

Empirical Models of Non-Transferable Utility Matching

Nikhil Agarwal and Paulo Somaini*

August 11, 2021

1 Introduction

Empirical models play a distinctive role in the study of matching markets. They provide a quantitative framework for measuring heterogeneity in preferences for schools ([Hastings et al., 2009](#)), comparing school assignment mechanisms ([Abdulkadiroglu et al., 2017](#); [Agarwal and Somaini, 2018](#); [Calsamiglia et al., 2020](#)), understanding preferences in the marriage market ([Chiappori et al., 2012](#)), and measuring the effects of market power in the medical residence match ([Agarwal, 2015](#)). The approach taken by these papers is based on first estimating the preferences of the agents in these markets and then using those estimates to make economic conclusions. This chapter provides a unified framework for analyzing agents' preferences in empirical matching models with non-transferable utility. Our objective is to provide a roadmap of the existing literature and highlight avenues for future research.

Unlike in the textbook model of consumer choice, in a matching market agents cannot choose with who to match with at posted prices. As examples, marriages are formed by mutual consent; schools do not admit all students willing to pay their tuition; public housing is allocated via waiting lists that use priorities; and employers and job applicants must compete with others in a labor market. Moreover, agents on both sides of the market have preferences over who to transact with, whereas firms rarely care about the identity of a consumer willing to meet their ask. Importantly, prices do not serve as the sole rationing device that equilibrates demand and supply.

Models of two-sided matching with non-transferable utility have three main features. First, the market has two sides – workers and firms, or schools and students – with agents on at least one side having preferences over agents on the other. This excludes one-sided matching problems, such as the matching of pairs of patients and living donors in kidney exchange markets. Second, agents form partnerships by mutual consent. And third, the terms of

*Agarwal: Department of Economics, MIT and NBER, email: agarwaln@mit.edu. Somaini: Stanford Graduate School of Business and NBER, email: soma@stanford.edu. With permission of the respective publishers, this chapter reuses portions from [Agarwal and Somaini \(2020\)](#) and adapts certain figures from [Agarwal and Somaini \(2018\)](#).

matching are fixed in advance. That is, either an agent on one side of the market cannot compensate an agent on the other side to enter into a partnership or such compensation, if any, is determined exogenously. This excludes models with transferable utility that are discussed in chapter ?? of this book.

Estimating preferences in matching markets depends crucially on the data that is available and institutional details about how the market is organized. We will consider two common data environments with the goal of identifying and estimating a random utility model further described in Section 2. In the first environment, the researcher has access only to data on realized matches. The second occurs when the researcher has data from a centralized assignment mechanism that elicits preferences in the form of rank-order lists of options. The appropriate approach to learning about the preferences of agents differs in these two environments, and the discussion in this chapter is split along these lines.

Learning about preferences when only data on final matches is available requires assuming a notion of equilibrium that results in these matches. The predominant concept used for this purpose is known as pairwise stability (see Chapter 3 of this book, and Roth and Sotomayor, 1992). If a match is stable, then one cannot find a pair of agents, one from each side, that would strictly prefer to match with each other relative to their current assignment. This concept is motivated by the idea that agents would have an incentive to break a match that is not stable. The challenge with using stability is that preferences on either side of the market could determine the final matches. Section 3 outlines various approaches for estimating preferences with such data.

In some cases, it is also possible to obtain more direct information on preferences when a centralized assignment mechanism is used to determine allocations. These mechanisms have had a big impact in numerous real-world contexts (Roth, 2018), including in school districts and residency programs around the world. The mechanisms require agents to indicate their preferences over being matched with agents on the other side of the market, which are then used as inputs in an algorithm that determines a final assignment. A researcher who knows the algorithm and has access to data on reported preferences has strictly more information than one with access only to data on realized allocations. A recent literature has developed methods for utilizing these data-reported preferences in a wide range of mechanisms. Estimating preferences is straightforward if participants report their true ordinal preferences but depends on the properties of the mechanism and behavioral assumptions in other cases. Section 4 provides a brief overview of these approaches.

A theme in this chapter is that the flexibility of the preference model that can be identified depends crucially on the available data. As we will discuss below, only data on final matches typically requires stronger independence assumptions than data on reported preferences. Similarly, our ability to estimate preferences in a one-to-one matching market will be more limited than in a many-to-one market, and estimating preferences of agents on both sides of the market is more demanding than only for one side. Our hope is to provide a guide to

these trade-offs on the current research frontier in a unified framework.

2 Empirical Model

We consider a two-sided matching market in which agents on one side, indexed by $i \in \mathcal{I} = \{1, \dots, I\}$, are assigned to agents on the other side, indexed by $j \in \mathcal{J} = \{1, \dots, J\}$. Agents $i \in \mathcal{I}$ may be matched with at most one agent in \mathcal{J} , but agents $j \in \mathcal{J}$ may be matched with $q_j \geq 1$ agents in \mathcal{I} . Agents on both sides may also remain unmatched.

Agents on each side of the market have preferences over the agents they match with. We will represent these preferences using a random utility model. Specifically, the utility of i from being matched with j is denoted by u_{ij} , and the utility of j from being matched with i is denoted by v_{ji} . In non-transferable utility models, these utilities cannot be changed by the agents in the market.¹ Moreover, the formulation implicitly assumes that the utility an agent receives from a given match does not depend on the other matches in the market. In particular, in the many-to-one context, the utility j receives from being matched with i does not depend on the other agents j is matched with.²

Our goal will be to identify and estimate the joint distributions of the vectors of random utilities $\mathbf{u}_i = (u_{i1}, \dots, u_{iJ})$ and $\mathbf{v}_i = (v_{1i}, \dots, v_{Ji})$ conditional on observable characteristics.³ Let u_{i0} and v_{j0} denote the utility of remaining unmatched.

We start by considering one side of the market. The most general form of 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 .⁴ 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., a measure of distance between i and j) that we will use as a numeraire, and represents a metric for utility.

¹Theoretical models of non-transferable utility are also closely related to settings that involve matching with contracts (see [Hatfield and Milgrom, 2005](#); [Hatfield et al., 2013](#)). We are not aware of empirical work that directly works with such models. The most closely related work is on models with imperfectly transferable utility, by [Galichon et al. \(2019\)](#).

²This assumption is sometimes referred to as “responsive preferences” (see [Roth and Sotomayor, 1992](#), Chapter 5).

³We refer the reader to [Matzkin \(2007\)](#) for the formal definition of identification that we employ in this chapter.

⁴Recent results in [Allen and Rehbeck \(2017\)](#) suggest that it may be possible to generalize this specification to allow utility to be non-linear but still separable in d_{ij} .

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.⁵

Identification and estimation of this model are usually studied in an environment in which the random variables for each i and each j are independent and identically distributed draws from some population distribution. Moreover, we will typically assume a conditional independence condition:

$$\epsilon_i \perp \mathbf{d}_i \mid \mathbf{z}_i, \mathbf{x}, (\xi_j)_{j=1}^J, \quad (2)$$

where $\mathbf{d}_i = (d_{i1}, \dots, d_{iJ})$ and $\mathbf{x} = (x_1, \dots, x_J)$. The independence condition (2) assumes that agent i 's unobserved taste shocks are independent of the vector of numeraire match-specific characteristics \mathbf{d}_i , conditional on the other observed characteristics of i , \mathbf{z}_i , and the vectors of observed and unobserved characteristics of agents on the other side of the market, $\mathbf{x}, (\xi_j)_{j=1}^J$, respectively. The assumption must be evaluated within each empirical application, and it is typically reasonable if \mathbf{x} is a sufficiently rich control.⁶

Example. (School Choice) In this canonical example, let i denote a student and j denote a school. The term u_{ij} is the utility that student i , or her parents, derive from being matched with school j . If we let d_{ij} be the distance from i 's residence to j 's location (as in [Abdulkadiroglu et al., 2017](#), for example), then student's or her parents' preferences can be summarized in terms of their "willingness to travel." In this example, the conditional independence assumption requires that the distance to school be 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 achievement, demographics, and socioeconomic 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.

While the analysis of identification can often allow for general functional forms, empirical methods will typically use additional parametric assumptions to ease the computational burden and to get statistically precise estimates with finite sample sizes. The most convenient functional forms depend on available data and the mechanism or setting being analyzed.

A commonly used parametric form encompassed by the above model assumes that

$$u_{ij} = \mathbf{x}'_j \beta + \mathbf{x}'_j \bar{\gamma} \mathbf{z}_i + \xi_j + \mathbf{x}'_j \gamma_i + \epsilon_{ij} - d_{ij}, \quad (3)$$

⁵The above specification also assumes that all agents dislike increases in d_{ij} . This restriction is not essential in many cases as discussed below, and the sign of this coefficient can be estimated.

⁶Relaxing this assumption is a fruitful avenue for future research. It likely requires augmenting the model to incorporate other sources of exogenous variation and specifying how it affects the data-generating process.

where γ_i and ε_{ij} are mean-zero, normally distributed random variables with variances to be estimated, and $\bar{\gamma}$ is a matrix conformable with \mathbf{x}'_j and \mathbf{z}_i . We denote by θ the vector of unknown parameters of the model, namely, $(\beta, \bar{\gamma}, \xi_1, \dots, \xi_J)$ and the parameters governing the distribution of ε_{ij} and γ_i . These functional form assumptions may be varied depending on the application.

This formulation is both tractable and flexible. It can capture many determinants of preferences and can be used to measure various aspects of preferences including quality differences across j due to observable and unobservable characteristics, preference heterogeneity across i due to observed and unobserved taste shifters, and also idiosyncratic preferences via ε_{ij} .

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](#)). 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

$$v_{ji} = v(\mathbf{x}_j, \mathbf{z}_i, \eta_i) - w_{ji}, \quad (4)$$

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 and strengthen the independence assumption (2) to

$$(\eta_i, \epsilon_i) \perp (\mathbf{d}_i, \mathbf{w}_i) \mid \mathbf{z}_i, \mathbf{x}, (\xi_j)_{j=1}^J, \quad (5)$$

where $\mathbf{w}_i = (w_{1i}, \dots, w_{Ji})$.

A few assumptions in the preference model are worth pointing out. First, there are no externalities. An agent's utility only depends on her own matches. This rules out preferences for attending a school with specific peers or working in a firm with specific colleagues, for example. The current approach to capturing such preferences is to include characteristics of the student body from prior years, but capturing such peer effects is an important avenue for future research; see [Section 5](#) for a more detailed discussion.⁷ 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 \(2018\)](#), who considers the possibility that preferences evolve after agents receive an initial assignment.

Third, the model assumes that unobservable characteristics are independent of observables. This assumption may be violated for several reasons, including if certain observed characteristics are chosen endogenously by agents in the market. For example, schools may invest in quality by hiring teachers in order to obtain a better set of students. Such endogenously selected characteristics pose empirical challenges because we cannot directly interpret the

⁷[Epple et al. \(2018\)](#) and [Allende \(2019\)](#) are two notable exceptions.

measured relationship between the observed characteristics and the preferences as causal. Additionally, counterfactual situations that change the allocation mechanism or the market's competition level may alter the incentives to invest in the characteristic. Methods that account for these sources of endogeneity when analyzing counterfactuals deserve further research.

In this chapter, we will analyze preferences in two different data environments. In the first environment, we will assume access only to information on final matches, and we will observe the identities of the agents in the market and their characteristics. We will use the assumption of stability defined in chapter 3 in our empirical approach. In the second environment, we will assume access to information on the preferences of each agent as reported to an assignment mechanism. These reports need not be truthful. Chapters 3 and ?? describe some commonly used mechanisms that we will study.

3 Analysis Using Final Matches and Stability

This section reviews different approaches when we observe only data on the realized matches from a single large matching market. Throughout, we will assume that these matches are stable (see Chapter). We distinguish between the following three types of markets. The first type is one-to-one matching with a large number of agents on each side. A canonical example is the market for marriages between women and men.⁸ The second type is few-to-one matching. Each agent in \mathcal{J} can match with only a few agents in \mathcal{I} . Again, we assume that both I and J are large. The market for medical residents is the canonical example of this type of matching market. Other entry-level labor markets with no salary bargaining can also fit this category if stability is a reasonable assumption. The third type is many-to-one matching. In this type, each agent in \mathcal{J} can match with many agents in \mathcal{I} . We assume that I is large but J is small. The leading example of this type of setting is the assignment of students to schools.

3.1 One-to-one Matching

Consider a market in which each agent in \mathcal{J} can match with at most one agent in \mathcal{I} . Assume that both I and J are large. The analysis of this model has been split into two types. 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 the case where preferences are heterogeneous such that two agents may have differing preferences over agents on the other side of the market. We discuss each of the two cases below.

⁸In other types of marriages, the matches cannot be described using a bipartite graph and are therefore not two-sided.

3.1.1 Double-Vertical Model

In this model, all agents on each 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 equation (2) by $\xi_j \perp \mathbf{x}_j$. This model omits both observed and unobserved sources of quality heterogeneity, resulting in a desirability 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.

Assume that we have access to data on the observable characteristics of the agents in a matching market. Therefore, we can identify and estimate the joint distribution $F_{X,Z}$ of the observable characteristics of matched agents. Since we are studying the marriage market between men and women, we follow 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

$$u(\mathbf{x}_j) = F_U^{-1}(F_V(v(\mathbf{z}_i) + \eta_i)) - \xi_j. \quad (6)$$

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 types that end up matching with women with the same distributions of observable characteristics. The same reasoning allows us to identify iso-attractiveness profiles for women.⁹ Chiappori et al. (2012) allow v_i to depend on two observable characteristics: body-mass index (BMI) and wages.

To formalize the intuition, suppose that for any function ϕ_x of observables \mathbf{x}_j there exist a function ϕ_v of the index v such that $\mathbb{E}(\phi_x(\mathbf{x}_j)|z_i) = \phi_v(v(z_i))$, 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. By differentiating both sides with respect to two components of z_i , we can measure the marginal rate of substitution

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

because the right-hand side is observed. Since the above argument provides only 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 that there is a characteristic that is known to be valued monotonically and is desirable.

A limitation of one-to-one matching data is that we are able to assess only 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.¹⁰ We will discuss how to identify the relative contribution of observables and unobservables when we discuss few-to-one matches below.

3.1.2 Heterogeneous Preferences

A strong restriction in the above model is that all agents have the same preferences over the agents on the other side of the market. To make progress on relaxing this assumption, Menzel (2015) revisited Dagsvik’s (2000) model in which the utilities are parametrized as follows:

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

⁹While Chiappori et al. (2012) consider more general preferences, their main results hold for the double-vertical model. More specific assumptions are required in the general case (see Chiappori et al., 2020).

¹⁰Observe, however, that the polar cases where η_i and ξ_j are both identically equal to zero for all i and j can be ruled out. This is because equation (6) reduces to $u(\mathbf{x}_j) = F_U^{-1}(F_V(v(z_i)))$. In this case, men with a given set of characteristics z_i match with women whose observables lie exactly on the iso-attractiveness curves described in Figure ??.

where $u_{i0} = 0 + \max_{k=1, \dots, J} \{\varepsilon_{i0,k}\}$, 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 standard Gumbel, or logit for short. 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 a fixed fraction of agents prefer to remain unmatched 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 types of agents matched with each other, denoted by $f(\mathbf{x}, \mathbf{z})$, has a very tractable functional form. Specifically, we get that

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

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 that i gets matched with j is the probability that j is i 's most preferred option in the set \mathcal{J}_i . And, similarly, i must be j 's most preferred option. Each of these probabilities are given by a logit-like formula in a large market, 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.¹¹

Another approach, due to [Sorensen \(2007\)](#), is to assume that matches depend only on the joint surplus $S(\mathbf{x}_j, \mathbf{z}_i) + \eta_{ij}$, but that the partners split this surplus via Nash bargaining after the match is formed. That is, side \mathcal{I} receives $\lambda S(\mathbf{x}_j, \mathbf{z}_i) + \lambda \eta_{ij}$ from a realized match and side \mathcal{J} receives $(1 - \lambda) S(\mathbf{x}_j, \mathbf{z}_i) + (1 - \lambda) \eta_{ij}$ for some $\lambda \in [0, 1]$. This model implies a constant ratio u_{ij}/v_{ij} for every pair of agents. Using the terminology of [Niederle and Yarov \(2009\)](#), 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 $S(\mathbf{x}_j, \mathbf{z}_i)$ directly.

These results stress the limitation of data from one-to-one matches, this time in a model with heterogeneous preferences. Namely, only some aggregate notion of joint surplus have been shown to be identified.¹² As in the double-vertical model, the difficulty still lies in trying to determine whether the preferences on side \mathcal{I} or on side \mathcal{J} are driving the observed matches.

¹¹This formula is remarkably similar to Choo and Siow's formula for transferable utility models (see Chapter ??). In both cases, the relative frequencies of matches between types \mathbf{x}, \mathbf{z} depends monotonically on $u(\mathbf{x}, \mathbf{z}) + v(\mathbf{x}, \mathbf{z})$.

¹²In early work, [Logan et al. \(2008\)](#) used a Bayesian approach to estimate heterogeneous preferences of both men and women over their partners. However, we are not aware of results that show identification in this model.

3.2 Few-to-one Matching

The above discussion demonstrates that only some features of preferences – either the portion of utility determined by observables or the sum of the surpluses – are identified given only data on final matches in a one-to-one matching market. However, answers to certain questions may require that the distribution of preferences on both sides be separately identified. For example, it is not sufficient to know only $u(\cdot)$ if we are interested in assessing the probability that a type \mathbf{x} is preferred to \mathbf{x}' . This probability depends on the full distribution of preferences.

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 result for the case of double-vertical preferences. As argued in [Section 3.1.1](#), 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 $u_j = u(\mathbf{x}_j) + \xi_j$ and $v_i = v(\mathbf{z}_i) + \eta_i$, 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 places only a single restriction on the behavior expressed in [equation \(6\)](#), but there are two unobservables in the model, ξ_j and η_i . In other words, the matches are determined 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. Examples of many-to-one markets include labor markets such as the iconic medical residency match ([Roth and Peranson, 1999](#)) as well as education markets such as college or school admissions.

As before, if preferences on both sides are vertical, matches are stable if and only if they exhibit perfect sorting. Formally, in such a market, consider a pair of students i and i' matched with the same college j . [Equation \(6\)](#) generalizes to

$$\begin{aligned} u(\mathbf{z}_i) &= F_U^{-1}(F_V(v(\mathbf{x}_j) + \xi_j)) - \eta_i \\ u(\mathbf{z}_{i'}) &= F_U^{-1}(F_V(v(\mathbf{x}_j) + \xi_j)) - \eta_{i'} \end{aligned} \tag{7}$$

As in the marriage market problem, the lack of perfect sorting based on observables indicates the presence of the errors η_i and ξ_j . Now, however, the composition of the incoming class in each program provides additional information about the contribution of each of the error terms. The expressions [\(7\)](#) 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). The method consists of the following steps. First, define a set of moments of the data m that we will try to match 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.¹³

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 \mathbf{x}_j and \mathbf{z}_i are column vectors; thus, $\mathbf{x}_j \mathbf{z}'_i$ is a matrix. Averaging this matrix over all matches yields

$$\frac{1}{I} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} 1 \{\mu(i) = j\} \mathbf{x}_j \mathbf{z}'_i.$$

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

The first set of moments summarizes aggregate sorting patterns based on observable characteristics, similar to the information used in Chiappori et al. (2012). The second and third sets of moments include additional information that is required to identify the contribution of each of the error terms.¹⁴

Estimates of this model using data on matches from the market for family medicine residents show that several non-salary observables and unobservables make significant contributions

¹³Train (2009) provides an overview of best practices.

¹⁴Sorting patterns and simulation-based estimation methods have also been used in Boyd et al. (2013) to estimate the preferences of teachers 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.

to the programs' desirability indices, implying that residents are willing to forego higher salaries for training in more desirable programs. This preference gives desirable programs market power, allowing them to levy an implicit tuition of over \$23,000 through a markdown in salaries.

3.3 Many-to-one Matching

We now consider settings in which agents on side \mathcal{J} can match with a large number of agents on side \mathcal{I} , when the number of agents on side \mathcal{J} is small. The canonical examples of such settings are school and college admissions. We will therefore refer to agents on side \mathcal{I} as students and to agents on side \mathcal{J} as schools.

There are two relevant types of data in these settings. The first is when the preferences or priorities used by schools to admit students are known. That is, the researcher can directly ascertain how two students will be ranked, possibly up to a random tie-breaker. For example, many school districts grant priority to students in their walkzone and to students who have siblings already enrolled, and many college systems prioritize students using only high school grades or entrance exam scores. In this case, the researcher only needs to estimate the preferences of the students for the schools.

The second case is more challenging as we need to estimate preferences on both sides of the market. 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 where only data on final matches is available, assuming that pairwise stability is satisfied. As before, this assumption requires justifications based on theory and institutional details 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. [Azevedo and Leshno \(2016\)](#) show 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 $\mathbf{v}_i = (v_{1i}, \dots, v_{Ji})$. The vector of cutoffs $\mathbf{p} = (p_1, \dots, p_J)$ is set so that the total number of students i with school j as their preferred option within the set $S(\mathbf{v}_i; \mathbf{p})$ does not exceed the capacity q_j at the school if $p_j > 0$. Moreover, in a market with an infinite number of students and a fixed number of schools, the stable match and the corresponding cutoffs p_j are unique.

3.3.1 Known Priorities: School Choice

Suppose we know each student’s eligibility score for each school, denoted by v_{ji} , up to a tie-breaker and the final assignment is stable. That is, $v_{ji} = v_j(\mathbf{z}_i, \eta_i)$, where the function $v_j(\cdot)$ and the distribution of the tie-breaker η_i are known. The cutoff scores p_j can be computed as the lowest eligibility score v_{ji} of a student that was matched to school j if the school does not have spare capacity. Otherwise, the cutoff p_j is equal to 0. The goal is then to estimate and identify the specification of preferences defined in equation (1).

This model is used by [Fack et al. \(2019\)](#) to study high school admissions in Paris, 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, [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. Figure 1 shows five regions of utilities denoted by Roman numerals. Each region implies different ordinal preferences except for region V, which pools the cases of $u_{i0} > u_{i1} > u_{i2}$ and $u_{i0} > u_{i2} > u_{i1}$. A student who is eligible for both schools will be assigned to school 1 if her utilities belong to either region I or II. Therefore, the share of students assigned to school 1 out of 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 Figure 2. In the later case, we can infer that $u_{i1} < u_{i0}$, which is lightly shaded. The share of students assigned to school 1 out of these students is an estimate of the total probability in regions I, II and, III of Figure 1.

These arguments are similar to those for standard consumer choice models but differ crucially in that not all students are assigned to their most preferred 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 scores \mathbf{v}_i requires 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 equation (1) are conditionally independent of eligibility given the observables included in the model. Formally, they require that

$$\epsilon_i \perp \mathbf{v}_i \mid \mathbf{z}_i, \mathbf{d}_i, \{\mathbf{x}_j, \xi_j\}_{j=1}^J. \quad (8)$$

This assumption may be a reasonable approximation if \mathbf{z}_i contains a rich set of student

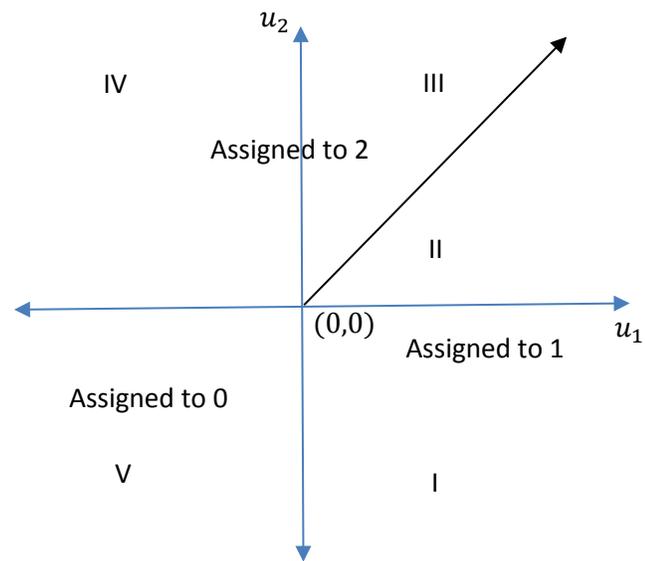


Figure 1: Stability – Both schools are feasible

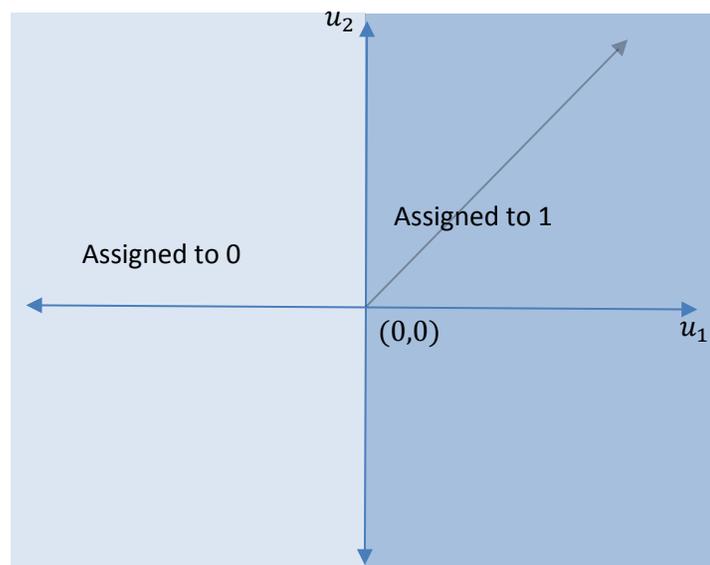


Figure 2: Stability: Only One School is Feasible

characteristics but can be violated, for example, if eligibility scores can be correlated with both unobserved student ability and unobserved preference parameters.

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 equations (2) and (8) 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

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

This expression enables estimation via maximum likelihood or other likelihood-based methods. Specific functional forms that are convenient for estimation are further discussed in Agarwal and Somaini (2020).

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, as a consequence, changes 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 Figure 1. This probability is identified in the two-school case if both schools are in the choice set or the assumption (8) 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, equations (1) and (2) 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.¹⁵

3.3.2 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 the leading example. The empirical challenge is not limited to estimating preferences for colleges. Because college preferences are unknown, it is not possible to make the same revealed preference arguments for students that we derived in the school choice context.

¹⁵It is also possible to develop the same identification argument using any other region in Figure 1. We choose region V because it is the negative orthant, which results in simpler expressions. Therefore, this model is over-identified.

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

- 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 i to college 1.

Agarwal and Somaini (in progress) show how to learn about preferences on both sides of the market simultaneously for the model described by equations (1) and (4). In the model discussed in Section 3.3.1, 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. (2020) yields a similar result under more stringent restrictions on equation (1).¹⁶

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 . Unfortunately, the location of v_{ji} is not identified when capacity is not known or when capacity limits are binding. This is because for students who strongly prefer college j but were not admitted, we can only deduce that $v_{ji} < p_j$. Hence, only the distribution of the difference $v_{ji} - p_j$ is identified.

Methods for estimating this model are still a subject of ongoing research. We refer the interested reader to He et al. (2020), Agarwal and Somaini (in progress), and Agarwal (2015).

4 Analysis Using Reported Preferences

Many centralized matching algorithms ask participants to submit a report that is used as an input to determine a match. Correspondingly, a well-developed literature has taken advantage of the additional information contained in these reports and derived methods to estimate agents' preferences based on these reports. Methods for analyzing preferences in this richer

¹⁶Specifically, He et al. (2020) 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 (in progress) 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 $g(\cdot)$.

data environment allow for more general preference models than the methods described in the previous section.

As in the many-to-one matching case, we call agents on side \mathcal{I} students and agents on side \mathcal{J} schools, since most of the applications studying rank-order data have been in the context of school choice. This focus corresponds to the widespread use of centralized mechanisms to assign students to schools. However, the methods discussed below are generally applicable to other settings where a researcher can obtain data on preferences.¹⁷

The discussion below is brief and focuses on the main differences from the previous section. We refer the reader to [Agarwal and Somaini \(2020\)](#) for a more thorough review, including a more in-depth discussion of various parametric models and methods for estimating preferences.

4.1 Truthful Reports

An important goal when designing assignment mechanisms is strategy-proofness ([Roth, 1982](#); [Abdulkadiroglu and Sonmez, 2003](#)). In such a mechanism, no student can benefit from submitting a list that does not rank schools in order of her true preferences. Strategy-proofness of a school choice mechanism can also enable an empirical strategy if agents understand it and follow this recommendation. Specifically, if agent i ranks j above j' , then we can infer that

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

It is less clear how to treat schools that are not ranked on the list. One approach is to assume that students rank all schools that are acceptable, i.e., preferable to the outside option. In this case, 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 [Figure 3](#) 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 or an empty list is submitted.

Observe that the rank-order lists provide richer information about preferences than in standard discrete choice models in which a consumer picks only her 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 [Figure 3](#), but we cannot distinguish between these two regions. The richer information in

¹⁷For 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. [Banerjee et al. \(2013\)](#) perform a similar analysis using a dataset of individuals placing and responding to matrimonial advertisements in a newspaper. They asked ad-placers to rank the letters they received and list the letters they are planning to follow up on.

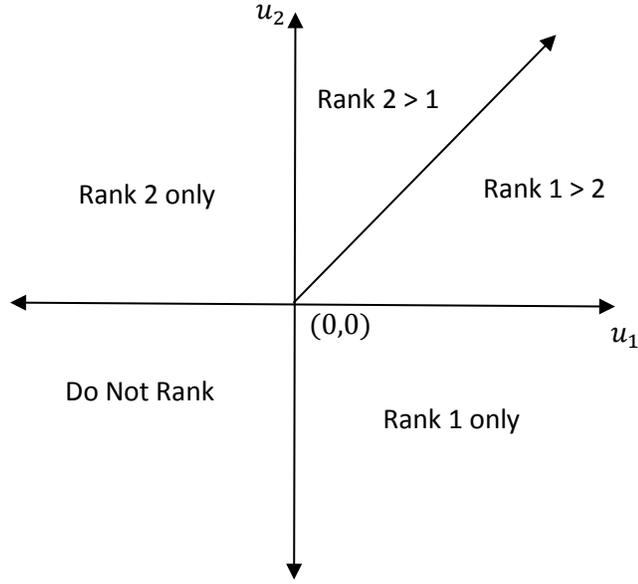


Figure 3: Revealed Preferences: Truthful Reports

ordered lists can help identify heterogeneity in preferences. In the school choice context, students often rank many more schools, allowing for very rich specifications for the distribution of utilities [Abdulkadiroglu et al. \(see 2017, for examples\)](#). Allowing for such heterogeneity is important for accurately estimating the value of improving assignments.

As in [Section 3.3.1](#), our goal is to identify the cdf F_{U^*} , i.e., 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 that the condition in [equation \(2\)](#) 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

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

An important point to note is that we can do away with the independence assumption in [equation \(8\)](#) required in [Section 3.3.1](#). This advantage is due to our ability to deduce whether or not $u_{ij} > u_{ij'}$ depends on the endogenous choice set of the agent as it did in [Section 3.3.1](#). This is the case because the model assumes that agents report preferences truthfully irrespective of the preferences of agents on the other side of the market.

4.2 Manipulable Mechanisms

There are many school districts that use non-strategy-proof mechanisms. The widely criticized but still commonly used Immediate Acceptance mechanism, for example, prioritizes students who rank a school higher, generating strategic incentives. To understand what can be learned from reports in manipulable mechanisms, it is useful to think of reports as actions in a game. Each action is associated with an expected payoff. If agents maximize expected utility, we can infer that the observed report yields the highest expected payoff. This approach assumes a considerable degree of sophistication as it requires agents to perform two cognitively demanding tasks. First, they have to be able to calculate the expected payoff for each possible report. Second, they have to maximize over all possible reports. We focus on the case where agents have rational expectations and can optimize, and conclude the section by 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, the collection of sets C_R for $R \in \mathcal{R}_{\mathcal{I}}$ partitions the space.¹⁸ Figure 4 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.

The discussion implicitly assumes that the vectors L_R for $R \in \mathcal{R}_{\mathcal{I}}$ are known to the researcher. In practice, they have to be estimated. Under rational expectations, these beliefs are objective assignment probabilities. Agarwal and Somaini (2018) show that almost all of the mechanisms used in practice can be described using a cutoff structure analogous to the one that applies to stable allocations. The distribution of these cutoffs in equilibrium determines the objective assignment probabilities. Thus, instead of estimating L_R , one can estimate the cutoff distribution instead, which is a lower-dimensional object. The cutoff structure is also useful to estimate 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

¹⁸More precisely, every $\mathbf{u} \in \mathbb{R}^J$ belongs to the interior of at most one of the sets in the collection and it 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.

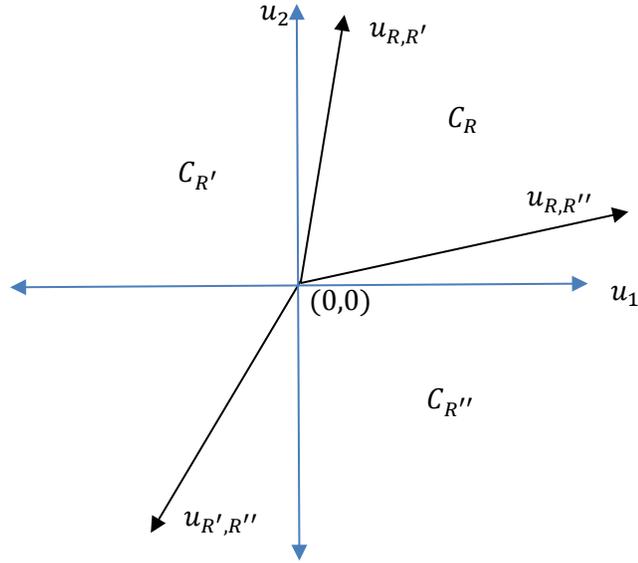


Figure 4: Manipulable Mechanisms

must be greater than the expected utility from any alternative report R . This expression forms the basis of estimation via maximum likelihood.

Several extensions have been based on this approach, but under alternative behavioral assumptions. [Kapor et al. \(2020\)](#) propose estimating \mathbf{L}_R by surveying agents. They surveyed families participating in the school choice mechanism in New Haven and found significant differences between elicited and objective assignment probabilities. [He \(2017\)](#) and [Hwang \(2016\)](#) do not impose all the conditions imposed by optimality. Instead, they derive a few intuitive conditions that reports have to satisfy and use the implied revealed preference relations to estimate preferences. Their approaches result in incomplete models of behavior that 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., the latter report their true ordinal preferences even if it is in their interest to report something else. A more detailed survey of methods for incomplete and mixture models is provided in [Agarwal and Somaini \(2020\)](#).

5 Applications, Extensions, and Open Questions

The empirical analysis of matching markets and related markets is still in a relatively early stage of development. This section highlights topics and areas of research that are ripe for study or have become active recently.

5.1 Applications

Comparing Mechanisms The theory on matching mechanisms has informed the implementation of coordinated mechanisms in various settings. Quantifying the benefits of centralization requires credible estimates of agents’ preferences using the methods proposed in this chapter. For example, [Abdulkadiroglu et al. \(2017\)](#) use the implementation of the New York City High School assignment system, which is based on the deferred acceptance algorithm, to quantify the welfare effects of coordinated school assignment. They find that, following the reform that centralized the assignment process, students were placed in more desirable schools and were more likely to enroll in their assigned school. Their analysis also compared the new DA-based system to the old system and to alternatives using a distance-metric utility function. 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 of the DA-based system, suggesting that the primary gains arise from coordinating assignments.

The methods developed here have also been used to compare various coordinated mechanisms. A number of papers have compared the deferred acceptance mechanism with alternatives. Most of the work has compared this celebrated mechanism to the immediate acceptance mechanism (also known as the Boston Mechanism). The common finding in this literature is that the Immediate Acceptance mechanism yields higher utility to students in the best case ([Agarwal and Somaini, 2018](#); [Calsamiglia et al., 2020](#)), but that this conclusion is not robust if students make mistakes ([Agarwal and Somaini, 2018](#); [Kapor et al., 2020](#)). We point the reader to [Agarwal and Somaini \(2020\)](#) for a more detailed review of this literature.

Rationing and Redistribution Regulated prices and capacity constraints can result in rationing and redistribution. Distributional concerns are particularly important in education ([Hastings et al., 2009](#); [Calsamiglia et al., 2020](#)), health care ([Agarwal et al., 2021](#)), and social assistance ([Waldinger, 2020](#)). The design of real-world allocation mechanisms needs to consider ethical and political constraints alongside traditional issues pertaining to efficiency and incentives. For example, school districts often implement quotas to equalize access to high-quality education, and publicly provided health care may need to ensure that all citizens receive adequate treatment. Preference estimates are a key tool for understanding the performance and distributional consequences of various mechanisms.

The tools described above are also useful when centralized assignment mechanisms are not used. For example, rationing also occurs if providers can choose who to treat based on their preferences. As in other matching markets, preferences on both sides of the market determine the final allocation.¹⁹ [Gandhi \(2020\)](#) takes this view when studying the market for nursing

¹⁹The key feature is that, unlike standard consumers, patients cannot choose their most preferred healthcare

homes in California. The paper finds that providers tend to discriminate against Medicaid-eligible patients who require lengthy anticipated stays because they are not profitable.

5.2 Extensions

There are several directions in which the methods discussed in this chapter can be extended. We discuss a few below:

Dynamics The models described in this chapter are static, in the sense that all agents arrive to the market simultaneously and match once and for all. However, there are environments like the markets for child care, public housing, and organ transplants in which 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. This decision is informative about agents' preferences in the same way that reports are for static allocation systems. Leveraging this insight, [Agarwal et al. \(2021\)](#) study the system that allocates deceased donor kidneys in the U.S., [Waldinger \(2020\)](#) public housing allocation; and [Reeling and Verdier \(2020\)](#) allocation of bear-hunting permits.

Externalities: Peer Effects and Competition A central assumption in the framework described above is that each agent has preferences only over the agents with whom they match. There is relatively little work on externalities, whereby the matches of others also affect an agent's payoffs. There are at least two important reasons why the matches of others may be important.

The first reason can be classified as peer effects. For example, students may derive utility from their classmates, and workers may have preferences over their co-workers. In this case, i 's preferences over j can depend on the set $\mu^{-1}(j)$. There is limited work on education markets that have 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.

The second reason is due to competitive effects. These are particularly important in industrial organization settings. For example, [Uetake and Watanabe \(2019\)](#) 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 that hold mineral

provider at posted prices. Instead, the provider can exercise discretion on who to treat. Therefore, insurance contracts may also result in rationing.

rights. In this model, the terms that an oil company can negotiate depend on their overall market presence. A challenge in these settings is to find an appropriate notion of stability that allows for externalities to be present.²⁰

6 Conclusion

Estimating preferences is a crucial first step toward understanding the effects of policy interventions in a matching market. Preference estimates enable both positive and normative analyses. Specifically, they can be used to predict how agents will behave after an intervention is implemented and how the allocation will change. They are also central to evaluating the welfare and distributional effects of such interventions.

However, standard tools for estimating consumer demand are not directly applicable in matching markets since prices do not clear the market. Instead, agents choose among an individualized set of options determined by the agents on the other side of the market that are willing to match with them. This choice set depends on the agent's desirability to the other side of the market. These features require the development of a new analytical toolset.

The appropriate toolset depends on whether the researcher has access to data only on final matches, or also to data on preferences submitted to an allocation mechanism. This chapter discussed methods for both cases, outlining the assumptions required in order to make progress in each environment.

A message of our chapter is that the data and institutional environment dictate whether a flexible preference model can be estimated. One-to-one matching markets require the most severe restrictions on preferences. The current set of results only justify estimating either homogeneous vertical preferences or some notion of aggregate surplus in models that allow some heterogeneity. More flexible models of preferences can be learned from markets with many-to-one matches. However, data on rank-order lists enables the estimation of the most flexible models of preferences.

While the methods discussed in this chapter have focused on two-sided matching, they also apply to many related environments. For example, they apply to environments where rationing occurs due to capacity constraints, including many education and healthcare markets, ranging from college admissions to the assignment of patients to nursing homes. Extensions of the two-sided matching models discussed here have also been used to describe the market for oil drilling leases and the allocation of organs to patients on a waiting list. We believe that there are many other applications and settings where insights from this rapidly growing literature can be applied in future research.

²⁰Conditions for existence of stable matchings with externalities is an active area of theoretical research. We refer the reader to [Pycia and Yenmez \(2019\)](#), [Fisher and Hafalir \(2016\)](#), and the references therein for some recent results.

References

- Abdulkadiroglu, Atila and Tayfun Sonmez**, “School Choice: A Mechanism Design Approach,” *The American Economic Review*, 2003, *93* (3), 729 – 747.
- , **Nikhil Agarwal**, and **Parag A. Pathak**, “The Welfare Effects of Coordinated School Assignment: Evidence from the NYC High School Match,” *American Economic Review*, 2017, *107* (12), 3635–89.
- Agarwal, Nikhil**, “An Empirical Model of the Medical Match,” *American Economic Review*, 2015, *105* (7), 1939–78.
- and **Paulo Somaini**, “Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism,” *Econometrica*, 2018, *86* (2), 391–444.
- and – , “Revealed Preference Analysis of School Choice Models,” *Annual Review of Economics*, aug 2020, *12* (1), 471–501.
- , **Itai Ashlagi**, **Michael A. Rees**, **Paulo Somaini**, and **Daniel Waldinger**, “Equilibrium Allocations Under Alternative Waitlist Designs: Evidence From Deceased Donor Kidneys,” *Econometrica*, 2021, *89* (1), 37–76.
- Akyol, Pelin and Kala Krishna**, “Preferences, selection, and value added: A structural approach,” *European Economic Review*, 2017, *91*, 89–117.
- Allen, Roy and John Rehbeck**, “Identification of Average Demand Models,” *Econometrica*, *forthcoming*, 2017.
- Allende, Claudia**, “Competition Under Social Interactions and the Design of Education Policies,” 2019.
- Azevedo, Eduardo M. and Jacob Leshno**, “A Supply and Demand Framework for Two-Sided Matching Markets,” *Journal of Political Economy*, 2016, *124* (5), 1235–1268.
- Banerjee, Abhijit, Esther Duflo, Maitreesh Ghatak, and Jeanne Lafortune**, “Marry for What? Caste and Mate Selection in Modern India,” *American Economic Journal: Microeconomics*, 2013, (2), 33–72.
- Becker, Gary S.**, “A Theory of Marriage: Part I,” *Journal of Political Economy*, jul 1973, *81* (4), 813–846.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff**, “Analyzing the determinants of the matching of public school teachers to jobs: Disentangling the preferences of teachers and employers,” *Journal of Labor Economics*, jan 2013, *31* (1), 83–117.

Bucarey, Alonso, “Who Pays for Free College? Crowding Out on Campus,” 2018.

Calsamiglia, Caterina, Chao Fu, and Maia Güell, “Structural Estimation of a Model of School Choices: the Boston Mechanism vs. Its Alternatives,” *Journal of Political Economy*, feb 2020, *128* (2), 642–680.

Chiappori, Pierre André, Sonia Oreffice, and Climent Quintana-Domeque, “Fatter attraction: Anthropometric and socioeconomic matching on the marriage market,” *Journal of Political Economy*, aug 2012, *120* (4), 659–695.

– , – , and – , “Erratum: Fatter attraction: Anthropometric and socioeconomic matching on the marriage market,” dec 2020.

Dagsvik, John K., “Aggregation in Matching Markets,” *International Economic Review*, feb 2000, *41* (1), 27–58.

Diamond, W. and N. Agarwal, “Latent indices in assortative matching models,” *Quantitative Economics*, 2017, *8* (3).

Epple, Dennis, Akshaya Jha, and Holger Sieg, “The superintendent’s dilemma: Managing school district capacity as parents vote with their feet,” *Quantitative Economics*, 2018, *9* (1), 483–520.

Fack, Gabrielle, Julien Grenet, and Yinghua He, “Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions,” *American Economic Review*, 2019, *109* (4), 1486–1529.

Fisher, James C.D. and Isa E. Hafalir, “Matching with aggregate externalities,” *Mathematical Social Sciences*, may 2016, *81*, 1–7.

Galichon, Alfred, Scott Duke Kominers, and Simon Weber, “Costly concessions: An empirical framework for matching with imperfectly transferable utility,” *Journal of Political Economy*, dec 2019, *127* (6), 2875–2925.

Gandhi, Ashvin, “Picking Your Patients: Selective Admissions in the Nursing Home Industry,” jun 2020.

Hastings, Justine S., Thomas J. Kane, and Douglas O. Staiger, “Heterogeneous Preferences and the Efficacy of Public School Choice,” 2009.

Hatfield, John William and Paul R. Milgrom, “Matching with contracts,” *American Economic Review*, sep 2005, *95* (4), 913–935.

– , **Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp**, “Stability and competitive equilibrium in trading networks,” *Journal of Political Economy*, oct 2013, *121* (5), 966–1005.

- He, Yinghua**, “Gaming the Boston School Choice Mechanism in Beijing,” 2017.
- , **Shruti Sinha**, and **Xiaoting Sun**, “Identification and Estimation in Many-to-One Two-sided Matching without Transfers,” 2020.
- Hitsch, Gunter J**, **Ali Hortacsu**, and **Dan Ariely**, “Matching and Sorting in Online Dating,” *American Economic Review*, mar 2010, *100* (1), 130–163.
- Hwang, Sam Il Myong**, “A Robust Redesign of High School Match,” 2016.
- Kapor, Adam**, **Christopher Neilson**, and **Seth Zimmerman**, “Heterogeneous Beliefs and School Choice Mechanisms,” *American Economic Review*, may 2020, *110* (5), 1274–1315.
- Logan, John Allen**, **Peter D Hoff**, and **Michael A Newton**, “Two-Sided Estimation of Mate Preferences for Similarities in Age, Education, and Religion,” *Journal of the American Statistical Association*, 2008, *103* (482), 559–569.
- Matzkin, Rosa L.**, “Nonparametric identification,” jan 2007.
- Menzel, Konrad**, “Large Matching Markets As Two-Sided Demand Systems,” *Econometrica*, 2015, *83* (3), 897–941.
- Narita, Yusuke**, “Match or Mismatch? Learning and Inertia in School Choice,” 2018.
- Niederle, Muriel** and **Leeat Yariv**, “Decentralized Matching with Aligned Preferences,” 2009.
- Pycia, Marek** and **M. Bumin Yenmez**, “Matching with Externalities,” 2019.
- Reeling, Carlson J** and **Valentin Verdier**, “Welfare effects of dynamic matching: An empirical analysis,” 2020.
- Roth, Alvin E.**, “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research*, nov 1982, *7* (4), 617–628.
- , “Marketplaces, Markets, and Market Design,” *American Economic Review*, jul 2018, *108* (7), 1609–1658.
- and **Elliott Peranson**, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 1999, *89* (4), 748–780.
- and **Marilda A. Oliveira Sotomayor**, *Two-Sided Matching*, Cambridge University Press, 1992.

Sorensen, Morten, “How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital,” *Journal of Finance*, 2007, 62 (6), 2725 – 2762.

Train, Kenneth E., *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009.

Uetake, Kosuke and Yasutora Watanabe, “Entry by Merger: Estimates from a Two-Sided Matching Model with Externalities,” 2019.

Vissing, Ashley, “One-to-Many Matching with Complementary Preferences: An Empirical Study of Market Concentration in Natural Gas Leasing,” 2018.

Waldinger, Daniel, “Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation,” 2020.

Index

Empirical model, 3

 random utility model, 3

 school choice, 4

Estimation using final matches

 many-to-one matching, 10

 one-to-one matching, 6

 stability, 6

Estimation using rank-order lists, 16

 manipulated reports, 19

 truthful reports, 17

Random utility model, 3

Stable matching, 6