just arrived. These decisions reveal agents' preferences in the same way that reports do for static allocation systems. Leveraging this insight, Agarwal et al. (2021) studied the system that allocates deceased donor kidneys in the U.S.; Waldinger (2021) analyzed public housing allocation; Verdier and Reeling (2021) analyzed bear-hunting licenses.

In other cases, both sides of the market have preferences and match over time. For example, Gandhi (2020) studies the market for nursing homes in which patients who need long-term care arrive over time and nursing homes decide whether to admit them or to hold a spot for a more profitable patient in the future. Similar considerations are important in Liu et al. (2021), who studies a peer-to-peer ride-sharing market in which drivers decide whether or not to give a passenger a ride or wait for the next passenger.

Peer effects

A central assumption in the models described above is that each agent has preferences only over who they match with. There is relatively little work on externalities, whereby others' matches also affect an agent's payoffs. One salient example oc-

curs when agents have preferences over their peer group. For example, students may derive utility from their classmates, and workers may have preferences over their co-workers. In this case, i 's preference for j can depend on the set $\mu^{-1}(j)$. Some work on education markets has addressed these issues (e.g. Epple et al., 2018; Allende, 2019). The typical approach here is to assume preferences for aggregate statistics of the composition of the student body in equilibrium.

Complex preferences

The methods described in this section mostly are developed for environments where agents require a single unit (e.g., a school placement, a job, a marriage partner,) or agents' preferences over multiple units are additive separable (e.g., a hospital hiring multiple residents). In the context of auctions, there is now an active literature on preference estimation in environments where bidders have more complex preferences over bundles of objects, such as substitutes and complements (see Cantillon and Pendorfer, 2006; Reguant, 2014, for example). A similar estimation problem presents itself in matching and assignment markets with multi-unit demand. A particularly challenging version of this problem presents itself in the combinatorial assignment context, where the choice set and therefore the space of preference types is high-dimensional (Budish and Kessler, 2021). Empirical approaches will need to consider both the problem of representing and estimating such preferences. As discussed in Section 4.3, this issue is also related to eliciting such preferences in practice.

Competition

One important reason why agents may care about matches other than their own is competitive effects, which are particularly important in settings concerning IO. For example, Uetake and Watanabe (2020) model an entry game in the banking industry using the tools of two-sided matching games in which a bank can enter a market by merging with an incumbent. In this entry game, payoffs are affected by the competitor banks that match in the market. Similarly, Vissing (2018) models the market for oil drilling leases as a matching game between oil companies and landlords who hold mineral rights. In this model, the terms that an oil company can negotiate depend on their overall market presence. A challenge in these settings lies in finding an appropriate notion of stability that allows for externalities.²⁸

A different approach to the formation of partnerships has been taken in IO and the study of vertical relationships in particular. Like in models of matching with transferable utility, these models analyze the relationships between upstream and downstream firms. Externalities, due to competition, are typically modeled through product market competition. Examples include the formation of networks between health insurance providers and hospitals (Ho, 2009), and relationships between content

²⁸ Conditions for the existence of stable matchings with externalities are an active area of theoretical research. We refer the reader to Pycia and Yenmez (2021), Fisher and Hafalir (2016), and references within for some recent results.

producers and cable-television service networks (Crawford and Yurukoglu, 2012). Chapter 9 provides a more thorough review of this literature.

4 Designing markets

This section discusses three primary goals in market design research. The first is to diagnose potential market failures and identify their causes. The second, complementary goal, is to evaluate markets and compare alternative designs in an effort to understand the qualitative and quantitative effects of various designs. The third goal is to propose new, improved designs, building on what we have learned from the first and second goals.²⁹ The discussion below is broader than the more specific focus on matching markets that was necessary to keep Sections 2 and 3 of this chapter self-contained.

4.1 Diagnosing market failures

Market design researchers often emphasize that seemingly small details about market rules can have a large impact on the ultimate performance of a market (see, for example, surveys by Roth (2002) and Klemperer (2002)). Put another way, flawed market rules can lead to market failures. In this sub-section, we give readers a brief tour of the kinds of market failures that have been documented. For each type of market failure, we give one or two leading examples, and then additional literature pointers. Some types of market failures that emerge in market design contexts will often be familiar to IO researchers.

Market power

Exercise of market power, leading to inefficiency, has been documented in a variety of market design settings in which some participants are “large,” suitably defined for the setting.

A recent example is the exertion of market power in the US Spectrum Incentive Auction (Milgrom and Segal, 2017, 2020). The auction design had the feature that it was strategy-proof for “small” sellers — defined as owners of at most a single television license per region — to reveal their true reservation value truthfully. However, as is commonly the case in multi-object auctions, sellers with multiple licenses could exert market power by withholding some of their supply. This is known as “supply reduction” (or, analogously, “demand reduction”) in the literature (Ausubel et al., 2005, 2014).

Many broadcast television licenses were indeed owned by relatively small firms, e.g., the local NBC affiliate, making this feature that the auction was strategy-proof

²⁹ Naturally, this classification is not exhaustive and several papers often have multiple objectives. Our hope is to organize what we see as the primary goals in market design research.

for small sellers desirable. However, after the auction was announced but prior to its taking place, some private equity firms aggregated broadcast television licenses from multiple sellers in a given region. This allowed them to then exert market power, by holding back some of the licenses from the auction. Doraszelski et al. (2019) estimate that payouts to station owners were higher by between 7-20% as a result of this exercise of market power.

This episode illustrates the importance of considering the “larger game” in which a particular market design takes place (Roth, 2018). It also echoes the guidance of Klemperer (2002) that in auction design, traditional IO issues such as market power, collusion, and entry can be as important as market design per se. Milgrom (2020) emphasizes that the inefficiency and increase in costs caused by this outside-the-game move, while worse than nothing and arguably preventable, should be considered relative to the magnitudes of the efficiency gains from the reallocation of spectrum and the revenue realized by the federal government.

Collusion

Centralized assignment mechanisms have sometimes been viewed with suspicion as a tool for facilitating collusion. In 2003, a group of former medical residents sued several medical associations and the National Residency Matching Program (NRMP) for fixing residency salaries at below competitive levels (Jung et al. versus Association of American Medical Colleges et al., 2002). The NRMP uses the Roth and Peranson (1999) algorithm, which is a variant of the Gale-Shapley Deferred Acceptance Algorithm, to place newly minted medical school students at residency programs. The plaintiffs argued that the match serves as a coordination device that is used to prohibit bargaining and to keep salaries low. In a brief submitted on behalf of the plaintiffs, Orley Ashenfelter argued the work done by medical residents is similar to that of nurse practitioners, suggesting that a perfectly competitive market would result in salaries for residents that are approximately \$40,000 higher.

A subtlety in this market, relative to other labor markets, is that hospitals set salaries uniformly across all of their residency positions, and do so before knowing which particular residents they will be matched with. Bulow and Levin (2006) analogize the restriction on salaries with having to offer the same salary for Barry Bonds (one of the greatest baseball players of all time, modulo subsequent concerns about performance-enhancing drug use) as for Mario Mendoza (a player who somehow became famous for mediocrity). They showed that this feature can depress and compress salaries relative to a competitive equilibrium.³⁰

The suit was dismissed following an exception to antitrust law enacted by Congress in 2004, but left open the question of whether the design of the match is responsible for low salaries. Empirical studies, however, suggest that the match may not

³⁰ It is worth emphasizing that the restriction on the salaries is not due to centralization per se. Crawford (2008) proposes a design reform that would allow both sides of the market to submit ranks that depend on various salary levels. Such a design could approximate competitive outcomes, although it may be burdensome to implement.

be directly responsible for low salaries. Niederle and Roth (2003a) show that salaries at medical fellowship markets that operate with and without the match are similar. Agarwal (2015) provides a rationale for this result, arguing that much of the salary gap arises due to market power issues that are unrelated to centralization. Medical residency programs have a limited number of accredited positions. Residents value training at these programs and are willing to take a salary cut. In particular, prestigious programs can reduce their salaries without substantially affecting the quality of residents they attract. Using estimates of residents' willingness to pay for training at various programs, Agarwal (2015) estimates that more than half of the salary gap can be explained by an implicit tuition for training at desirable programs. This gap would remain even in the absence of a match.

There have also been numerous studies of collusive behavior in auction settings. To name a few, Friedman (1991) famously alleged collusion in the market for US treasury bonds, which ultimately led to a market-design change from pay-as-bid to uniform-price; Porter and Zona (1993) demonstrate bid rigging in highway procurement auctions; and Klemperer (2002) describes some famous examples in the 1990s spectrum auctions.

Rent seeking

Budish et al. (2015) study rent seeking in the context of high-frequency trading and modern financial markets. They show that the design of modern financial exchanges, in which time is continuous and orders are processed serially (i.e., one-at-a-time) creates arbitrage opportunities that are a pure contest in speed. For example, if the S&P 500 futures contract goes up in price, this will lead to an arbitrage opportunity to buy S&P 500 exchange traded funds. These arbitrage opportunities in turn induce a socially-wasteful rent seeking competition. In equilibrium, the marginal firm's expenditure on speed leaves them with zero economic profit, just as in a standard rent-seeking tournament.

Leslie and Sorensen (2014) study rent seeking in the context of the market for concert tickets. In that market, ticket brokers often race to purchase underpriced tickets, with the intent of reselling in the secondary market. Leslie and Sorensen build a structural model in which they can quantify both the rent-seeking aspects of ticket resale and the potential for efficiency gains from reallocation. Bhavie and Budish (2018) study the mid-2000's introduction of primary-market auctions to the event ticket market by Ticketmaster, and, by matching primary-market auction data to secondary-market resale values from eBay, find that the auctions eliminated the scope for broker rents (as auctions should). Interestingly, however, the primary-market auctions were abandoned by Ticketmaster, who around the same time made a big push to enter the secondary-market — and hence get a piece of the resale rents themselves (Budish, 2019). Other useful references on this market, which has long captured the attention of economists, include Courty (2003, 2019); Sweeting (2012); and Krueger (2019).

Other recent studies that relate to rent seeking and market design include: Hakimov et al. (2019) on black markets for appointment slots, and Budish (2018) on the rent-seeking tournament by bitcoin miners.

Participation and entry

An important requirement for successful design is that agents have incentives to participate in the market. Otherwise, potentially valuable transactions will not be realized. In fact, Bulow and Klemperer (1996) show that securing an additional bidder's participation in a standard auction without a reserve price increases revenue more than using the optimal auction. The message is that participation is more important than optimal design.

The need to encourage participation has guided design in other contexts as well. Centralized matching mechanism design implicitly acknowledges this goal via emphasis on stable matching mechanisms. A mechanism that satisfies this property reduces the incentives for agents in the market to match outside the mechanism or to disobey proposed assignments. But satisfying this property is not always without costs as relaxing this requirement can often enable more efficient matches if other aspects of agent behavior can be held fixed.³¹ Nonetheless, stability is considered an important property in part because empirical observation suggests that markets that do not produce stable outcomes often collapse because agents start transacting outside the system (Roth, 1991, 2002, 2008).

Market fragmentation

Non-participation or partial participation can have particularly stark consequences because a fragmented marketplace can make it hard to find all potential transactions. Ashlagi and Roth (2014) and Agarwal et al. (2019) analyze this problem's causes and consequences in the living donor kidney exchange market. This market enables transplants for patients with kidney failure who have willing but incompatible living donors. Kidney exchanges are facilitated by individual hospitals or one of three large multi-hospital platforms. Thus, it is more efficient to co-ordinate exchanges via the large multi-hospital platforms because they can explore a greater set of potential transactions.

Unfortunately, a design flaw in the mechanism used by the platforms is partially responsible for the fragmented market structure. Specifically, the design gives incentives to the hospitals to retain some of their easiest-to-match patient-donor pairs and match them with hard-to-match pairs. This internal matching is instead of registering both pairs with a platform, causing fragmentation. The flaw in the design can be remedied by instituting incentives via priorities for patients at hospitals that participate to a greater extent. In addition to this problem, fragmentation may also result from unreimbursed fixed and variable costs of participating in a platform. Agarwal et

³¹ For example, in a school choice model, Kesten (2010) shows there are no strategy-proof mechanisms that are both stable and efficient.

al. (2019) estimate that proper co-ordination can result in more than a 30% increase in the total number of kidney exchange transplants. Realizing these gains would require addressing fragmentation by using either a mandate or both reimbursement reform and a redesign of the platform's algorithm. Both approaches bring up important and subtle issues that are left for future research.

Similar market fragmentation has also been documented and analyzed in the context of Chilean higher education (Kapor et al., 2020b). A centralized admissions system is used by a large number of prestigious colleges in Chile but a significant portion of the higher education sector does not participate on the platform. Frictions in the market result in inefficient clearing because admissions across the two paths to colleges are not well coordinated. Using estimates of a model that combines preferences and student achievement effects, they show the benefits of eliminating this market fragmentation.

Investment and innovation

Chapter 13 (this volume) provides a survey of the literature on incentives for innovation. One recent study that relates specifically to innovation in the context of market design is Budish et al. (2020). It uses theory and data to study the industrial organization of modern electronic stock exchanges, and uses the model to understand the incentives of exchanges to innovate on their market designs in ways that reduce some of the harms associated with high-frequency trading. One conclusion is that private innovation incentives diverge from social incentives, because exchanges earn economic profits in the status quo from selling inputs into the arms race for trading speed. Innovating to address the negative aspects of high-frequency trading would undermine this source of economic rent.

As Budish et al. (2020) point out, in many of the market design implementations that are well known to the literature — spectrum auctions, the medical match, school choice, kidney exchange, course allocation — the entities making adoption decisions are governments or non-profit organizations. In this context, where the key players are for-profit financial exchanges, private innovation incentives will play a role in whether socially optimal market designs are implemented. This intersection of IO, market design and innovation incentives seems a fruitful space for future research.

A different question concerns the incentives of participants in a designed marketplace to undertake costly investments, e.g., in increasing their values or reducing their costs. There is a classic result in the context of Vickrey auctions that, under some conditions, private investment incentives are socially efficient (Rogerson, 1992). The reason is that the Vickrey auction implements the socially efficient allocation in an incentive-compatible manner by giving each agent their marginal contribution to social welfare as private surplus (e.g., in a single-object private-values auction, the highest bidder gets as surplus the difference between their value and the second highest-value, which exactly aligns private and social incentives to increase one's private value). A recent paper of Hatfield et al. (2021) extends this idea to a more general mechanism design environment that includes many auction and matching problems as special cases.

Search costs

The baseline theoretical and empirical models assume away costs of participating in centralized marketplaces. However, real-world markets are characterized by frictions that are costly to overcome. For example, students and their parents spend time and effort learning about schools in their district. And, schools and employers exert effort in evaluating applicants. These frictions can result in suboptimal outcomes that can be important to manage. He and Magnac (2020) argue that increases in evaluation costs associated with making it easier for students to submit a large number of applications may outweigh the improvements in match quality. One way to mitigate this externality is to levy appropriately calibrated application fees.

In other cases, applicants may be applying to too few options because of costly search. This is especially true if evaluating and processing applications impose negligible costs, as is the case in many centralized school choice systems. Using detailed data from the application process and a survey conducted right after applications are submitted, Arteaga et al. (2021) show that over-optimism about admission chances results in sub-optimally low search and too few school options being ranked. As a remedy to the resulting market failures, they demonstrate the value of matching platforms with personalized warnings for applicants whose tentative rank-order list leaves substantial estimated risk of non-placement.

Strategic misreporting

Most market designs encountered in the literature have the property that, if participants report their preferences truthfully, the market will produce an efficient allocation, suitably defined given the context.

However, many market designs used in the world induce participants to strategically misreport their preferences, sometimes significantly, and this can lead the market to produce allocations that are efficient with respect to the reported preferences, but inefficient with respect to the underlying truthful preferences.

An early suggestive example in the matching literature is Roth (1991), which shows that medical specialties that used a matching algorithm called the priority match, as opposed to deferred acceptance, often abandoned their match, suggesting inefficiency. Also suggestive was early evidence on the Boston mechanism for school choice, for example Pathak and Sonmez (2008), which provided various forms of anecdotal evidence about some parents strategically misreporting their preferences while other parents participated more naively. In the next subsection, we discuss empirical work that has documented how agents behave when confronted with a manipulable school choice system. This work also quantifies the effects of this manipulation by comparing this mechanism to strategy-proof alternatives.

Concrete quantification of the harms from strategic preference misreporting is provided in Budish and Cantillon (2012). The context of that study is course allocation at Harvard Business School (HBS). HBS uses a mechanism colloquially known as a “snake draft”, in which students take turns choosing courses one at a time, with the order reversing in each round (thereby “snaking” up and down the

list of students, see Section 2.3.6). That mechanism provides incentive to strategically misreport, so as not to waste early round draft choices on courses that could be safely obtained in later rounds. The authors have both the actual (potentially strategic) reported preferences, as well as data on students' underlying truthful preferences from an administrative survey. These data are used to directly document that students strategically misreport their preferences in the direction suggested by the theory, and to calculate the effects of strategic misreporting.

4.2 Evaluating and comparing designs

In order to understand how market design may affect outcomes, it is important to compare existing designs to each other and to alternatives. The theoretical literature has found that it is often hard to simultaneously satisfy all the properties a designer might like to (see Vulkan et al., 2013 and Roth, 2018). For cases where comparing two mechanisms may be theoretically ambiguous, it is valuable to quantify these mechanisms' effects on various dimensions, for example, in terms of participant welfare.

The body of work discussed below highlights that detailed market rules interact with empirical research in three different ways. First, formally adopted mechanisms often yield data that enable empirical approaches. Second, the appropriate empirical approach depends on the rules adopted in the market. Third, this work can shed light on the quantitative effects of various rules and identify potential reforms.

School choice

Perhaps the most widely studied empirical problem is school choice design (see Pathak, 2017 and Agarwal and Somaini, 2020 for more complete surveys). This emphasis is partly because administrative data from school districts has been more readily available, but also because methods for school choice settings are comparatively better developed. The methodological advances in this area have been enabled by a relatively straightforward link between the type of preference data available (rank-order lists) and single-unit discrete choice demand models. Moreover, a detailed understanding of student priorities and the assignment process is helpful in developing approaches that are well-tailored to the empirical setting.

School choice reforms have aimed to coordinate admissions via formal matching mechanisms in order to increase allocative efficiency. But understanding the quantitative effects of centralizing admissions requires data from a system without a formal choice process, and such data are hard to come by. Abdulkadiroglu et al. (2017b) use the implementation of the New York City High School assignment system to quantify the welfare effects of centralized school assignment. They find that, following the reform that centralized the assignment process, students were more likely to enroll in their assigned school and exits from the public school system fell. Their analysis also compared the new DA-based system to the old system and alternatives motivated by matching theory. On a scale ranging from a no-choice neighborhood assignment to the utilitarian optimal, the new system realized 80% of the potential gains, whereas

the old system achieved one-third at most. Other ordinal mechanisms studied in the theoretical literature were within a few percentage points, suggesting that the primary gains arise from coordinating assignments. Thus, centralization and coordination are of first-order importance relative to the differences between well-designed alternatives.

A related issue in designing school choice mechanisms is comparing manipulable and non-manipulable systems. Although the theoretical literature advocates strongly for the strategy-proof Deferred Acceptance mechanism, commonly used manipulable alternatives such as the Immediate Acceptance mechanism cannot be ruled out based on theory alone (Abdulkadiroglu et al., 2011). The empirical literature has shed light on the comparison between these two mechanisms. Student welfare has been compared using various agent behavior models, ranging from equilibrium play (Agarwal and Somaini, 2018) and models that allow for mistaken reports (Calsamiglia et al., 2020; Agarwal and Somaini, 2018; He, 2017; Hwang, 2014) to models of heterogeneous beliefs estimated using survey data (Kapor et al., 2020a). These papers largely find that the average student welfare is higher under IA if students' behavior is described by equilibrium play, but the difference is small at best (Agarwal and Somaini, 2018; Kapor et al., 2020a). Survey evidence also suggests that many students are mistaken about their admission chances or the mechanism used, which can further weigh against IA (He, 2017; Kapor et al., 2020a). Thus, mistakes and agent behavior are important factors in a design's effects.

Practical experience with implementing school choice mechanisms and the data that have been generated has also revealed new design issues that need evaluation. For example, districts have implemented school choice menus to manage transportation costs (Shi, 2015). Districts have also used multi-stage clearing processes with a restricted the number of choices in order to simplify the process and to provide information about school competition (Ajayi and Sidibe, 2021; Lufade, 2019). Understanding these designs' effects requires careful empirical work. More broadly, a mechanism and priorities should be designed keeping the objective function of the school district and any constraints on the allocation in mind. Indeed, Shi (2021) develops an approach for this objective that uses estimates of preferences as an input into the design problem.

Another issue uncovered via practical experience is that students in large districts find it daunting to evaluate and rank many schools (Arteaga et al., 2021). This can result in numerous students remaining unmatched and some appealing their initial assignments after learning more about their assigned school (Narita, 2018). Managing such mismatch requires careful thinking about after-market design and the process through which students acquire information. Appropriately designing a mechanism with such issues in mind requires a deeper understanding of how the implementation of a mechanism affects incentives for acquiring information and engaging in search (Immorlica et al., 2020).

Dynamic assignment mechanisms

An important class of markets are ones in which agents and objects that need to be allocated arrive over time, rather than simultaneously as in school choice. Such

markets include public housing (Waldinger, 2021), deceased donor organs (Agarwal et al., 2021), hunting licenses (Verdier and Reeling, 2021), and foster care (Robinson-Cortés, 2019). In most of these cases, objects must be assigned to agents without full knowledge of which agents or objects will be available for assignment in the future either because waiting is costly or storage is not possible. For the designer and the empirical researcher, this adds another challenge in addition to the considerations in static settings discussed above.

Recent methodological research has made some progress on these issues. A challenge with designing practical mechanisms is that it is rarely feasible to elicit agents' preferences over all potential objects and their assignment time. This curse of dimensionality results in much coarser information that a mechanism can use to allocate objects. For example, on the organ waitlist, patients can either accept or decline an organ offer. In the latter case, they are exercising the option of waiting for another offer. Consequently, unlike school choice, administrative data from these systems also do not allow for a direct comparison between any two objects. The empirical approaches mentioned above adapt models of dynamic discrete choice (Pakes, 1986; Rust, 1987; Hotz and Miller, 1993) and dynamic games (Bajari et al., 2007) in order to infer the comparison between two objects. Like school choice, these methods also leverage the detailed understanding of the assignment process that a formal mechanism affords.

Commonly studied designs often involve waitlists with agents who are offered objects in a priority order as they arrive. Each agent can be assigned at most one object, and they may either accept or decline an offer in the hope that a preferable object will be offered in the future. A number of papers have studied how to prioritize agents on this list. For example, one could organize the waitlist according to a first-come first-served principle or a last-come first-served protocol. Theoretical results suggest that the comparison between these two simple mechanisms is sensitive to the nature of preferences in the market. Consider that if preferences are largely idiosyncratic, then a first-come first-served protocol is preferable (Bloch and Cantala, 2017). This mechanism induces agents at the top of the list to be selective and only accept objects that are a good match for them. However, if objects are vertically differentiated, then a last-come first-served system is preferable because it reduces selectivity and minimizes waste (Su and Zenios, 2004). This conclusion is special to the model with vertically differentiated preferences because who is assigned the object is inconsequential. In other cases, mechanisms other than these two polar extremes can be preferable (Leshno, 2019).

These findings suggest the need to understand the distribution of preferences when recommending a design. The detailed administrative data collected during formal assignment processes provide a path to achieving these ends. This recent body of empirical work has used such data to build an empirical approach to addressing the design problem. For example, Agarwal et al. (2021) study the assignment of deceased donor kidneys to patients with kidney failure. They use data on the decisions made by agents in this market to estimate preferences as a function of patient-donor characteristics. The results are then used to design priority rules for the waitlist to

maximize patient welfare. In a related study of the assignment system for public housing, Waldinger (2021) explores how leaning on choices versus on priorities affects the trade-off between allocative and targeting efficiency.

While patients waiting for a kidney and public housing applicants are typically assigned to only one object from the list, in certain contexts agents demand assignments of multiple goods over time. A theoretical literature studies how optimal mechanisms in these contexts significantly differ from waitlists and usually involve explicit incentives to decline an object (see Guo and Horner, 2018, and references therein). A real-world example of such a design is the allocation of bear hunting licenses. Many states use a lottery system in which an agent who is not assigned a license during a particular season gains priority in the next season (Verdier and Reeling, 2021).

Alternative approaches to evaluating mechanisms

The discussion above focuses on studies that are based on methods described in Section 3 of this chapter. These studies involve using data generated by one mechanism to project counterfactual outcomes in another using a structural model of preferences and behavior. Yet in some instances, empirical analysis is possible by either directly describing mechanism-generated data or using surveys. These studies often compare outcomes from two different designs observed in the same market or cases when a counterfactual mechanism is straightforward to simulate. Examples of these studies include Niederle and Roth (2003b), which measures mobility in the gastroenterology fellowship market that abandoned a centralized matching procedure between 1997 and 1999; Budish and Cantillon (2012), which studies a course allocation program by surveying participants to obtain truthful preference data and simulate counterfactuals; and Prendergast (2021), which describes how trade patterns changed after food banks adopted a new market design. The next subsection describes some of these applications in greater detail.

Policy evaluation using data from centralized marketplaces

Centralized marketplaces are also an ideal laboratory for studying questions beyond the immediate purview of market design. They are unique repositories of detailed and administrative data on market participants and their outcomes. Moreover, the rules governing these markets are well-understood.

A small empirical literature has taken advantage of this confluence of data and a grasp of institutional detail to address questions central to health and education economics. To name a few examples, Hastings et al. (2009) investigates the importance of socio-economic heterogeneity in preferences in explaining inequality in enrollment at good schools; Abdulkadiroglu et al. (2017a) uses randomization inherent in a school choice mechanism to evaluate school quality in Denver; and Agarwal (2017) compares policies (financial incentives and quantity regulations) for increasing the supply of medical residents in underserved rural areas. Needless to say, it is important to pay careful attention to the market's design when formulating an empirical strategy for policy evaluation. Just as the effects of market power are intermediated

through the design of the market, so is the effect of policy reforms such as incentives. Therefore, understanding how a market's design shapes the effect of a policy can enable novel empirical strategies.

Evaluating non-utilitarian objectives

Most of the aforementioned work evaluates mechanisms based on traditional (utilitarian) notions of agent welfare. However, the need for market designers to devise mechanisms that can be deployed in the real world has also motivated research that directly evaluates other objectives (not necessarily utilitarian) that are important to policymakers. For example, public housing authorities value targeting allocations to the neediest, which may be at odds with allocative efficiency (Waldinger, 2021). Similarly, organ allocation authorities value survival outcomes or transplanting the urgently sick (Agarwal et al., 2020). In addition, distributional issues across demographics are often central to the types of problems considered (Dworczak et al., 2021). On the empirical side, Tanaka et al. (2020) studies the effects of centralized assignment on meritocratic admissions and long-term student outcomes. An open area of future research is designing mechanisms that directly incorporate these considerations.

An empirical analysis that evaluates market performance using non-utilitarian outcomes requires methods for measuring the effects of changing assignments on these outcomes. For example, we need to estimate how alternative school assignments affect student achievement, or how alternative organ waitlist designs will affect patient survival. The empirical challenge is that agent choices may be correlated with outcomes (Roy, 1951). As is well-known, the resulting selection problems can bias observational estimates. Agarwal et al. (2020) combine a model of choices and outcomes in an assignment mechanism to estimate treatment effects as a function of observed and unobserved agent and object heterogeneity. The approach tailors methods from a large literature on estimating treatment effects using quasi-experimental variation (Imbens and Angrist, 1994; Heckman and Vytlačil, 2005; Heckman and Navarro, 2007) to the context of an assignment mechanism. Using this approach, Agarwal et al. (2020) estimate that improved assignments can increase patient survival on the deceased donor waiting list by several years. Insights from these methods have also been used by Kapor et al. (2020b) to show that improving the coordination of college admissions in Chile lead to increased on-time graduation rates.

4.3 Proposing new market designs

A third objective of market design research, as the term suggests, is to propose new market designs. Often, diagnosing failures of an existing market (per Section 4.1), and evaluation of the performance of existing market designs (per Section 4.2), are key inputs into this step. This section briefly describes some economist-proposed market designs for real-world settings.

Medical match

Each year, the National Residency Matching Program (NRMP) matches approximately 25,000 newly minted graduates of medical schools to training positions at hospitals called residencies. Both residents and residency programs are unique, differing in quality and idiosyncratic fit, much like many other labor markets or educational markets. But, unlike the vast majority of labor markets, the institution is centralized. Each side of the market – residents and programs – submits a rank-ordered preference list of agents on the other side to a centralized clearinghouse that uses an algorithm to determine the final matches.

Amazingly, the NRMP first implemented the Gale-Shapley matching algorithm in the 1950s — well before Gale and Shapley did their research (by all accounts, Gale and Shapley were not aware of this usage of the algorithm in practice). This independent discovery of the algorithm, by practitioners as opposed to game theorists, is remarkable, and perhaps speaks to the algorithm's intrinsic beauty and appeal.

Roth and Xing (1994) report on the nature of the medical job market in the 1940s, prior to the adoption of a centralized matching algorithm. The market was decentralized and residency programs vied for the best medical students by offering some positions earlier than their competitors in order to gain an edge. This process unraveled to the point that some students were offered positions several years in advance of their graduation. Such early matching can be inefficient for a variety of reasons. As just one example, medical students' preferences over the various fields of medicine may not be fully formed before students are able to do rotations across specialties, which is part of medical school training.³²

Roth (1984) first reported on the connection between Gale and Shapley (1962) and the NRMP, and, moreover, noted some potential problems that were starting to emerge with the algorithm. As women were admitted to medical schools in greater numbers, starting in the 1960s and especially the 1970s, it became more and more common for medical residents to be married couples, seeking a pair of positions. But, the Gale-Shapley algorithm treated each worker as their own individual. Roth and Peranson (1999) proposed a modification of the Gale-Shapley algorithm which allowed for married couples to participate in the match in a way that, with high probability, would yield a stable match (see Ashlagi et al., 2014; Kojima et al., 2013).

Notably, for IO readers, direct empirical evidence on the welfare gains of the match has been somewhat elusive. Some indirect evidence comes from the fact that the algorithm has been maintained in the market for so long, and that specialties that adopted different, non-stable algorithms, have tended to abandon those non-stable algorithms in favor of Gale-Shapley or Roth-Peranson (Roth, 2002). There is also indirect evidence from an interesting case study in gastroenterology, which, for idiosyncratic reasons, went back and forth between utilizing a centralized match and a decentralized market (Niederle and Roth, 2003b), and from the context of matching

³² This would be like having to decide whether or not to be an IO economist before taking the second-year field courses in IO.

college football teams to bowl games (Fréchette et al., 2007). But, there is no direct evidence that speaks to the overall social welfare effects of the centralized match as researchers have not yet found a way to get enough grip on the counterfactual.

School choice

In an influential paper, Abdulkadiroglu and Sonmez (2003) introduced matching students to schools as a market design problem; documented that some school districts were using the “Boston mechanism” (aka Immediate Acceptance algorithm) which, as discussed above in Section 2.3.2, is strategically manipulable; and proposed variations on the deferred acceptance algorithm and top-trading cycles algorithm as a theoretically superior alternative. Since that paper, the theoretical literature on school choice has been very active, and many school districts have adopted deferred acceptance (the adoption of top trading cycles has been more mixed, as discussed shortly). A useful survey is Pathak (2011).

Here we emphasize a few observations about this literature, and the role it is played in market design reform. As discussed above in Section 4.2, several papers have used modern structural IO methods to compare mechanisms. One design takeaway from these papers is that the *equilibrium efficiency* differences between well-designed candidate mechanisms are rather small. This in turn then suggests that factors aside from equilibrium efficiency per se are likely to be central in design decisions.

Two factors in particular seem to be central. First, the *complexity of equilibrium behavior* in the Immediate Acceptance mechanism, and the associated fairness costs of strategic mistakes, appear to have been the decisive factor in economist recommendations to move away from that mechanism. Recall that survey evidence from Kapor et al. (2020a) directly documents the difficulty some families had with figuring out how to play equilibrium. Yet, the Immediate Acceptance mechanism remains widely used across the world.

Second, the key difference between deferred acceptance (which has been widely adopted) and top trading cycles (which has not been widely adopted) is the *interpretation of priority* awarded to students based on whether they live nearby to the school (“walk zone priority”) or have a sibling who has attended the school (“sibling priority”). The top trading cycles algorithm treats such priorities as property rights that can be traded to secure a school the student likes better. The deferred acceptance algorithm treats schools’ rank-orderings as either a non-tradable property right of students, or as a preference of the school (or both, they are different English descriptions of the same underlying mathematics). The fact that most school districts, when presented with deferred acceptance and top trading cycles as economist-blessed market design alternatives, have opted for deferred acceptance over top trading cycles, strikes us as evidence that most public school districts do not interpret such priorities as tradable property rights. Additionally, it may be that a subset of schools have preferences over which students enroll, and respecting these preferences is important to secure their participation in the school choice mechanism.

A last point we emphasize, consistent with the previous, is that the evidence in Abdulkadiroglu et al. (2017b) suggests that it is far more important for welfare to move from an uncoordinated system to a sensibly-designed mechanism, than it is to move from a sensibly-designed mechanism to the optimal mechanism. This result echoes the argument of Klemperer (2002) in the context of spectrum auctions.

Course allocation

Universities often place limits on the number of students who can enroll in a particular class. This gives rise to a challenging allocation problem: how can a university ensure that students receive schedules of courses that efficiently reflect their preferences, in a manner that is perceived as fair, and under a procedure that is incentive compatible? The series of papers that led to Wharton and other universities adopting Budish's A-CEEI mechanism for course allocation usefully illustrates how market design reforms often take multiple papers using multiple different methodological approaches, in this case involving theory, data, experiments, and computation (Roth, 2002).

First, Budish and Cantillon (2012) used theory and data to study the draft mechanism for course allocation used at Harvard Business School. Draft mechanisms are common in practice (e.g., for professional sports leagues) but their theoretical and empirical properties are not well understood. Budish and Cantillon (2012) obtained administrative data from HBS that contained both the preferences students actually reported to the draft (potentially strategically misreported) as well as their underlying truthful preferences, from a survey. This data showed that students strategically misreported their preferences, as suggested they would by theory (not wasting earlier draft picks on courses that a student prefers but knows to be unpopular). By simulating student outcomes under hypothetical truthful play of the draft, and comparing it to student outcomes under actual strategic play of the draft, the authors could show that strategic misreporting harmed overall student welfare — underscoring the value of strategy-proofness. Yet, the comparison of the manipulable and manipulated draft mechanism, to the theoretical counterfactual random serial dictatorship mechanism — which is both strategy-proof and ex-post Pareto efficient — showed the draft to be much better for student welfare. Budish and Cantillon (2012) thus indicated that the theory literature would have to find some way to move beyond dictatorships, which in turn would involve some compromise away from either strategy-proofness or ex-post Pareto efficiency, given the relevant impossibility theorems.

Next, Budish (2011) proposed the A-CEEI mechanism, based in part off of the empirical lessons from Budish and Cantillon.³³ A-CEEI is *approximately* but not exactly ex-post Pareto efficient, approximately strategy-proof, and satisfies criteria of outcome fairness called envy bounded by a single good and maximin share (now referred to as EF1 and MMS in the fair division literature). Using the HBS data from

³³ This section reports on papers in the order they were first completed, not the order in which they were ultimately published.

Budish and Cantillon (2012), and a computational algorithm for A-CEEI developed by Othman et al. (2010), Budish (2011) shows that the A-CEEI mechanism outperforms the draft and the dictatorship on measures of both efficiency and fairness. We note that whereas in the school choice literature discussed above, efficiency differences between mechanisms are fairly small, in course allocation they are large. The intuition is that in multi-unit allocation there is a lot more scope for realizing efficiency gains from trade than in allocation problems in which each agent consumes just a single unit.

Budish and Kessler (2021) then ran an experimental test of A-CEEI at the Wharton School at the University of Pennsylvania. The main goal of the experiment was to assess whether market participants could report their preferences to the A-CEEI mechanism “accurately enough” for the mechanism’s efficiency and fairness advantages to manifest in practice. This question necessitated a novel style of laboratory experiment, in which subjects participated based on their real preferences for real schedules in a simulated (i.e., unincentivized) environment, as opposed to an environment in which subjects are endowed with artificial preferences. The evidence in Budish and Kessler (2021) suggested that students at Wharton could report their preferences with sufficient accuracy, but that preference-reporting mistakes did have a measurable effect on the market design’s performance in the lab.

Last, Budish et al. (2017) report on the actual implementation of the A-CEEI mechanism at Wharton. This involved finalizing the preference-reporting language based on the evidence from Budish and Kessler (2021), improving the algorithm for discovering A-CEEI prices relative to the initial version in Othman et al. (2010), and various other kinds of work involved in bringing a market design from theoretical idea to practical implementation.

Food banks

Prendergast (2017; 2021) reports on the implementation of a pseudomarket mechanism for Feeding America, the national umbrella organization responsible for food banks across the United States. The context for this market design research is that Feeding America frequently receives donations of food (e.g., a truckload of cereal) that it has to allocate to one of the many food banks across the United States. Previously, Feeding America made allocation decisions using human judgment — making phone calls to various food banks, to ascertain who had an especially high need for the donation (e.g., based on other food on hand at that food bank at the moment, from other sources), while also keeping track of fairness considerations (e.g., spreading donations fairly across the country).

Together with many colleagues, Prendergast implemented a market based on an internal currency to replace the prior process based on human judgments. This market allocated a new internal currency to food banks across the country in proportion to estimated local need. Then, when donations came in to Feeding America, the food banks could bid for the donation. The highest bidder for a given donation would then receive this donation (the food bank would be responsible for all transportation costs), and the artificial currency that was bid would then be electronically distributed

to all of the other food banks, in proportion to their needy population. In this way, the overall level of the currency in the system stayed constant over time, and food banks could learn, over time, the approximate costs of certain kinds of donations. For a variety of fascinating reasons, proteins like chicken and shelf-stable foods like pasta and cereal were richly priced in the market, whereas produce and dairy, which perish quickly, and junk food with poor nutritional value, were very cheap.

Prendergast (2021) documents that the pseudomarket mechanism, dubbed the Choice System, lead to a more equitable allocation of food, more efficient sorting of donations to nearby recipients (thereby economizing on transport costs), and, perhaps most surprisingly, vertical sorting of food banks across the country, whereby food banks with lots of other sources of donation (e.g., in wealthy cities) used the Choice System to get a small amount of highly-priced items (e.g., chicken) whereas food banks with fewer other sources of donation used the Choice System to get a much larger amount of food overall, in essence “trading” their shares of chicken for large amounts of other foods.

Financial exchanges

Budish et al. (2015) propose a new market design for financial exchanges, frequent batch auctions (FBA), to replace the continuous limit order book (CLOB) that is widely used for the trading of stocks, futures, and other financial assets around the world. Budish et al. (2015) use both theory and data to critique the CLOB design and make a case for the FBA design. Empirically, they show that market correlations break down at fine-enough time intervals (e.g., fractions of a second), leading to frequent, fairly obvious arbitrage opportunities. These arbitrage opportunities in turn induce an arms race for trading speed. The theory model shows that the arms race for trading speed in a sense is never ending, since it is *relative*, not *absolute*, trading speed that is essential for capturing fleeting arbitrage opportunities. More subtly, the model shows that the rents from these arbitrage opportunities come at the expense of market participants who are providing liquidity to the rest of the market (i.e., who stand willing to either buy or sell, at a spread). In equilibrium, liquidity providers pass on this cost to end investors. Thus, there is an equivalence in the model between the rents in the arms race for speed, the investments in speed, and the cost to end investors.

The paper then uses theory to propose an alternative market design, frequent batch auctions, that eliminates the rents from fleeting arbitrages, and in so doing ends the arms race for trading speed and improves market liquidity. Intuitively, in a continuous market, even a one-microsecond speed advantage (i.e., 0.000001 seconds) is enough to win the race. Whereas in a batch auction market, even if the batch auctions occur as often as every tenth or even thousandth of a second, a one-microsecond speed advantage is basically meaningless. Competition is on price, rather than the last epsilon of trading speed.

Aquilina et al. (2021) use a novel form of financial market data, obtained using a specific regulatory authority of the UK Financial Conduct Authority (where Aquilina and O’Neill are based), to measure the extent of the speed-based arbitrage opportuni-

ties suggested by the theoretical model of Budish et al. (2015). Specifically, the novel data is called “message data”, which, crucially, contains messages sent by market participants who *lose*, not just win, a race for a particular trade. Aquilina et al. (2021) show that, for stocks in the UK FTSE 100 index (roughly analogous to the S&P 500 index in the United States), there is about one race per minute (537 per symbol per day), and that races constitute over 20% of all trading volume. In the modal race, the difference in time between the winner and the loser is between 0 and 15 millionths of a second (i.e., less than 0.000015 seconds). Races are for small amounts per race, but, because of the volumes, they add up to significant sums. The authors compute that speed-based arbitrage constitutes about 33% of the literature’s main measure of the cost of liquidity, called price impact, that implementing FBAs would reduce the market’s cost of liquidity by 17%, and that speed-based arbitrages are worth about \$5 billion per year in equities markets alone.

Budish et al. (2020) study the incentives for exchanges to innovate on their market design, e.g., by adopting FBAs. If FBAs reduce the cost of liquidity by 17%, not to mention make the market computationally simpler (avoiding much of the complexity of managing systems that can manage and adjudicate microsecond-level speed races), how come exchanges have not rushed to adopt? The answer, Budish et al. (2020) suggest, is a divergence between private and social incentives for innovation (see Chapter 13 (this volume)). Specifically, incumbent exchanges earn significant economic profits from selling speed into the speed race — e.g., selling the right to co-locate one’s own computers next to the exchange’s computers — and would cut off this source of revenue if they adopted a market design that stopped the arms race for speed.

Thus, market design reform in this setting may require a regulatory intervention, as opposed to coming about purely from private sector incentives. Notably, in many of the other market design reforms noted above — school choice, NRMP, course allocation, Food Banks — private and social incentives to adopt a new, better market design were *aligned* as opposed to *divergent*. A broader lesson for the market design literature, especially in its intersection with IO, is to pay attention not just to the ideal design of market institutions, but to the incentives of those with power to implement such market institutions.

5 Conclusion

An exciting aspect of market design research is that it involves close interaction with the real world. Marketplaces and the rules that govern them come in all shapes and sizes. Market designers seek to understand which rules matter. Given the myriad of rules, it is often unclear which ones are important, a point that has been emphasized before (see Roth 2002; Klemperer 2004). This analysis requires understanding the context in which the rules operate, how they interact with the behavior of market participants, and what works and what does not.

Not only does this engagement with real-world marketplaces improve our knowledge of markets, but it also allows us to improve markets themselves. While the starting point is often existing market institutions, it can also begin with recognizing a gap – a problem that clearly needs new market institutions. Design proposals can be based on mechanisms with well-understood theoretical properties, designs imported from other settings, or novel rules. Both past experience and theoretical analysis play an important role in this process.

The experience acquired through the practice of market design itself creates knowledge. Moreover, as markets mature, new issues are often uncovered that need to be addressed, making the endeavor of market design an iterative process.³⁴ Achieving this goal may require a researcher to step outside the traditional bounds of the academic process. But when the opportunity presents itself, a body of research that has fleshed out practically relevant issues is essential to bringing the knowledge of market designers to bear on reforms.

In this way, theory, data, and practice play complementary roles in market design. Theory allows us to identify the qualitative effects of different designs, isolate key tradeoffs, hypothesize the effects of different designs, and identify open issues. Empirical work helps generate empirical regularities that need explanation, quantify economic effects, test hypotheses, and address theoretically ambiguous issues. Practical work with designing markets and domain expertise sheds light on new issues, helps bring data to bear on open questions, and reveals new designs that need further analysis. Throughout, market designers pay specific attention to institutional details in order to guide research towards understanding and improving real-world markets.

This focus on real-world design and the synthesis of approaches echoes Roth (2002), which famously described the “economist as engineer.” Roth argued that game theory would need to be complemented with other methodological tools such as computational simulations and experimental evidence for effective market design. Just as an engineer relies on theory and simulations when designing a bridge, the ultimate goal of a market designer is to take what we have learned and institute new designs that remedy market failure.

We argue that insights from IO add to this engineering approach to design economics in at least two ways. First, our chapter illustrates that another central tool in the market design toolbox is econometric models of market primitives. These models can be estimated using data in order to simulate counterfactuals, a core method in empirical IO. These tools are closely synced with both the theory and practice of market design. Second, we argue that a holistic study of markets requires analyzing both market design and market power. With the notable exception of auction markets and a handful of examples discussed in this chapter, these two issues have largely

³⁴ Roth (2008) describes the evolution of kidney exchange markets from their inception. The challenges faced by the market changed with its development, and a number of issues required new solutions. Even at a relatively advanced stage, there are new opportunities for expanding kidney exchange that require novel market design solutions (Roth, 2018; Ashlagi and Roth, 2021).

been studied independently. We expect that there is much to learn from the synthesis of ideas across these fields.

References

- Abdulkadiroglu, Atila, Sonmez, Tayfun, 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66 (3), 689.
- Abdulkadiroglu, Atila, Sonmez, Tayfun, 2003. School choice: a mechanism design approach. *The American Economic Review* 93 (3), 729–747.
- Abdulkadiroglu, Atila, Angrist, Joshua D., Narita, Yusuke, Pathak, Parag A., 2017a. Research design meets market design: using centralized assignment for impact evaluation. *Econometrica* 85 (5), 1373–1432.
- Abdulkadiroglu, Atila, Agarwal, Nikhil, Pathak, Parag A., 2017b. The welfare effects of coordinated school assignment: evidence from the NYC high school match. *The American Economic Review* 107 (12), 3635–3689.
- Abdulkadiroglu, Atila, Pathak, Parag A., Roth, Alvin E., 2005. The New York city high school match. *The American Economic Review* 95 (2), 364–367.
- Abdulkadiroglu, Atila, Pathak, Parag A., Roth, Alvin E., 2009. Strategy-proofness versus efficiency in matching with indifferences: redesigning the NYC high school match. *The American Economic Review* 99 (5), 1954–1978.
- Abdulkadiroglu, Atila, Che, Yeon-Koo, Yasuda, Yosuke, 2011. Resolving conflicting preferences in school choice: the Boston mechanism reconsidered. *The American Economic Review* 101 (1), 399–410.
- Agarwal, Nikhil, 2015. An empirical model of the medical match. *The American Economic Review* 105 (7), 1939–1978.
- Agarwal, Nikhil, 2017. Policy analysis in matching markets. *The American Economic Review* 107 (5), 246–250.
- Agarwal, Nikhil, Somaini, Paulo J., 2021a. Demand Analysis under Latent Choice Constraints: an Application to the US Dialysis Market.
- Agarwal, Nikhil, Somaini, Paulo, 2018. Demand analysis using strategic reports: an application to a school choice mechanism. *Econometrica* 86 (2), 391–444.
- Agarwal, Nikhil, Somaini, Paulo, 2020. Revealed preference analysis of school choice models. *Annual Review of Economics* 12 (1), 471–501.
- Agarwal, Nikhil, Somaini, Paulo, 2021b. Empirical models of non-transferable utility matching. In: Echenique, Federico, Immorlica, Nicole S., Vazirani, Vijay (Eds.), *Online and Matching-Based Market Design*.
- Agarwal, Nikhil, Hodgson, Charles, Somaini, Paulo, 2020. Choices and Outcomes in Assignment Mechanisms: the Allocation of Deceased Donor Kidneys. NBER Working Paper 28064.
- Agarwal, Nikhil, Ashlagi, Itai, Azevedo, Eduardo, Featherstone, Clayton R., Karaduman, Ömer, 2019. Market failure in kidney exchange. *The American Economic Review* 109 (11), 4026–4070.
- Agarwal, Nikhil, Ashlagi, Itai, Rees, Michael A., Somaini, Paulo, Waldinger, Daniel, 2021. Equilibrium allocations under alternative waitlist designs: evidence from deceased donor kidneys. *Econometrica* 89 (1), 37–76.
- Ajayi, Kehinde, Sidibe, Modibo, 2021. School Choice Under Imperfect Information.

- Akkus, Oktay, Cookson, J. Anthony, Hortaçsu, Ali, 2016. The determinants of bank mergers: a revealed preference analysis. *Management Science* 62 (8), 2241–2258.
- Akyol, Pelin, Krishna, Kala, 2017. Preferences, selection, and value added: a structural approach. *European Economic Review* 91, 89–117.
- Allende, Claudia, 2019. Competition Under Social Interactions and the Design of Education Policies.
- Aquilina, Matteo, Budish, Eric, O’neill, Peter, 2021. Quantifying the high-frequency trading “arms race”. *The Quarterly Journal of Economics*. Forthcoming.
- Arteaga, Felipe, Kapor, Adam J., Neilson, Christopher A., Zimmerman, Seth D., 2021. Smart Matching Platforms and Heterogeneous Beliefs in Centralized School Choice.
- Artemov, Gregory, Che, Yeon-Koo, He, Yinghua, 2021. Strategic Mistakes: Implications for Market Design Research.
- Ashlagi, Itai, Roth, Alvin E., 2014. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics* 9 (3), 817–863.
- Ashlagi, Itai, Roth, Alvin E., 2021. Kidney exchange: an operations perspective. *Management Science*. Forthcoming.
- Ashlagi, Itai, Braverman, Mark, Hassidim, Avinatan, 2014. Stability in large matching markets with complementarities. *Operations Research* 62 (4), 713–732.
- Aue, Robert, Klein, Thilo, Ortega, Josue, 2020. What Happens when Separate and Unequal School Districts Merge?
- Ausubel, L.M., Cramton, P., Pycia, M., Rostek, M., Weretka, M., 2014. Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies* 81 (4), 1366–1400.
- Ausubel, Lawrence M., Aperjis, Christina, Baranov, Oleg, 2017. Market Design and the FCC Incentive Auction.
- Ausubel, Lawrence M., Cramton, Peter, Milgrom, Paul R., 2005. The clock-proxy auction: a practical combinatorial auction design. In: Cramton, Peter, Shoham, Yoav, Steinberg, Richard (Eds.), *Combinatorial Auctions*. MIT Press, Cambridge, pp. 115–138. Chapter 5.
- Azevedo, Eduardo M., Budish, Eric, 2019. Strategy-proofness in the large. *The Review of Economic Studies* 86 (1), 81–116.
- Azevedo, Eduardo M., Leshno, Jacob, 2016. A supply and demand framework for two-sided matching markets. *Journal of Political Economy* 124 (5), 1235–1268.
- Azevedo, Eduardo M., Hatfield, John William, 2018. Existence of Equilibrium in Large Matching Markets with Complementarities.
- Bajari, Patrick, Lanier Benkard, C., Levin, Jonathan, 2007. Estimating dynamic models of imperfect competition. *Econometrica* 75 (5), 1331–1370.
- Becker, Gary S., 1973. A theory of marriage: part I. *Journal of Political Economy* 81 (4), 813–846.
- Becker, Gary S., 1974. A theory of marriage: part II. *Journal of Political Economy* 82 (2, Part 2), S11–S26.
- Beggs, S., Cardell, S., Hausman, J., 1981. Assessing the potential demand for electric cars. *Journal of Econometrics* 17 (1), 1–19.
- Berry, Steven T., Levinsohn, James, Pakes, Ariel, 1995. Automobile prices in market equilibrium. *Econometrica* 63 (4), 841–890.
- Berry, Steven T., Levinsohn, James, Pakes, Ariel, 2004. Differentiated products demand systems from a combination of micro and macro data: the new car market. *Journal of Political Economy* 112 (1), 68–105.
- Bhave, Aditya, Budish, Eric, 2018. Primary-Market Auctions for Event Tickets: Eliminating the Rents of “Bob the Broker”? Working Paper.

- Bloch, Francis, Cantala, David, 2017. Dynamic assignment of objects to queuing agents. *American Economic Journal: Microeconomics* 9 (1), 88–122.
- Block, H., Marshall, J., 1960. Random orderings and stochastic theories of responses. In: Olkin, I., Ghurye, S., Hoefding, W., Madow, W.G., Mann, H.B. (Eds.), *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*. Stanford University Press.
- Bogomolnaia, Anna, Moulin, Hervé, 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100 (2), 295–328.
- Bordon, Paola, Fu, Chao, 2015. College-major choice to college-then-major choice. *The Review of Economic Studies* 82 (4), 1247–1288.
- Boyd, Donald, Lankford, Hamilton, Loeb, Susanna, Wyckoff, James, 2013. Analyzing the determinants of the matching of public school teachers to jobs: disentangling the preferences of teachers and employers. *Journal of Labor Economics* 31 (1), 83–117.
- Brams, Steven J., Taylor, Alan D., 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Brams, Steven J., Straffin, Philip D., 1979. Prisoners' dilemma and professional sports drafts. *The American Mathematical Monthly* 86 (2), 88.
- Bucarey, Alonso, 2018. Who Pays for Free College? Crowding Out on Campus.
- Budish, Eric, 2011. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119 (6), 1061–1103.
- Budish, Eric, 2012. Matching “versus” mechanism design. *ACM SIGecom Exchanges* 11 (2), 4–15.
- Budish, Eric, 2018. The Economic Limits of Bitcoin and the Blockchain. NBER Working Paper 24717.
- Budish, Eric, 2019. How to Fix the Market for Event Tickets.
- Budish, Eric, Cantillon, Estelle, 2012. The multi-unit assignment problem: theory and evidence from course allocation at Harvard. *The American Economic Review* 102 (5), 2237–2271.
- Budish, Eric, Kessler, Judd B., 2021. Can market participants report their preferences accurately (enough)? *Management Science*. Forthcoming.
- Budish, Eric, Cachon, Gérard P., Kessler, Judd B., Othman, Abraham, 2017. Course match: a large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research* 65 (2), 314–336.
- Budish, Eric, Cramton, Peter, Shim, John, 2015. The high-frequency trading arms race: frequent batch auctions as a market design response. *The Quarterly Journal of Economics* 130 (4), 1547–1621.
- Budish, Eric, Lee, Robin, Shim, John, 2020. A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market? NBER Working Paper 25855.
- Budish, Eric, Che, Yeon Koo, Kojima, Fuhito, Milgrom, Paul, 2013. Designing random allocation mechanisms: theory and applications. *The American Economic Review* 103 (2), 585–623.
- Bulow, Jeremy, Levin, Jonathan, 2006. Matching and price competition. *The American Economic Review* 96 (3), 652–668.
- Bulow, Jeremy, Klemperer, Paul, 1996. Auctions versus negotiations. *The American Economic Review* 86 (1), 180–194.
- Calsamiglia, Caterina, Güell, Maia, 2018. Priorities in school choice: the case of the Boston mechanism in Barcelona. *Journal of Public Economics* 163, 20–36.
- Calsamiglia, Caterina, Fu, Chao, Güell, Maia, 2020. Structural estimation of a model of school choices: the Boston mechanism vs. its alternatives. *Journal of Political Economy* 128 (2), 642–680.

- Cantillon, Estelle, Pesendorfer, Martin, 2006. Combination Bidding in Multi-Unit Auctions.
- Carvalho, José Raimundo, Magnac, Thierry, Xiong, Qizhou, 2019. College choice, selection, and allocation mechanisms: a structural empirical analysis. *Quantitative Economics* 10 (3), 1233–1277.
- Castillo, Juan Camilo, Ahuja, Amrita, Athey, Susan, Baker, Arthur, Budish, Eric, Chipty, Tasneem, Glennerster, Rachel, Duke Kominers, Scott, Kremer, Michael, Larson, Greg, Lee, Jean, Prendergast, Canice, Snyder, Christopher M., Tabarrok, Alex, Tan, Brandon Joel, Więcek, Witold, 2021. Market design to accelerate COVID-19 vaccine supply. *Science* 371 (6534), 1107–1109.
- Chen, Yan, Sonmez, Tayfun, 2006. School choice: an experimental study. *Journal of Economic Theory* 127 (1), 202–231.
- Chiappori, Pierre-André, Salanié, Bernard, 2016. The econometrics of matching models. *Journal of Economic Literature* 54 (3), 832–861.
- Chiappori, Pierre André, Oreffice, Sonia, Quintana-Domeque, Climent, 2012. Fatter attraction: anthropometric and socioeconomic matching on the marriage market. *Journal of Political Economy* 120 (4), 659–695.
- Choo, Eugene, Siow, Aloysius, 2006. Who marries whom and why. *Journal of Political Economy* 114 (1), 175–201.
- Courty, Pascal, 2003. Some economics of ticket resale. *The Journal of Economic Perspectives* 17 (2), 85–97.
- Courty, Pascal, 2019. Ticket resale, bots, and the fair price ticketing curse. *Journal of Cultural Economics* 43 (3), 345–363.
- Crawford, Gregory S., Yurukoglu, Ali, 2012. The welfare effects of bundling in multichannel television markets. *The American Economic Review* 102 (2), 643–685.
- Crawford, Vincent P., 2008. The flexible-salary match: a proposal to increase the salary flexibility of the national resident matching program. *Journal of Economic Behavior & Organization* 66 (2), 149–160.
- de Haan, Monique, Gautier, Pieter A., Oosterbeek, Hessel, van der Klaauw, Bas, 2018. The Performance of School Assignment Mechanisms in Practice.
- Demange, Gabrielle, Gale, David, Sotomayor, Marilda, 1986. Multi-item auctions. *Journal of Political Economy* 94 (4), 863–872.
- Diamond, W., Agarwal, N., 2017. Latent indices in assortative matching models. *Quantitative Economics* 8 (3).
- Doraszelski, Ulrich, Seim, Katja, Sinkinson, Michael, Wang, Peichun, 2019. Ownership Concentration and Strategic Supply Reduction. NBER Working Paper 23034.
- Dubins, L.E., Freedman, D.A., 1981. Machiavelli and the Gale-Shapley algorithm. *The American Mathematical Monthly* 88 (7), 485–494.
- Dworczak, Piotr, Kominers, Scott Duke, Akbarpour, Mohammad, 2021. Redistribution through markets. *Econometrica* 89 (4), 1665–1698.
- Einav, Liran, Farronato, Chiara, Levin, Jonathan, 2016. Peer-to-peer markets. *Annual Review of Economics* 8 (1), 615–635.
- Epple, Dennis, Jha, Akshaya, Sieg, Holger, 2018. The superintendent’s dilemma: managing school district capacity as parents vote with their feet. *Quantitative Economics* 9 (1), 483–520.
- Eraslan, Hulya, Fox, Jeremy T., He, Yinghua, Pirozhenko, Yakym, 2021. Measuring the Welfare Gains from Cardinal-Preference Pseudomarkets in School Choice.
- Erdil, Aytok, Ergin, Haluk, 2008. What’s the matter with tie-breaking? Improving efficiency in school choice. *The American Economic Review* 98 (3), 669–689.

- Fack, Gabrielle, Grenet, Julien, He, Yinghua, 2019. Beyond truth-telling: preference estimation with centralized school choice and college admissions. *The American Economic Review* 109 (4), 1486–1529.
- Fisher, James C.D., Hafalir, Isa E., 2016. Matching with aggregate externalities. *Mathematical Social Sciences* 81, 1–7.
- Fox, Jeremy T., 2010. Identification in matching games. *Quantitative Economics* 1 (2), 203–254.
- Fox, Jeremy T., 2018. Estimating matching games with transfers. *Quantitative Economics* 9 (1), 1–38.
- Fox, Jeremy T., Bajari, Patrick, 2013. Measuring the efficiency of an FCC spectrum auction. *American Economic Journal: Microeconomics* 5 (1), 100–146.
- Fox, Jeremy T., Yang, Chenyu, Hsu, David H., 2018. Unobserved heterogeneity in matching games. *Journal of Political Economy* 126 (4), 1339–1373.
- Fréchet, Guillaume R., Roth, Alvin E., Ünver, M. Utku, 2007. Unraveling yields inefficient matchings: evidence from post-season college football bowls. *The Rand Journal of Economics* 38 (4), 967–982.
- Friedman, Milton, 1991. How to sell government securities. *The Wall Street Journal*.
- Gale, David, Shapley, Lloyd S., 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69 (1), 9–15.
- Galichon, Alfred, Salanie, Bernard, 2020. *Cupid's Invisible Hand: Social Surplus and Identification in Matching Models*.
- Galichon, Alfred, Salanie, Bernard, 2021. Structural estimation of matching markets with transferable utility. In: Echenique, Federico, Immorlica, Nicole S., Vazirani, Vijay (Eds.), *Online and Matching-Based Market Design*.
- Galichon, Alfred, Kominers, Scott Duke, Weber, Simon, 2019. Costly concessions: an empirical framework for matching with imperfectly transferable utility. *Journal of Political Economy* 127 (6), 2875–2925.
- Gandhi, Ashvin, 2020. *Picking Your Patients: Selective Admissions in the Nursing Home Industry*.
- Gelman, Andrew, Rubin, Donald B., 1992. Inference from iterative simulation using multiple sequences. *Statistical Science* 7 (4), 457–472.
- Graham, Bryan S., 2011. Econometric methods for the analysis of assignment problems in the presence of complementarity and social spillovers. In: Benhabib, Jess, Bisin, Alberto, Jackson, Matthew O. (Eds.), *Handbook of Social Economics*, vol. 1B. North-Holland, Amsterdam, pp. 965–1052. Chapter 19.
- Graham, Bryan S., 2014. Errata in “Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers”.
- Guo, Yingni, Horner, Johannes, 2018. *Dynamic Mechanisms Without Money*. Working Paper.
- Hakimov, Rustamdjan, Heller, C.-Philipp, Kübler, Dorothea, Kurino, Morimitsu, 2019. *How to Avoid Black Markets for Appointments with Online Booking Systems*. Working Paper.
- Hassidim, Avinatan, Romm, Assaf, Shorrer, Ran I., 2020. The limits of incentives in economic matching procedures. *Management Science* 67 (2), 951–963.
- Hastings, Justine S., Kane, Thomas J., Staiger, Douglas O., 2009. *Heterogeneous Preferences and the Efficacy of Public School Choice*.
- Hatfield, John William, Milgrom, Paul R., 2005. Matching with contracts. *The American Economic Review* 95 (4), 913–935.
- Hatfield, John William, Kojima, Fuhito, Kominers, Scott Duke, 2021. *Strategy-Proofness, Investment Efficiency, and Marginal Returns: an Equivalence*.

- He, Yinghua, 2017. Gaming the Boston School Choice Mechanism in Beijing.
- He, Yinghua, Magnac, Thierry, 2020. Application Costs and Congestion in Matching Markets.
- He, Yinghua, Sinha, Shruti, Sun, Xiaoting, 2021. Identification and Estimation in Many-to-One Two-sided Matching without Transfers.
- Heckman, James J., Vytlacil, Edward, 2005. Structural equations, treatment effects, and econometric policy evaluation. *Econometrica* 73 (3), 669–738.
- Heckman, James J., Navarro, Salvador, 2007. Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics* 136 (2), 341–396.
- Hitsch, Gunter J., Hortacısu, Ali, Ariely, Dan, 2010. Matching and sorting in online dating. *The American Economic Review* 100 (1), 130–163.
- Ho, Katherine, 2009. Insurer-provider networks in the medical care market. *The American Economic Review* 99 (1), 393–430.
- Hotz, V. Joseph, Miller, Robert A., 1993. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies* 60 (3), 497–529.
- Hwang, Sam Il Myong, 2014. A Robust Redesign of High School Match.
- Hylland, Aanund, Zeckhauser, Richard, 1979. The efficient allocation of individuals to positions. *Journal of Political Economy* 87 (2), 293–314.
- Imbens, Guido W., Angrist, Joshua D., 1994. Identification and estimation of local average treatment effects. *Econometrica*.
- Immorlica, Nicole, Mahdian, Mohammad, 2015. Incentives in large random two-sided markets. *ACM Transactions on Economics and Computation* 3 (3), 1–25.
- Immorlica, Nicole S., Leshno, Jacob D., Lo, Irene Y., Lucier, Brendan J., 2020. Information Acquisition in Matching Markets: the Role of Price Discovery.
- Jung, et al. versus Association of American Medical Colleges, et al., 2002. Class Action Complaint, No. 1:02-CV-00873.
- Kapor, Adam, Neilson, Christopher, Zimmerman, Seth, 2020a. Heterogeneous beliefs and school choice mechanisms. *The American Economic Review* 110 (5), 1274–1315.
- Kapor, Adam, Karnani, Mohit, Neilson, Christopher, 2020b. Aftermarket Frictions and the Cost of Off-Platform Options in Assignment Mechanisms.
- Kesten, Onur, 2010. School choice with consent. *The Quarterly Journal of Economics* 125 (3), 1297–1348.
- Klemperer, Paul, 2002. What really matters in auction design. *The Journal of Economic Perspectives* 16 (1), 169–189.
- Klemperer, Paul, 2004. *Auctions: Theory and Practice*. Princeton University Press, Princeton, NJ.
- Kojima, Fuhito, Pathak, Parag A., 2009. Incentives and stability in large two-sided matching markets. *The American Economic Review* 99 (3), 608–627.
- Kojima, Fuhito, Pathak, Parag A., Roth, Alvin E., 2013. Matching with couples: stability and incentives in large markets. *The Quarterly Journal of Economics* 128 (4), 1585–1632.
- Krueger, Alan B., 2019. *Rockonomics: A Backstage Tour of What the Music Industry Can Teach Us About Economics and Life*. Currency, New York.
- Laffont, Jean-Jacques, Tirole, Jean, 1993. *A Theory of Incentives in Procurement and Regulation*. MIT Press.
- Leshno, Jacob, 2019. Dynamic Matching in Overloaded Waiting Lists.
- Leshno, Jacob D., Lo, Irene, 2021. The cutoff structure of top trading cycles in school choice. *The Review of Economic Studies* 88 (4), 1582–1623.
- Leslie, P., Sorensen, A., 2014. Resale and rent-seeking: an application to ticket markets. *The Review of Economic Studies* 81 (1), 266–300.

- Levin, Jonathan, 2013. The economics of Internet markets. In: Acemoglu, Daron, Arellano, Manuel, Dekel, Eddie (Eds.), *Advances in Economics and Econometrics: Tenth World Congress, Volume I: Economic Theory*. Cambridge University Press, Cambridge, pp. 48–75. Chapter 2.
- Li, Shengwu, 2017. Obviously strategy-proof mechanisms. *The American Economic Review* 107 (11), 3257–3287.
- Liu, Tracy, Wan, Zhixi, Yang, Chenyu, 2021. The Efficiency of a Dynamic Decentralized Two-Sided Matching Market.
- Logan, John Allen, Hoff, Peter D., Newton, Michael A., 2008. Two-sided estimation of mate preferences for similarities in age, education, and religion. *Journal of the American Statistical Association* 103 (482), 559–569.
- Luflade, Margaux, 2019. The value of information in centralized school choice systems.
- Manski, Charles F., 1975. Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics* 3 (3), 205–228.
- Manski, Charles F., 1977. The structure of random utility models. *Theory and Decision* 8 (3), 229–254.
- Matzkin, Rosa L., 2007. Nonparametric identification. In: Heckman, James, Leamer, Edward (Eds.), *Handbook of Econometrics*, vol. 6B. North Holland, Amsterdam, pp. 5307–5368. Chapter 73.
- McCulloch, Robert, Rossi, Peter E., 1994. An exact likelihood analysis of the multinomial probit model. *Journal of Econometrics* 64 (1–2), 207–240.
- McFadden, Daniel, 1973. Conditional logit analysis of qualitative choice behavior. In: Zarembka, Paul (Ed.), *Frontiers in Econometrics*. Academic Press, New York, pp. 105–142. Chapter 4.
- Menzel, Konrad, 2015. Large matching markets as two-sided demand systems. *Econometrica* 83 (3), 897–941.
- Milgrom, Paul, 2020. The Market Design Community and the Broadcast Incentive Auction: Fact-Checking Glen Weyl’s and Stefano Feltri’s False Claims.
- Milgrom, Paul, Segal, Ilya, 2020. Clock auctions and radio spectrum reallocation. *Journal of Political Economy* 128 (1), 1–31. <https://doi.org/10.1086/704074>.
- Milgrom, Paul R., Segal, Ilya R., 2017. Designing the US incentive auction. In: Bichler, Martin, Goeree, Jacob K. (Eds.), *Handbook of Spectrum Auction Design*. Cambridge University Press, pp. 803–812. Chapter 37.
- Miralles, Antonio, 2008. School Choice: the Case for the Boston Mechanism.
- Narita, Yusuke, 2018. Match or Mismatch? Learning and Inertia in School Choice.
- Niederle, Muriel, Roth, Alvin E., 2003a. Relationship between wages and presence of a match in medical fellowships. *JAMA, Journal of the American Medical Association* 290 (9).
- Niederle, Muriel, Roth, Alvin E., 2003b. Unraveling reduces mobility in a labor market: gastroenterology with and without a centralized match. *Journal of Political Economy* 111 (6), 1342–1352.
- Niederle, Muriel, Yariv, Leeat, 2009. Decentralized Matching with Aligned Preferences.
- Othman, Abraham, Sandholm, Tuomas, Budish, Eric, 2010. Finding approximate competitive equilibria: efficient and fair course allocation. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, vol. 1, pp. 873–880.
- Pakes, Ariel, 1986. Patents as options: some estimates of the value of holding European patent stocks. *Econometrica* 54 (4), 755.
- Pathak, Parag A., 2011. The mechanism design approach to student assignment. *Annual Review of Economics* 3 (1), 513–536.

- Pathak, Parag A., 2017. What really matters in designing school choice mechanisms. In: Honoré, Bo, Pakes, Ariel, Piazzesi, Monika, Samuelson, Larry (Eds.), *Advances in Economics and Econometrics, 11th World Congress of the Econometric Society*, vol. 1. Cambridge University Press, Cambridge, pp. 176–214. Chapter 6.
- Pathak, Parag A., Sethuraman, Jay, 2011. Lotteries in student assignment: an equivalence result. *Theoretical Economics* 6 (1), 1–17.
- Pathak, Parag A., Sonmez, Tayfun, 2008. Leveling the playing field: sincere and sophisticated players in the Boston mechanism. *The American Economic Review* 98, 1636–1652.
- Petrin, Amil, 2002. Quantifying the benefits of new products: the case of the minivan. *Journal of Political Economy* 110 (4), 705–729.
- Porter, Robert H., Zona, J. Douglas, 1993. Detection of bid rigging in procurement auctions. *Journal of Political Economy* 101 (3), 518–538.
- Prendergast, Canice, 2017. How food banks use markets to feed the poor. *The Journal of Economic Perspectives* 31 (4), 145–162.
- Prendergast, Canice, 2021. The allocation of food to food banks. *Journal of Political Economy*. Forthcoming.
- Pycia, Marek, Yenmez, M. Bumin, 2021. Matching with Externalities.
- Rees-Jones, Alex, 2018. Suboptimal behavior in strategy-proof mechanisms: evidence from the residency match. *Games and Economic Behavior* 108, 317–330.
- Reguant, Mar, 2014. Complementary bidding mechanisms and startup costs in electricity markets. *The Review of Economic Studies* 81 (4), 1708–1742.
- Reny, Philip J., 2017. Assignment problems. *Journal of Political Economy* 125 (6), 1903–1914.
- Robinson-Cortés, Alejandro, 2019. Who Gets Placed Where and Why? An Empirical Framework for Foster Care Placement. Working Paper.
- Rogerson, William P., 1992. Contractual solutions to the hold-up problem. *The Review of Economic Studies* 59 (4), 793.
- Roth, A.E., 1991. A natural experiment in the organization of entry-level labor markets: regional markets for new physicians and surgeons in the United Kingdom. *The American Economic Review* 81 (3), 415–440.
- Roth, A.E., Sonmez, T., Unver, M.U., 2004. Kidney exchange. *The Quarterly Journal of Economics* 119 (2), 457–488.
- Roth, Alvin, Xing, Xiaolin, 1994. Jumping the gun: imperfections and institutions related to the timing of market transactions. *The American Economic Review* 84 (4), 992–1044.
- Roth, Alvin E., 1982a. Incentive compatibility in a market with indivisible goods. *Economics Letters* 9 (2), 127–132.
- Roth, Alvin E., 1982b. The economics of matching: stability and incentives. *Mathematics of Operations Research* 7 (4), 617–628.
- Roth, Alvin E., 1984. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy* 92 (6), 991–1016.
- Roth, Alvin E., 1985. The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory* 36 (2), 277–288.
- Roth, Alvin E., 2002. The economist as engineer: game theory, experimentation, and computation as tools for design economics. *Econometrica* 70 (4), 1341–1378.
- Roth, Alvin E., 2008. What have we learned from market design? *The Economic Journal* 118 (527), 285–310.
- Roth, Alvin E., 2018. Marketplaces, markets, and market design. *The American Economic Review* 108 (7), 1609–1658.

- Roth, Alvin E., Peranson, Elliott, 1999. The redesign of the matching market for American physicians: some engineering aspects of economic design. *The American Economic Review* 89 (4), 748–780.
- Roth, Alvin E., Sotomayor, Marilda A. Oliveira, 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, Cambridge.
- Roth, Alvin E., Rothblum, Uriel G., 1999. Truncation strategies in matching markets-in search of advice for participants on JSTOR. *Econometrica* 67 (1), 21–43.
- Roth, Alvin E., Xing, Xiaolin, 1997. Turnaround time and bottlenecks in market clearing: decentralized matching in the market for clinical psychologists. *Journal of Political Economy* 105 (2), 284–329.
- Roth, Alvin E., Sönmez, Tayfun, Utku Ünver, M., 2007. Efficient kidney exchange: coincidence of wants in markets with compatibility-based preferences. *The American Economic Review* 97 (3), 828–851.
- Roth, Alvin E., Sönmez, Tayfun, Ünver, Utku M., 2005. Pairwise kidney exchange. *Journal of Economic Theory* 125 (2), 151–188.
- Roy, A.D., 1951. Some thoughts on the distribution of earnings. *Oxford Economic Papers*.
- Rust, J., 1987. Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher. *Econometrica* 55 (5), 999–1033.
- Shapley, L.S., Shubik, M., 1971. The assignment game I: the core. *International Journal of Game Theory* 1 (1), 111–130.
- Shapley, Lloyd, Scarf, Herbert, 1974. On cores and indivisibility. *Journal of Mathematical Economics* 1 (1), 23–37.
- Shi, Peng, 2015. Guiding school-choice reform through novel applications of operations research. *Interfaces* 45 (2), 117–132.
- Shi, Peng, 2021. Optimal priority-based allocation mechanisms. *Management Science*. Forthcoming.
- Shorrer, Ran I., Sovago, Sandor, 2019. Obvious Mistakes in a Strategically Simple College-Admissions Environment.
- Sorensen, Morten, 2007. How smart is smart money? A two-sided matching model of venture capital. *The Journal of Finance* 62 (6), 2725–2762.
- Su, Xuanming, Zenios, Stefanos, 2004. Patient choice in kidney allocation: the role of the queueing discipline. *Manufacturing & Service Operations Management* 6 (4), 280–301.
- Sweeting, Andrew, 2012. Dynamic pricing behavior in perishable goods markets: evidence from secondary markets for major league baseball tickets. *Journal of Political Economy* 120 (6), 1133–1172.
- Tanaka, Mari, Narita, Yusuke, Moriguchi, Chiaki, 2020. Meritocracy and Its Discontent: Long-run Effects of Repeated School Admission Reforms. Working Paper.
- Train, Kenneth E., 2009. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Uetake, Kosuke, Watanabe, Yasutora, 2020. Entry by Merger: Estimates from a Two-Sided Matching Model with Externalities.
- van der Vaart, A.W., 2000. *Asymptotic Statistics*. Cambridge University Press.
- Verdier, Valentin, Reeling, Carson, 2021. Welfare effects of dynamic matching: an empirical analysis. *The Review of Economic Studies*. Forthcoming.
- Vissing, Ashley, 2018. One-to-Many Matching with Complementary Preferences: an Empirical Study of Market Concentration in Natural Gas Leasing.
- Vulkan, Nir, Roth, Alvin E., Neeman, Zvika (Eds.), 2013. *The Handbook of Market Design*, 1 ed. Oxford University Press, Oxford.

- Waldinger, Daniel, 2021. Targeting in-kind transfers through market design: a revealed preference analysis of public housing allocation. *The American Economic Review* 111 (8), 2660–2696.
- Whinston, Michael Dennis, 2006. *Lectures on Antitrust Economics*. MIT Press, Cambridge.