Hedonic and Matching Models

draft of a *Handbook of Econometrics* chapter[1]

Alfred Galichon[2] **Bernard Salanié**[3]

November 13, 2024

[1]This research has benefited from support from the European Research Counsel under the consolidator grant No. 866274. The authors are indebted to Antoine Jacquet for helpful comments and to Isabella Giancola for excellent research assistance.

[2]New York University and Sciences Po, alfred.galichon@nyu.edu.
[3]Columbia University, bsalanie@columbia.edu.

Contents

Introduction

The term "hedonic" is widely used in economics and in the applied statistics literature, where it may refer to quite different concepts. In the broader sense, "price hedonics" is a tool to explore the relationship between the prices of differentiated products and their characteristics. As such, it has a long history, with contributions dating as far back as Waugh (1928). The development of the characteristic-based approach to consumer theory by Becker (1965) and Lancaster (1966) and the pioneering empirical work of Griliches (1961) brought it back to the fore. Since then "hedonic regressions" have often been used to adjust for quality shifts or technological upgrades in price indices[1].

Simple price regressions can only constitute reduced-form equations of an equilibrium system, however; and their estimated parameters combine demand and supply factors. In a fundamental contribution, Rosen (1974) introduced *hedonic equilibrium models* in which consumers buy differentiated varieties of a good from sellers. Consumers vary according to their

[1] In the US, the Bureau of Labor Statistics now uses hedonic adjustments for 34 of its 273 item categories. These include electric appliances, electronic equipment, and also clothing items (`https://www.bls.gov/cpi/quality-adjustment/home.htm`, consulted June 28, 2022).

willingness to pay for varieties with different characteristics; and producers have idiosyncratic production costs for each variety. These differences in utilities and costs induce a sorting between consumer and producer types in competitive equilibrium; this is mediated by a *hedonic price function* which maps the characteristics of an alternative into its equilibrium price.

This chapter will focus on such *competitive equilibrium* hedonic models and on related models of matching. The term "competitive" reflects a substantial restriction. While it greatly simplifies the analysis, markets for differentiated products may be far from perfect competition. Then markups need to be incorporated in the analysis, as argued by Pakes (2003).

The analysis in most of this chapter will also assume that agents' utilities and costs are quasilinear in prices, and that agents can transfer utility (or another numéraire) one-for-one, without cost or friction. This is usually called the perfectly transferable utility (TU) setting. Some of our results extend to utilities which are not quasilinear in prices; we will describe these imperfectly transferable utility (ITU) extensions in later sections. However, we shall always assume that transfers are possible. We will thus be leaving aside recent developments in the econometrics of matching with (perfectly) non-transferable utility (NTU), for which we refer the reader to the recent survey by Agarwal and Somaini (2023).

It may not be obvious a priori why hedonic markets and matching markets belong in the same chapter. As Roth (2015, p.4) eloquently put it, matching is about "how we get the many things we choose in life that also must choose us". To cite some leading examples: marriage partners (Becker (1973, 1974)) choose each other; so do hospitals and residents (Roth (1984));

colleges and students (Gale and Shapley (1962)); and employers and employees (Kelso and Crawford (1982)). This two-sided choice problem is *a priori* absent from hedonic markets: hedonic demand is just a particular way of representing consumer choice between passive objects. Similarly, hedonic supply has sellers choose which objects they bring to market.

To illustrate this difference, compare the house sale market and the house rental market. The seller of a house only cares about the price they receive, as the relationship with the buyer ends once the transaction is concluded. The sale market can be described quite naturally as a hedonic market. On the other hand, in the rental market the landlord cares about more than just the rental price; various characteristics of the other party matter to her, such as the ability of the tenant to pay in time and to maintain the housing unit. The rental market is therefore better described as a matching market[2].

In spite of these differences, the connection between hedonic and matching models is deep and fruitful. Fundamentally, the objects traded in a hedonic market can be seen as virtual intermediaries that match buyers and sellers. Their prices determine the equilibrium utilities in a corresponding matching market. We will show that the analogy is exact if we endow objects with a well-chosen (if virtual) objective function. Hedonic equilibrium, which requires that demand equal supply for each object, generates a matching of buyers and sellers via the objects they exchange in equilibrium. This analogy can be pushed further. The choice of an object by a consumer is a

[2] This is related to the distinction between private values and common values that plays a crucial role in hidden information models.

simple form of matching in which the object is passive; just as the choice of selling a particular object is matching a seller with a passive object. Hedonic equilibrium thus breaks down into two simple matching problems, which are linked by the equilibrium condition that reconciles the aggregate demand for each object to its aggregate supply. It is therefore not so surprising that hedonic models are closely related to matching models.

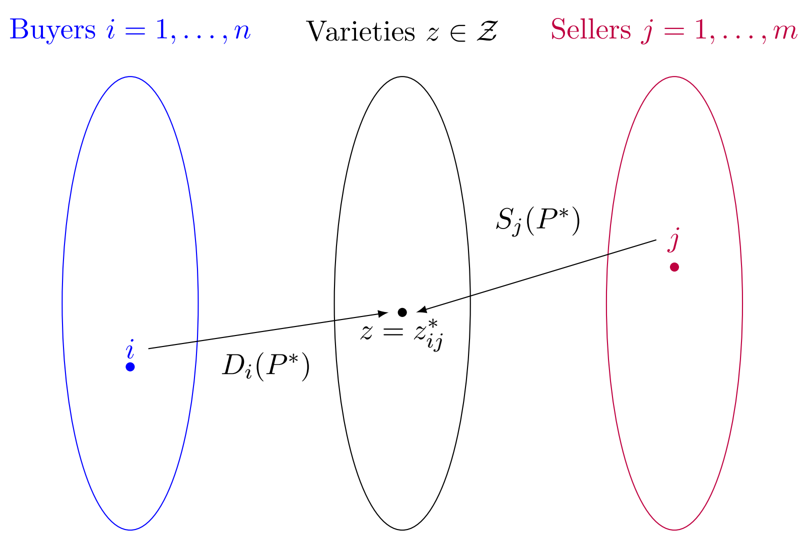

Figure 1: A Stylized Hedonic Market

In order to convey the intuition more precisely, we now present a stylized hedonic market. We consider the market for varieties z of an indivisible good. Given a price function for varieties $z \to P^*(z)$, each buyer i buys zero or one unit of the good; and each seller j has a production capacity of one unit . Buyers choose which variety of the good they buy, if any, and sellers choose which variety they produce, if any. Figure 1 shows a trade where a buyer i buys a unit of a variety z^*_{ij} from a seller j.

Let us now assume that the preferences of all buyers and sellers are quasi-linear, with a common and constant marginal utility of income: buyer i chooses the $z = D_i(P^*)$ that maximizes her net utility $U_i(z) - P^*(z)$, and seller j chooses the variety $z_j = S_j(P^*)$ that maximizes her profit $P^*(z) - C_j(z)$. From an efficiency perspective, it is natural to consider the variety that maximizes the joint surplus of the buyer and seller, that is the sum $\Phi_{ij}(z) \equiv U_i(z) - C_j(z)$ of consumer surplus plus seller profit. The value of the maximized surplus $\Phi^*_{ij} = \max_z \Phi_{ij}(z)$ is a monetary equivalent that captures the value of the trade. Furthermore, it is also natural to match buyers and sellers in a way that maximizes the total value of these joint surpluses. This leads to modeling hedonic equilibrium as maximizing the total surplus. Figure 2 illustrates the matching model that corresponds to the hedonic market of Figure 1.

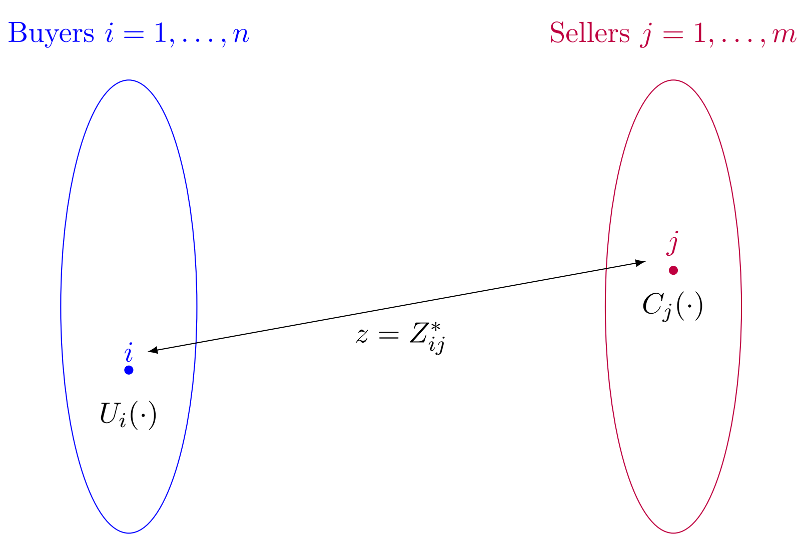

Figure 2: A Stylized Matching Market

In any trade, the price paid by the buyer to the seller determines how the joint surplus from the transaction is shared between the two parties. In bilateral monopoly, the mutually acceptable prices only need to ensure that both the buyer and the seller are above their reservation utility. The situation is very different in a competitive market, as agents have the possibility to seek another counterpart. In equilibrium, no seller-buyer pair could set up a new, mutually advantageous trade. As competitive equilibrium lies in the core, no such blocking pair may exist: this is the *stability* condition introduced by Gale and Shapley (1962).

Now any seller-buyer pair that does trade in equilibrium generates exactly the joint surplus; therefore *in equilibrium*, the sum of the equilibrium utilities of all buyers and sellers must equal the sum of the surplus from all trades. Since stability requires that the sum of utilities of any potential buyer-seller pair be at least as large as the surplus they could get by trading, in equilibrium the sum of all utilities must be the smallest possible among all stable matchings.

To summarize: minimizing the sum of utilities under the stability requirement defines a linear programming problem. It is closely related by duality to another linear optimization problem, which maximizes the total surplus under the feasibility (scarcity) constraints. The maximization of the total surplus, whose decision variables are the quantities traded, is traditionally called the *primal*, while total utility minimization under stability constraints, whose decisions variables are prices, is called the corresponding *dual*. The Duality Theorem for linear programming implies that the prices (and utilities) that solve the dual are the Lagrange multipliers associated with the feasibility constraints of the primal, while the traded quantities that maximize the total surplus in the primal are the multipliers of the stability constraints in the dual.

This powerful duality result relies heavily on the assumption that utility is perfectly transferable. It applies both in matching markets and in hedonic markets (reinterpreting a match as an assignment of a buyer, a seller, and the variety they trade.) As we will see, both frameworks can be interpreted as an *optimal transportation problem* whose solution yields the equilibrium prices, trades and utilities in the hedonic market, and the equilibrium matches and utilities in the matching market.

Section 1 derives more rigorously this relationship between hedonic equilibrium and stable matchings when utility is perfectly transferable. We start with some recent results in classic one-sided discrete choice problems; then we use them to analyze hedonic equilibria and stable matchings.

In most empirical work, the analyst only observes part of the information that is available to the agents. We will assume in this chapter that all agents are perfectly informed of all that is payoff-relevant to them. On the other hand, the econometrician only has access to a more limited set of variables; the rest is *unobserved heterogeneity.* In hedonic markets, it consists of unobserved characteristics of products, and how they enter the buyers' preferences and the sellers' costs. In matching markets, some characteristics of the partners are unobserved, as are idiosyncratic components of their preferences.

The explicit incorporation of unobserved heterogeneity into our models increases their plausibility by allowing agents who are indistinguishable from the econometrician's point of view to make different choices. It brings in new problems, however. Even in one-sided choice problems, little progress can be made without restricting the specification of unobserved heterogeneity. This difficulty is magnified in two-sided markets. Our chapter adopts the dominant approach in this field, which assumes a form of *separability* in the way observed and unobserved characteristics interact in payoff functions. Section 2.2.1 defines separability and the restrictions it entails: it rules out interactions between unobserved characteristics in the surplus of a trade (or a match). It still allows for "matching on unobservables" in a restricted sense, as well as unrestricted matching on observable characteristics.

The great value of separability lies in that it greatly enhances the tractability of this class of models. As we will see, it takes us from the realm of linear programming to convex programming. Moreover, in large markets the objects of interest are strictly convex and smooth, unlike the piecewise linear world of linear programming[3]. This allows us to apply the powerful tools of convex analysis. In particular, utilities (or prices) and matches (or trades) are linked in a bijective manner by *convex duality*.

Convex duality should be familiar to anyone who took a graduate microeconomics class, in the form of Shephard's Lemma in consumer theory or Hotelling's Lemma in producer theory. In demand systems and in hedonic and matching markets, it yields simple identification formulæ. These formulæ, which we derive in Section 2, allow the analyst to recover the underlying utility functions and/or the joint surplus from observed data.

Several methods can be used to estimate and test these models, most

[3]For a visual analogy, imagine going from a polygon to a circle.

of them directly inspired by the identification formulæ. We describe non-parametric and parametric estimators in Section 3, along with extensions to more complex matching markets. We also summarize there some of the main empirical applications.

Computational issues loom large in this field, especially in the context of parametric estimation. Section 4 builds on the connections we established with well-tested methods in applied mathematics; it shows how to adapt existing algorithms to solve for equilibrium in hedonic and matching markets.

While the past decade has seen remarkable progress in both methods and applications in this subfield, many issues remain unresolved. We end this chapter with promising research directions as we see them.

Chapter 1

Specification

Hedonic and matching models are two-sided equilibrium models, where agents on each side have preferences. Nevertheless, we start in this section with only one side of a hedonic market. This will allow us to introduce the objects and methods that lie at the core of our approach to two-sided markets: linear programming, optimal transport, and duality. We then show how when utilities are quasi-linear, hedonic equilibrium breaks down into two simple one-sided problems. This extends to bipartite matching markets when utility is perfectly transferable, provided that the joint surplus is separable. We conclude with an extension to matching with imperfectly transferable utility.

A word on notation: we denote a measure ν on a set S as $\nu(ds)$. This allows us to deal both with discrete and with continuous sets, as for any measurable subset T of S,

$$\nu(T) = \int \mathbb{1}(s \in T)\, \nu(ds).$$

If ν has a density f in s, then $\nu(ds) = f(s)ds$; if ν has a mass p at t, then locally $\nu(ds) = p\delta_t ds$, where $\delta_t(s) = \delta(s-t)$ and δ is the Dirac mass.

1.1 One-sided choice

As in the introduction, each agent chooses which variety of a good to purchase or sell. To simplify the exposition, we will focus on buyers choosing a "product," as in much of the empirical industrial organization literature. Our results apply equally well to the other side of the market (sellers deciding which varieties they will produce).

Both agents and products are heterogeneous. Each buyer has a *full type* which we denote $\tilde{x}$. The econometrician knows that $\tilde{x}$ is distributed according to a measure $\tilde{n}\,(d\tilde{x})$ on a set $\tilde{\mathcal{X}}$. In addition, she observes a part of the full type $\tilde{x}$ which we call the *observed type* x. Without loss of generality, we will often write $\tilde{x} = (x, \varepsilon)$ where ε is a random variable with a conditional distribution that derives from $\tilde{n}(d\tilde{x}|x)$; and we denote $n(dx)$ the distribution of the observed type x over its set of possible values $\mathcal{X}$.

We will assume in this section that each product's characteristics are fully observed by the econometrician: it takes a value z in a set $\mathcal{Z}$. The set $\mathcal{Z}$ of alternatives is typically a subset of a finite-dimensional vector space. It is often discrete, and even finite: the buyer may opt for an apple or for a pear. Sometimes it is better to model it as continuous, as with the square footage of a house.

The price of a product of full (and observed) type z is $P\,(z)$; the prices of all products that have non-zero market share are fully observed by the econometrician. There usually is a "zero option"; we denote it as 0, and we assume that $P(0) = 0$. All agents take the price function as given.

In some contexts, it would be useful to allow the quantity of variety z to take other values than zero or one. Formally, we could redefine z to be a vector of quantities, one for each variety. Again depending on the context, the price function might be required to be linear in quantities. In a further extension, each seller could be allowed to produce several varieties. While we will not pursue these extensions here, Section 3.2.3 considers corresponding variants of the matching model; see also Section 4.1 for a reinterpretation in terms of flows on a network.

1.1.1 Utility and demand

Let $U(\tilde{x}, z, p)$ denote the net utility of a buyer $\tilde{x}$ if he purchases product z at a price p. $U(\tilde{x}, z, p)$ is generally assumed to be continuous, and decreasing in p. In leading applications, its variations are restricted further. We now turn to four classic examples.

Example 1 (Additively separable Random Utility)**.** *In this setting, additively separable random utility (ARUM) rules out interactions between the price p and the unobserved type of the buyer. This translates into*

$$U(\tilde{x}, z, p) \equiv V(x, z, p) + \zeta(\tilde{x}, z). \tag{1.1}$$

The reason for this restriction will become clear when we discuss identification of hedonic demand models in Section 2.1.

Example 2 (Quasilinear Utility)**.** *In the quasilinear case, the marginal utility of money is constant and identical for all buyers:*

$$U\left(\tilde{x}, z, p\right) \equiv \bar{U}\left(\tilde{x}, z\right) - p. \tag{1.2}$$

Note that it is also additively separable as in Example 1, as the disutility of p is constant.

Example 3 (Logit case)**.** *Suppose that $\mathcal{Z}$ is discrete and $\boldsymbol{\varepsilon}$ is a vector of iid Gumbel variables, distributed independently of x. Let $\tilde{U}\left(\tilde{x}, z, p\right) = U\left(x, z\right) - p + \varepsilon_z$. This defines the multinomial logit discrete choice model; it is quasilinear and therefore additively separable.*

Example 4 (The Tinbergen model)**.** *An alternative subcase of Example 2 is the quadratic-in-characteristics case in which*

$$U\left(\tilde{x}, \boldsymbol{z}, p\right) \equiv \boldsymbol{z}'\boldsymbol{a}\left(x\right) + \boldsymbol{z}'\boldsymbol{B}\boldsymbol{z}/2 - p + \boldsymbol{z}'\boldsymbol{\varepsilon}, \tag{1.3}$$

where $\boldsymbol{a}(x)$ is a vector function, $\boldsymbol{B}$ is a symmetric matrix, and $\tilde{x} = (x, \boldsymbol{\varepsilon})$.

This model was pioneered by Tinbergen (1956); its properties are analyzed by Epple (1987) and Ekeland, Heckman, and Nesheim (2004). It allows the marginal willingness to pay for a characteristic

$$\boldsymbol{a}(x) + \boldsymbol{B}\boldsymbol{z} + \boldsymbol{\varepsilon}$$

to vary with the full type of the buyer via $\boldsymbol{a}(x)$ and $\boldsymbol{\varepsilon}$, and to depend on $\boldsymbol{z}$ itself (uniformly across buyers).

1.1.1.1 Individual Demand

Faced with the price function $P(\cdot)$, the buyer will purchase a product that maximizes $U(\tilde{x}, z, P(z))$ over $z \in \mathcal{Z}$. Equivalently, let $\tilde{\pi}^D(dz|\tilde{x}; P)$ denote the conditional choice probability of a buyer of type $\tilde{x}$ over the products in $\mathcal{Z}$ given the price function P. The support of this measure must be included in the set of maximizers over z, which we denote $Z^D(\tilde{x}; P)$:

$$\text{supp } \tilde{\pi}^D(\cdot|\tilde{x}; P) \subseteq Z^D(\tilde{x}; P) \equiv \arg\max_{z \in \mathcal{Z}} U(\tilde{x}, z, P(z)). \tag{1.4}$$

While this formulation may seem abstract, it will have clear advantages when we aggregate over buyers and especially when we move to equilibrium models in Section1.2.

1.1.1.2 Aggregate Demand

In empirical applications, the data often only contains the sum of the individual demands of all consumers of a given observed type x for a price function P. Now consumer behavior, as represented by the choice probability $\tilde{\pi}^D(\cdot|\tilde{x}; P)$, *transports* the distribution of the full types $\tilde{x}$ of this subset of buyers to a measure on the space $\mathcal{Z}$. This measure is simply

$$\pi^D(dz|x; P) = \int_{\tilde{\mathcal{X}}} \tilde{\pi}^D(dz|\tilde{x}; P)\, \tilde{n}(d\tilde{x}|x). \tag{1.5}$$

Since the econometrician can only observe the subtypes x of the buyers, the data will only inform her on this set of conditional measures π^D. We call π^D the *conditional demand* associated with the price function P. Integrating over x gives the *aggregate demand* q^D:

$$q^D(dz;P) = \int_{\mathcal{X}} \pi^D(dz|x;P)\, n(dx). \tag{1.6}$$

To summarize:

Definition 1 (Aggregate Demand)**.** *A hedonic demand side* $(U, \tilde{n})$ *consists of a set of preferences* U *and a distribution* $\tilde{n}$ *over full types of buyers. For any given price function* P*, it generates three measures over* $\mathcal{Z}$*:*

- *the individual demand* $\tilde{\pi}^D(\cdot|\tilde{x};P)$*,*
- *the conditional demand* $\pi^D(\cdot|x;P)$ *as in* (1.5)*,*
- *and the aggregate demand* $q^D(\cdot;P)$ *as in* (1.6)*.*

1.1.1.3 Differentiable Models

The differentiable case is of particular interest. Assume that $\mathcal{Z}$ is an open subset of $\mathbb{R}^d$, that $U(\tilde{x}, \boldsymbol{z}, p)$ is continuously differentiable with respect to $\boldsymbol{z}$ and p, and that the price function $P(\cdot)$ is also continuously differentiable. Then the function $\boldsymbol{z} \rightarrow U(\tilde{x}, \boldsymbol{z}, P(\boldsymbol{z}))$ is also continuously differentiable on $\mathcal{Z}$. Its interior maxima must satisfy the first order conditions

$$\nabla_{\boldsymbol{z}} U(\tilde{x}, \boldsymbol{z}, P(\boldsymbol{z})) + \frac{\partial U}{\partial p}(\tilde{x}, \boldsymbol{z}, P(z)) \nabla P(\boldsymbol{z}) = 0. \tag{1.7}$$

If the set $Z^D(\tilde{x};P)$ is a singleton $\{\boldsymbol{z}^D(\tilde{x};P)\}$ for all buyers of observed type x, then the conditional demand is given by

$$\pi^D(B|x;P) = \int_{\tilde{\mathcal{X}}} \mathbb{1}\left(\boldsymbol{z}(\tilde{x};P) \in B\right) \tilde{n}(d\tilde{x}|x)$$

for any measurable set $B \subseteq \mathcal{Z}$.

Example 2 (continued) *In the quasilinear model, assume that* $\mathcal{Z} = \mathbb{R}^d$ *and* $\bar{U}(\tilde{x}, \boldsymbol{z})$ *is continuously differentiable with respect to* $\boldsymbol{z}$*. The first order conditions* (1.7) *equalize marginal utilities of characteristics and (marginal) prices:*

$$\nabla_{\boldsymbol{z}} \bar{U}(\tilde{x}, \boldsymbol{z}) = \nabla P(\boldsymbol{z}).$$

Example 4 (continued) *In the quadratic-in-characteristics Tinbergen model, equation (1.7) becomes*

$$\boldsymbol{a}(x) + \boldsymbol{B}\boldsymbol{z} + \boldsymbol{\varepsilon} = \nabla P(\boldsymbol{z}).$$

Under appropriate conditions, this defines $Z^D(\tilde{x};P)$*. The set of buyers who choose products in a set* B *is*

$$\{\tilde{x} = (x, \boldsymbol{\varepsilon}) \mid \boldsymbol{a}(x) + \boldsymbol{\varepsilon} = \nabla P(\boldsymbol{z}) - \boldsymbol{B}\boldsymbol{z} \;\; \textit{for some} \;\; z \in B\}.$$

1.1.2 The Quasilinear Case

We assume in this subsection that utilities are quasi-linear, as in Example 2:

$$U(\tilde{x}, z, p) = \bar{U}(\tilde{x}, z) - p.$$

We already used the idea that given a price function P, consumer behavior "transports" the distribution of consumer types $\tilde{n}$ over $\tilde{\mathcal{X}}$ into a distribution of aggregate hedonic demand $q^D(\cdot; P)$ over $\mathcal{Z}$. It does so via the individual demand $\tilde{\pi}^D(\cdot|\cdot; P)$; this defines a joint measure

$$\mu^D(d\tilde{x}, dz; P) = \tilde{\pi}^D(dz|\tilde{x}t; P)\ \tilde{n}(d\tilde{x})$$

over the product space $\tilde{\mathcal{X}} \times \mathcal{Z}$. By construction, the marginal distributions of $\mu^D(\cdot, \cdot; P)$ over $\tilde{\mathcal{X}}$ and $\mathcal{Z}$ are $\tilde{n}$ and $q^D(\cdot; P)$. As the set of joint measures with given margins will play a central role in this chapter, we introduce a notation for it:

Definition 2 (Joint measures with given margins)**.** *Let A and B be two measurable spaces. For any two positive measures a on A and b on B, we denote $\mathcal{M}(a, b)$ the set of positive measures on $A \times B$ with margins a and b.*

The set $\mathcal{M}(a, b)$ is clearly non-empty and convex; beyond that, it can be very high-dimensional. If for instance A and B are subsets of the real line, each choice of copula would generate an element of $\mathcal{M}(a, b)$. In the hedonic demand model, $\mu^D(\cdot, \cdot; P) \in \mathcal{M}\left(\tilde{n}, q^D(\cdot; P)\right)$ is generated by the buyers' choices of products for a given price function P.

1.1.2.1 Allocative efficiency

The concept of a Product (with a capital P) will prove useful here. A Product z has a utility p if it is bought at a price p. In the quasilinear case, utility is perfectly transferable between the buyer $\tilde{x}$ and the Product z. As the buyer gets net utility $\bar{U}(\tilde{x}, z) - P(z)$ and the Product gets a (virtual) utility of $P(z)$, the transaction generates a *virtual surplus* of $\bar{U}(\tilde{x}, z)$. Integrating over any joint distribution μ over $\tilde{\mathcal{X}} \times \mathcal{Z}$ yields an *aggregate virtual surplus*[1]

$$\mathcal{W}(\mu) \equiv \int_{\tilde{\mathcal{X}} \times \mathcal{Z}} \bar{U}(\tilde{x}, z)\mu(d\tilde{x}, dz).$$

In economic terms, the choice of μ realizes an assignment of objects z to buyers $\tilde{x}$. This must respect two scarcity constraints: the distribution of buyers, and the distribution of the characteristics of the Products. These two constraints define *feasible* measures μ. Given perfectly transferable utility, allocative efficiency requires that the assignment maximize the aggregate virtual surplus over all feasible joint measures.

Now suppose that we are given a family of preferences $\bar{U}$, a distribution of full types of buyers $\tilde{n}$, and a distribution of product characteristics q on $\mathcal{Z}$. Feasibility requires that each buyer $\tilde{x}$ trade one Product: $\mu(d\tilde{x}, \mathcal{Z}) \equiv \tilde{n}(d\tilde{x})$; and that each Product z be traded exactly once: $\mu(\tilde{\mathcal{X}}, dz) \equiv q(dz)$. These constraints are equivalent to the requirement that $\mu \in \mathcal{M}(\tilde{n}, q)$. By integration, they imply that $\tilde{n}$ and q have the same total mass: $q(\mathcal{Z}) = \tilde{n}(\tilde{\mathcal{X}})$.

This discussion shows that given $\bar{U}, \tilde{n}$, and q such that $q(\mathcal{Z}) = \tilde{n}(\tilde{\mathcal{X}})$, the efficient allocations solve the following program:

$$\begin{aligned} \max_{\mu} \quad & \int_{\tilde{\mathcal{X}} \times \mathcal{Z}} \bar{U}(\tilde{x}, z)\mu(d\tilde{x}, dz) \\ s.t. \quad & \mu \in \mathcal{M}\left(\tilde{n}, q\right). \end{aligned} \tag{1.8}$$

This is the primal formulation of an *optimal transport problem* with surplus

[1] Assuming that $\bar{U}(\tilde{x}, \emptyset) = P(\emptyset) = 0$.

$\bar{U}$ and with margins $(\tilde{n}, q)$. Box 1.1 gives a brief introduction to optimal transport; we will only use elementary results.

Note that the primal program is isomorphic to a house matching problem: reinterpret each Product as a house, and $\bar{U}(\tilde{x}, z)$ as the housing utility that buyer $\tilde{x}$ gets from a house with characteristics z. Allocative efficiency requires that houses be allocated to buyers in a way that maximizes total housing utility.

We not only seek to determine a joint measure μ that achieves allocative efficiency; we also want to find a price function P that supports it, in the sense that

$$\frac{\mu(d\tilde{x}, dz)}{\tilde{n}(d\tilde{x})} \equiv \tilde{\pi}^D(dz|\tilde{x}; P) :$$

the conditional demand under P coincides with the measure μ. The duality theorem of optimal transport gives us the solution. It states that under minimal conditions, the value of the primal program coincides with that of the *dual program*:

$$\begin{aligned} \min_{\tilde{u}, P} \quad & \int_{\tilde{\mathcal{X}}} \tilde{u}(\tilde{x}) \tilde{n}(d\tilde{x}) + \int_{\mathcal{Z}} P(z) q(dz) \\ s.t. \quad & \tilde{u}(\tilde{x}) + P(z) \geq \bar{U}(\tilde{x}, z), \end{aligned} \tag{1.9}$$

where the inequality must hold for all $(\tilde{x}, z)$. Moreover, the multipliers of the constraints in the primal problem solve the dual problem, and vice versa.

The intuition of this result is simple. Remember that the feasibility constraint $\mu \in \mathcal{M}(\tilde{n}, q)$ requires that

$$\mu(d\tilde{x}, \mathcal{Z}) = \tilde{n}(d\tilde{x}) \text{ for all } \tilde{x}$$
$$\mu(\tilde{\mathcal{X}}, dz) = q(dz) \text{ for all } z.$$

Let $\tilde{u}(\tilde{x})$ (resp. $P(z)$) be the Lagrange multipliers for the first (resp. the second) line of constraints. Since μ cannot be negative, the Karush-Kuhn-Tucker (KKT) conditions imply that

$$\bar{U}(\tilde{x}, z) \leq \tilde{u}(\tilde{x}) + P(z),$$

with equality where $\mu(d\tilde{x}, dz) > 0$. This is just saying that the net utility of the buyer of any product z, which equals $\bar{U}(\tilde{x}, z) - P(z)$, cannot exceed $\tilde{u}(\tilde{x})$; and that it must equal it if the buyer does end up with this product. More succinctly, $\tilde{u}(\tilde{x}) = \max_{z \in \mathcal{Z}}(\bar{U}(\tilde{x}, z) - P(z))$ is attained on the support of μ. If the price function is P, we end up with an efficient allocation at which the net utilities of buyers are given by the function $\tilde{u}$.

Conversely, denote μ the multipliers for the inequalities that constrain the dual. The KKT conditions imply $\tilde{n}(d\tilde{x}) = \int_{z \in \mathcal{Z}} \mu(d\tilde{x}, dz)$ and $q(dz) = \int_{\tilde{x} \in \tilde{\mathcal{X}}} \mu(d\tilde{x}, dz)$; that is, $\mu \in \mathcal{M}(\tilde{n}, q)$. Moreover, the inequality must be binding on the support of μ:

$$\tilde{u}(\tilde{x}) = \max_{z \in \mathcal{Z}}(\bar{U}(\tilde{x}, z) - P(z))$$

and the maximum is attained on the support of μ. Therefore the solutions to the dual and their multipliers coincide with, respectively, the multipliers and the solutions of the primal.

1.1.2.2 Assortative Matching

As Becker (1973) showed, allocative efficiency in itself may have strong consequences on matching patterns. Assume that $\tilde{\mathcal{X}}$ and $\mathcal{Z}$ are subsets of the real line, and that the surplus U is *strictly supermodular*: if $\tilde{x}' \geq \tilde{x}$ and $z' \geq z$ with at least one strict inequality, then

$$U(\tilde{x}', z') + U(\tilde{x}, z) > U(\tilde{x}, z') + U(\tilde{x}', z).$$

This would be natural if for instance z is a quality index and $\tilde{x}$ measures a preference for higher quality. Then any μ that is allocatively efficient must be *positively assortative*: higher-$\tilde{x}$ buyers are assigned higher-z products. To see this, suppose for instance that $\tilde{n}$ and q have no mass point[2]. Then any feasible μ has a density; denote it m. Suppose that $\tilde{x}' > \tilde{x}$, that $m(\tilde{x}, z) > 0$, and that (contrary to positively assortative matching) buyer $\tilde{x}'$ buys a lower-tier product $z' < z$ with positive probability: $m(\tilde{x}', z') > 0$. Letting $\underline{m} > 0$ be the minimum of $m(\tilde{x}, z)$ and $m(\tilde{x}', z')$, move a density $\underline{m}$ from $(\tilde{x}, z)$ to $(\tilde{x}', z)$, and the same $\underline{m}$ from $(\tilde{x}', z')$ to $(\tilde{x}, z')$. The resulting joint measure is non-negative and is still feasible as the marginal densities are preserved. The total surplus changes by

$$\underline{m} \times \left(\bar{U}(\tilde{x}', z) + \bar{U}(\tilde{x}, z') - \bar{U}(\tilde{x}, z) - \bar{U}(\tilde{x}', z') \right),$$

[2] The proof is essentially unchanged when they have discrete support.

Box 1.1: The optimal transport problem

The generic optimal transport problem is defined by two measures F on a set A and G on a set B, and a surplus function R from $A \times B$ to $\mathbb{R}$. It consists in maximizing the integral

$$\int_{A\times B} R(a,b)H(da,db) \tag{1.10}$$

over all joint measures H with margins F and G—that is, $H \in \mathcal{M}(F,G)$. This is a linear programming problem, whose dual consists of finding real-valued functions ϕ and ψ that minimize

$$\int_A \varphi(a)F(da) + \int_B \psi(b)G(db) \tag{1.11}$$

subject to the constraint $\varphi(a) + \psi(b) \geq R(a,b)$.

Under suitable assumptions on the surplus function R, both the primal problem (1.10) and the dual (1.11) have solutions, and their values coincide. Moreover, complementary slackness holds: if $H \in \mathcal{M}(F,G)$ and $\varphi(a) + \psi(b) \geq R(a,b)$, then H and (φ,ψ) are respectively solutions to the primal and the dual problems if and only if H assigns probability one to the set of (a,b) values such that $\varphi(a) + \psi(b) = R(a,b)$.

When A and B are finite sets of respective sizes I and J, the problem boils down to finding a vector $H_{ij} \geq 0$ which maximizes $\sum_{i=1}^{I}\sum_{j=1}^{J} R_{ij}H_{ij}$ under the constraints

$$\sum_{j=1}^{J} H_{ij} = F_i \quad \text{for } i = 1,\dots,I, \text{ and } \sum_{i=1}^{I} H_{ij} = G_j \quad \text{for } i = j,\dots,J.$$

It is a finite-dimensional linear programming problem, whose dual consists in minimizing $\sum_{i=1}^{I} F_i\varphi_i + \sum_{j=1}^{J} G_j\psi_j$ subject to $\phi_i + \psi_j \geq R_{ij}$ for all i and j. As is well-known, the primal problem has a solution, which is generically unique; the dual problem has a set of solutions that is defined by a system of linear inequalities.

If A and B are not finite sets, the existence and uniqueness of a solution are not as straightforward: in fact, the original instance of the optimal transportation problem in Monge (1781) was not fully solved until the 1980s. There is a rich mathematical literature on the topic. For a standard reference on the mathematical theory of optimal transport, see Villani (2003); for a more accessible introduction to the topic and its applications to economics, we refer the reader to Galichon (2016). We discuss various numerical solutions of the optimal transport problem in section 4.2.1.

which is positive since $\tilde{x}' > \tilde{x}$, $z' < z$, and $\bar{U}$ is strictly supermodular. It follows that μ cannot be allocatively efficient.

One can go further: the support of the allocatively efficient μ is the graph of the strictly increasing map $z = T(\tilde{x})$ defined by $\tilde{n}([\tilde{x}, \infty) \cap \tilde{X}) = q([z, \infty) \cap \mathcal{Z})$. If we denote $\tilde{N}$ and Q the cumulative distribution functions of $\tilde{n}$ and q, this is simply the equality $\tilde{N}(\tilde{x}) = Q(z)$; and $T = Q^{-1} \circ \tilde{N}$.

1.1.2.3 An Equivalence Theorem for One-sided Choice

To summarize our results so far, consider a given population of buyers. If a buyer of type $\tilde{x}$ decides to purchase a product of type z, then $(\tilde{x}, z)$ is in the support of the allocatively efficient μ, which solves program (1.8). Conversely, if μ is allocatively efficient (it solves the optimal transport problem (1.8)), then there exists a function P such that μ is generated by the hedonic demand associated with price function P. Moreover, the function P, along with $\tilde{u}(\tilde{x}) = \max_{z\in\mathcal{Z}}(\bar{U}(\tilde{x}, z) - P(z))$, solves the dual program (1.9).

More formally, we state the following theorem, which adapts results in Galichon and Salanié (2022):

Theorem 5 (Equivalence Theorem for Quasilinear One-sided Choice)**.** *Fix the distribution of buyer characteristics $\tilde{n}$ and a family of quasilinear preferences $\bar{U}$ for these buyers, and consider a price function P.*

Let $\tilde{\pi}^D(dz|\tilde{x}; P)$ be the individual demand of buyer $\tilde{x}$ for the price function P and $q^D(dz; P)$ the corresponding aggregate demand, as defined in paragraph 1.1.1.2. Define $\tilde{u}(\tilde{x}) = \max_{z\in\mathcal{Z}} \left\{\bar{U}(\tilde{x}, z) - P(z)\right\}$, the indirect utility of consumer $\tilde{x}$ under P. Then *(i) The joint measure* $\mu(d\tilde{x}, dz) := \tilde{\pi}^D(dz|\tilde{x}; P)\tilde{n}(d\tilde{x})$ *solves the primal optimal transport problem* (1.8) *with surplus* $\bar{U}$ *and margins* $(\tilde{n}, q^D(\cdot; P))$*.*

(ii) The pair $(\tilde{u}, P)$ *solves the dual optimal transportation problem* (1.9)*, and the aggregate surplus and the sum of utilities (indirect buyer utilities and virtual Product utilities) coincide:*

$$\int_{\tilde{\mathcal{X}}\times\mathcal{Z}} \bar{U}(\tilde{x}, z)\mu(d\tilde{x}, dz) = \int_{\tilde{\mathcal{X}}} \tilde{u}(\tilde{x})\tilde{n}(d\tilde{x}) + \int_{\mathcal{Z}} P(z)q(dz). \qquad (1.12)$$

(iii) Conversely, take any measure q *on* $\mathcal{Z}$ *with total mass* $q(\mathcal{Z}) = \tilde{n}(\tilde{x})$*. If* $\mu \in \mathcal{M}(\tilde{n}, q)$ *solves the primal optimal transport problem* (1.8) *and* $(\tilde{u}, P)$ *solves the corresponding dual* (1.9)*, then the conditional distribution* $\mu(\cdot|\tilde{x})$ *coincides with the individual demand* $\tilde{\pi}^D(\cdot|\cdot; P)$ *for the price function* P*, and* q *is the aggregate demand* $q^D(\cdot; P)$*.*

Theorem 5 relies on the Monge–Kantorovich duality theorem, which is fundamental in the theory of optimal transport. Note that the dual problem can be formulated in terms of the price function P only: the P that rationalizes a distribution q on product characteristics $\mathcal{Z}$ must minimize

$$\int_{\tilde{\mathcal{X}}} \max_{z\in\mathcal{Z}} \left(U(\tilde{x}, z) - P(z)\right) \tilde{n}\left(d\tilde{x}\right) + \int_{\mathcal{Z}} P\left(z\right) q\left(dz\right). \qquad (1.13)$$

Example 3 (continued) *To illustrate these results, let us return to the logit example. Given a price function* P*, standard formulæ give the conditional demand of buyers of observed type* x*:*

$$\pi^D(z|x; P) = \frac{\exp(\bar{U}(x, z))}{\sum_{z'\in\mathcal{Z}} \exp(\bar{U}(x, z'))}$$

and their "inclusive value" is

$$E\left(\tilde{u}(\tilde{x})|x\right) = \int_{\varepsilon \in \mathbb{R}^{\mathcal{Z}}} \max_{z \in \mathcal{Z}} \left\{ \tilde{U}\left(x, \boldsymbol{\varepsilon}, z\right) - P\left(z\right) \right\} \tilde{n}\left(d\varepsilon|x\right) = \log \sum_{z \in \mathcal{Z}} \exp\left(U\left(x, z\right) - P\left(z\right)\right).$$

Taking expectations over x and substituting in (1.13) *implies that the price function P that rationalizes a product distribution q minimizes*

$$\int_{x \in \mathcal{X}} n\left(dx\right) \log\left(\sum_{z \in \mathcal{Z}} \exp\left(U\left(x, z\right) - P\left(z\right)\right)\right) + \sum_{z \in \mathcal{Z}} P\left(z\right) q\left(z\right).$$

The first order conditions yield the usual formulæ for aggregate demand under P:

$$q\left(z\right) = q^{D}(z; P) = \int_{x \in \mathcal{X}} n\left(dx\right) \frac{\exp\left(U\left(x, z\right) - P\left(z\right)\right)}{\sum_{z' \in \mathcal{Z}} \exp\left(U\left(x, z'\right) - P\left(z'\right)\right)}.$$

1.1.3 Beyond the Quasilinear Case

If the buyer's utility $U(\tilde{x}, z, p)$ is not additive in p, we cannot define the joint surplus from a trade and allocative efficiency loses its appeal. It is still possible to characterize hedonic demand (Bonnet, Galichon, Hsieh, O'hara, and Shum, 2022).

Given a price function P, buyer $\tilde{x}$ maximizes $U(\tilde{x}, z, P(z))$ and attains utility $\tilde{u}\left(\tilde{x}\right) = \max_{z \in \mathcal{Z}} U\left(\tilde{x}, z, P\left(z\right)\right)$. This leads as before to a joint distribution

$$\mu(d\tilde{x}, dz) := \tilde{\pi}^{D}(dz|\tilde{x}; P)\, \tilde{n}(d\tilde{x})$$

over buyers and purchased products. Let $q = q^{D}(\cdot; P)$ denote the aggregate demand distribution of the purchased products. These three objects $(\mu, \tilde{u}, q)$

are linked by the following conditions:

$$\begin{cases} \mu \in \mathcal{M}(\tilde{n}, q) \\ \tilde{u}(\tilde{x}) \geq U(\tilde{x}, z, P(z)) \text{ for all } \tilde{x} \in \tilde{\mathcal{X}} \text{ and } z \in \mathcal{Z} \\ \text{If } \tilde{u}(\tilde{x}) > U(\tilde{x}, z, P(z)), \text{ then } (\tilde{x}, z) \text{ is not in the support of } \mu. \end{cases}$$

Conversely, if the triple $(\mu, \tilde{u}, P)$ satisfies these three conditions, it can be shown that $\mu(\cdot|\tilde{x})$ coincides with the individual demand $\tilde{\pi}^D(\cdot|\cdot; P)$ and q is the aggregate demand $q^D(\cdot; P)$ for the price function P.

From a mathematical point of view, this is an *equilibrium transport problem*; in the notation of Box 1.2, we have $\phi = \tilde{u}, \psi = P$, and $S(\tilde{x}, z, p, q) = p - U(\tilde{x}, z, q)$.

1.2 Hedonic Equilibrium

We now bring together supply and demand. The modeling of the supply side mirrors that of the demand side. Each seller has a full type $\tilde{y} \in \tilde{\mathcal{Y}}$ and derives a utility (or profit) $V(\tilde{y}, z, p)$ of selling an object of characteristics z at price p. Given a price function P, the individual supply of seller $\tilde{y}$ is the conditional probability $\pi^S(\cdot|\tilde{y}; P)$. The support of this distribution is the set of maximizers of the seller's utility:

$$\text{supp}(\pi^S(\cdot|\tilde{y}; P) \subseteq \arg\max_{z \in \mathcal{Z}} V(\tilde{y}, z, P(z)) \equiv Z^S(\tilde{y}; P).$$

Box 1.2: The equilibrium transport problem

We start from the same notation as in Box 1.1. Let p and q denote the "payoffs" of a and b. Payoffs p and q are feasible if $S(a,b,p,q) \leq 0$, where S is a function that is increasing in p and in q. We interpret this inequality as constraining the transfers between types a and b if they are matched. More precisely, (H,φ,ψ) solves the *equilibrium transport problem* defined by (F,G,S) iff

1. $H \in \mathcal{M}(F,G)$
2. $S(a,b,\varphi(a),\psi(b)) \geq 0$ for all $a \in A, b \in B$
3. $\Pr_H(S(a,b,\varphi(a),\psi(b)) = 0) = 1$.

As noted in Galichon (2016, Definition 10.1), optimal transport is the special case where

$$S(a,b,p,q) = p + q - R(a,b). \tag{1.14}$$

Let $\tilde{m}$ denote the measure over full types $\tilde{y}$. The *hedonic supply* for prices P is the joint measure μ^S over seller and product characteristics given by

$$\mu^S(d\tilde{y},dz) \equiv \pi^S(dz|\tilde{y};P)\,\tilde{m}(d\tilde{y}),$$

and its margin over $\mathcal{Z}$ is the aggregate supply $q^S(\cdot;P)$.

Proceeding as with hedonic demand, we know that the joint measure μ^S has margins $q^S(\cdot|P)$ and $\tilde{m}$, and that its support only includes pair $(\tilde{x},z)$ such that $z \in Z^S(\tilde{y};P)$. The definition of hedonic equilibrium for utility functions U and V and distributions of full types $\tilde{n}$ and $\tilde{m}$ follows directly.

1.2.1 Definition

Definition 3. *A hedonic equilibrium is given by a price function P, a hedonic demand μ^D, and a hedonic supply μ^S such that:*

1. *the margin of μ^D on $\tilde{\mathcal{X}}$ is $\tilde{n}$: $\int_{z \in \mathcal{Z}} \mu(d\tilde{x}, dz) = \tilde{n}(d\tilde{x})$;*
2. *the margin of μ^S on $\tilde{\mathcal{Y}}$ is $\tilde{m}$: $\int_{z \in \mathcal{Z}} \mu^S(d\tilde{y}, dz) = \tilde{m}(d\tilde{y})$;*
3. *the margins of μ^D and μ^S on $\mathcal{Z}$ coincide:*

$$\int_{\tilde{x} \in \tilde{\mathcal{X}}} \mu(d\tilde{x}, dz) = \int_{\tilde{y} \in \tilde{\mathcal{Y}}} \mu^S(d\tilde{y}, dz);$$

4. *each buyer $\tilde{x}$ consumes their preferred product given P:*

   the support of $\mu^D(\cdot|\tilde{x}; P)$ is contained in $Z^D(\tilde{x}; P)$

5. *each seller $\tilde{y}$ supplies their preferred product given P:*

   the support of $\mu^S(\cdot|\tilde{y}; P)$ is contained in $Z^S(\tilde{y}; P)$.

The common margin referred to in definition 3 is the distribution $q(\cdot; P)$ of the characteristics of products traded in equilibrium. By construction, $q(\cdot; P) = q^D(\cdot; P) = q^S(\cdot; P)$.

When utilities on both sides of the market are quasilinear, we can apply Theorem 5 to demand and to supply separately. We could pick a price function P and solve the two primal problems to obtain $q^D(\cdot; P)$ and $q^S(\cdot; P)$. Alternatively, we could choose a measure q such that $q(\mathcal{Z}) = \tilde{n}(\tilde{x}) = \tilde{m}(\tilde{y})$, and solve the two dual problems to obtain price functions P^D and P^S. Equilibrium obtains in the first approach when q^S and q^D coincide; in the second approach, it is achieved when P^D and P^S coincide.

To put it differently: a hedonic market brings together two one-sided matching problems—one between buyers and Products, and one between sellers and Products. Either matching problem generates a relationship between a measure of Products traded and a price function; equilibrium obtains when these objects are the same in the two matching problems[3].

1.2.2 Solving the Quadratic Example

To illustrate this, let us consider a simple instance of the quadratic model of Example 4. To simplify it even further, we assume that buyers and sellers only have a scalar, unobserved type: $\tilde{x}$ is simply ε and $\tilde{y}$ is η. We assume that the masses of buyers and sellers both equal one and we denote their cdfs as F_ε and F_η. We write the utility functions as respectively $U(\varepsilon, z, p) \equiv \varepsilon z + Bz^2/2 - p$ and $V(\eta, z, p) \equiv \eta z + Cz^2/2 + p$.

For any given price function P, buyer ε chooses a product $z^D(\varepsilon)$ that solves $\varepsilon + Bz = P'(z)$, and seller η chooses a product $z^S(\eta)$ that solves $\eta + Cz = -P'(z)$. The second-order conditions imply that $B < P''(z) < -C$.

Given the complementarity of type and product, both z^D and z^S are increasing functions; as a consequence, so is the function $\eta = M(\varepsilon)$ that matches a seller to a buyer. Moreover, we can eliminate the marginal price from the definitions of z^D and z^S to obtain $M(\varepsilon) = -\varepsilon - (B + C)z^D(\varepsilon)$.

Hedonic equilibrium is achieved for a price function P that induces a

[3] As we will explain in Section 4.1, this connection between hedonic equilibrium models and matching models can also be seen by viewing them as instances of network flow problems.

matching $M = (z^S)^{-1} \circ z^D$ such that

$$F_\eta(M(\varepsilon)) = F_\varepsilon(\varepsilon) \text{ for all } \varepsilon. \tag{1.15}$$

To simplify this even more, we now assume that F_ε and F_η belong to the same location-scale family of distributions, so that $F_\eta(t) \equiv F((t-m)/\sigma)$ for some m and σ. Then (1.15) holds if and only if $M(\varepsilon) \equiv m + \sigma\varepsilon$. If follows that

$$z^D(\varepsilon) = -\frac{M(\varepsilon)+\varepsilon}{B+C} = -\frac{m+(\sigma+1)\varepsilon}{B+C}.$$

Changing variables again, we obtain $z = -(m+(\sigma+1)(P'(z)-Bz))/(B+C)$ and finally

$$P'(z) = \frac{(B\sigma - C)z - m}{\sigma+1}.$$

so that the equilibrium price function is quadratic. Its level depends on the value of the outside options.

The dual approach builds on (1.13): the equilibrium trade measure q on $\mathcal{Z}$ must be such that the same price function P minimizes both

$$\int \max_{z\in\mathcal{Z}} \left(z\varepsilon + Bz^2/2 - P(z)\right) F_\varepsilon(d\varepsilon) + \int P(z)q(dz)$$

and

$$\int \max_{z\in\mathcal{Z}} \left(z\eta + Cz^2/2 + P(z)\right) F_\eta(d\eta) - \int P(z)q(dz).$$

Now consider an infinitesimal variation $\delta P(z)$ in the price function. By the envelope theorem, it contributes $-\delta P(z^D(\varepsilon))$ to each $\max_{z\in\mathcal{Z}}$ in the first problem, and $\delta P(z^S(\eta))$ to each $\max_{z\in\mathcal{Z}}$ in the second problem. In the end, both

$$-\int \delta P(z^D(\varepsilon))F_\varepsilon(d\varepsilon) + \int \delta P(z)q(dz)$$

and

$$\int \delta P(z^S(\eta))F_\eta(d\eta) - \int \delta P(z)q(dz)$$

must be approximately zero for each such price variation. This can only happen if at each $z = z^D(\varepsilon) = z^S(\eta)$ we have $-F_\varepsilon(d\varepsilon) + q(dz) = 0 = F_\eta(d\eta) - q(dz)$. Since $\varepsilon = P'(z) - Bz$ and $\eta = -P'(z) - Cz$,

$$d\varepsilon = (P''(z) - B)dz \quad \text{and} \quad d\eta = -(P''(z) + C)dz.$$

Under our assumption that $F_\eta(t) \equiv F_\varepsilon((t - m)/\sigma))$, we must have

$$P''(z) - B \equiv -\frac{P''(z) + C}{\sigma}.$$

Therefore P'' is a constant and the equilibrium function must be quadratic. It is easy to check that the leading coefficient is the same as in the primal approach.

1.3 Two-sided matching models

The results in the previous section readily adapt to the matching model of Becker (1973) and Shapley and Shubik (1972): the one-to-one bipartite matching market with perfectly transferable utility. The term “bipartite” here means that (just as with a trade in the hedonic market) the two partners in a match are drawn from separate subpopulations of agents. Since Becker’s original application was to the marriage market, we follow tradition and we refer to potential partners as "men" and "women"[4].

1.3.1 Preliminaries

We use the same notation for the agents as with hedonic markets: a man has full type $\tilde{x} = (x, \varepsilon)$ and a woman has a full type $\tilde{y} = (y, \eta)$. We will often use the shortcut $\tilde{x} \in x$, say, to state that a man of type $\tilde{x}$ has observable type x. We assume that each man, and each woman, observes their own full type and the full type of any potential partner.

A *match* can be three things: a couple $(\tilde{x}, \tilde{y})$; a single man $\tilde{x}$; or a single woman $\tilde{y}$. It will be convenient to denote $\tilde{\mathcal{X}}_0 = \{0\} \cup \tilde{\mathcal{X}}$ and $\tilde{\mathcal{Y}}_0 = \{0\} \cup \tilde{\mathcal{Y}}$, where 0 applies to singles ("matched with 0"). We define $\mathcal{X}_0$ and $\mathcal{Y}_0$ in the same way. A *matching* is a joint measure over the set $\tilde{\mathcal{A}} = \left(\tilde{\mathcal{X}}_0 \times \tilde{\mathcal{Y}}\right) \cup \left(\tilde{\mathcal{X}} \times \tilde{\mathcal{Y}}_0\right)$. It is *feasible* if the margins of the joint measure do not excess $\tilde{n}$ and $\tilde{m}$ respectively.

Transfers within a match are costless, unlimited, and have a constant marginal utility of 1 for each partner. As a consequence, utilities within a match are only required to add up to a number which we call the *joint utility*: if a match between a man of full type $\tilde{x}$ and a woman of full type $\tilde{y}$ forms, then their respective utility levels $\tilde{U}$ and $\tilde{V}$ need only satisfy the constraint

$$\tilde{U} + \tilde{V} = \tilde{\Phi}(\tilde{x}, \tilde{y}).$$

Let us denote $\tilde{\Phi}(\tilde{x}, 0)$ (resp. $\tilde{\Phi}(0, \tilde{y})$) the utility man $\tilde{x}$ (resp. woman $\tilde{y}$) would

[4]We assume here that marriages are heterosexual and monogamous. We will discuss same-sex marriage and more generally non-bipartite models in Section 3.2.3.3.

attain if single. The *joint surplus* created by this match is the excess of the joint utility over the sum of the utilities the man and the woman achieve as singles:

$$\tilde{S}(\tilde{x},\tilde{y}) = \tilde{\Phi}(\tilde{x},\tilde{y}) - \tilde{\Phi}(\tilde{x},0) - \tilde{\Phi}(0,\tilde{y}).$$

Any match must guarantee to each partner as much utility as they could obtain on their own or with an alternative, willing partner. Consider any other woman, of full type $\tilde{y}'$, and denote $\tilde{v}(\tilde{y}')$ the utility she obtains in her current match. To match her, man $\tilde{x}$ must leave to her at least a utility $\tilde{v}(\tilde{y}')$. As a consequence, a $\tilde{x}$ man cannot obtain more than $\tilde{\Phi}(\tilde{x},\tilde{y}') - \tilde{v}(\tilde{y}')$ with this alternative partner. If

$$\tilde{U} \geq \max_{\tilde{y}'} \left(\tilde{\Phi}(\tilde{x},\tilde{y}') - \tilde{v}(\tilde{y}')\right),$$

then man $\tilde{x}$ has no incentive to look for a different partner than woman $\tilde{y}$. Denoting $\tilde{u}(\tilde{x})$ the utility he gets from his match, this can be written more symmetrically as $\tilde{u}(\tilde{x}) + \tilde{v}(\tilde{y}') \geq \tilde{\Phi}(\tilde{x},\tilde{y}')$. As this must apply to all potential partners on both sides, we have

$$\tilde{u}(\tilde{x}) + \tilde{v}(\tilde{y}) \geq \tilde{\Phi}(\tilde{x},\tilde{y}) \text{ for all } (\tilde{x},\tilde{y}). \tag{1.16}$$

In addition, no individual can obtain less utility than they would by remaining single:

$$\tilde{u}(\tilde{x}) \geq \tilde{\Phi}(\tilde{x},0) \text{ for all } \tilde{x}; \text{ and } \tilde{v}(\tilde{y}) \geq \tilde{\Phi}(0,\tilde{y}) \text{ for all } \tilde{y}. \tag{1.17}$$

This implies in particular that a match cannot form if its joint surplus is negative, as this would make at least one of the two partners worse off than staying single.

The set of inequalities in (1.16) and (1.17) spell out the requirements of *stability* for a one-to-one, bipartite matching with perfectly transferable utility.

To summarize:

Definition 4 (Matchings with transferable utility)**.** *A one-to-one, bipartite* matching market with perfectly transferable utility *consists of two measures $\tilde{n}$ over $\tilde{\mathcal{X}}$ and $\tilde{m}$ over $\tilde{\mathcal{Y}}$ and a joint utility function $\tilde{\Phi}$ from $\tilde{\mathcal{A}} = (\tilde{\mathcal{X}} \times \tilde{\mathcal{Y}}_0) \cup (\tilde{\mathcal{X}}_0 \times \tilde{\mathcal{Y}})$ to $\mathbb{R}$.*

A match *is a 2-uple $(\tilde{x}, \tilde{y})$ in $\tilde{\mathcal{A}}$. A* matching *is a measure over the set $\tilde{\mathcal{A}}$; it records the frequency with which each possible pair is formed. A matching is* feasible *if and only if each individual appears once and only once. It is* stable *if and only if there exist two functions $\tilde{u} : \tilde{\mathcal{X}} \to \mathbb{R}$ and $\tilde{v} : \tilde{\mathcal{Y}} \to \mathbb{R}$ such that*

- *the inequalities* (1.16) *and* (1.17) *hold for all $\tilde{x}$ and $\tilde{y}$;*
- *if the matching assigns positive mass to $(\tilde{x}, \tilde{y})$ pairs, then* (1.16) *is an equality;*
- *if the matching assigns positive mass to $\tilde{x}$ (resp. $\tilde{y}$) remaining single, then $\tilde{u}(\tilde{x}) = \tilde{\Phi}(\tilde{x}, 0)$ (resp. $\tilde{v}(\tilde{y}) = \tilde{\Phi}(0, \tilde{y})$).*

1.3.2 Primal and Dual

As the notation suggests, under perfectly transferable utility we can directly transpose most of our results on one-sided discrete choice models in section 1.1.2. To any matching $\tilde{\mu}$, we can associate the support of the distribution $\tilde{\mu}$, which is a subset $\mathcal{S}$ of $\tilde{\mathcal{A}}$. A measure which is supported by $\mathcal{S}$ is feasible if and only if its margins are $\tilde{n}$ and $\tilde{m}$; in the notation of section 1.1.2, $\tilde{\mu} \in \mathcal{M}(\tilde{n}, \tilde{m})$.

By construction, to a feasible matching corresponds a feasible measure $\tilde{\mu}$. It generates a *social welfare* given by the sum of joint utilities over matches:

$$\mathcal{W}(\tilde{\mu}) = \int_{\tilde{\mathcal{A}}} \tilde{\Phi}(\tilde{x}, \tilde{y})\ \tilde{\mu}(d\tilde{x}, d\tilde{y}). \tag{1.18}$$

Just as in section 1.1.2, the measure associated to a stable matching maximizes the social welfare $\mathcal{W}$ over the set of feasible measures: that is, it solves the optimal transport problem

$$\begin{aligned} &\max \mathcal{W}(\tilde{\mu}) \\ &\text{s.t.} \int_{\tilde{\mathcal{Y}}_0} \tilde{\mu}(d\tilde{x}, d\tilde{y}) = \tilde{n}(d\tilde{x}) \ \text{ for all } \tilde{x} \in \tilde{\mathcal{X}} \\ &\text{and} \int_{\tilde{\mathcal{X}}_0} \tilde{\mu}(d\tilde{x}, d\tilde{y}) = \tilde{m}(d\tilde{y}) \ \text{ for all } \tilde{y} \in \tilde{\mathcal{Y}}. \end{aligned} \tag{1.19}$$

Let $\alpha(\tilde{x})$ and $\beta(\tilde{y})$ be the Lagrange multipliers associated to the constraints at the optimum. If neither $\tilde{x}$ nor $\tilde{y}$ is zero, the first order conditions give

$$\alpha(\tilde{x}) + \beta(\tilde{y}) \geq \tilde{\Phi}(\tilde{x}, \tilde{y}),$$

with equality in the support of $\tilde{\mu}$ — that is, when $\tilde{x}$ and $\tilde{y}$ are matched. In addition, $\alpha(\tilde{x}) \geq \tilde{\Phi}(\tilde{x}, 0)$ and $\beta(\tilde{y}) \geq \tilde{\Phi}(0, \tilde{y})$, and these are equalities for singles. Following Section 1.1.2, we can identify α and β to $\tilde{u}$ and $\tilde{v}$, the utilities obtained at the stable matching.

To the primal problem (1.19), we associate the dual. It consists of *minimizing* the sum of utilities:

$$\begin{aligned} & \min \left(\int_{\tilde{\mathcal{X}}} \tilde{u}(\tilde{x}) \tilde{n}(d\tilde{x}) + \int_{\tilde{\mathcal{Y}}} \tilde{v}(\tilde{y}) \tilde{m}(d\tilde{y}) \right) \qquad (1.20) \\ \text{s.t.} \quad & \tilde{u}(\tilde{x}) + \tilde{v}(\tilde{y}) \geq \tilde{\Phi}(\tilde{x}, \tilde{y}) \ \text{ for all } \ (\tilde{x}, \tilde{y}) \\ & \text{and } \tilde{u}(\tilde{x}) \geq \tilde{\Phi}(\tilde{x}, 0) \ \text{ for all } \ \tilde{x} \\ & \text{and } \tilde{v}(\tilde{y}) \geq \tilde{\Phi}(0, \tilde{y}) \ \text{ for all } \ \tilde{y}. \end{aligned}$$

It is easy to check that if $(\tilde{u}, \tilde{v})$ are associated to a stable matching with measure $\tilde{\mu}$, then they must solve the dual program (1.20); and the Lagrange multipliers at the minimum constitute the measure $\tilde{\mu}$. The proofs of these results are essentially identical to those of Section 1.1.2.

1.3.3 Separability

Remember that the available data only contains information about observed types x and y, whose distributions are

$$n(dx) = \int \tilde{n}(d\tilde{x}|x) \ \text{ and } \ m(dy) = \int \tilde{m}(d\tilde{y}|y).$$

The econometrician can only hope to estimate the distributions n over $\mathcal{X}$ and m over $\mathcal{Y}$, and the joint measure

$$\mu(dx, dy) = \int_{\tilde{\mathcal{A}}} \tilde{\mu}(d\tilde{x}, d\tilde{y})\tilde{n}(d\tilde{x}|x)\tilde{m}(d\tilde{y}|y).$$

This is obviously not enough to recover the joint utility $\tilde{\Phi}(\tilde{x}, \tilde{y})$, unless we restrict how the unobserved types ε and η enter $\tilde{\Phi}$. Following Choo and Siow (2006), most of the literature has adopted a *separability* assumption. In order to introduce it, it is useful to think of the joint utility of a match in analysis-of-variance terms. Since $\tilde{x} = (x, \varepsilon)$ and $\tilde{y} = (y, \eta)$, four classes of terms may contribute to the joint utility of the match:

1. the additive terms, that do not interact $\tilde{x}$ and $\tilde{y}$ within $\tilde{\Phi}$;

2. the terms that interact x and y;

3. the terms that interact x and $\tilde{y}$, or y and $\tilde{x}$;

4. the terms that contain interactions between the unobserved components ε and η.

The first group helps explain the singlehood rates across categories. As Becker (1973) emphasized, complementarity in the joint surplus drives matching patterns. Groups 2 to 4 above represent these complementarities with more and more complexity. The assumptions and/or data required to identify these terms grow as we move from group 2 to group 4. In particular, identifying the distribution of ε, that of η, and the form of their interaction in the 3-way terms is very challenging[5]. For this reason, the literature has

[5] We will return to this point in Section 3.2.3.1.

converged towards the class of models that rules out the terms in group 4. This amounts to assuming that the joint surplus is *separable*, in the following sense.

Assumption 1 (Separability in Matching Markets)**.** *The joint utility of a match between $\tilde{x} = (x, \varepsilon)$ and $\tilde{y} = (y, \eta)$ is*

$$\tilde{\Phi}(\tilde{x}, \tilde{y}) = \Phi(x, y) + \zeta(\tilde{x}, y) + \xi(x, \tilde{y}).$$

We normalize $\Phi(x, 0) = \Phi(0, y) = 0$ for all x and y, so that remaining unmatched gives $\tilde{x}$ (resp. $\tilde{y}$) a utility $\zeta(\tilde{x}, 0)$ (resp. $\xi(0, \tilde{y})$).

Here $\Phi(x, y)$ subsumes the 0-way terms and the 1-way terms, while $\zeta(\tilde{x}, y)$ and $\xi(x, \tilde{y})$ are the 2-way terms.

Since we normalized the value of Φ to zero for singles, it represents the part of the joint surplus that is generated by interactions between the observable types; it will be the central parameter of interest that we seek to estimate. We will call it the *surplus function* in what follows.

It is easy to imagine examples in which separability fails. If for instance the joint utility of a match is higher when the partners share some unobservable characteristic such as hair color, then separability cannot be expected to hold. As always, the more interesting question is how damaging it is to assume separability when it does not hold. By analogy with standard results on omitted variables in linear models, one might hope that such misspecifications would have mild consequences when the unobservable traits that interact are distributed independently of the observable traits. Simulations reported in Chiappori, Nguyen, and Salanié (2019) are encouraging in this respect. They show, in particular, that it takes large injections of this type of non-separability to seriously bias estimates of complementarities between observable characteristics. More generally, separability is less demanding when using richer datasets as the set of unobservable traits shrinks.

Finally, let us note that separability is a property of the *joint utility* of a match, not of the pre-transfer utilities of the partners[6]. While it clearly holds if the utilities of the partners are $\bar{U}(\tilde{x}, y) - t$ and $\bar{V}(x, \tilde{y}) + t$ after they match and transfer utility t to each other, this specification is by no means necessary[7].

Separability was one of the assumptions introduced by Choo and Siow (2006). It was named by Chiappori, Salanié, and Weiss (2017), who derived its crucial implication for identification: it reduces the matching problem to a much lower dimension.

Theorem 6 (Chiappori, Salanié, and Weiss (2017),Galichon and Salanié (2022))**.** *Under Assumption 1, there exist two scalar functions U and V defined over $\mathcal{A} = (\mathcal{X} \times \mathcal{Y}_0) \cup (\mathcal{X}_0 \times \mathcal{Y})$, with $U(x, 0) = V(0, y) = 0$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$, such that at a stable matching:*

(i) a man of full type $\tilde{x} = (x, \varepsilon)$ will match with a woman of an observable type y that maximizes $U(x, y) + \zeta(\tilde{x}, y)$;

(ii) a woman of full type $\tilde{y} = (y, \eta)$ will match with a man of an observable type x that maximizes $V(x, y) + \xi(x, \tilde{y})$;

[6] Which Mailath, Postlewaite, and Samuelson (2013) call the *premunerations*.

[7] It is not even necessary to adopt the perspective we used in this survey, in which we associated ε to the man and η to the woman; see Galichon and Salanié (2022, Appendix C.1) for an example.

(iii) $U(x,y) + V(x,y) \geq \Phi(x,y)$ *for all* (x,y)*, with equality if the set of matches with partners* $\tilde{x} \in x$ *and* $\tilde{y} \in y$ *is non-empty.*

Proof. Fix a stable matching and consider a man of full type $\tilde{x} = (x, \varepsilon)$. Given that the women's utilities are $\tilde{v}(\cdot)$, he can choose to remain single and get utility $\zeta(\tilde{x}, 0)$ or to form a match with a woman of full type $\tilde{y}$ and obtain utility $\tilde{\Phi}(\tilde{x}, \tilde{y}) - \tilde{v}(\tilde{y})$. Therefore he will obtain utility

$$\tilde{u}(\tilde{x}) = \max\left(\zeta(\tilde{x}, 0), \max_{\tilde{y}}\left(\tilde{\Phi}(\tilde{x}, \tilde{y}) - \tilde{v}(\tilde{y})\right)\right). \tag{1.21}$$

The term $\tilde{\Phi}(\tilde{x}, \tilde{y}) - \tilde{v}(\tilde{y})$ can be rewritten

$$\Phi(x,y) + \zeta(\tilde{x}, y) + \xi(x, \tilde{y}) - \tilde{v}(\tilde{y})$$

so that

$$\max_{\tilde{y}}\left(\tilde{\Phi}(\tilde{x}, \tilde{y}) - \tilde{v}(\tilde{y})\right) = \max_{y}\left(\Phi(x,y) + \zeta(\tilde{x}, y) - \min_{\tilde{y} \in y}\left(\tilde{v}(\tilde{y}) - \xi(x, \tilde{y})\right)\right).$$

The minimum in this expression only depends on x and y; let us denote it $\bar{V}(x,y)$. Then (1.21) becomes

$$\tilde{u}(\tilde{x}) = \max\left(\zeta(\tilde{x}, 0), \max_{y}\left(\Phi(x,y) - \bar{V}(x,y) + \zeta(\tilde{x}, y)\right)\right).$$

Remember that $\Phi(x, 0) = 0$; we can extend the function $\bar{V}$ to $\bar{V}(x, 0) = 0$, so that

$$\tilde{u}(\tilde{x}) = \max_{y \in \mathcal{Y}_0}\left(\Phi(x,y) - \bar{V}(x,y) + \zeta(\tilde{x}, y)\right). \tag{1.22}$$

A similar argument proves that

$$\tilde{v}(\tilde{y}) = \max_{x \in \mathcal{X}_0} \left(\Phi(x,y) - \bar{U}(x,y) + \xi(x,\tilde{y}) \right) \tag{1.23}$$

with $\bar{U}(x,y) = \min_{\tilde{x} \in x} (\tilde{u}(\tilde{x}) - \zeta(\tilde{x},y))$ and $\bar{U}(0,y) = 0$. Now define the functions $U = \Phi - \bar{V}$ and $V = \Phi - \bar{U}$. Then equations (1.22) and (1.23) imply that for all $\tilde{x} \in x$ and $\tilde{y} \in y$,

$$(\tilde{u}(\tilde{x}) - \zeta(\tilde{x},y)) + (\tilde{v}(\tilde{y}) - \xi(x,\tilde{y})) \geq U(x,y) + V(x,y).$$

Minimizing over $\tilde{x}$ and $\tilde{y}$ gives $U(x,y) + V(x,y) \geq \Phi(x,y)$. Finally, if $\tilde{x}$ and $\tilde{y}$ are matched then $\tilde{\Phi}(\tilde{x},\tilde{y}) = \tilde{u}(\tilde{x}) + \tilde{v}(\tilde{y})$ implies that $U(x,y) + V(x,y) = \Phi(x,y)$. ∎

The intuition of this result is straightforward. The term $\zeta(\tilde{x},y)$ can be seen as a contribution that $\tilde{x}$ brings to all matches he could establish with women with observable type y. Just as a worker who is \$1 more productive than another in every job will get a \$1 higher wage in equilibrium, a man with a higher $\zeta(\tilde{x},y)$ will reap its value at any stable matching that puts positive mass on such matches. Separability is crucial here: without Assumption 1, the contribution of any man to the joint surplus of a given match would also depend on the unobservable traits η of the woman and this argument would fail.

An immediate consequence of Theorem 6 is that if we know the functions U and V for a stable matching, then we can solve two independent discrete choice problems:

$$\tilde{u}(\tilde{x}) = \max_{y \in \mathcal{Y}_0} (U(x,y) + \zeta(\tilde{x}, y)) \tag{1.24}$$

and

$$\tilde{v}(\tilde{y}) = \max_{x \in \mathcal{X}_0} (V(x,y) + \xi(x, \tilde{y})). \tag{1.25}$$

The values of these problems are the utilities the individuals obtain at a stable matching; and the maximizers define the observable types of their possible equilibrium partners. This will allow us to reuse much of the tools we discussed in Section 1.1.

Assuming separability greatly reduces the complexity of the matching problem: our unknown now are the functions U and V, which are defined on the space of observable types rather than on the space of full types. This is especially valuable when types are discrete, as the problem becomes finite-dimensional and the functions U and V can be represented by matrices. If moreover every type of match in $\mathcal{A}$ occurs, then buy Theorem 6.(iii) we have $V \equiv \Phi - U$ and the problem simplifies further.

1.3.4 Matching with imperfectly transferable utility

With perfectly transferable utility, the utilities U and V are jointly attainable by $\tilde{x}$ and $\tilde{y}$ if $U + V \leq \tilde{\Phi}(\tilde{x}, \tilde{y})$. This is a very specific case; if transfers are costly, limited, or have different marginal values for the partners, we need a more general way to describe the set of feasible utilities. We follow Galichon, Kominers, and Weber (2019) and introduce the *distance function* for this purpose.

Given a set F included in $\mathbb{R}^2$, we define the distance of a point (a, b) to

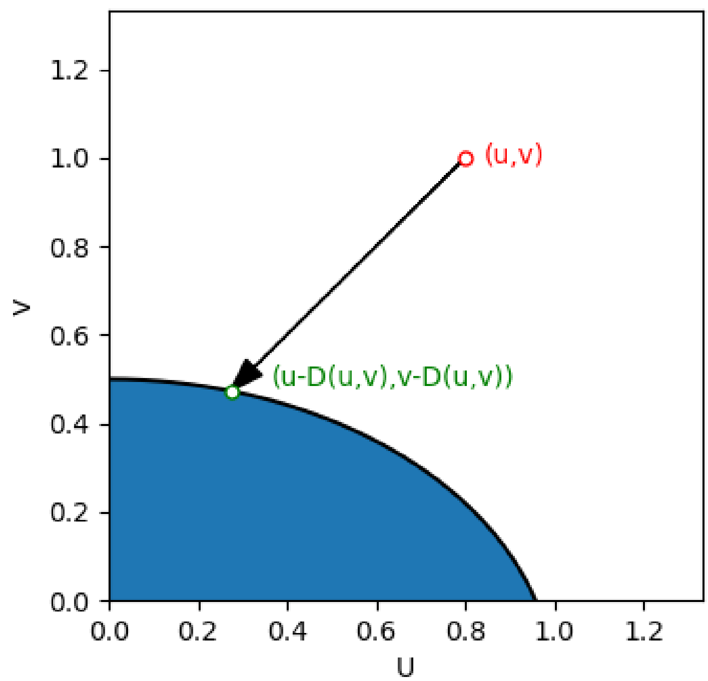

Figure 1.3: Distance function

F as the number (positive or negative) that must be subtracted from (a, b) on the diagonal in order to reach the set F:

$$D_F(a, b) = \min_{t \in \mathbb{R}} \{t \mid (a - t, b - t) \in F\}.$$

Given a potential match $(\tilde{x}, \tilde{y})$, we denote $\tilde{\mathcal{F}}$ the set of all feasible pairs of utility claims (U, V). Then the distance $D_{\tilde{F}}(U, V)$ is the smallest amount of utility that must be subtracted from both prospective partner's claims U and V so that these claims become feasible. By construction, (U, V) is feasible if and only if $D_{\tilde{F}}(U, V) \leq 0$, and $D_{\tilde{F}}(U, V) = 0$ if and only if

(U, V) is on the efficient frontier of $\tilde{\mathcal{F}}$. As in Galichon, Kominers, and Weber (2019), we make three natural assumptions about the feasible sets:

Assumption 2. *Fix $\tilde{x} \neq 0$ and $\tilde{y} \neq 0$; the set $\tilde{\mathcal{F}}$ of feasible payoffs for the match $(\tilde{x}, \tilde{y})$ is such that $D_{\tilde{F}}$ is an* admissible distance function*:*

(i) $D_{\tilde{F}}$ is continuous on all of $\mathbb{R}^2$;

(ii) $D_{\tilde{F}}$ is isotone*: if $u \leq u'$ and $v \leq v'$, then $D_{\tilde{F}}(u, v) \leq D_{\tilde{F}}(u', v')$;*

(iii) Given $(u, v) \in \mathbb{R}^2$, $D_{\tilde{F}}(u + t, v)$ and $D_{\tilde{F}}(u, v + t)$ become positive for t large enough.

The first assumption means that the feasible sets are non-empty and closed. The second one is a free disposal property. The last one is a scarcity requirement that ensures that one side of the relationship cannot obtain an arbitrarily high payoff without imposing an arbitrarily high cost to the partner on the other side. Note that if utility is perfectly transferable, all of these assumptions hold: each set $\tilde{F}$ is a half plane bounded by an infinite line with slope -1, and it is easy to see that the distance function is just the (signed) orthogonal distance to this line: $D_{\tilde{F}}(U, V) = (U + V - \tilde{\Phi}(\tilde{x}, \tilde{y}))/2$.

Not only is the distance function a very convenient tool to express feasibility: it also defines blocking pairs. Suppose that $\tilde{x}$ and $\tilde{y}$ have payoffs (U, V), and let $\tilde{F}$ be the set of feasible payoffs in their match. If $D_{\tilde{F}}(U, V) < 0$, these two individuals could form a match and achieve equal gains of $t = -D_{\tilde{F}}(U, V) > 0$ utils. If on the other hand $D_{\tilde{F}}(U, V) \geq 0$, under Assumption 2 there is no way for them to achieve a greater utility by forming a match.

The following definition summarizes this discussion; it extends Definition 4.

Definition 5 (Matchings with imperfectly transferable utility)**.** *A one-to-one, bipartite* matching market *with imperfectly transferable utility consists of two full type measures $\tilde{n}$ over $\tilde{\mathcal{X}}$ and $\tilde{m}$ over $\tilde{\mathcal{Y}}$; utilities when single $\tilde{\Phi}(\tilde{x},0)$ and $\tilde{\Phi}(0,\tilde{y})$; and a function $d(\tilde{x},\tilde{y},\tilde{u},\tilde{v})$ from $\tilde{\mathcal{X}} \times \tilde{\mathcal{Y}} \times \mathbb{R} \times \mathbb{R}$ to $\mathbb{R}$ such that for every $(\tilde{x} \neq 0, \tilde{y} \neq 0)$ and corresponding set of feasible payoffs $\tilde{F}$, the distance function $D_{\tilde{F}}(U,V) \equiv d(\tilde{x},\tilde{y},U,V)$ is admissible.*

A matching $\tilde{\mu}$ is feasible *if it has margins $\tilde{n}$ and $\tilde{m}$. It is* stable *if and only if there exist two functions $\tilde{u} : \tilde{\mathcal{X}} \to \mathbb{R}$ and $\tilde{v} : \tilde{\mathcal{Y}} \to \mathbb{R}$ such that*

- *$d(\tilde{x},\tilde{y},\tilde{u}(\tilde{x}),\tilde{v}(\tilde{y})) \geq 0$ for all $\tilde{x}$ and $\tilde{y}$;*
- *if $d(\tilde{x},\tilde{y},\tilde{u}(\tilde{x}),\tilde{v}(\tilde{y})) > 0$, then $(\tilde{x},\tilde{y})$ is not in the support of $\tilde{\mu}$;*
- *if $\tilde{x}$ (resp. $\tilde{y}$) is single, then $\tilde{u}(\tilde{x}) = \tilde{\Phi}(\tilde{x},0)$ (resp. $\tilde{v}(\tilde{y}) = \tilde{\Phi}(0,\tilde{y})$).*

The separability assumption of paragraph 1.3.3 can be extended to the imperfectly transferable context.

Assumption 3. *The function d of Definition 5 is* separable *if and only if there exists a family of admissible distance functions $(D_{xy})_{x\neq 0, y\neq 0}$ and functions $\zeta(\tilde{x},y)$ and $\xi(x,\tilde{y})$ such that for all $\tilde{x} \in x \neq 0, \tilde{y} \in y \neq 0$ and corresponding set of feasible payoffs $\tilde{F}$, the distance function $D_{\tilde{F}}$ satisfies the identity*

$$D_{\tilde{F}}(U,V) = D_{xy}(U - \zeta(\tilde{x},y), V - \xi(x,\tilde{y})).$$

To illustrate these concepts, let us suppose that in a $(\tilde{x}, \tilde{y})$ match, the efficient boundary of the set of feasible payoffs $\tilde{F}$ is a parameterized curve whose shape only depends on (x, y), translated by the usual random terms:

$$w \in \mathbb{R} \to (U_{xy}(w), V_{xy}(w)) + (\zeta(\tilde{x}, y), \xi(x, \tilde{y}))$$

where U_{xy} is increasing and V_{xy} is decreasing. Then the model is separable. To see this, fix (U, V) and let $t \in \mathbb{R}$ be the solution to the scalar equation $U_{xy}^{-1}(U - t) + V_{xy}^{-1}(V - t) = 0$. One can easily check that this t defines the distance $D_{xy}(U, V)$ of Assumption 3.

Chapter 2

Identification

We can now use the concepts and results in Section 1 to present identification results. We keep the same notation throughout. Our first task is to show how the distribution of preferences on one side of the market can be recovered from observed choices. This leads us to introduce a new tool, the *generalized entropy*, which we then use to derive identification results in separable matching markets. Finally, we bring demand and supply together in a hedonic market, with a variety of identifying approaches.

2.1 Identifying Demand

Let us examine the demand side of the market—the same approach applies to supply, with obvious changes. Empirical studies of demand on a market typically seek to estimate the distribution of preferences $U(\tilde{x}, z, p)$ for each observed type x. We assume that the analyst observes a sample of individuals $i = 1, \ldots, N$; what we called in Section 1.1 their observed types x_i; and the product z_i they purchased, along with its price p_i. Given an infinite sample, this identifies, in our notation, the distribution $n(dx)$ of observed types, the price function $z \to P(z)$, and the distribution $\mu^D(dz|x) \equiv \pi^D(dz|x; P)$ for this price function P. A fortiori and by integrating, we recover $q^D(dz) \equiv \int \mu^D(dz|x)n(dx)$.

Much of our analysis will take the set of products $\mathcal{Z}$ to be finite; with a slight abuse of notation, we use the notation $\pi^D(z; x, P)$ and $q^D(z)$ for the choice probabilities in this case.

2.1.1 Buyer Choice

We start by computing the market shares associated to given utility parameters. When $\mathcal{Z}$ is finite, we will work with the set of preferences defined in Assumption 4. This is by far the most common choice in applications; still, we note here that it does not generalize readily when $\mathcal{Z}$ is infinite. We will return to this point.

Assumption 4 (Quasilinear Preferences over a Finite Product Set)**.**

1. *$\mathcal{Z}$ is a finite set of products $(z_1, \ldots, z_J)$;*

2. *the preferences of the buyers are quasilinear, as in Example 2:*

$$U(\tilde{x}, z, p) = \bar{U}(x, z) - p + \zeta(\tilde{x}, z);$$

3. *conditional on x, the distribution of the random vector $(\zeta(\tilde{x}, z_j))_{j=1,\ldots,J}$ is a probability distribution $\mathbb{P}_x$ on $\mathbb{R}^J$; and the vectors $\boldsymbol{\zeta}_i = (\zeta_{ij})_{j=1,\ldots,J}$ are iid across all i such that $x_i = x$.*

As part 3 of Assumption 4 may seem opaque, we now turn to a celebrated example.

2.1.1.1 Logit demand

The most common specification adds to Assumption 4 by requiring that for every x, the random vector $\boldsymbol{\zeta}$ be drawn from a standard type I extreme value distribution. This is simply Example 3, with slightly different notation:

Example 7 (Multinomial Logit)**.** *Suppose that the preferences of the buyers satisfy Assumption 4 and that the distribution of the random vector* $(\zeta(\tilde{x}, z_j))_{j=1,\dots,J}$ *conditional on* x *is i.i.d. standard type I extreme value (also known as the Gumbel distribution). More concretely, for every buyer* $i = 1, \dots, n$*, the values* $\zeta(\tilde{x}_i, z_j) \equiv \zeta_{ij}$ *for* $j = 1, \dots, J$ *are iid draws from a standard type I-EV distribution; and the vectors* $\boldsymbol{\zeta}_i = (\zeta_{ij})_{j=1,\dots,J}$ *are iid across* i*.*

We already computed the corresponding choice probabilities in Section 1.1.2: for an observed type x,

$$\pi^D(z_j|x;P) = \frac{\exp(\bar{U}(x, z_j))}{\sum_{k=1}^J \exp(\bar{U}(x, z_k))}$$

and the market share over the population is

$$q^D(z_j) = \int_{x \in \mathcal{X}} n(dx) \frac{\exp(\bar{U}(x, z_j))}{\sum_{k=1}^J \exp(\bar{U}(x, z_k))}.$$

2.1.1.2 Other quasilinear models of demand

Even within the class of quasilinear separable models, the multinomial logit is restrictive in several ways. It imposes a parameter-free distribution; and it does not allow for heteroskedasticity or for any form of correlation of taste shocks across products[1]. In addition, it exhibits the Independence of Irrelevant Alternatives property, which is often violated in the data.

We now return to the more general Assumption 4. Since the set of products here is finite, we simplify notation by denoting

- $\boldsymbol{\zeta}^x$ the random vector $(\zeta(\tilde{x}, z_j))_{j=1,\dots,J}$ for given x;
- $\boldsymbol{\mu}^x = (\mu_j^x)_{j=1,\dots,J}$ the vector whose j-th component is the conditional probability $\mu(z_j|x)$ that a buyer of observed type x chooses product j;
- $\boldsymbol{U}(x; P)$ the vector formed by the $(U_j(x; P))_{j=1,\dots,J}$, where $U_j(x; P) \equiv \bar{U}(x, z_j) - P(z_j)$;
- and we denote $E_{\mathbb{P}_x} f(\boldsymbol{\zeta})$ the expectation $\int f(\boldsymbol{\zeta})\mathbb{P}_x(d\boldsymbol{\zeta})$ of $f(\boldsymbol{\zeta})$ under $\mathbb{P}_x$.

A buyer of type $\tilde{x}$, faced with products $(z_j)_{j=1,\dots,J}$ and the price function P, will choose the product that maximizes $U_j(x; P) + \zeta_j^x$. Taking expectations over $\boldsymbol{\zeta}^x$, buyers of observed type x obtain on average the utility

$$G_x(\boldsymbol{U}(x; P)) \equiv E_{\mathbb{P}_x} \max_{j=1,\dots,J} \left(U_j(x; P) + \zeta_j^x \right).$$

[1] A multinomial probit with a flexible variance-covariance matrix would; it is much less attractive for computational reasons, however.

This G_x function is often called the "Emax" function, as the expectation of the maximum utility a buyer can achieve. We will use it frequently in this chapter. Since G_x is the expectation of a maximum of functions that are linear in $\boldsymbol{U}(x;P)$, it is a convex function of its arguments. As such, it has a subgradient everywhere, and is differentiable almost everywhere. By the classic Daly–Zachary–Williams result[2], the probability that a buyer with observed type x chooses product j is in its subgradient:

$$\mu_j^x \in \partial G_x(\boldsymbol{U}(x;P)). \tag{2.1}$$

In the logit case, a well-known formula gives

$$G_x(\boldsymbol{U}(x;P)) = \log \sum_{j=1}^{J} \exp\left(U_j(x,P)\right),$$

which is everywhere differentiable, so that

$$\mu_j^x = \frac{\exp\left(U_j(x,P)\right)}{\sum_{k=1}^{J} \exp\left(U_k(x,P)\right)}, \tag{2.2}$$

as expected.

These formulæ show why the case when $\mathcal{Z}$ is infinite requires special precautions: as the number of products goes to infinity, the value of the Emax function goes to infinity. There are several ways to circumvent this difficulty. Menzel (2015) lets the standard errors of the $\boldsymbol{\zeta}$ terms shrink at the appropriate rate. Dagsvik (1994) followed by Dupuy and Galichon (2014) assume a specific form of frictions in the way partners can meet. We will

[2] See Daly and Zachary (1978) and Williams (1977).

describe the latter in much more detail in Section 3.2.3.4.

2.1.2 Inverting Demand

The probabilities $\boldsymbol{\mu}^x$ are observed in the data; to recover the preferences of the buyers, we need to invert the inclusion in 2.1. To do so, we use *convex duality.* Let G_x^* denote the Legendre–Fenchel transform of the convex function G_x: for nonnegative $\boldsymbol{q} = (q_1, \ldots, q_J)$ such that $\sum_{j=1}^J q_j = 1$,

$$G_x^*(\boldsymbol{q}) = \max_{\boldsymbol{u} \in \mathbb{R}^J} \left(\boldsymbol{q} \cdot \boldsymbol{u} - G_x(\boldsymbol{u}) \right). \tag{2.3}$$

Note that restricting $\boldsymbol{q}$ to the simplex is necessary: it is easy to see that for any scalar t we have $G_x(\boldsymbol{u}+t) \equiv G_x(\boldsymbol{u})+t$, so that the maximum would be infinite if we did not impose that $\sum_j q_j = 1$. Therefore the vector $\boldsymbol{q}$ must be a probability over $\{1, \ldots, J\}$. Galichon and Salanié (2022) call the function G_x^* the *generalized entropy of choice.*

The function G_x^* is a maximum of functions that are linear in q, and is therefore convex. By the envelope theorem, its subgradient at $\boldsymbol{q}$ is the set of vectors $\boldsymbol{u}$ that achieve the maximum. Since G_x is convex, these vectors are given by the first-order condition $\boldsymbol{q} \in \partial G_x(\boldsymbol{u})$.

Comparing with (2.1) shows that

$$\boldsymbol{U}(x; P) \in \partial G_x^*(\boldsymbol{\mu}^x). \tag{2.4}$$

This inclusion calls for three comments. First, if $\boldsymbol{u} \in \partial G_x^*(\boldsymbol{q})$ then for any scalar t we also have $\boldsymbol{u} + t \in \partial G_x^*(\boldsymbol{q})$. This is simply reflecting the fact that $G_x(\boldsymbol{u}+t) = G_x(\boldsymbol{u}) + t$; it implies that at best utilities are only defined up to an additive constant, as in any discrete choice model. Normalizing their location by (for instance) $U_1(x;P) = 0$ for all x fixes the constant. Moreover, the subgradient of G_x is a singleton when G_x is continuously differentiable. If G_x is also strictly convex, then the subgradient of G_x^* is a line $\{\boldsymbol{u}\} + \mathbb{R}\mathbf{1}_J$, where $\mathbf{1}_J$ is a vector of ones in $\mathbb{R}^J$ that reflects the indeterminacy of the general level of utilities. We will define the following two important assumptions:

Assumption 5. *For each $x \in \mathcal{X}$, the probability distribution $\mathbb{P}_x$ has no mass point: $\mathbb{P}_x(\{\boldsymbol{e}\}) = 0$ for any $\boldsymbol{e} \in \mathbb{R}^J$.*

and

Assumption 6. *For each $x \in \mathcal{X}$, the probability distribution $\mathbb{P}_x$ has full support.*

One can show that:

Theorem 8. *Assumption 5 implies that G_x is differentiable: for any $\boldsymbol{u}$,*

$$\partial G_x(\boldsymbol{u}) = \left\{ \frac{\partial G_x}{\partial \boldsymbol{u}}(\boldsymbol{u}) \right\}.$$

Assumption 6 implies that that G_x is strictly convex on any subspace defined by $U_k = c$ for some $k \in \{1, ..., J\}$ and some constant number c. This in turn implies that for every $\boldsymbol{\mu}$, $\partial G_x^*(\boldsymbol{\mu})$ contains only one point $\boldsymbol{U}$ such that $U_k = c$.

To illustrate this, we return to the logit model of Example 7. Since the distribution of the $\boldsymbol{\zeta}^x$ is parameter-free, our goal is simply to recover an estimate of the function $\bar{U}$. With this specification, equation (2.2) gives

$$U_j(x; P) = \log \mu_j^x + \log \sum_{k=1}^{J} \exp(U_k(x; P));$$

making the Emax term on the right-hand side equal to some arbitrary number c_x identifies all $U_j(x; P)$. Then the utility function $\bar{U}$ is identified non-parametrically on $\mathcal{Z}$ by

$$\bar{U}(x, z_j) = P(z_j) + \log \mu_j^x + c_x.$$

To put it differently, we identify the differences

$$\bar{U}(x, z_j) - \bar{U}(x, z_k) = P(z_j) - P(z_k) + \log \mu_j^x / \mu_k^x$$

for all $j, k = 1, \ldots, J$.

To summarize our results:

Theorem 9 (Identification in the Hedonic Demand Model)**.** *Under Assumption 4, denote P the observed price function and take any observed type x.*

1. *For any choice of the distribution $\mathbb{P}_x$, the differences in the utilities $\boldsymbol{U}(x; P)$ are partially identified by*

$$\boldsymbol{U}(x; P) \in \partial G_x^*(\boldsymbol{\mu}^x).$$

For any j and $k \neq j$, the difference $\bar{U}(x, z_j) - \bar{U}(x, z_k)$ is in the set

$$\partial_j G_x^*(\boldsymbol{\mu}^x) - \partial_k G_x^*(\boldsymbol{\mu}^x) + P(z_j) - P(z_k).$$

2. *if the distribution $\mathbb{P}_x$ yields a strictly convex and differentiable G_x, then the differences in the utilities $\boldsymbol{U}(x; P)$ are point-identified; for any j and $k \neq j$,*

$$\bar{U}(x, z_j) - \bar{U}(x, z_k) = \frac{\partial G_x^*}{\partial \mu_j^x}(\boldsymbol{\mu}^x) - \frac{\partial G_x^*}{\partial \mu_k^x}(\boldsymbol{\mu}^x) + P(z_j) - P(z_k).$$

3. *if $\mathbb{P}_x$ satisfies Assumptions 5 and 6, then 2. obtains.*

These results depend on observing the prices of all varieties. It is clear from Theorem 9 that partial identification of the price function translates directly into partial identification of the utilities. Even more importantly, the identification is conditional on the distribution $\mathbb{P}_x$ being fully known. If the analyst only assumes that $\mathbb{P}_x$ belongs to a family of distributions, this indeterminacy again translates into partial identification of the differences of the utilities. We will give an example in Section 2.1.3.

2.1.3 Implementing demand inversion

Sometimes, as in Example 7, the functions G_x and G_x^* are available in closed form. It is easy to check that given the multinomial logit structure,

$$G_x^*(\boldsymbol{q}) = \sum_{j=1}^{J} q_j \log q_j,$$

the entropy of the probability distribution $\boldsymbol{q}$. To make it slightly more general, assume that $\mathbb{P}_x$ has a scale parameter σ_x. Then $G_x(\boldsymbol{u}) = \sigma_x \log \sum_{j=1}^J \exp(u_j/\sigma_x)$ and $G_x^*(\boldsymbol{q}) = \sigma_x \sum_{j=1}^J q_j \log q_j$, whose subgradient is

$$\partial G_x^*(\boldsymbol{q}) = \sigma_x \{\log \boldsymbol{q}\} + \mathbb{R}\mathbf{1}_J.$$

Therefore differences of utilities are identified by $U_j(x;P) - U_k(x;P) = \sigma_x \log(\mu_j^x/\mu_k^x)$ if σ_x is known a priori. If not, then only ratios of differences of utilities are identified: we could for instance normalize location and scale by $U_1(x;P) = 0$ and $U_J(x;P) = \log(\mu_J^x/\mu_1^x)$ to set $\sigma_x = 1$ and $U_j(x;P) = \log(\mu_j^x/\mu_1^x)$.

This illustrates a general point: in discrete choice models, even under our separability assumption, preferences are only identified nonparametrically if the distribution $\mathbb{P}_x$ of their individual-specific component (our vector $\boldsymbol{\zeta}^x$) is "known"—which usually means "assumed". Since we only observe the J numbers in $\boldsymbol{\mu}^x$ for each x, and their sum is fixed, we cannot hope to identify more than the J values of the vector $\boldsymbol{U}$, up to an additive constant. The alternatives are to use a parametric specification for the preferences $\bar{U}(x,z)$ and the distribution of $\boldsymbol{\zeta}(x,\tilde{z})$; to restrict the variation in preferences across x; and more generally, to combine additional data and exclusion restrictions. We will return to this point in more detail in Section 3.1.

2.1.4 Demand and inverse demand beyond the quasi-linear case

It is possible to extend the previous results to preferences $U(\tilde{x}, z, p)$ that are neither quasilinear nor separable, even when $\mathcal{Z}$ is not a finite set. To see this, remember that the support of the individual demand $\tilde{\pi}^D(\cdot|\tilde{x}; P)$ of a full type $\tilde{x}$ for the price function P is

$$\mathcal{S}(\tilde{x}; U, P) \equiv \arg\max_{z \in \mathcal{Z}} U(\tilde{x}, z, P(z)).$$

Now fix an observed type x and let $\tilde{n}(d\tilde{x}|x)$ the distribution of full types whose observed types is x, as in Section 1.1. Take a measurable subset B of $\mathcal{Z}$. Let $\tilde{\mathcal{X}}_B^-$ (resp. $\tilde{\mathcal{X}}_B^+$) denote the set of $\tilde{x}$ such that $\mathcal{S}(\tilde{x}; U, P) \subset B$ (resp. $\mathcal{S}(\tilde{x}; U, P) \cap B \neq \emptyset$). While $\tilde{\mathcal{X}}_B^-$ consists of the full types whose preferred choices are all in B, the larger set $\tilde{\mathcal{X}}_B^+$ also includes full types some of whose preferred choices are in B. The measure $\mu^x(B)$ of the set of observed types who choose a product in B clearly cannot be smaller than that of $\tilde{\mathcal{X}}_B^-$; and it cannot be greater than that of $\tilde{\mathcal{X}}_B^+$. A theorem of Strassen (1965) can be used to prove that this is in fact an exact characterization of the demand μ^x: we can replace the identification inclusion (2.4) with $\mu^x \in Q_x$, where Q_x is the set of such measures ν over $\mathcal{Z}$ such that for all measurable subsets B of $\mathcal{Z}$,

$$\tilde{n}(\tilde{\mathcal{X}}_B^-|x) \leq \nu(B) \leq \tilde{n}(\tilde{\mathcal{X}}_B^+|x). \tag{2.5}$$

This type of characterization has been used in the context of partial identification, see e.g. Galichon and Henry (2011).

If the preferred set of products $\mathcal{S}(\tilde{x}; U, P)$ of almost every full type $\tilde{x}$ with observed type x is a singleton under the price function P, then the sets $\tilde{\mathcal{X}}_B^-$ and $\tilde{\mathcal{X}}_B^+$ coincide and μ^x is uniquely determined.

Since the $\tilde{\mathcal{X}}_B$ sets depend on the family of preferences U, the inequalities (or equalities) in (2.5) can be used to partially identify U. The difficulty of this demand inversion very much depends on the specification of U. To take a trivial example, let us modify the logit model of Example 7 by replacing $\bar{U}(x,z) - P(z)$ with some $\bar{U}(x,z,P(z))$. Then (2.5) boils down to requiring that

$$\mu^x(z_j) = \frac{\exp(\bar{U}(x, z_j, P(z_j))}{\sum_{k=1}^J \exp(\bar{U}(x, z_k, P(z_k))}$$

for all products z_j; inversion gives $\bar{U}(x, z_j, P(z_j)) = \log \mu^x(z_j) + c_x$ for an arbitrary family of numbers c_x.

2.1.5 Endogenous Prices

It is worth stressing here the similarities and differences with the way demand is modeled in empirical industrial organization. Typically, empirical IO models choices between a finite number of varieties, with heterogeneous consumer preferences over a vector of product characteristics $\boldsymbol{C}(z)$. Using our notation, a generic Berry, Levinsohn, and Pakes (1995) specification would have $\tilde{x}_i = (x_i, \boldsymbol{\varepsilon}_i)$ and

$$U(\tilde{x}_i, z_j, p_j) = \boldsymbol{C}(z_j) \cdot (\boldsymbol{a}(x_i) + \boldsymbol{\varepsilon}_i) - p_j + \xi_j + v_{ij}. \tag{2.6}$$

Here the vector $\boldsymbol{v}_i = (v_{ij})_{j=1,\dots,J}$ consists of J iid draws from a type I-EV distribution, and the vectors $\boldsymbol{v}_i$ are iid across i. The “product effects” ξ_j

are unobserved and usually correlated with p_j. The goal of the analyst is to estimate the vector function $\boldsymbol{a}$ and the distribution of $\boldsymbol{\varepsilon}$ conditional on x. This is typically done by imposing a set of moment conditions $E(\xi_j|W_j) = 0$.

The first two terms in (2.6) map directly into the preferences we are using:

$$\bar{U}(x,z) = \boldsymbol{C}(z_j) \cdot \boldsymbol{a}(x_i)$$
$$\zeta_j^x = \boldsymbol{C}(z_j) \cdot \boldsymbol{\varepsilon}_i.$$

On the other hand, IO includes the additional term $(\xi_j + v_{ij})$ to represent the effect of unobserved product characteristics. While the methods presented in this chapter cannot incorporate the idiosyncratic terms (v_{ij}), it is easy to accommodate the unobserved market-wide product effects ξ_j. All we need to do is change $U_j(x;P)$ to $U_j(x;P) + \xi_j$ to get

$$\boldsymbol{U}(x;P) + \boldsymbol{\xi} \ \in \ \partial G_x^*(\boldsymbol{\mu}^x).$$

Suppose for simplicity that G_x^* is strictly convex and differentiable, and fix the general level of utilities. Then the set of moment conditions $E(\xi_j|W_j) = 0$ can be rewritten as

$$E\left(\frac{\partial G_x^*}{\partial \mu_j^x}(\boldsymbol{\mu}^x) - U_j(x;P) \,\middle|\, W_j\right) = 0 \ \text{ for } \ j = 1, \dots, J,$$

which identifies $\boldsymbol{U}(x;P)$ under completeness conditions—in practical terms, if the moment conditions are informative enough.

To put it differently, our specification is very close to the "pure characteristics" model of Berry and Pakes (2007).

2.2 Identification in Separable Matching Models

The primitives of a separable matching problem consist of

- the joint distribution $\tilde{n}$ of $\tilde{x} = (x, \varepsilon)$; we denote n its margin on x, and $\mathbb{P}_x$ the distribution of ε conditional on x;
- the joint distribution $\tilde{m}$, with margin m on y and a conditional distribution $\mathbb{Q}_y$ of η given y;
- the functions Φ, ζ, and ξ.

We saw in Section 1.3.2 that any matching problem with perfectly transferable utility can be reframed as an optimal transportation problem. Consider the dual version, as stated in (1.20): we minimize the sum of utilities

$$\int_{\tilde{\mathcal{X}}} \tilde{u}(\tilde{x})\tilde{n}(d\tilde{x}) + \int_{\tilde{\mathcal{Y}}} \tilde{v}(\tilde{y})\tilde{m}(d\tilde{y})$$

under the stability constraints. Theorem 6 implies that at the optimum , there exists a function U such that

$$\int \tilde{u}(\tilde{x})\tilde{n}(d\tilde{x}) = \int \max_{y \in \mathcal{Y}_0}(U(x,y)+\zeta(x,\varepsilon,y))n(dx)\mathbb{P}_x(d\varepsilon) = \int G_x(U(x,\cdot))n(dx)$$

where $G_x(U(x,\cdot))$ is the "Emax" function for men of observable type x:

$$G_x(U(x,\cdot)) = \int \max_{y \in \mathcal{Y}_0}(U(x,y) + \zeta(x,\varepsilon,y))\mathbb{P}_x(d\varepsilon).$$

This is exactly the function we introduced in Section 1.1.2.3 for one-sided choice problems; again, we normalize it by $U(x, 0) = 0$ for all x. Define $H_y(V(\cdot, y))$ in a similar way for women of observable type y. With this notation, the social welfare is the sum of Emax utilities

$$\mathcal{W} = \int G_x(U(x,\cdot))n(dx) + \int H_y(V(\cdot, y))m(dy). \tag{2.7}$$

Like the functions G_x and H_y, it is convex; all we have to do is minimize it with respect to U and $V \geq \Phi - U$, then define $\tilde{u}$ and $\tilde{v}$ by (1.24) and (1.25), e.g.

$$\tilde{u}(\tilde{x}) = \max_{y \in \mathcal{Y}_0}(U(x,y) + \zeta(x, \varepsilon, y).$$

By construction, the functions $\tilde{u}$ and $\tilde{v}$ satisfy the stability conditions.

To make this more concrete, let us suppose that the sets of observed types are finite and all (x, y) matching cells are non-empty. We will return to continuous types in Section 3.2.3.4. Assuming full support is a minor technical convenience.

Assumption 7 (Discrete types)**.** *The sets of observed types are $\mathcal{X} = \{1, \ldots, X\}$ and $\mathcal{Y} = \{1, \ldots, Y\}$; no matching cell (x, y) in $\mathcal{A} = (\mathcal{X} \times \mathcal{Y}_0) \cup (\mathcal{X}_0 \times \mathcal{Y})$ is empty.*

With discrete types, $\boldsymbol{\Phi}, \boldsymbol{U}$ and $\boldsymbol{V}$ are $X \times Y$ matrices. We accordingly adapt the notation by using subscripts instead of function arguments: e.g. type x (resp. y) has a mass n_x (resp. m_y); the matrix $\boldsymbol{U}$ has rows $\boldsymbol{U}_{x\cdot}$ and columns $\boldsymbol{U}_{\cdot y}$. The social welfare in (2.7) becomes

$$\sum_{x=1}^{X} n_x G_x(\boldsymbol{U}_{x\cdot}) + \sum_{y=1}^{Y} m_y H_y(\boldsymbol{V}_{\cdot y}).$$

The full support assumption has two consequences. First, $\boldsymbol{V} = \boldsymbol{\Phi} - \boldsymbol{U}$ by Theorem 6. Second, all functions G_x and H_y are differentiable at the optimum. Then minimizing the social welfare with respect to $\boldsymbol{U}$ gives first-order conditions

$$n_x \frac{\partial G_x}{\partial U_{xy}}(\boldsymbol{U}_{x\cdot}) = m_y \frac{\partial H_y}{\partial V_{xy}}(\boldsymbol{\Phi}_{\cdot y} - \boldsymbol{U}_{\cdot y}). \tag{2.8}$$

By the Daly-Zachary-Williams theorem, the left-hand side of (2.8) is the number of matches of men of observable type x with women of observable type y. These first-order conditions simply state that at the minimum of (2.7), it must equal the number of matches of women of observable type y with men of observable type x. Therefore the matrix $\boldsymbol{U}$ supports a stable matching. In turn, the common value μ_{xy} of the two sides is the number of matches in the (x, y) cell of observable types; and the utilities are distributed as per (1.24) and (1.25).

These results are proved rigorously in Galichon and Salanié (2022). While useful, they still do not give a straightforward way to solve for $\boldsymbol{U}$. In the discrete example, the first-order conditions are a system of $X \times Y$ nonlinear equations in as many unknowns. We could minimize the social welfare directly, as it is globally convex. This can still be costly if there are many possible observable types on each side of the match. Fortunately, the theory of *convex duality* allows us to go even further and to derive a closed-form solution.

First note that under Assumption 7, a *feasible matching* is simply a matrix $\boldsymbol{\mu}$ whose elements μ_{xy} are positive and sum up to no more than the margins $\boldsymbol{n} = (n_1, \ldots, n_X)$ on rows and to no more than the margins $\boldsymbol{m} = (m_1, \ldots, m_Y)$ on columns: for all x and y,

$$\sum_{t=1}^{Y} \mu_{xt} \leq n_x \text{ and } \sum_{z=1}^{X} \mu_{zy} \leq m_y.$$

Note that this defines implicitly $\mu_{x0} = n_x - \sum_{y=1}^{Y} \mu_{xy}$ (resp. $\mu_{0y} = m_y - \sum_{z=1}^{X} \mu_{zy}$) as the number of men of type x (resp. of women of type y) who end up single in the matching $\boldsymbol{\mu}$.

For any feasible matching, we define the conditional matching probabilities $\mu_{y|x} = \mu_{xy}/n_x$ and $\mu_{x|y} = \mu_{xy}/m_y$; the feasibility constraints become $\sum_{y=1}^{Y} \mu_{y|x} \leq 1$ and $\sum_{x=1}^{X} \mu_{x|y} \leq 1$. Under our full support assumption, these inequalities are strict and all conditional probabilities are positive.

Now remember from equation (2.3) that we can associate to the convex function G_x its convex (Legendre-Fenchel) transform by

$$G_x^*(\boldsymbol{\mu}_{\cdot|x}) = \max_{\boldsymbol{U}_{x\cdot}} \left(\sum_{y=1}^{Y} \mu_{y|x} U_{xy} - G_x(\boldsymbol{U}_{x\cdot}) \right)$$

for all $\boldsymbol{\mu}_{\cdot|x}$ such that $\sum_{y=1}^{Y} \mu_{y|x} < 1$.

Applying the envelope theorem to the above program gives the set of conditions

$$\frac{\partial G_x^*}{\partial \mu_{y|x}}(\boldsymbol{\mu}_{\cdot|x}) = U_{xy},$$

which associate to any feasible matching $\boldsymbol{\mu}$ the matrix $\boldsymbol{U}$ that rationalizes it:

$$\frac{\partial G_x}{\partial U_{xy}}(\boldsymbol{U}_{\cdot|x}) = \mu_{xy}.$$

This applies in particular to the equilibrium matching; writing

$$\begin{aligned} U_{xy} &= \frac{\partial G_x^*}{\partial \mu_{y|x}}(\boldsymbol{\mu}_{\cdot|x}) \\ V_{xy} &= \frac{\partial H_y^*}{\partial \mu_{x|y}}(\boldsymbol{\mu}_{\cdot|y}) \end{aligned}$$

and adding up identifies Φ_{xy} as

$$\Phi_{xy} = \frac{\partial G_x^*}{\partial \mu_{y|x}}(\boldsymbol{\mu}_{\cdot|x}) + \frac{\partial H_y^*}{\partial \mu_{x|y}}(\boldsymbol{\mu}_{\cdot|y}). \tag{2.9}$$

If $\boldsymbol{U}_{x\cdot}$ rationalizes $\boldsymbol{\mu}_{\cdot|x}$, then $G_x(\boldsymbol{U}_{x\cdot}) + G_x^*(\boldsymbol{\mu}_{\cdot|x}) = \sum_{y=1}^{Y} \mu_{y|x} U_{xy}$. Multiplying by n_x gives $n_x G_x(\boldsymbol{U}_{x\cdot}) + n_x G_x^*(\boldsymbol{\mu}_{x\cdot}/n_x) = \sum_{y=1}^{Y} \mu_{xy} U_{xy}$; summing over x and adding the corresponding equations on the y side, we obtain

$$\begin{aligned} \sum_{x=1}^{X}\sum_{y=1}^{Y} \mu_{xy}\Phi_{xy} = &\left(\sum_{x=1}^{X} n_x G_x(\boldsymbol{U}_{x\cdot}) + \sum_{y=1}^{Y} m_y H_y(\boldsymbol{V}_{\cdot y})\right) \\ &+ \left(\sum_{x=1}^{X} n_x G_x^*(\boldsymbol{\mu}_{x\cdot}/n_x) + \sum_{y=1}^{Y} m_y H_y^*(\boldsymbol{\mu}_{\cdot y}/m_y)\right). \end{aligned}$$

The first term on the right hand side is the value of social welfare $\mathcal{W}$ (see (2.7)). The second term is the opposite of the *generalized entropy* $\mathcal{E}$, defined in Galichon and Salanié (2022) as the concave function

$$\mathcal{E}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}) = -\sum_{x=1}^{X} n_x G_x^*(\boldsymbol{\mu}_{x\cdot}/n_x) - \sum_{y=1}^{Y} m_y H_y^*(\boldsymbol{\mu}_{\cdot y}/m_y).$$

Given this definition, the closed-form identification formula (2.9) can be rewritten as

$$\boldsymbol{\Phi} = -\frac{\partial \mathcal{E}}{\partial \boldsymbol{\mu}}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}). \tag{2.10}$$

These are simply the first-order conditions for the maximization of the social welfare

$$\mathcal{W}(\boldsymbol{\Phi}, \boldsymbol{n}, \boldsymbol{m}) = \max_{\boldsymbol{\mu}} \left(\sum_{x=1}^{X} \sum_{y=1}^{Y} \mu_{xy} \Phi_{xy} + \mathcal{E}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}) \right). \tag{2.11}$$

Suppose we scale up the distribution of the error terms by a factor $\sigma > 0$. It is easy to see that

$$G_x^\sigma(\boldsymbol{U}_{x\cdot}) = E \max_{y=0,1,\dots,Y} \big(U_{xy} + \sigma \zeta_{xy}(\varepsilon) \big) = \sigma G_x^1(\boldsymbol{U}_{x\cdot}/\sigma)$$

and its convex transform $(G_x^\sigma)^*$ is simply $\sigma (G_x^1)^*$. When σ is very small, matching only involves observables; the generalized entropy function becomes negligible and the social welfare reduces to the first term in (2.11). Matching on unobservables adds another contribution to the social welfare $\mathcal{W}$, via the generalized entropy term.

2.2.1 Gaining Identification Power

It is important to understand what (2.9) achieves. It identifies the matrix $\mathbf{\Phi}$ as a function of the observed matching patterns $\boldsymbol{\mu}$ in so far as we can compute the derivatives of the Legendre-Fenchel transforms G_x^* and H_y^*. Since the latter depend on G_x and on H_y, which in turn integrate over the distribution of the unobservable types, this gives us nonparametric identification conditional on knowing (or assuming) the distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$ of the 2-way terms $\zeta_{xy}(\varepsilon)$ and $\xi_{xy}(\eta)$. There is just not enough information in the data on matching patterns to identify more, or even to test the model: for any assumed distributions of the ζ and ξ terms, for any observed matching patterns $\boldsymbol{\mu}$, there exists a (unique) $\mathbf{\Phi}$ that rationalizes $\boldsymbol{\mu}$.

There are (at least) two possible responses. One is to parameterize or otherwise restrict the specification of $\mathbf{\Phi}$, and to reuse degrees of freedom in the specification of ζ and ξ. The other response is to add more data; this has mostly been done by pooling data from several matching markets and assuming that some elements of the specification are constant across markets. These approaches are clearly non-exclusive. We start with the former, and we will describe the latter in Section 3.2.3.1.

2.2.2 The Logit Model

The most popular specification of the separable model is the multinomial logit of Choo and Siow (2006). In addition to separability (Assumption 1) and discrete types (Assumption 7), Choo and Siow assumed that the error terms have the familiar structure that underlies one-sided multinomial discrete choice models.

Assumption 8 (Logit Specification)**.**

$$\zeta_{xy}(\varepsilon) = \varepsilon_y \quad and \quad \xi_{xy}(\eta) = \eta_x$$

for all (x, y) *in* $\mathcal{A}$*, where the vectors* $(\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_Y)$ *and* $(\eta_0, \eta_1, \ldots, \eta_X)$ *are independent draws from the standardized type-I extreme value distribution.*

The calculations become very simple and familiar. Consider an infinitely large matching market. Then the "Emax" functions take the familiar "log-sum-exp" form; for instance,

$$G_x(\boldsymbol{U}_{x\cdot}) = \log\left(1 + \sum_{t=1}^{Y} \exp(U_{xt})\right)$$

where we used $\exp(U_{x0}) = \exp(0) = 1$. The conditional matching patterns associated to a given U are

$$\mu_{y|x} = \frac{\exp(U_{xy})}{1 + \sum_t \exp(U_{xt})}$$

and the Legendre-Fenchel transforms are

$$G_x^*(\boldsymbol{\mu}_{\cdot|x}) = \sum_{y=1}^{Y} \mu_{y|x} \log \mu_{y|x} + \mu_{0|x} \log \mu_{0|x},$$

defining $\mu_{0|x} = 1 - \sum_{y=1}^{Y} \mu_{y|x}$. In this case, the generalized entropy $\mathcal{E}$ is simply the standard entropy

$$\mathcal{E}(\boldsymbol{\mu};\boldsymbol{n},\boldsymbol{m}) = -\sum_{x=1}^{X}\sum_{y=1}^{Y}\mu_{xy}\log\frac{\mu_{xy}^2}{n_x m_y} - \sum_{x=1}^{X}\mu_{x0}\log\frac{\mu_{x0}}{n_x} - \sum_{y=1}^{Y}\mu_{0y}\log\frac{\mu_{0y}}{m_y}.$$

Equation (2.9) gives the very simple *Choo and Siow formula*

$$\Phi_{xy} = \log\frac{\mu_{xy}^2}{\mu_{x0}\mu_{0y}}. \tag{2.12}$$

2.2.3 Other Specifications

Since it has no free distributional parameter, the logit specification circumvents the identification issues mentioned in Section 2.2.1. On the other hand, it suffers from the usual problems of the multinomial logit: it has very constrained comparative statics, and relabeling the types has spurious effects[3].

By construction, the unobservable shocks ε_y and ε_t in the logit specification are independent if $y \neq t$. Even with discrete variables, natural notions of distance between y and t often emerge. In such cases, a more reasonable specification would allow for a form of "local" correlation by that metric. The more general results in Section 2.2 can accommodate such models, and more.

A first, simple change is to keep the type-I distributional form and the independence of the draws in Assumption 8, but to allow for scale factors

[3]See Galichon and Salanié (2022, Appendix C.2), and Galichon and Salanié (2019) for a longer discussion.

that may depend on the type:

$$\zeta_{xy}(\varepsilon) = \sigma_x \varepsilon_y \text{ and } \xi_{xy}(\eta) = \tau_y \eta_x$$

where $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$ are vectors of positive numbers. The identification formulæ become

$$\Phi_{xy} = \log \frac{\mu_{xy}^{\sigma_x+\tau_y}}{\mu_{x0}^{\sigma_x}\mu_{0y}^{\tau_y}}.$$

We give formulæ for the two-level nested logit and for a mixed logit in Galichon and Salanié (2022, Section 2). The logit structure is not essential, however; for Generalized Extreme Value specifications, see Galichon and Salanié (2022, Appendix B).

2.2.4 Beyond Transferable Utility

Galichon, Kominers, and Weber (2019) showed how formula (2.9) extends to the imperfectly transferable utility specification of Section 1.3.4. Recall that in this class of separable models, the frontier of the set of feasible utilities for a match $(\tilde{x}, \tilde{y})$ is the set of $(\tilde{U}, \tilde{V})$ pairs such that

$$D_{xy}\left(\tilde{U} - \zeta(\tilde{x}, y), \tilde{V} - \xi(\tilde{y}, x)\right) = 0, \tag{2.13}$$

where D_{xy} is an admissible distance function[4].

Under these assumptions, we still have

$$\tilde{U} = \frac{\partial G_x^*}{\partial \mu_{y|x}}\left(\mu_{\cdot|x}\right) + \zeta(\tilde{x}, y) \text{ and } \tilde{V} = \frac{\partial H_y^*}{\partial \mu_{x|y}}\left(\mu_{\cdot|y}\right) + \xi(x, \tilde{y}), \tag{2.14}$$

[4]See Assumptions 2 and 3.

where G_x^* and H_y^* are the Legendre-Fenchel transforms of the Emax functions that correspond to the distributions of the $\boldsymbol{\zeta}$ and $\boldsymbol{\xi}$, respectively.

Substituting (2.14) in (2.13) gives, for all (x, y),

$$D_{xy}\left(\frac{\partial G_x^*}{\partial \mu_{y|x}}\left(\mu_{\cdot|x}\right), \frac{\partial H_y^*}{\partial \mu_{x|y}}\left(\mu_{\cdot|y}\right)\right) = 0. \tag{2.15}$$

This set of equations provides identifying information on the function D_{xy}, given known/assumed distributions for $\boldsymbol{\zeta}$ and $\boldsymbol{\xi}$.

In the logit case, we know that

$$U_{xy} = \log \frac{\mu_{xy}}{\mu_{x0}} \text{ and } V_{xy} = \log \frac{\mu_{xy}}{\mu_{0y}}. \tag{2.16}$$

Using the property $D_{xy}\left(u + c, v + c\right) = D_{xy}\left(u, v\right) + c$, formula (2.15) gives

$$\log \mu_{xy} + D_{xy}\left(-\log \mu_{x0}, -\log \mu_{0y}\right) = 0. \tag{2.17}$$

Note that if utility is perfectly transferable, $D_{xy}(u, v) = (u + v - \Phi_{xy})/2$ and we recover the Choo and Siow formula (2.12).

2.3 Identification in hedonic models

Our exposition of identification in hedonic models will follow the thread that runs from Tinbergen (1956) through Rosen (1974) to Ekeland, Heckman, and Nesheim (2004) and Heckman, Matzkin, and Nesheim (2010). We will also discuss the contributions of Bajari and Benkard (2005) and of Epple (1987) and Epple and Sieg (1999).

We will assume throughout that all characteristics of the products are observed, so that the full type z and the observed type z coincide. We will return to this assumption in Section 2.3.4. On the other hand, agent full types have both observed and unobserved components, which we denote $\tilde{x} = (x, \varepsilon)$ and $\tilde{y} = (y, \eta)$.

2.3.1 The identification problem: Rosen's approach and its shortcomings

Rosen's pioneering paper illustrated the essence of the identification problem: since different agents trade different products, the slope of the price function P for a given product z reflects the marginal utilities of the agents who trade that product. Only in very special cases will it give the information necessary to identify the underlying utility functions.

This is easily seen in the version of the quadratic model of Example 4 that we solved in Section 1.2, where the hedonic equilibrium price function was also quadratic:

$$P(z) = k + lz + mz^2/2.$$

Given an infinite sample, we easily identify the function P and recover the coefficients k, l, and m. The first-order condition of buyer's ε problem yields

$$P'(z) = l + mz = Bz + \varepsilon. \tag{2.18}$$

This can be seen as a regression of the known (or identified) $P'(z_i)$ on z_i, where i ranges over buyers. Rosen (1974) initially proposed to jointly estimate two such regressions (one on each side of the market). As pointed out by a number of papers from Brown and Rosen (1982) to Ekeland, Heckman, and Nesheim (2004), this approach has serious shortcomings. Since buyer i selects their preferred product z_i on the basis of their unobserved characteristics ε_i, the regressor and the error term are correlated [5]. Over the last twenty years, a series of papers have revived this issue and proposed several strategies to tackle it.

2.3.2 Rank-based Identification

One strategy, advocated particular by Heckman, Matzkin, and Nesheim (2010), relies on *quantile identification*. In our quadratic example, the percentile of z in its distribution coincides with the percentiles of both the buyer's and seller's unobserved taste ε and η. Under suitable assumptions, this extends to a broader classes of models, after conditioning on observed types x and y.

Heckman, Matzkin, and Nesheim (2010) assume that the unobserved variation in agent types is scalar; independently distributed of the observed type; and satisfies a single-crossing condition. More precisely, assume for simplicity that the set of products $\mathcal{Z}$ is a closed interval of the real line $[\underline{z}, \overline{z}]$.

We focus on the demand side, where the buyer utility is quasilinear: we denote it $U(x, \varepsilon, z) - p$. We assume that the unobserved type ε is a scalar distributed independently of x, and that the marginal utility of the product characteristic z is an increasing function of ε. Then the choice probability

[5] In the example of Section 1.2, the correlation is positive and induces a positive bias.

$\pi^D(\cdot|\tilde{x}; P)$ is a Dirac mass on a point $z^D(x, \varepsilon)$ such that

$$\partial_z U(\tilde{x}, z) = P'(z);$$

$z^D(x, \varepsilon)$ is an increasing function of ε, and the aggregate demand has $q^D(B) = \Pr(z^D(\tilde{x}) \in B)$ for any measurable subset of $\mathcal{Z}$.

The data identifies the price function P and the joint distribution of product and buyer types (z, x). Since we know that $z^D(x, \varepsilon)$ is increasing in ε, and ε is independent of x, the usual quantile inversion argument allows us to write that for any scalar t,

$$\Pr(z \leq t|x) = \Pr(\varepsilon \leq D^1(x, t))$$

where $D^1(x, \cdot)$ is the inverse of the increasing function $\varepsilon \to z^D(x, \varepsilon)$. While the left-hand side of this equation is directly identified from the data, its right-hand side depends on two unknown quantities: the distribution of ε and the function D^1. This was to be expected: in the absence of any functional or distributional restriction, we need to impose a normalization. We could let ε be uniformly distributed over $[0, 1]$, for instance, which allows us to interpret it as a quantile. Then we obtain

$$D^1(x, z) = F_{Z|X}(z|x),$$

and by direct inversion the demand z^D. A similar argument allows us to recover the supply z^S.

For many applications, and in particular when we seek welfare evaluations, we will want to identify the utility functions themselves. Unfortunately, knowing the price function and the demand and supply functions is not enough. Consider the first-order condition for buyer $\tilde{x}$:

$$\partial_z U(\tilde{x}, z^D(\tilde{x})) = P'(z^D(\tilde{x})).$$

The marginal price function P' is known, and we just identified the function z^D. But this only identifies the marginal utility $\partial_z U$ on the graph of the demand function z^D: it gives us no information[6] on $\partial_z U(\tilde{x}, z)$ if $z \neq z^D(\tilde{x})$. The only way to extrapolate its values is to restrict the form of the function U somehow.

Imposing more structure on U may also allow to relax the assumption that the distribution of ε is known. In this vein, Ekeland, Heckman, and Nesheim (2004) make further assumptions on the form U in z and require continuous observable characteristics x. Let buyer $\tilde{x}$'s utility function be $za(x) + \varepsilon z - B(z) - p$. The first-order conditions then become

$$P'(z) = a(x) + \varepsilon - B'(z),$$

which we can rewrite as $d(z) - a(x) = \varepsilon$ with $d \equiv P' + B'$. If the function d is increasing in z, Ekeland, Heckman, and Nesheim (2004) note that

$$F_{Z|X}(z|x) = \Pr(\varepsilon \leq d(z) - a(x)) = F_\varepsilon(d(z) - a(x)).$$

[6] We did impose that U_z increase with ε for given x; this is very coarse information, however.

Taking the derivatives with respect to x and z gives

$$\partial_x F_{Z|X}(z|x) = -f_\varepsilon(d(z) - a(x))\, a'(x)$$
$$\partial_z F_{Z|X}(z|x) = f_\varepsilon(d(z) - a(x))\, d'(z).$$

As a result, the ratio $\partial_z F_{Z|X}(z|x)/\partial_x F_{Z|X}(z|x) = -d'(z)/a'(x)$. It is easy to see that this identifies the functions d and a up to location parameters and to a common scale parameter; since P' is directly identified, this also identifies the function B'. This strategy obviously cannot be applied to discrete-valued observed types x, however.

2.3.3 The optimal transportation approach

Chernozhukov, Galichon, Henry, and Pass (2021) extend Heckman, Matzkin, and Nesheim (2010)'s identification strategy by leveraging the tight connection between hedonic models and the optimal transportation problem. We focus here on the class of hedonic markets described in paragraph 1.2, where agents are endowed with quasilinear utilities and unobserved heterogeneity is additively separable:

$$U(\tilde{x}, z; P) = \bar{U}(x, z) - P(z) + \zeta(\tilde{x}, z) \;\text{ and }\; V(\tilde{y}, z; P) = \bar{V}(y, z) - P(z) + \xi(\tilde{y}, z).$$

As in previous sections, we assume that the joint distribution $F_{X,Z}$ of (x, z) is identified from the data. We adopt two different approaches, depending on whether prices are observed or not.

2.3.3.1 Observed Prices

If the price function $P(\cdot)$ is identified from the data, we can identify preferences of buyers and of sellers separately, for any given distribution of the unobserved heterogeneity.

To see this, take the demand side for instance. The Daly-Williams-Zachary theorem tells us that the derivatives of the Emax function coincide with the choice probabilities. In the notation of Section 1.1,

$$\frac{\partial G_x}{\partial \bar{U}(x,z)} = \pi^D(z|x;P).$$

Fix x; these equations are the first-order conditions of the problem

$$\min_{U^x} \left(\int \max_{z \in \mathcal{Z}} \left\{ U^x(z) - P(z) + \zeta(\tilde{x}, z) \right\} \tilde{n}(d\tilde{x}|x)) - \int_{\mathcal{Z}} U^x(z) F_{Z|X}(dz|x) \right) \tag{2.19}$$

where the function U^x gives $\bar{U}(x,z) = U^x(z)$.

If we know (or assume we know) *exactly* the function $\zeta(x,\cdot,\cdot)$ and the conditional distribution $\tilde{n}(\cdot|x)$ at a type x, then minimizing (2.19) directly gives us the utility function U^x, up to a location parameter $c(x)$. Any unknown parameters in $\zeta(x,\cdot,\cdot)$ or $\tilde{n}(\cdot|x)$ can only be identified if further restrictions are imposed on the preferences $\bar{U}$. For a simple illustration, let us consider a supermodular example:

Example 10. *We follow here Matzkin (2003)'s identification strategy. Suppose that both z and ε are scalar, with continuous distributions $F_{Z|X}$ and $F_{\varepsilon|X}$ on $\mathbb{R}$, and that for all x, $(\varepsilon, z) \to \zeta(x, \varepsilon, z)$ is supermodular. The preferred choice of $\tilde{x} = (x, \varepsilon)$ is the value of z that maximizes $\bar{U}(x,z) - P(z) +$*

$\zeta(x,\varepsilon,z)$. *Since* ζ *is supermodular, this* z *is an increasing function* $z^D(x,\varepsilon)$ *of* ε*; it solves*

$$\frac{\partial \bar{U}}{\partial z}(x,z) = P'(z) - \frac{\partial \zeta}{\partial z}(x,\varepsilon,z). \tag{2.20}$$

Since $\varepsilon \to z^D(x,\varepsilon)$ *is increasing, it has an inverse which we denote* $\varepsilon = D^1(x,z)$. *Moreover, we must have* $F_{\varepsilon|X}(D^1(x,z)|x) = F_{Z|X}(z|x)$. *Substituting into* (2.20)*, we get*

$$\frac{\partial \bar{U}}{\partial z}(x,z) = P'(z) - \frac{\partial \zeta}{\partial z}(x, F^{-1}_{\varepsilon|X}(F_{Z|X}(z|x)|x), z);$$

this integrates into

$$\bar{U}\,(x,z) = c\,(x) + P(z) - \int_0^z \frac{\partial \zeta}{\partial z}\left(F^{-1}_{\varepsilon|X}\left(x, F_{Z|X}\,(t|x)\,|x\right), t\right) dt,$$

where $c\,(x)$ *is an arbitrary number.*

Even in this scalar, supermodular case, we need to make additional assumptions on the function ζ. If for instance we impose that unobserved heterogeneity in preferences only shift the marginal utility of z, we could let $\zeta(\varepsilon,z) \equiv \varepsilon z$. This would give us

$$U\,(x,z) = c\,(x) - \int_0^z F^{-1}_{\varepsilon|X}\left(F_{Z|X}\,(t|x)\,|x\right) dt;$$

further assumptions on $F_{\varepsilon|X}$ would add identifying power.

Alternatively, we could start from (2.19) and specify a parametric form. Let β collect the unknown parameters of $\bar{U}$, ζ, and $\tilde{n}(\cdot|x)$. Then we would identify β as the minimizer (or the set of minimizers) of the following expression:

$$\int_{\mathcal{X}} n(dx) \left(\int_{\tilde{\mathcal{X}}} u^{\beta}(\tilde{x}) \tilde{n}^{\beta}(d\tilde{x}|X=x) - \int_{\mathcal{Z}} \bar{U}^{\beta}(x,z) F_{Z|X=x}(dz|x) \right),$$

where $u^{\beta}(\tilde{x}) = \max_{z\in\mathcal{Z}} \left(\bar{U}^{\beta}(x,z) - P(z) + \zeta(\tilde{x},z) \right)$ if $\tilde{x}$ has observed type x.

2.3.3.2 Unobserved Prices

Let us now turn to the more interesting case when the data does not have any information on prices. For simplicity, we assume that the sets of observed types and products $\mathcal{X}, \mathcal{Y}$, and $\mathcal{Z}$ are finite and we adapt the notation accordingly. We denote $\mathcal{Z}_0 = \mathcal{Z} \cup \{0\}$ the set of products augmented by the exit option, for which $P(0) = \bar{U}(x,0) = \bar{V}(y,0) = 0$. We assume that the econometrician observes the quantities purchased and sold $\mu(z|x)$ and $\mu(z|y)$ for all x, y, and z.

The Emax operators of a buyer of type x and a seller of type y are respectively

$$\begin{aligned} G_x\left(\bar{\boldsymbol{U}}\right) &= \mathbb{E}_{\mathbb{P}_x} \max_{z\in\mathcal{Z}_0} \left(\bar{U}_{xz} - P_z + \varepsilon_z^x\right) \\ H_y\left(\bar{V}\right) &= \mathbb{E}_{\mathbb{Q}_y} \max_{z\in\mathcal{Z}_0} \left(\bar{V}_{yz} + P_z + \eta_z^y\right) \end{aligned}$$

Using Legendre-Fenchel transforms and duality, we get

$$\begin{aligned} \bar{U}_{xz} - P_z &= \frac{\partial G_x^*}{\partial \mu_{z|x}} \left(\mu_{\cdot|x}\right) \\ \bar{V}_{yz} + P_z &= \frac{\partial H_y^*}{\partial \mu_{z|y}} \left(\mu_{\cdot|y}\right), \end{aligned}$$

and we can eliminate the unobserved prices to obtain an analog of the non-parametric identification formula of Section 2.9:

$$\bar{U}_{xz} + \bar{V}_{yz} = \frac{\partial G_x^*}{\partial \mu_{z|x}} \left(\mu_{.|x} \right) + \frac{\partial H_y^*}{\partial \mu_{z|y}} \left(\mu_{.|y} \right).$$

This shows that even though the prices are not observed, we can identify $\bar{U}_{xz}$ and $\bar{V}_{yz}$ up to a common z fixed effect, provided that we know (or assume) the distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$ that determine the shape of the functions G_x^* and H_y^*.

2.3.4 Dealing with Unobserved Product Characteristics

We assumed so far that the econometrician observes all characteristics of the products that are payoff-relevant for the agents. This is an unpalatable assumption: as mentioned in Section 2.1.5, the empirical industrial organization literature, for instance, gives a large role to unobserved "product effects". These terms can be seen as the utility value of product characteristics that the agents observe but the econometrician does not. Let us now assume that the product characteristics have two components: $\tilde{z} = (z, \xi)$, where the characteristics z are in the data but the unobserved characteristics ξ are not. We denote $U(\tilde{x}, \tilde{z}, p)$ the utility of a consumer $\tilde{x}$ who buys a variety $\tilde{z}$ at price p. The equilibrium price function now must depend on both components of $\tilde{z}$: we have $p = P(z, \xi)$. Since ξ is unobserved, the econometrician can only estimate the distribution of $P(z, \xi)$ conditional on z.

While identifying the U preferences from the data may seem hopeless, Bajari and Benkard (2005) show three sets of assumptions under which it can be done. The common underlying idea is to combine monotonicity and independence assumptions to identify the value of ξ for any product whose characteristics z are observed, essentially reducing the identification problem to the case when all product characteristics are observed.

To see how this works, assume that ξ is a scalar characteristic and that the utility U is increasing in ξ. Now consider the consumers who buy a product with observed characteristics z. Since U decreases with the price p, none of these consumers will buy (z, ξ_1) in preference to (z, ξ_2) if $\xi_1 < \xi_2$ and $P(z, \xi_1) \geq P(z, \xi_2)$. Therefore for given z, the equilibrium price function $\xi \mapsto P(z, \xi)$ must be increasing.

The next step is to see that if we assume moreover that ξ is distributed independently of z, then the conditional quantiles of the price identify ξ. The argument is very similar to that in Section 2.3.2. Denote $P^1(z, p)$ the inverse function of $\xi \to P(z, \xi) = p$. We have

$$\begin{aligned} F_{p|Z}(p|z) &= \Pr(P(z,\xi) \leq p | Z = z) \\ &= \Pr(\xi \leq P^1(z,p)) \\ &= P^1(z,p) \end{aligned}$$

where the second line uses independence and the third line normalizes the distribution of ξ to be $U[0, 1]$, again without loss of generality.

Since the distribution of price conditional on z is observed, we identify the function $P^1(z, p)$ and therefore its inverse P (at least on the range of variation of these variables). Now take any product $\tilde{z} = (z, \xi)$, with its price

$P(z, \xi)$. By the above equations, we have

$$F_{p|Z}(P(\tilde{z})|z) = P^1(z, P(\tilde{z})) = \xi$$

so that we recover the unobserved characteristics of this product.

Bajari and Benkard (2005) also discuss two variants of this idea. The first one uses the price variation between different "option packages" on the same product. The second one relaxes the independence of observed and unobserved characteristics by introducing instruments.

Chapter 3

Inference

We now turn to estimation and testing. This section follows the structure of the previous ones: we start with individual demand; then we move to matching models, and finally to hedonic equilibrium. Minimum-distance estimation will be a recurrent theme, as it leads to practical estimation method that do not require solving for the stable matching or equilibrium. We also present a variety of other approaches to inference, however.

3.1 Inference on Demand

Let us first consider individual demand in the very common case when the utility function is assumed to be quasi-linear and separable:

$$U(\tilde{x}, z, p) = \bar{U}(x, z) - p + \zeta(\tilde{x}, z).$$

We also assume that z only takes J values. In the language of the econometrics of discrete choice, we have a multinomial Additive Random Utility Model, that is quasi-linear in price. There are of course many well-known methods to estimate such models. Rather than reviewing them, we will use this opportunity to introduce the minimum distance estimator that plays a central role in this section.

Suppose that a parameter vector $\boldsymbol{\theta}$ is used to index both the preferences $\bar{U}^{\boldsymbol{\theta}}(x,z)$ and the distribution $\mathbb{P}_x^{\boldsymbol{\theta}}$ of the vector $\boldsymbol{\zeta}^x = \zeta(\tilde{x},\cdot)$ conditional on x. Then the model is fully parametric. Let us denote, as in Section 1.1, the Emax function by G_x, its Legendre-Fenchel transform by G_x^*; $\boldsymbol{\mu}^x = \mu(\cdot|x)$ the choice probabilities of buyers with observed type x; and $\boldsymbol{U}^{\boldsymbol{\theta}}(x;P)$ the vector of the values of $\bar{U}^{\boldsymbol{\theta}}(x,z) - P(z)$ for all $z \in \mathcal{Z}$. Then the function G_x^* depends on $\boldsymbol{\theta}$ and the identification condition (2.4) becomes:

$$\boldsymbol{U}^{\boldsymbol{\theta}}(x;P) \in \partial G_x^*(\boldsymbol{\mu}^x;\boldsymbol{\theta}).$$

If the model is well-specified, the true parameter vector $\boldsymbol{\theta}_0$ satisfies this inclusion for all values of x; and any parameter vector θ that does is consistent with the data.

In the simplest (and very commonly used) case when each G_x^* is strictly convex and differentiable, all inclusions are equalities:

$$\boldsymbol{d}_x(\boldsymbol{\mu};\boldsymbol{\theta}) \equiv \boldsymbol{U}^{\boldsymbol{\theta}}(x;P) - \nabla_{\boldsymbol{\mu}} G_x^*(\boldsymbol{\mu}^x;\boldsymbol{\theta}) = \mathbf{0}.$$

The identified set Θ^I can be estimated by setting up a minimum distance estimation problem based on the following *mixed hypothesis*:

$$\text{there exists } \boldsymbol{\theta} \text{ such that for all } x,\ \boldsymbol{d}_x(\boldsymbol{\mu}^x;\boldsymbol{\theta}) = \boldsymbol{0}. \tag{3.1}$$

Any consistent estimator $\hat{\boldsymbol{\mu}}$ of $\boldsymbol{\mu}$ can then be used to obtain a consistent estimator of $\boldsymbol{\theta}$.

Suppose for simplicity that $\mathcal{X}$ is finite (or take a finite subset of X values of x) and stack the vectors $(\boldsymbol{d}_x(\hat{\boldsymbol{\mu}}^x;\boldsymbol{\theta}))_{x\in\mathcal{X}}$ into a vector $\boldsymbol{D}(\hat{\boldsymbol{\mu}};\boldsymbol{\theta}) \in \mathbb{R}^{J\times X}$. Then for any positive definite matrix $\boldsymbol{W}$, solving

$$\min_{\boldsymbol{\theta}} \boldsymbol{D}(\hat{\boldsymbol{\mu}};\boldsymbol{\theta})'\ \boldsymbol{W}\ \boldsymbol{D}(\hat{\boldsymbol{\mu}};\boldsymbol{\theta}) \tag{3.2}$$

gives a consistent estimator of Θ^I. More precisely, the set $\hat{\Theta}^I_n$ of values of θ that achieve the minimum in (3.2) converges to Θ^I for the Hausdorff distance as n goes to infinity.

With a finite choice set, the vector $\hat{\boldsymbol{\mu}}_n = (\hat{\mu}_{x,n})_{x\in\mathcal{X}}$ is asymptotically normal with asymptotic variance-covariance matrix $\boldsymbol{M}$:

$$\sqrt{n}\,(\hat{\boldsymbol{\mu}}_n - \boldsymbol{\mu}) \to N(\boldsymbol{0},\boldsymbol{M}),$$

with $M_{xx'} = \mu_x\left(\mathbb{1}(x = x') - \mu_{x'}\right)$.

Under point-identification, the set $\hat{\Theta}^I_n$ always consists of a single point $\hat{\boldsymbol{\theta}}_n$. General results in minimum distance estimation readily apply: the estimator $\hat{\boldsymbol{\theta}}_n$ is asymptotically normal; its asymptotic variance is minimized by choosing $\boldsymbol{W}$ to be a consistent estimator of

$$\boldsymbol{W}^* = \left(\boldsymbol{D}_{\boldsymbol{\mu}}\ \boldsymbol{M}\ \boldsymbol{D}'_{\boldsymbol{\mu}}\right)^{-1}$$

where $\boldsymbol{D}_{\boldsymbol{\mu}}$ is the Jacobian of $\boldsymbol{D}$ with respect to $\boldsymbol{\mu}$ at $(\boldsymbol{\mu}, \boldsymbol{\theta}_0)$. For this choice of $\boldsymbol{W}$,

$$\sqrt{n}\left(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0\right) \to N(\boldsymbol{0}, \boldsymbol{V})$$

with

$$\boldsymbol{V} = \left(\boldsymbol{D}_{\boldsymbol{\theta}}\ \boldsymbol{W}^*\ \boldsymbol{D}'_{\boldsymbol{\theta}}\right)^{-1}$$

with $\boldsymbol{D}_{\boldsymbol{\theta}}$ the Jacobian of $\boldsymbol{D}$ with respect to $\boldsymbol{\theta}$ at $(\boldsymbol{\mu}, \boldsymbol{\theta}_0)$. Moreover, the minimum value of the objective function (3.2) then is asymptotically distributed as a χ^2 with $J \times X - K$ degrees of freedom, where K is the dimension of $\boldsymbol{\theta}$. This gives a straightforward specification test.

3.2 Inference in Matching Models with Transferable Utility

We now turn to matching markets, where the data the econometrician observes typically consists of a list of matches along with some characteristics of the partners in each match: "who matches whom" in the jargon of this subfield. This kind of data boils down to a set of numbers, each of which counts the observed matches where the partners have given observed types. In the marriage market, the data could record the number of marriages between a fireman and a midwife. In the labor market, it could be the number of midsize furniture manufacturers that employ three high-school graduates as cabinetmakers.

Sometimes more information is available. The labor market is an important example where transfers may be observed: the data often records wages paid, and perhaps some characteristics of the tasks of the worker. Proxy measures for the "outcomes" of a match may also be recorded. One can think of student grades in a teacher-students match, or worker performance. Section 3.2.3 will explain how such additional information, when available, can help identify and estimate the model.

As in Sections 1.3 and 2.2, most of this section will focus on the one-to-one, bipartite model—the heterosexual marriage market of Becker (1973). We will also explain in Section 3.2.3 how some of the results extend to other matching markets with transferable utility. We assume throughout that utility is perfectly transferable: at no cost, without limits, one for one. Section 3.2.4 discusses imperfectly transferable utility.

Our discussion of identification in TU matching models will emphasize the class of "separable" models, from Choo and Siow (2006) to Chiappori, Salanié, and Weiss (2017) and our own work: Galichon and Salanié (2017, 2022, 2024). We will also describe the class of "index" models pioneered by Chiappori, Oreffice, and Quintana-Domeque (2012). Finally, we discuss the maximum-score approach of Fox (2010, 2018) and his coauthors, with its recent developments that link it to theoretical results by Azevedo and Hatfield (2018).

Estimation first requires taking a stand on sampling variation. The data has two sources of sampling variation. The first one is reflected in estimates

$\hat{\boldsymbol{\mu}}$ of the matching patterns $\boldsymbol{\mu}$, and of the margins $\boldsymbol{n}$ and $\boldsymbol{m}$. The second one relates to the evaluation of the Emax functions.

Let us first consider the sampling unit. It can be the individual, or it can be the match. To pursue the example of the marriage market, the data may come from a survey of individuals, or from a survey of households. The probability of observing a couple in cell (x, y), for instance, depends both on the model for the population, as described in the previous sections, and on the sampling scheme. This has consequences for the likelihood function of the observations. Sampling schemes always matter, of course; still, this general point needs to be paid even more careful attention when working on matching data.

As for the Emax functions, remember that when we defined G_x and H_y, we took expectations with respect to the distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$ of the unobservables. Stability conditions really refer to the empirical distributions $\tilde{\mathbb{P}}_x$ and $\tilde{\mathbb{Q}}_y$, as perceived by the agents. There are two difficulties here. First, the population comprises some potential partners who are not in the sample. Second, each individual may only be aware of a limited number of potential partners. In principle, we should account for these two issues when applying the identification formulæ. This is hard to do in practice (see Menzel (2015) for a proposed solution in a non-transferable utility model of matching). If agents match within a large population, one may want to appeal to a large markets assumption and neglect this source of variability. This has been the dominant approach in the literature.

To simplify exposition, we assume here that sampling was done at match level and that markets are large enough that we can neglect the sampling variation in $\mathbb{P}_x$ and $\mathbb{Q}_y$. We observe a sample of N households, and some characteristics of both partners.

3.2.1 Nonparametric Estimation of the Surplus Function

Assume for now that observed types x and y are discrete. From the data, we obtain the number $\hat{\mu}_{xy}$ of households that are matches between observed types x and y; $\hat{\mu}_{x0}$, the number of single men of observed type x, and $\hat{\mu}_{0y}$, the number of single women of observed type y. This results in *margins*

$$\hat{n}_x = \sum_y \hat{\mu}_{xy} + \hat{\mu}_{x0}$$
$$\hat{m}_y = \sum_x \hat{\mu}_{xy} + \hat{\mu}_{0y}.$$

The distributions of these estimators of matching patterns and of the margins are given by the standard formulæ of discrete choice problems. In the simplest case, if the N households are drawn with equal probabilities from an infinite population characterized by true matching patterns $\boldsymbol{\mu}$, then for $(x, y) \in \mathcal{A}$ and $(z, t) \in \mathcal{A}$[1]:

$$\text{cov}(\hat{\mu}_{xy}, \hat{\mu}_{zt}) = \frac{1}{N}\mu_{xy}(\mathbb{1}(x = z, y = t) - \mu_{zt}).$$

If the distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$ are parameter-free and have continuous and full support, one can use the empirical analog of (2.9) directly to obtain

[1] Recall that $\mathcal{A} = (\mathcal{X} \times \mathcal{Y}_0) \cup (\mathcal{X}_0 \times \mathcal{Y})$.

an estimate of the surplus matrix:

$$\hat{\Phi}_{xy} = -\frac{\partial \mathcal{E}}{\partial \mu_{xy}}(\hat{\boldsymbol{\mu}}; \hat{\boldsymbol{n}}, \hat{\boldsymbol{m}}) = -\frac{\partial G_x^*}{\partial \mu_{y|x}}(\hat{\boldsymbol{\mu}}_{\cdot|x}) - \frac{\partial H_y^*}{\partial \mu_{x|y}}(\hat{\boldsymbol{\mu}}_{\cdot|y}).$$

This assumes of course that the generalized entropy $\mathcal{E}$ corresponding to these error distributions can be computed, at least numerically.

The distribution of the resulting estimator follows directly from that of the estimated matching patterns $\hat{\boldsymbol{\mu}}$. Like them, it is $\sqrt{N}$-consistent and asymptotically normal. The logit model is a leading example; Choo and Siow (2006) used (2.12) to estimate the surplus function. Note that the term "nonparametric" only applies here to the matrix $\mathbf{\Phi}$, as the distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$ are imposed by the analyst.

3.2.2 Parametric Estimation of the Surplus Function

If the error distributions are not fully-specified, then the model is underidentified unless restrictions are imposed on the specification of the surplus function $\mathbf{\Phi}$ and/or more data is used (see Section 3.2.3 for the latter). We maintain our assumption that observed types are discrete; and we specify a fully parametric model $(\mathbf{\Phi}^{\boldsymbol{\theta}}, (\mathbb{P}_x^{\boldsymbol{\theta}}), (\mathbb{Q}_y^{\boldsymbol{\theta}}))$, with a true parameter vector $\boldsymbol{\theta}_0$. In typical applications, the dimension of $\boldsymbol{\theta}$ is much smaller than $X \times Y$, and the model is identified and testable. We will assume that to be the case.

For any given value of $\boldsymbol{\theta}$, we can use one of the algorithms described in Section 4.2 to solve the matching problem for the equilibrium values $\boldsymbol{\mu}^{\boldsymbol{\theta}}$ of the matching patterns given the surplus function $\mathbf{\Phi}^{\boldsymbol{\theta}}$, the error distributions $\mathbb{P}_x^{\boldsymbol{\theta}}$ and $\mathbb{Q}_y^{\boldsymbol{\theta}}$, and the observed margin distributions $\hat{n}_x$ and $\hat{m}_y$.

Note that this results in a number of households

$$N^{\boldsymbol{\theta}} = \sum_{x=1}^{X} \sum_{y=1}^{Y} \boldsymbol{\mu}_{xy}^{\boldsymbol{\theta}} + \sum_{x=1}^{X} \boldsymbol{\mu}_{x0}^{\boldsymbol{\theta}} + \sum_{y=1}^{Y} \boldsymbol{\mu}_{0y}^{\boldsymbol{\theta}}$$

that will in general differ from the observed number of households N, as $\boldsymbol{\theta}$ predicts more or fewer matches than are observed in the sample.

3.2.2.1 Maximum Likelihood Estimation

The most obvious way to estimate a parametric separable matching model is maximum likelihood. Under our assumptions the likelihood function of the sample is

$$\log L(\boldsymbol{\theta}) = \sum_{x=1}^{X} \sum_{y=1}^{Y} \hat{\mu}_{xy} \log \frac{\mu_{xy}^{\boldsymbol{\theta}}}{N^{\boldsymbol{\theta}}} + \sum_{x=1}^{X} \hat{\mu}_{x0} \log \frac{\mu_{x0}^{\boldsymbol{\theta}}}{N^{\boldsymbol{\theta}}} + \sum_{y=1}^{Y} \hat{\mu}_{0y} \log \frac{\mu_{0y}^{\boldsymbol{\theta}}}{N^{\boldsymbol{\theta}}}.$$

The maximum likelihood estimator $\hat{\boldsymbol{\theta}}^{MLE}$ given by the maximization of $\log L$ is consistent, asymptotically normal, and asymptotically efficient under the usual set of assumptions.

3.2.2.2 Moment Matching

Instead of maximizing the likelihood of a fully specified model, the analyst may want to only impose some moment conditions. This approach is particularly attractive when the model is *semilinear* :

Definition 6 (Semilinear Model)**.** *The parameter vector $\boldsymbol{\theta}$ consists of $(\boldsymbol{\alpha}, \boldsymbol{\beta})$. The parameters in $\boldsymbol{\alpha}$ only affect the distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$; those in $\boldsymbol{\beta}$ only affect the joint surplus, which consists of a linear expansion over a family* *of K known basis functions $\boldsymbol{\phi}$: for all (x, y),*

$$\Phi^{\boldsymbol{\beta}}_{xy} = \sum_{k=1}^{K} \beta_k \phi_{xy,k}.$$

This *semilinear* case is of special interest as it opens the door to very simple estimation methods. In particular, it seems natural to want to match the observed *comoments* $\hat{C}_k = \sum_{x=1}^{X} \sum_{y=1}^{Y} \hat{\mu}_{xy} \phi_{xy,k}$ with their simulated counterparts:

$$C^{\boldsymbol{\theta}}_k = \sum_{x=1}^{X} \sum_{y=1}^{Y} \mu^{\boldsymbol{\theta}}_{xy} \phi_{xy,k} \quad \text{for } k = 1, \ldots, K.$$

As an illustration, suppose that the k-th basis function is a dummy variable that equals one when $x = y$: $\phi_{xy,k} = \mathbb{1}(x = y)$. Then the k-th comoment counts the number of couples where the observed types of the man and the woman coincide. By construction, equating $\boldsymbol{C}^{\boldsymbol{\theta}}$ to $\hat{\boldsymbol{C}}$ gives as many estimating equations as there are parameters in $\boldsymbol{\beta}$. It can be used when the error distributions are parameter-free, or for fixed values of their parameters $\boldsymbol{\alpha}$. In the latter case, we obtain estimators $\hat{\boldsymbol{\beta}} = \hat{\boldsymbol{B}}(\boldsymbol{\alpha})$, and one can (for instance) maximize the profile likelihood over $(\hat{\boldsymbol{B}}(\boldsymbol{\alpha}), \boldsymbol{\alpha})$.

Galichon and Salanié (2022) show that for fixed $\boldsymbol{\alpha}$, the estimator $\hat{\boldsymbol{B}}(\boldsymbol{\alpha})$ can be obtained by maximizing the function

$$\sum_{k=1}^{K} \beta_k \hat{C}_k - \mathcal{W}^{\boldsymbol{\alpha}}(\boldsymbol{\Phi}^{\boldsymbol{\beta}}; \hat{\boldsymbol{n}}, \hat{\boldsymbol{m}})$$

where $\mathcal{W}^{\boldsymbol{\alpha}}$ is the social welfare defined in (2.11) when the entropy is computed using the distributions $\mathbb{P}_x^{\boldsymbol{\alpha}}$ and $\mathbb{Q}_y^{\boldsymbol{\alpha}}$.

Since $\mathcal{W}^{\boldsymbol{\alpha}}$ is convex in $\boldsymbol{\Phi}$ and $\boldsymbol{\Phi}^{\boldsymbol{\beta}}$ is linear in $\boldsymbol{\beta}$, this function is globally concave in $\boldsymbol{\beta}$ and its maximizer can be obtained very easily. Using (2.11), the first-order conditions at the maximum are

$$\hat{C}_k = \sum_{x=1}^{X}\sum_{y=1}^{Y} \frac{\partial \mathcal{W}^{\boldsymbol{\alpha}}}{\partial \Phi_{xy}} \frac{\partial \Phi_{xy}^{\boldsymbol{\beta}}}{\partial \beta} = \sum_{x=1}^{X}\sum_{y=1}^{Y} \mu_{xy}^{\boldsymbol{\theta}} \phi_{xy,k}$$

which indeed match observed and simulated comoments.

While the moment matching estimator is not efficient in general, it is consistent and inference can be conducted on its parameters in the standard way.

3.2.2.3 Minimum-distance Estimation

Both maximum likelihood and moment matching require solving for the equilibrium matching patterns $\boldsymbol{\mu}$ (or evaluating the social welfare $\mathcal{W}$) for many values of the parameter vector $\boldsymbol{\theta}$. We will describe in Section 4.2 our IPFP algorithm, which provides a fast way of doing so for many models. Still, one may want to circumvent solving for equilibrium during the estimation process. We present here a minimum-distance method that achieves this goal.

Given a parametric specification, our Equation (2.9) can be rewritten as

$$\boldsymbol{\Phi}^{\boldsymbol{\theta}_0} + \frac{\partial \mathcal{E}^{\boldsymbol{\theta}_0}}{\partial \boldsymbol{\mu}}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}) = \boldsymbol{0},$$

where $\mathcal{E}^{\boldsymbol{\theta}}$ is the generalized entropy function that is generated by the distributions $\mathbb{P}_x^{\boldsymbol{\theta}}$ and $\mathbb{Q}_y^{\boldsymbol{\theta}}$. This suggests a minimum-distance estimator $\hat{\boldsymbol{\theta}}$ that is obtained by minimizing a norm of the vector

$$\boldsymbol{D}(\boldsymbol{\theta}) \equiv \boldsymbol{\Phi}^{\boldsymbol{\theta}} + \frac{\partial \mathcal{E}^{\boldsymbol{\theta}}}{\partial \boldsymbol{\mu}}(\hat{\boldsymbol{\mu}}; \hat{\boldsymbol{n}}, \hat{\boldsymbol{m}}) :$$

we only need to choose a positive-definite matrix $\boldsymbol{S}$ and to minimize $\boldsymbol{D}(\boldsymbol{\theta})'\boldsymbol{S}\boldsymbol{D}(\boldsymbol{\theta})$.

If the model is identified, this estimator is consistent for any choice of $\boldsymbol{S}$. As in Section 3.1, there is a choice of the weighting matrix $\boldsymbol{S}$ that minimizes the asymptotic variance of the estimator. Here it is

$$\boldsymbol{S}^* = \left(\boldsymbol{H}\boldsymbol{M}\boldsymbol{H}'\right)^{-1},$$

where $\boldsymbol{M}$ is the asymptotic variance-covariance matrix of $\hat{\boldsymbol{\mu}}$ and $\boldsymbol{H}$ is the Hessian of the generalized entropy $\mathcal{E}^{\boldsymbol{\theta}_0}$ at $\boldsymbol{\mu}$. If the efficient weighting matrix $\boldsymbol{S}^*$ is chosen then the value of the objective function at the optimum $\boldsymbol{D}(\boldsymbol{\theta})'\boldsymbol{S}^*\boldsymbol{D}(\boldsymbol{\theta})$ can be used as a specification test: under the null hypothesis of correct specification, it is distributed as a χ^2 with $(X \times Y - \dim \boldsymbol{\theta})$ degrees of freedom.

The minimum-distance estimator takes a remarkably simple form if the model is semilinear[2] and the gradient of the generalized entropy $\mathcal{E}^{\boldsymbol{\alpha}}$ is linear-affine in the parameters $\boldsymbol{\alpha}$:

$$\frac{\partial \mathcal{E}^{\boldsymbol{\alpha}}}{\partial \boldsymbol{\mu}}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}) = \boldsymbol{e}_0(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}) + \sum_{l=1}^{L} \boldsymbol{e}_l(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m})\alpha_l.$$

Then the efficient minimum-distance estimator can be obtained by quasi-

[2]See Definition 6.

generalized least squares. In a first step, any weighting matrix $\boldsymbol{S}$ can be used; GLS gives a consistent estimator $\tilde{\boldsymbol{\theta}}$. This allows us to compute the efficient weighting matrix $\boldsymbol{S}^*$, which is used in a second-step GLS. If moreover the gradient of the generalized entropy is parameter-free (so that $L = 0$ above), the efficient weighting matrix $\boldsymbol{S}^*$ is simply the inverse of the variance of $\boldsymbol{e}_0(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m})$, which does not depend on $\boldsymbol{\theta}$. The minimum-distance estimator then is a one-step GLS estimator.

The semilinear/linear case is not as rare as one might think: the logit model is parameter-free, and both the heteroskedastic logit model and the nested logit model generate generalized entropy functions that are linear in the parameters. We refer the reader to Galichon and Salanié (2024) for the formulæ, and to Salanié (2023) for their Python implementation.

The minimum-distance estimator has two minor drawbacks. To derive the efficient weighting matrix $\boldsymbol{S}^*$, it is necessary to compute the Hessian of the generalized entropy. This is often a very sparse matrix, and it can be done numerically[3]. Moreover, the estimator relies on (2.9), which only holds when $\hat{\boldsymbol{\mu}} \gg \boldsymbol{0}$. Galichon and Salanié (2024) discuss several solutions.

3.2.2.4 Estimating the Semilinear Logit Model

Since the semilinear logit model has proved to be so popular, we present here another estimator that circumvents solving for equilibrium. We prove in Galichon and Salanié (2024) that in this model, the moment-matching estimator is formally identical to the maximum-likelihood estimator of a Poisson regression with two-way fixed effects.

[3]Note that it only needs to be done once.

To see this, we define an augmented parameter vector $\boldsymbol{\beta} = (\boldsymbol{\theta}, \boldsymbol{a}, \boldsymbol{b})$, where $\boldsymbol{a}$ has X components and $\boldsymbol{b}$ has Y. We stack the observed matching patterns $\hat{\boldsymbol{\mu}}$ into a vector of $X \times Y$ elements in row-major order; and we add the X values μ_{x0} and the Y values μ_{0y} to form a vector of $(X \times Y + X + Y)$ elements. For each basis function $k = 1, \ldots, K$, we stack the matrix $\boldsymbol{\phi}_{\cdot\cdot,k}$ in a vector in row-major order, and we create an $(X \times Y, K)$ matrix $\bar{\boldsymbol{\phi}}$ from these K vectors. Using row-major order again, define a vector $\boldsymbol{w}$ with $(X \times Y + X + Y)$ elements and

$$w_{xy} = 2, w_{x0} = 1, w_{0y} = 1 \; \forall \; x, y.$$

Finally, let

$$\boldsymbol{\mathcal{X}} = \begin{pmatrix} \bar{\boldsymbol{\phi}}/2 & -\frac{1}{2} I_X \otimes 1_Y & -\frac{1}{2} 1_X \otimes I_Y \\ 0_{X \times K} & -I_X & 0_{Y \times X} \\ 0_{Y \times K} & 0_{X \times Y} & -I_Y \end{pmatrix}.$$

Let $(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{a}}, \tilde{\boldsymbol{b}})$ be the estimates of the coefficients of the explanatory variables $\boldsymbol{\mathcal{X}}$ in the pseudo-maximum likelihood of the Poisson regression with dependent variables $\hat{\boldsymbol{\mu}}$ and weights $\boldsymbol{w}$. Galichon and Salanié (2024) prove that

- $\tilde{\boldsymbol{\beta}}$ coincides with the moment-matching estimator $\hat{\boldsymbol{\beta}}$;

- $\hat{\boldsymbol{u}} = \tilde{\boldsymbol{a}} + \log \hat{\boldsymbol{n}}$ and $\hat{\boldsymbol{v}} = \tilde{\boldsymbol{b}} + \log \hat{\boldsymbol{m}}$ are consistent estimators of the expected utilities of the groups of men and women.

3.2.2.5 Maximum-score Methods

In a series of papers starting with Fox (2010), Fox has developed an empirical approach to matching with transferable utility that relies on a selecting set of "matching inequalities." The intuition behind it is simple when we assume away unobserved heterogeneity. Suppose for now that all matches between men and women of given observed types generate the same surplus Φ_{xy}: all ζ and ξ terms are identically zero. Recall that the observed matching maximizes the social welfare under the scarcity constraint. If we observe two matches (x, y) and (x', y'), reshuffling partners cannot increase the sum of the surpluses from these two matches:

$$\left(\Phi_{xy} + \Phi_{x'y'}\right) - \left(\Phi_{xy'} + \Phi_{x'y}\right) \geq 0.$$

To simplify the notation, we borrow from Chiappori, Salanié, and Weiss (2017) and we denote $\mathcal{D}^{xy,x'y'}(\boldsymbol{\Phi})$ the double difference above. Let $\Phi_{xy} \equiv \boldsymbol{\phi}'_{xy}\boldsymbol{\theta}$ for some set of K basis functions $\boldsymbol{\phi}$, as in the semilinear model of Definition 6. For any two observed matches i and j, define the K-dimensional vector

$$\boldsymbol{\Delta}_{ij} \equiv \mathcal{D}^{x_i y_i, x_j y_j}(\boldsymbol{\phi}).$$

Then for every $i < j$, we must have $\boldsymbol{\Delta}'_{ij}\boldsymbol{\theta} \geq 0$. These inequalities define a convex polytope, the partially identified set for $\boldsymbol{\theta}$. Under reasonable conditions, this set should be small.

Taking this intuition to a model that includes matching on unobservable types is more challenging. Any generalization of these inequalities must be stated in terms of probabilities. In a one-sided choice model, we know that the probability of choosing an alternative increases with the mean utility of this alternative. This suggests that the double-difference $\mathbf{\Delta}_{ij}$ defined above should be linked to a double-difference of increasing functions of the matching patterns. In fact, this is clearly true in the logit model of Section 2.2.2, as (2.12) implies that

$$\mathcal{D}^{xy,x'y'}(\mathbf{\Phi}) = 2\,\mathcal{D}^{xy,x'y'}(\log \boldsymbol{\mu}). \tag{3.3}$$

In particular,

$$\mathcal{D}^{xy,x'y'}(\mathbf{\Phi}) > 0 \ \text{ iff } \ \mathcal{D}^{xy,x'y'}(\log \boldsymbol{\mu}) > 0. \tag{3.4}$$

Fox (2010) called this the *rank-order property.*

Unfortunately, in general all we can say for separable models is that

$$\mathcal{D}^{xy,x'y'}(\mathbf{\Phi}) = -\mathcal{D}^{xy,x'y'}\left(\frac{\partial \mathcal{E}}{\partial \boldsymbol{\mu}},\right),$$

which can be a very complicated function of the observed matching patterns. However, Graham (2011, 2014) proved that if for each value of x (resp. for each value of y), all error terms $\zeta(\tilde{x}, y)$ (resp. all error terms $\xi(x, \tilde{y})$) are iid, then the rank-order property holds. Fox (2018) showed that Graham's result extends to more general (separable) models of matching; among other things, he weakened the iid requirement to exchangeability.

While the set of inequalities in (3.4) is less informative than (3.3), it is valid in a much larger set of models. It covers for instance the heteroskedastic logit models used in Chiappori, Salanié, and Weiss (2017). On the other hand, it does not allow for nested logit or mixed logit structures, which violate the exchangeability requirement.

Now assume a semilinear $\mathbf{\Phi}$ again and consider the scalar function

$$F(\boldsymbol{\theta}) \equiv \sum_{i,j} \mathbb{1}\left(\mathbf{\Delta}'_{ij}\boldsymbol{\theta} > 0\right)$$

where i and j range over the set of observed matches. This double sum can also be rewritten as

$$F(\boldsymbol{\theta}) = \sum_{x,y}\sum_{x',y'}\sum_{i,j} \mathbb{1}(x_i = x, y_i = y)\mathbb{1}(x_j = x', y_j = y')\mathbb{1}\left(\left(\mathcal{D}^{xy,x'y'}(\boldsymbol{\phi})\right)'\boldsymbol{\theta} > 0\right).$$

In this form, it is easy to see that the expected value of $F(\theta)$ over a sample with observed matching patterns $\hat{\mu}$ is

$$\sum_{x,y}\sum_{x',y'} \hat{\mu}_{xy}\hat{\mu}_{x'y'}\mathbb{1}\left(\left(\mathcal{D}^{xy,x'y'}(\boldsymbol{\phi})\right)'\boldsymbol{\theta} > 0\right).$$

We know from (3.4) that $\left(\mathcal{D}^{xy,x'y'}(\phi)\right)'\theta > 0$ if and only if $\mathcal{D}^{xy,x'y'}(\log \boldsymbol{\mu}^{\boldsymbol{\theta}}) > 0$, where $\boldsymbol{\mu}^{\boldsymbol{\theta}}$ are the equilibrium matching patterns for parameter values $\boldsymbol{\theta}$. Much like in Manski (1975), maximizing $F(\boldsymbol{\theta})$ gives a set-valued estimator of $\boldsymbol{\theta}$ that converges to a set that includes the true parameter value.

This approach has two benefits. First, it applies to many-to-one or many-to-many matching with little modification. Second, it allows the researcher to select the inequality conditions she wants to impose for estimation: the sums over x, y and x', y' need not include all possible values. This may allow for more robust inference. On the other hand, the maximum-score method minimizes a discontinuous function and only yields a set-valued estimator; and it only applies to models with exchangeable error distributions.

3.2.3 Extensions

The methodology described in Sections 2.2 and 3.2 can be extended in several directions.

3.2.3.1 Using Richer Data

We have assumed so far that the econometrician only observed the matching patterns, "who matches whom"; and that she only observed one large market, say the marriage market in 2020 in California.

Other important variables can sometimes be observed: they could be transfers (like wages in labor markets) or proxies for the realized joint surplus. Observing transfers at the match level brings very useful information; on the other hand, it has little identifying power if it can only be done at the (x, y) cell level (see (Salanié, 2015)). If elements of the realized joint surplus can be observed (grades in a student-teacher match, for instance), then they also can bring useful information (see (Graham, Imbens, and Ridder, 2014)).

To continue with the example of marriage in the US, most marriages involve partners who live in the same state. This suggests pooling data on all fifty states, modeled as independent marriage markets, and imposing that some elements of the specification are constant across states. The same idea can apply to different age cohorts, as in Chiappori, Salanié, and Weiss (2017). Studying changes in marriage patterns by education groups for the cohorts born between 1943 and 1972, they start with a logit model that only allows for two-way interactions between the educations and the two partners and the cohort of the husband:

$$\Phi^c_{xy} = \Phi^0_{xy} + \gamma^c_x + \delta^c_y.$$

This simple model cannot account for the observed changes in marriage patterns among white individuals; it is strongly rejected by the data. It is necessary to include a linear trend term $\Phi^1_{xy} \times c$ that represents changes in the joint surplus across cohorts for given education pairs. Note that both models are massively overidentified, as the number of parameters is much lower than the number of marriage patterns $\hat{\mu}^c_{xy}$.

Fox, Yang, and Hsu (2018) show that one can in fact use variation across many markets to identify the distributions of the error terms (in our notation, $\mathbb{P}_x$ and $\mathbb{Q}_y$). To convey some of the underlying intuition, recall that in separable models,

$$\Phi_{xy} + \frac{\partial \mathcal{E}}{\partial \mu_{xy}}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}) = 0$$

and that the generalized entropy function $\mathcal{E}$ depends on the $\mathbb{P}_x$ and $\mathbb{Q}_y$ distributions. Suppose that the matrix $\boldsymbol{\Phi}$ is known and varies across markets (as do the margins $\boldsymbol{n}$ and $\boldsymbol{m}$), while the unknown distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$ are the same in all markets. Then the shape of the observed mapping from $(\boldsymbol{\Phi}, \boldsymbol{n}, \boldsymbol{m})$ to $\boldsymbol{\mu}$ across markets is informative on the underlying error distributions. This intuition extends to cases when the matrix $\boldsymbol{\Phi}$ is unknown but its variations across markets are restricted.

Many-to-one and many-to-many matching markets also contain information that helps identify the weights of different variables in the joint surplus. To see this, consider worker-employer matches, as a stylized many-to-one market. Suppose that employers care only about the productivity of each employee, and that they observe workers' education and past job experience. Everything else equal, the range of variation of education between employees of a given firm will be greater if employers value job experience more. This is useful identifying information to an econometrician who only observes education and not job experience. This insight was used by Agarwal (2015) in non-transferable utility setting and formalized by Diamond and Agarwal (2017).

Observing transitions between matches can also be very helpful, especially if transfers or proxies for the surplus are observed. The labor market is a prominent example. In their pioneering work with matched employer-employee data, Abowd, Kramarz, and Margolis (1999) assumed that job mobility was exogenous. More recent contributions, such as Bonhomme, Lamadon, and Manresa (2019), allow for worker sorting across firms. When transfers are not observed, as in marriage markets, changes in labor supply or consumption patterns within households can serve as a proxy—see Goussé, Jacquemet, and Robin (2017).

3.2.3.2 Matching without Singles

Sometimes the data only contain realized matches: singles are not observed. The adding up constraints become

$$\sum_{y=1}^{Y} \mu_{xy} = n_x \text{ for all } x \; ; \; \sum_{x=1}^{X} \mu_{xy} = m_y \text{ for all } y,$$

and the total numbers of men and women must be equal:

$$\sum_{x=1}^{X} n_x = \sum_{y=1}^{Y} m_y.$$

Suppose that the data-generating process has a joint surplus matrix $\mathbf{\Phi}$. If we add arbitrary numbers a_x and b_y to its elements, the total joint surplus for a *feasible* matching $\boldsymbol{\mu}$ must increase by

$$\sum_{x,y} \mu_{xy}(a_x + b_y) = \sum_{x} a_x n_x + \sum_{y} b_y m_y.$$

Since this takes the same value for all feasible matchings $\boldsymbol{\mu}$, the stable matching cannot change. As a consequence, a_x and b_y cannot be identified. To put it differently, only the double differences of the $\mathbf{\Phi}$ matrix can be identified. If for instance the joint surplus is linear in the parameters as in Definition 6:

$$\Phi_{xy} = \boldsymbol{\phi}_{xy} \cdot \boldsymbol{\beta},$$

then the coefficients of those basis functions that do not interact x and y are not identified.

Another way to see this is that since

$$G_x(\boldsymbol{U}_{x\cdot}) = E \max_{y \in \mathcal{Y}} (U_{xy} + \varepsilon_y),$$

for any $\boldsymbol{\mu}_{\cdot|x}$ such that $\sum_{y=1}^{Y} \mu_{y|x} = 1$ the value of $\left(\sum_{y=1}^{y} \mu_{y|x} U_{xy} - G_x(\boldsymbol{U}_{x\cdot})\right)$ is unchanged in all translations $\boldsymbol{U}_{x\cdot} \to \boldsymbol{U}_{x\cdot} + a_x$. As a consequence, the subgradient of the Legendre-Fenchel transform at $\boldsymbol{\mu}_{\cdot|x}$ is a whole line, and the generalized entropy $\mathcal{E}$ also has a subgradient that is invariant by this group of transformations.

More precisely, the identifying inclusion

$$\boldsymbol{\Phi} \in \partial \mathcal{E}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m})$$

can be rewritten as the equality

$$\boldsymbol{D}_2 \boldsymbol{\Phi} = \boldsymbol{D}_2 \, \partial_{\boldsymbol{\mu}} \mathcal{E}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m})$$

where $\boldsymbol{D}_2$ is the double differencing $(X \times Y, X \times Y)$ matrix: for a vector $\boldsymbol{c} = (c_{xy})$,

$$(\boldsymbol{D}_2 \boldsymbol{c})_{xy} = c_{xy} - \frac{1}{Y} \sum_{t=1}^{Y} c_{xt} - \frac{1}{X} \sum_{z=1}^{X} c_{zy} + \frac{1}{XY} \sum_{z=1}^{X} \sum_{t=1}^{Y} c_{zt}.$$

Minimum Distance Estimation without Singles The matrix $\boldsymbol{D}_2$ only has rank $(X-1) \times (Y-1)$. To avoid zero eigenvalues in the variance that is used in the optimal weighting matrix, one should premultiply $\boldsymbol{D}_2$ by a $((X-1)(Y-1), XY)$ matrix $\boldsymbol{A}$ such that $\boldsymbol{A}\boldsymbol{D}_2$ is invertible. The identifying equations become

$$\boldsymbol{A}\boldsymbol{D}_2 \, \boldsymbol{\Phi} = \boldsymbol{A}\boldsymbol{D}_2 \, \partial_{\boldsymbol{\mu}} \mathcal{E}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}).$$

The minimum distance estimator can be applied to the resulting mixed hypothesis, as in Section 3.2.2. If the efficient weighting matrix is used, the χ^2 specification test statistic has $(X+Y-1)$ fewer degrees of freedom than when singles are observed.

Poisson Estimation without Singles The Poisson-GLM estimator of Section 3.2.2.4 is easily adapted to the Choo and Siow homoskedastic model without singles. The weights are now $w_{xy} = 2$ and the $\boldsymbol{X}$ matrix simplifies to

$$\boldsymbol{\mathcal{X}} = \begin{pmatrix} \bar{\boldsymbol{\phi}}/2 & -\frac{1}{2} I_X \otimes 1_Y & -\frac{1}{2} 1_X \otimes I_Y \end{pmatrix}.$$

The first-order conditions of the Poisson estimator are

$$\sum_{x,y} \boldsymbol{X}_{xy} \exp((\boldsymbol{\phi}_{xy}\boldsymbol{\gamma} - u_x - v_y)/2) = \sum_{x,y} \boldsymbol{X}_{xy}\boldsymbol{\gamma}.$$

Even if all basis functions ϕ^k_{xy} interact x and y, only the sums of the expected utilities $u_x + v_y$ are identified in the absence of singles. Therefore one normalization condition must be imposed when estimating the model. If we fix say $u_1 = 0$, we can drop the first column of $\boldsymbol{X}$. All other utility estimates must be interpreted accordingly.

3.2.3.3 Roommate Matching

Same-sex marriage has become legal in more and more countries across the world. More generally, many one-to-one matching situations do not restrict partners to belong to two separate subpopulations, as bipartite matching does. Gale and Shapley (1962) used the term “roommate matching” to refer to such markets. They showed that a stable matching may fail to exist: since anyone can match with anyone, the number of potential blocking coalitions may just be too large. While their model excluded transfers, later literature showed that this result extends to the TU world. Chiappori, Galichon, and Salanié (2019) showed that in large roommate matching markets, a stable matching always exists. Their paper also demonstrates a simple way to reformulate any such market as a bipartite market; identification and estimation results can then be translated from the latter to the former. This was extended to models with continuous observed types by Ciscato, Galichon, and Goussé (2020), building on the methods of Dupuy and Galichon (2014) described in Section 3.2.3.4 below. Using data on Californian households from 2008 to 2012, they find that same-sex couples display less preference for assortative matching than different-sex couples with respect to ethnic background, and more for education.

3.2.3.4 Continuous Observed Characteristics

As is well-known, continuous-choice models are not a natural limit of discrete-choice models. A multinomial logit with m alternatives, for instance, yields an expected utility that increases like $\log(m)$, so that the Emax functions are not well-defined when there is an infinity of alternatives. Dagsvik (1994) proposed a clever solution to this conundrum in a matching framework. Consider a man $\tilde{x} = (x, \varepsilon)$. In a separable model, we know that in equilibrium there exists a function $U : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ such that this man will choose the observed type y of his partner by maximizing $U(x, y) + \zeta(x, y, \varepsilon)$. Suppose that each individual can only match with partners they randomly meet; Dagsvik (1994) showed that if meetings are generated as the points of a Poisson process with a well-chosen intensity, this results in formulæ that are the continuous analog of those that Choo and Siow (2006) obtained for the discrete logit model of Section 2.2.2.

More precisely, let the full types $\tilde{y}_m = (y_m, \eta_m)$ of women for $m \geq 1$ be draws from a Poisson process on the set of full types with intensity $dy \times \exp(-\eta)$. Then it can be shown that the probability density on matches for the average man with an observed type x is

$$\mu(y|x) = \frac{\exp(U(x,y))}{1 + \int \exp(U(x,t))dt},$$

and that this man stays single with probability

$$\mu(0|x) = \frac{1}{1 + \int \exp(U(x,t))dt}.$$

Expected utilities are finite and given by $u(x) = -\log \mu(0|x) = \log(1 + \int \exp(U(x,t))dt)$, which is the direct equivalent of the G_x function in the discrete version of the model. Dupuy and Galichon (2014) showed how the techniques described in previous sections extend naturally to this continuous logit model. They used a Dutch survey to estimate the effect of the "big five" personality traits on marriage patterns. Bojilov and Galichon (2016) derived closed-form formulæ for the special case when in addition $\Phi(x,y)$ is quadratic and the observed types x and y are normally distributed.

Lindenlaub (2017) used a quadratic-Gaussian model with bivariate observed types to disentangle the effects of skill-biased and task-biased technological change. Her approach differs from that adopted in this survey: her model has agents and firms matching on observables only, and she introduces error terms at the estimation stage.

Guadalupe, Rappoport, Salanié, and Thomas (2024) build on Dupuy and Galichon (2014) to estimate a model with both discrete and continuous characteristics. The resulting formulæ are natural extensions of the fully continuous case: if observed types y have discrete components y_d and continuous components y_c, then

$$\mu(y_d, y_c|x) = \frac{\exp(U(x,(y_d,y_c)))}{1+\sum_{t_d}\int \exp(U(x,(t_d,t_c)))dt_c}.$$

They use this specification to study mergers between European firms; in their application, the discrete characteristics are industry and country, and the continuous characteristics are productivity and size.

To the best of our knowledge, the literature has not come up with a continuous separable model that goes beyond the logit structure, with its appealing simplicity but also its restrictive properties.

3.2.3.5 Index Models

The great appeal of the separable approach is that by ruling out interactions between unobserved characteristics of agents on the two sides of the market, it makes it much easier to identify and estimate matching and hedonic models. An alternative approach assumes that agents of a given (observed) type all use the same low-dimensional index to choose between alternatives. As this is not a very appealing assumption when types are discrete, we now assume that they are continuous; we denote $\boldsymbol{x} = (x_1, \dots, x_J)$ and $\boldsymbol{y} = (y_1, \dots, y_K)$ the observed types, with full types $\tilde{x} = (\boldsymbol{x}, \varepsilon)$ and $\tilde{y} = (\boldsymbol{y}, \eta)$.

Chiappori, Oreffice, and Quintana-Domeque (2012, 2020) pioneered the study of index models. They assumed the following:

Definition 7 (Matching by Indices)**.**

(i) all women value the observed types of men with the same scalar index of marital attractiveness $\bar{x} = g(\boldsymbol{x})$, and all men value the observed types of women with the same scalar index $\bar{y} = h(\boldsymbol{y})$. The joint utility formed by a match of $\tilde{x} = (\boldsymbol{x}, \varepsilon)$ and $\tilde{y} = (\boldsymbol{y}, \eta)$ only depends on the values of these two indices and on the unobserved types:

$$\tilde{\Phi}(\tilde{x}, \tilde{y}) = \Phi(\bar{x}, \bar{y}, \varepsilon, \eta);$$

(ii) ε (resp. η) is independent of $\boldsymbol{x}$ conditional on $\bar{x}$ (resp. of $\boldsymbol{y}$ conditional on $\bar{y}$).

Note that both parts of this assumption are important: if for instance (ii) failed, agents on one side of the market would value unobserved characteristics differently.

Some models in this class are separable; a very simple example would be

$$\Phi(\bar{x}, \bar{y}, \varepsilon, \eta) = \bar{x}\bar{y} + \varepsilon + \eta.$$

Diamond and Agarwal (2017) focussed on a different case: the "doubly vertical" subclass for which $\Phi(\bar{x}, \bar{y}, \varepsilon, \eta) = \bar{\Phi}(\bar{x}+\varepsilon, \bar{y}+\eta)$ with a supermodular $\bar{\Phi}$, and ε (resp. η) independent of $\boldsymbol{x}$ (resp. $\boldsymbol{y}$).

Chiappori, Oreffice, and Quintana-Domeque (2012) showed that in models that satisfy (i) and (ii), the distribution of $\bar{x}$ conditional on $\boldsymbol{y}$ *across* *couples* only depends on $\bar{y}$, and vice-versa. In more compact form,

$$\bar{x} \perp\!\!\!\perp \boldsymbol{y} \mid \bar{y} \text{ and } \bar{y} \perp\!\!\!\perp \boldsymbol{x} \mid \bar{x}.$$

If we knew one of the two indices, say $\bar{x}$, we could estimate the other index $\bar{y}$ by running a nonparametric regression of $\bar{x}_c$ on $\boldsymbol{y}_c$ across the observed couples $c = 1, \ldots, C$. In general, we know neither $\bar{x}$ nor $\bar{y}$; we can only regress $\boldsymbol{x}$ on $\boldsymbol{y}$, or the reverse. This is a problem since as shown by Chiappori, Oreffice, and Quintana-Domeque (2020), *all* matchings that map $\bar{x}$ to $\bar{y}$ in the same way are stable. More concretely, if a matching with couples $(\boldsymbol{x}_c, \boldsymbol{y}_c)_{c=1}^C$ is stable, then any other matching with couples $(\boldsymbol{x}'_c, \boldsymbol{y}'_c)_{c=1}^C$ such that

$$g(\boldsymbol{x}'_c) = g(x_c) \text{ and } h(\boldsymbol{y}'_c) = h(y_c) \text{ for all } c = 1, \ldots, C$$

is also stable.

As shown by Guadalupe, Rappoport, Salanié, and Thomas (2024), the continuous logit model is one favorable case in which regressing $\boldsymbol{x}$ on $\boldsymbol{y}$ will identify the index $\bar{y}$. To see this, remember that in the logit model

$$\mu(\boldsymbol{x}, \boldsymbol{y}) = \sqrt{\mu(\boldsymbol{x}, 0)\mu(0, \boldsymbol{y})} \exp(\Phi(\boldsymbol{x}, \boldsymbol{y})/2).$$

Suppose that $\Phi(\boldsymbol{x}, \boldsymbol{y}) = \bar{\Phi}(\bar{x}, \bar{y})$. The conditional expectation of $\boldsymbol{x}$ given $\boldsymbol{y}$, across observed couples, is

$$E(\boldsymbol{x}|\boldsymbol{y}) = \frac{\int \boldsymbol{x}\mu(\boldsymbol{x}, \boldsymbol{y})d\boldsymbol{x}}{\int \mu(\boldsymbol{x}, \boldsymbol{y})d\boldsymbol{x}} = \frac{\int \boldsymbol{x}\sqrt{\mu(\boldsymbol{x}, 0)} \exp(\bar{\Phi}(\bar{x}, \bar{y})/2)d\boldsymbol{x}}{\int \sqrt{\mu(\boldsymbol{x}, 0)} \exp(\bar{\Phi}(\bar{x}, \bar{y})/2)d\boldsymbol{x}}$$

which only depends on $\boldsymbol{y}$ via $\bar{y}$. Therefore for any component $j = 1, \ldots, J$, we have

$$E(x_j|\boldsymbol{y}) = f_j(\bar{y})$$

for some scalar function f_j. This generates testable implications as the "marginal rates of substitution"

$$\frac{\frac{\partial E(x_j|\boldsymbol{y})}{\partial y_k}}{\frac{\partial E(x_j|\boldsymbol{y})}{\partial y_l}} = \frac{\frac{\partial \bar{y}}{\partial y_k}}{\frac{\partial \bar{y}}{\partial y_l}}$$

must be the same for all $j = 1, \ldots, J$. These equations identify the index $\bar{y}$ up to an invertible transformation that is economically irrelevant.

Dupuy and Galichon (2014) developed inference procedures for a different version of the semilinear logit model. Suppose that $\Phi(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x}'\boldsymbol{A}\boldsymbol{y} = \sum_{i=1}^{n}\sum_{j=1}^{m} A_{ij}x_i y_j$, where A is an unknown "affinity matrix"; and that the error terms have the structure described in Section 3.2.3.4. If $\boldsymbol{A}$ is a rank one matrix, then it can be written as $\boldsymbol{A} = \boldsymbol{u}\boldsymbol{v}'$, where $\boldsymbol{u}$ and $\boldsymbol{v}$ are two column vectors defined up to a multiplicative scalar. In that case, $g(\boldsymbol{x}) = \boldsymbol{x}'\boldsymbol{u}$ and $h(\boldsymbol{y}) = \boldsymbol{v}'\boldsymbol{y}$ can be interpreted as one-dimensional attractiveness indices. More generally, if $\boldsymbol{A}$ is of rank p then $\Phi(\boldsymbol{x}, \boldsymbol{y})$ can be written as a sum of products of indices: $\sum_{k=1}^{p}(\boldsymbol{x}'\boldsymbol{u}_k)(\boldsymbol{v}_k'\boldsymbol{y})$. Dupuy and Galichon (2014) show how $\boldsymbol{A}$ can be estimated, and how to test for its rank. Alternatively, if the rank of $\boldsymbol{A}$ is imposed it can be estimated using a rank-constrained procedure (see (Dupuy, Galichon, and Sun, 2019)).

3.2.3.6 Many-to-one and Many-to-many Matching

When each match may consist of more than two partners, there may not exist a stable matching. This has been known since the work of Kelso and Crawford (1982) on matching with transfers in the labor market. The underlying intuition is most simply seen by thinking of convergence towards a stable matching. If hiring a worker makes another worker more valuable to a firm, this complementarity may hinder convergence, just as in a Walrasian tâtonnement process or in simultaneous ascending auctions. Many papers since then have extended this finding and made it more precise.

Assuming away complementarities is one strategy, which is not always very credible. Recent literature (e.g Fox, 2018) instead switched focus from stable outcomes to competitive equilibria. In large markets, a competitive equilibrium typically exists and is efficient (Azevedo and Hatfield, 2018). As such, it maximizes social welfare under the scarcity constraints—which is exactly what we have been doing in Section 1.3.2.

So far, the maximum-score methods developed by Fox (2010, 2018) have dominated applications in one-to-many and many-to-many markets. Bajari and Fox (2013) modeled the stability of the FCC spectrum license auctions in the mid 1990s. The license seller (the US government) has very restricted preferences; still, this was a matching game as resale among participants is permitted, and in fact quite an active secondary market exists. Bajari and Fox (2013) assumed that the allocation of licenses was pairwise stable: an exchange of two licenses by winning bidders could not have raised the sum of the valuations of the two bidders. Their estimates suggest that given large complementarities between licenses, packaging them by (large) region would have raised the allocative efficiency of the auction.

Fox (2018) applied these methods to a many-to-many market: matching between car parts suppliers and assemblers. The unit of observation is a "trade": a supplier selling a car part to an assembler. Each supplier may sell various car parts to different assemblers, and vice-versa. Since car parts need to fit together and there are economies of scope, the valuation function of each firm depends on its whole portfolio of trades and is not additively separable across its trades. With hundreds of suppliers, eleven assemblers, and more than 30,000 car parts, Fox (2018) could generate a huge number of matching inequalities. He chose to randomly sample among them.

The methods we developed for a bipartite market under separability extend readily to some other markets; we only develop one example here. Suppose that each match on the labor market must associate a firm $\tilde{x} = (x, \varepsilon)$ and at most one worker in each of p tasks. We denote a worker in task $k = 1, \dots, p$ as $\tilde{y}^k = (y^k, \eta^k)$, and we assume that this worker cannot work in another task. In a separable model, the joint utility from a match of firm $\tilde{x}$ and workers $\tilde{\boldsymbol{y}} = (\tilde{y}^1, \dots, \tilde{y}^p)$ is

$$\Phi(x, \boldsymbol{y}) + \varepsilon(\boldsymbol{y}) + \sum_{k=1}^{p} \eta^k(x, \boldsymbol{y}^{-k}),$$

where $\boldsymbol{y} = (y^1, \dots, y^p)$ and $\boldsymbol{y}^{-k}$ denotes $\boldsymbol{y}$ minus its k-th component—that is, the observed types of workers in all tasks except task k. This separability assumption allows the value-added of a task-k worker to vary both with the observed types of firms and with the observed types of workers in other tasks in that firm. We allow for unfilled tasks $y^k = 0$ in the obvious way.

It is easy to extend Theorem 6 to this $(p+1)$-partite separable model. There exist functions $U(x, \boldsymbol{y})$ and $V^k(x, \boldsymbol{y})$ for $k = 1, \ldots, p$ that add up to $\Phi(x, \boldsymbol{y})$ such that in equilibrium:

- firm $f \in x$ matches with workers of observed types $\boldsymbol{y}$ that jointly achieve the maximum of $U(x, \boldsymbol{y}) + \varepsilon^f(\boldsymbol{y})$;
- a worker $\tilde{y}^k$ in task k matches with a firm of observed type x and workers of observed types $\boldsymbol{y}^{-k}$ that jointly achieve the maximum of $V^k(x, \boldsymbol{y}) + \eta^k(x, \boldsymbol{y}^{-k})$.

As in Section 2.2, one can define Emax functions G_x and H_{y^k}, their Legendre-Fenchel transforms, and a generalized entropy. The identification equation (2.9) becomes

$$\Phi(x, y) = \frac{\partial G_x^*}{\partial \mu(\boldsymbol{y}|x)}(\mu(\cdot|x)) + \sum_{k=1}^{p} \frac{\partial H_{y^k}^*}{\partial \mu(x, \boldsymbol{y}^{-k}|y^k)}(\mu(\cdot|y^k)).$$

All of our results in Section 3.2 adapt straightforwardly to this one-to-many setting.

In the logit model for instance, the Choo and Siow formula (2.12) becomes

$$\Phi(x, \boldsymbol{y}) = \log \frac{\mu_{xy}^{1+p}}{\mu_{x0} \prod_{k=1}^{p} \mu_{0y^k}}.$$

Corblet (2022) uses this variant of the logit model to examine how workers in different age and education groups sort across firms in different industries on the Portuguese labor market. In her application, what we denoted k is a worker age-education group, rather than a task, and firms choose how many workers of each group they want to hire. These two models have a very similar structure. Let y^k representing the number of workers in group k that the firm hires. Then the joint utility is

$$\Phi(x, \boldsymbol{y}) + \varepsilon(\boldsymbol{y}) + \sum_{k=1}^{p} y^k \eta^k(x, \boldsymbol{y}^{-k});$$

a worker in group k maximizes $V_k(x, \boldsymbol{y}^{-k}) + \eta^k(x, \boldsymbol{y}^{-k})$, where

$$U(x, \boldsymbol{y}) + \sum_{k=1}^{p} y^k V_k(x, \boldsymbol{y}^{-k}) = \Phi(x, \boldsymbol{y});$$

and in the logit model, the identification formula is

$$\Phi(x, \boldsymbol{y}) = \log \frac{\mu(x, \boldsymbol{y})}{\mu(x, 0)} + \sum_{k=1}^{p} y^k \log \frac{y^k \mu(x, \boldsymbol{y})}{\mu(0, k)}.$$

3.2.3.7 Skill Bundling

Heckman and Scheinkman (1987) emphasized that workers should really be seen as bundles of skills that cannot be sold separately on the labor market. As a consequence, firms are constrained in the total vectors of skills they can acquire. To make this more precise, suppose that a worker of type i can be described as a bundle $\tilde{\boldsymbol{y}} = (\tilde{y}^1, \ldots, \tilde{y}^S)$, where $\tilde{y}^s$ is that worker's endowment of skill s. Skill bundling implies that a firm can only acquire restricted combinations of skills. To take an extreme example, firms in a particular industry may only use worker skill $s = 2$; yet that skill will only come bundled with other skills that have zero value in this industry.

Let us denote $\tilde{F}_j(a^1, \ldots, a^S)$ the value of the output of the representative firm in industry j when it uses a^s units of skill $s = 1, \ldots, S$; we assume that F_j is increasing, continuously differentiable and concave. Given skill bundling, each a^s is a weighted sum $\sum_i \mu_{ij} \tilde{y}_i^s$, where $\mu_{ij} = 1$ is the number of workers of type i employed by industry j.

In a competitive equilibrium, industry j offers a vector of wages $(w_{ij})_i$. If workers of type i face disutility d_{ij} from working in industry j, their equilibrium utility is simply

$$u_i = \max_j (w_{ij} - d_{ij}).$$

The labor demands $\boldsymbol{\mu}_{\cdot j}$ of industry j maximize its profit

$$F_j \left(\sum_i \mu_{ij} \tilde{\boldsymbol{y}}_i \right) - \sum_i \mu_{ij} w_{ij}$$

under the constraints that each $\mu_{ij} \geq 0$ (neglecting integer constraints for simplicity).This gives the first-order conditions

$$\sum_s \frac{\partial F_j}{\partial a_{js}} (a_{j1}, \ldots, a_{jS}) \; \tilde{y}_i^s \leq w_{ij},$$

with equality if $\mu_{ij} > 0$. This shows that if industry j does employ workers of type i, then their wage can be written as

$$w_{ij} = \sum_s y_i^s \bar{w}_j^s.$$

Therefore each industry j uses a vector of skill-specific wage rates $\bar{\boldsymbol{w}}_j$ to compensate workers. These wage rates equal the marginal productivities of the corresponding skill in that industry.

If skills could be "unbundled" and hired separately, productive efficiency implies that all industries should use the same skill-specific wage rates: $\bar{w}_j^s \equiv \bar{w}^s$ for all industries that use skill s. Heckman and Scheinkman (1987) show that with skill bundling, there may exist equilibria in which different industries apply different skill-specific wage rates: $w_j^s \neq w_k^s$. These occur in "boundary cases" in which the set of productively optimal allocations of worker skills to firms hits the boundary of the set of feasible allocations, given the restrictions implied by skill bundling.

It is tempting to bring this model inside the matching framework. The joint surplus from a match between the representative firm in industry j and numbers $(\mu_{ij})_i$ of the different types of workers is the value of the output minus the sum of workers' disutilities:

$$F_j\left(\sum_i \mu_{ij}\tilde{\boldsymbol{y}}_i\right) - \sum_i \mu_{ij} d_{ij}$$

The stable matching maximizes the sum of these surpluses across industries, under the usual constraints on the number of workers of each type: $\sum_j \mu_{ij} \leq n_i$. If the production functions F_j are not linear, the total surplus is not a linear function of the matching patterns $\boldsymbol{\mu}$. This makes the analysis of this model more intricate. Choné and Kramarz (2022) use the techniques of "weak optimal transport" (see Gozlan, Roberto, Samson, and Tetali, 2017) to study the properties of the set of stable matchings. They show that technological innovations such as platforms, by making skill unbundling easier, benefit workers with specialized skills more than "generalists"; moreover, specialized firms tend to specialize even further. Skans, Choné, and Kramarz (2021) validates these predictions on the Swedish labor market.

3.2.4 Imperfectly Transferable Utility

Recall that in separable matching models with imperfectly transferable utility, equation (2.15) holds:

$$D_{xy}\left(\frac{\partial G_x^*}{\partial \mu_{y|x}}\left(\mu_{\cdot|x}\right), \frac{\partial H_y^*}{\partial \mu_{x|y}}\left(\mu_{\cdot|y}\right)\right) = 0.$$

Suppose that we parameterize the distance function D_{xy} and the distribution $\mathbb{P}_x$ and $\mathbb{Q}_y$ with some vector of unknown parameters $\boldsymbol{\theta}$. We can write the following mixed hypothesis:

$$\text{there exists } \boldsymbol{\theta} \text{ such that for all } (x,y),\ D_{xy}^{\boldsymbol{\theta}}\left(\frac{\partial G_x^{*,\boldsymbol{\theta}}}{\partial \mu_{y|x}}\left(\mu_{\cdot|x}\right), \frac{\partial H_y^{*,\boldsymbol{\theta}}}{\partial \mu_{x|y}}\left(\mu_{\cdot|y}\right)\right) = 0.$$

Given estimates $\hat{\boldsymbol{\mu}}$, this can be used as the basis for a minimum distance estimator of $\boldsymbol{\theta}$, just as in Section 3.2.2.

Let us illustrate this on the logit model, which only has parameters in the distance function. Using (2.17), the mixed hypothesis is

$$\text{there exists } \boldsymbol{\theta} \text{ such that for all } (x,y),\ \log \mu_{xy} + D_{xy}^{\boldsymbol{\theta}}\left(-\log \mu_{x0}, -\log \mu_{0y}\right) = 0.$$

Suppose that utility is close to perfectly transferable: the utility possibility frontier is defined by the set of U, V such that $\exp(\lambda(U - \alpha_{xy}^{\boldsymbol{\theta}})) + \exp(\lambda(V - \gamma_{xy}^{\boldsymbol{\theta}})) = 2$ for a small unknown $\lambda > 0$. We let $\Phi_{xy}^{\boldsymbol{\theta}} = \alpha_{xy}^{\boldsymbol{\theta}} + \gamma_{xy}^{\boldsymbol{\theta}}$.

Since we assumed that λ is small, we approximate $\exp(\lambda t) \simeq 1+\lambda t+\lambda^2 t^2/2$. Simple calculations show that up to a first-order approximation in λ,

$$D^{\boldsymbol{\theta}}_{xy}(u,v) = \frac{u+v-\Phi_{xy}}{2}+\frac{\lambda}{4}(\left(u-\alpha^{\boldsymbol{\theta}}_{xy}\right)^2+\left(v-\gamma^{\boldsymbol{\theta}}_{xy}\right)^2-\left(u-\alpha^{\boldsymbol{\theta}}_{xy}\right)\left(v-\gamma^{\boldsymbol{\theta}}_{xy}\right)).$$

The first term on the right-hand side corresponds to the case of perfectly transferable utility; the second term reflects the effect of the costs of utility transfers. In the end the mixed hypothesis can be approximated by

$$\begin{aligned} \log \frac{\mu^2_{xy}}{\mu_{x0}\mu_{0y}} &= \alpha_{xy}+\gamma_{xy} \\ &+\frac{\lambda}{2}\left(\begin{array}{c} \left(\log\frac{\mu_{xy}}{\mu_{x0}}-\alpha^{\boldsymbol{\theta}}_{xy}\right)^2+\left(\log\frac{\mu_{xy}}{\mu_{0y}}-\gamma^{\boldsymbol{\theta}}_{xy}\right)^2 \\ -\left(\log\frac{\mu_{xy}}{\mu_{x0}}-\alpha^{\boldsymbol{\theta}}_{xy}\right)\left(\log\frac{\mu_{xy}}{\mu_{0y}}-\gamma^{\boldsymbol{\theta}}_{xy}\right) \end{array}\right), \end{aligned}$$

where we recognize the Choo and Siow formula (2.12) when $\lambda = 0$. Given a parameterization of α, γ, and λ, this can be estimated by minimum-distance for instance. One could also run a test for the hypothesis $\lambda = 0$ of perfectly transferable utility.

Maximum likelihood estimation is also possible. Just as in Section 3.2.2, the log-likelihood of the sample is

$$\sum_{xy\in\mathcal{X}\times\mathcal{Y}} \hat{\mu}_{xy}\ln\mu^{\theta}_{xy}+\sum_{x\in\mathcal{X}}\hat{\mu}_{x0}\ln\mu^{\theta}_{x0}+\sum_{y\in\mathcal{Y}}\hat{\mu}_{0y}\ln\mu^{\theta}_{0y}-\hat{N}\ln N^{\theta},$$

where the number of households at the stable matching $N^{\boldsymbol{\theta}}$ is given by

$$N^{\boldsymbol{\theta}} = \sum_{x\in\mathcal{X}}\mu^{\boldsymbol{\theta}}_{x0}+\sum_{y\in\mathcal{Y}}\mu^{\boldsymbol{\theta}}_{0y}+\sum_{xy\in\mathcal{X}\times\mathcal{Y}}\mu^{\boldsymbol{\theta}}_{xy}.$$

This has a nice interpretation in the ITU model with logit heterogeneity. For a given parameter vector $\boldsymbol{\theta}$, the numbers of matches can be defined by

$$\begin{cases} \mu_{xy}^{\boldsymbol{\theta}} = \exp\left(-D_{xy}^{\boldsymbol{\theta}}\left(u_x, v_y\right)\right) \\ \mu_{x0}^{\boldsymbol{\theta}} = \exp\left(-u_x\right) \\ \mu_{0y}^{\boldsymbol{\theta}} = \exp\left(-v_y\right), \end{cases}$$

where u_x and v_y are the average utilities of each observed type at the stable matching. The first order conditions of the MLE can be rewritten as

$$\begin{aligned} \frac{\hat{\mu}_{x0}}{\hat{N}} + \sum_{y\in\mathcal{Y}} \frac{\hat{\mu}_{xy}}{\hat{N}} \frac{\partial D_{xy}^{\boldsymbol{\theta}}\left(u_x, v_y\right)}{\partial u_x} &= \tfrac{\mu_{x0}^{\theta}}{N^{\theta}} + \textstyle\sum_{y\in\mathcal{Y}} \tfrac{\mu_{xy}^{\theta}}{N^{\theta}} \tfrac{\partial D_{xy}^{\boldsymbol{\theta}}(u_x,v_y)}{\partial u_x} \\ \frac{\hat{\mu}_{0y}}{\hat{N}} + \sum_{x\in\mathcal{X}} \frac{\hat{\mu}_{xy}}{\hat{N}} \frac{\partial D_{xy}^{\boldsymbol{\theta}}\left(u_x, v_y\right)}{\partial v_y} &= \tfrac{\mu_{0y}^{\theta}}{N^{\theta}} + \textstyle\sum_{y\in\mathcal{Y}} \tfrac{\mu_{xy}^{\theta}}{N^{\theta}} \tfrac{\partial D_{xy}^{\boldsymbol{\theta}}(u_x,v_y)}{\partial v_y} \\ \sum_{xy\in\mathcal{X}\times\mathcal{Y}} \frac{\hat{\mu}_{xy}}{\hat{N}} \frac{\partial D_{xy}^{\boldsymbol{\theta}}\left(u_x, v_y\right)}{\partial \boldsymbol{\theta}_k} &= \textstyle\sum_{xy\in\mathcal{X}\times\mathcal{Y}} \tfrac{\mu_{xy}^{\theta}}{N^{\theta}} \tfrac{\partial D_{xy}^{\boldsymbol{\theta}}(u_x,v_y)}{\partial \boldsymbol{\theta}_k}. \end{aligned}$$

The interpretation of these equations is clear: the MLE ensures that certain moments predicted by the parameter value $(u, v, \boldsymbol{\theta})$ match their empirical counterparts. The first set of equations matches the observed and fitted moments of $\partial D_{xy}^{\boldsymbol{\theta}}/\partial u_x$; this is the partial derivative of the distance function with respect to the average utility of agents of type x. With perfectly transferable utility, this quantity is $1/2$ uniformly. When utility is imperfectly transferable, it can be shown that the vector $(\partial D_{xy}^{\boldsymbol{\theta}}/\partial u_x, \partial D_{xy}^{\boldsymbol{\theta}}/\partial v_y)$ is tangent to the feasible utility set for (x, y) couples.

3.3 Inference in Hedonic Equilibrium Models

We focus here on the class of hedonic markets of Section 2.3.3, where agents are endowed with quasilinear utilities, with and additively separable unobserved heterogeneity and finite sets of observed types and products. We denote the utilities

$$U(\tilde{x}, z; P) = \bar{U}_{xz} - P_z + \zeta_z^{\tilde{x}} \text{ and } V(\tilde{y}, z; P) = \bar{V}_{yz} + P_z + \xi_z^{\tilde{y}}$$

and we let $\hat{\mu}_{z|x}$ and $\hat{\mu}_{z|y}$ denote the estimated choice probabilities of buyers and sellers.

If prices are observed, the function P is identified from the data and the most direct way to estimate the model is simply to estimate separately the models for individual demand and for individual supply, using for instance the results in Section 3.1.

If on the other hand the data has no information on prices, we can base a minimum-distance estimator on the formula that we derived in Section 2.3.3.2:

$$\bar{U}_{xz} + \bar{V}_{yz} = \frac{\partial G_x^*}{\partial \boldsymbol{\mu}(\cdot|x)}(\mu_{z|x}) + \frac{\partial H_y^*}{\partial \boldsymbol{\mu}(\cdot|y)}(\mu_{z|y}).$$

Let us parameterize the functions $\bar{U}$ and $\bar{V}$ and the distributions $\mathbb{P}_x$ and $\mathbb{Q}_y$ with an unknown vector $\boldsymbol{\theta}$ (and normalize their location). Then we can estimate $\boldsymbol{\theta}$ by minimizing the norm of a vector with components

$$\bar{U}_{xz}^{\boldsymbol{\theta}} + \bar{V}_{yz}^{\boldsymbol{\theta}} = \frac{\partial G_x^{*\boldsymbol{\theta}}}{\partial \mu_{z|x}}\left(\mu_{.|x}\right) + \frac{\partial H_y^{*\boldsymbol{\theta}}}{\partial \mu_{z|y}}\left(\mu_{.|y}\right)$$

for all or a subset of (x, y, z) 3-uples, given a consistent estimator $\hat{\boldsymbol{\mu}}$ of $\boldsymbol{\mu}$.

In a logit model, the right-hand side is simply

$$\log \frac{\hat{\mu}_{xz}\hat{\mu}_{yz}}{\hat{\mu}_{x0}\hat{\mu}_{y0}}.$$

Assume moreover that the buyer and seller utility functions are linear: $\bar{U}^{\boldsymbol{\theta}}_{xz} = \sum_{k=1}^{K} \phi^k_{xz}\theta_k$ and $\bar{V}^{\boldsymbol{\theta}}_{yz} = \sum_{k=1}^{K} \psi^K_{yz}\theta_k$ (letting e.g. $\boldsymbol{\psi}^k = 0$ if $\boldsymbol{\phi}^k$ only makes sense for buyers). Then the mixed hypothesis is

$$\text{there exists } \boldsymbol{\theta} \text{ such that for all } x, y, z,\ \sum_{k=1}^{K} \left(\phi^k_{xz} + \psi^k_{yz}\right)\theta_k = \log \frac{\mu_{xz}\mu_{yz}}{\mu_{x0}\mu_{y0}}.$$

Just as in Section 3.2.2, $\boldsymbol{\theta}$ can be estimated by GLS. The logarithm can cause difficulties in hedonic models, where many (x, y, z) cells may be empty. The Poisson estimator of Section 3.2.2 circumvents this issue. One can prove the following result:

Theorem 11. *For given $\boldsymbol{\theta}$, consider the convex optimization problem*

$$\max_{\boldsymbol{u},\boldsymbol{v},\boldsymbol{p}} \left\{ \begin{array}{c} \sum_{\substack{x\in\mathcal{X} \\ z\in\mathcal{Z}}} U^{\boldsymbol{\theta}}_{xz}\hat{\mu}_{xz} + \sum_{\substack{y\in\mathcal{Y} \\ z\in\mathcal{Z}}} V^{\boldsymbol{\theta}}_{yz}\hat{\mu}_{yz} - \sum_x n_x u_x - \sum_y m_y v_y \\ - \sum_{\substack{x\in\mathcal{X} \\ z\in\mathcal{Z}}} n_x \exp\left(U^{\boldsymbol{\theta}}_{xz} - u_x - p_z\right) - \sum_{x\in\mathcal{X}} n_x \exp\left(-u_x\right) \\ - \sum_{\substack{y\in\mathcal{Y} \\ z\in\mathcal{Z}}} m_y \exp\left(V^{\boldsymbol{\theta}}_{yz} - v_y + p_z\right) - \sum_{y\in\mathcal{Y}} m_y \exp\left(-v_y\right) \end{array} \right\}.$$

Let $F(\boldsymbol{\theta})$ denote its value, and define $\hat{\boldsymbol{\theta}}$ to be a minimizer of the function F.

Then

(i) $\hat{\boldsymbol{\theta}}$ is a consistent and asymptotically normal estimator of $\boldsymbol{\theta}$.

(ii) in a semilinear model, F is a convex function and $\hat{\boldsymbol{\theta}}$ can be obtained

by running a Poisson regression of the stacked vector

$$\hat{\boldsymbol{\mu}} = \left((\hat{\mu}_{xz})'_{x,z}, (\hat{\mu}_{yz})'_{y,z}, (\hat{\mu}_{x0})'_{x}, (\hat{\mu}_{y0})'_{y} \right)$$

on the rows of the matrix

$$\boldsymbol{M} = \begin{pmatrix} \boldsymbol{\phi} & -\mathbf{1}_{|\mathcal{X}|} \otimes \mathbf{I}_{|\mathcal{Z}|} & -\mathbf{I}_{|\mathcal{X}|} \otimes \mathbf{1}_{|\mathcal{Z}|} & \mathbf{0}_{|\mathcal{X}||\mathcal{Z}| \times |\mathcal{Y}|} \\ \boldsymbol{\psi} & \mathbf{1}_{|\mathcal{Y}|} \otimes \mathbf{I}_{|\mathcal{Z}|} & \mathbf{0}_{|\mathcal{Y}||\mathcal{Z}| \times |\mathcal{X}|} & -\mathbf{I}_{|\mathcal{Y}|} \otimes \mathbf{1}_{|\mathcal{Z}|} \\ \mathbf{0}_{|\mathcal{X}| \times K} & \mathbf{0}_{|\mathcal{X}| \times |\mathcal{Z}|} & -\mathbf{I}_{|\mathcal{X}|} & \mathbf{0}_{|\mathcal{X}| \times |\mathcal{Y}|} \\ \mathbf{0}_{|\mathcal{Y}| \times K} & \mathbf{0}_{|\mathcal{Y}| \times |\mathcal{Z}|} & \mathbf{0}_{|\mathcal{Y}| \times |\mathcal{X}|} & -\mathbf{I}_{|\mathcal{Y}|} \end{pmatrix}.$$

Let $\hat{\boldsymbol{\lambda}}$ be the estimated coefficients—that is, the solution of

$$\max_{\boldsymbol{\lambda}} \left(\mathbf{1}' \exp\left(\boldsymbol{M} \boldsymbol{\lambda} \right) - \hat{\boldsymbol{\mu}}' \right);$$

then $\hat{\boldsymbol{\theta}}$ consists of the first K components of $\hat{\boldsymbol{\lambda}}$.

Chapter 4

Computation

Suppose we are given a complete specification of a hedonic market: the distributions of buyer and seller types, their preferences, and the set of products. This could for instance result from applying the identification and estimation procedures of Sections 2.3 and 3.3 to data. We may then want to compute the equilibrium prices and allocations for different counterfactuals. While we solved for equilibrium in a simple quadratic instance of a hedonic market in Section 1.2.2, the method we applied does not extend easily to more complex models. We show here how solving for the hedonic equilibrium can be reformulated as a network flow problem, for which fast algorithms are readily available.

In separable matching models, we obtained formula (2.10):

$$\boldsymbol{\Phi}_{xy} = -\frac{\partial \mathcal{E}}{\partial \mu_{xy}}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}).$$

Ir gives a very convenient way to identify the joint utility matrix $\boldsymbol{\Phi}$ when the distributions $(\mathbb{P}_x)$ and $(\mathbb{Q}_y)$ are known, given the matching patterns $\boldsymbol{\mu}$ and the margins $\boldsymbol{n}$ and $\boldsymbol{m}$. It is often necessary to solve in the other direction, however: given the specification $(\boldsymbol{\Phi}, (\mathbb{P}_x), (\mathbb{Q}_y))$ of a separable model, and the margins $\boldsymbol{n}$ and $\boldsymbol{m}$, what are the stable matching patterns $\boldsymbol{\mu}$? We will present the IPFP algorithm and several alternative computational approaches.

4.1 Solving for Hedonic Equilibria

We focus here on the quasi-linear case, when buyer and seller utilities are given by

$$U(\tilde{x}, z, P(z)) = \bar{U}(\tilde{x}, z) - P(z)$$
$$V(\tilde{y}, z, P(z)) = \bar{V}(\tilde{y}, z) + P(z)$$

so that a sale of one unit of z by $\tilde{x}$ to $\tilde{y}$ creates a joint surplus $\bar{U}(\tilde{x}, z) + \bar{V}(\tilde{y}, z)$. The hedonic equilibrium consists of measures μ^D on $\tilde{\mathcal{X}} \times \mathcal{Z}$ and μ^S on $\tilde{\mathcal{Y}} \times \mathcal{Z}$ that maximize the total joint surplus

$$S(\mu, \mu^S) := \int_{\tilde{\mathcal{X}} \times \mathcal{Z}} \bar{U}(\tilde{x}, z)\mu^D(d\tilde{x}, dz) + \int_{\tilde{\mathcal{Y}} \times \mathcal{Z}} \bar{V}(\tilde{y}, z)\mu(d\tilde{y}, dz)$$

under three types of constraints:

C1: there exists a price function $P(z)$ such that buyers only buy products that maximize their net utility, and sellers only sell products that maximize theirs. That is, for almost every $\tilde{x}$ (resp. $\tilde{y}$) the support of the conditional measure $\mu^D(\cdot|\tilde{x})$ (resp. $\mu^S(\cdot|\tilde{y})$) on $\mathcal{Z}$ is contained in the set of z values that maximize $\bar{U}(\tilde{x}, z) - P(z)$ (resp. $\bar{V}(\tilde{y}, z) + P(z)$).

C2: each buyer buys at most, and each seller sells at most, one unit of product; so that

$$\int_{\mathcal{Z}} \mu^D(dz|\tilde{x}) \leq 1 \text{ for almost all } \tilde{x}$$

and

$$\int_{\mathcal{Z}} \mu^S(dz|\tilde{y}) \leq 1 \text{ for almost all } \tilde{y}.$$

C3: demand equals supply for each product:

$$\int_{\tilde{X}} \mu^D(d\tilde{x}, dz) = \int_{\tilde{Y}} \mu^S(d\tilde{y}, dz) \text{ for all } z.$$

The objective function S and the constraints **C2** and **C3** are linear in the measures μ^D and μ^S. Unfortunately, the constraints in **C1** are not; this is what makes solving for equilibrium a non-trivial task in hedonic markets[1].

Following a suggestion from Maurice Queyranne (Queyranne (2011)), we will reformulate this constrained optimization program as a minimum cost network flow problem. This has the great advantage that many powerful algorithms are available to solve such problems[2].

To do so, we construct a simple network whose nodes are the sellers, the products, and the buyers. Sellers and buyers are only linked indirectly, via their links to products. The sale of a product z by $\tilde{y}$ to $\tilde{x}$ is represented as a

[1] Note in passing that under **C1**, the conditional measures $\mu^D(\cdot|\tilde{x})$ and $\mu^S(\cdot|\tilde{y})$ coincide with the measures $\tilde{\pi}^D(\cdot|\tilde{x}; P)$ and $\tilde{\pi}^S(\cdot|\tilde{y}; P)$ that we defined in Section 1.2.

[2] See Ahuja, Magnanti, and Orlin (1993) or Williamson (2019).

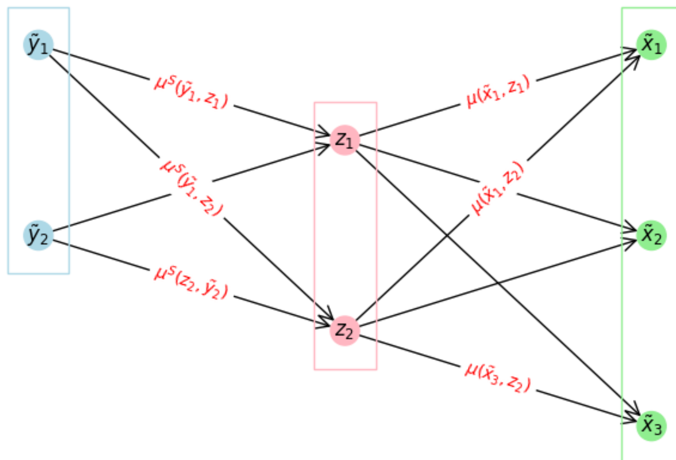

Figure 4.1: Hedonic Equilibrium as a Network Flow

Only some of the values of the flow are indicated.

unit mass that *flows* into the network at $\tilde{y}$, transits on the $\tilde{y} \to z$ link, then on the $z \to \tilde{x}$ link, and finally exits the network at $\tilde{x}$. Figure 4.1 illustrates this; to simplify the picture, we only show some of the values of the flows. In addition, we leave out flows to and from $z = 0$, which correspond to buyers and sellers who fail to transact with probability one.

In the language of network flow models, the seller nodes are *sources* and the buyer nodes are *destinations*. Each source $\tilde{y}$ is endowed with $\tilde{m}(\tilde{y})$ units of mass, and each destination $\tilde{x}$ has a capacity limit of $\tilde{n}(\tilde{x})$. These limits correspond to the **C2** constraints above. In addition, masses of products must be balanced: the flows into a product node z must equal the flows out of this node. These *mass balance conditions* are exactly the **C3** constraints.

All that is left is to translate the objective function into the cost of the network flow. This is achieved simply by assigning to a unit flow from $\tilde{y}$ to z (resp. from z to $\tilde{x}$) a cost $-\bar{V}(\tilde{y}, z)$ (resp. $-\bar{U}(\tilde{x}, z)$), so that the total cost of the network flow is

$$-\int_{\tilde{\mathcal{Y}}\times\mathcal{Z}} \bar{V}(\tilde{y}, z)\mu^S(d\tilde{y}, dz) - \int_{\tilde{\mathcal{X}}\times\mathcal{Z}} \bar{U}(\tilde{x}, z)\mu^D(d\tilde{x}, dz),$$

which is exactly the opposite of the joint surplus.

Finally, note that it would be easy to extend this analogy to hedonic models where agents can trade varieties in less constrained quantities. Denote $\bar{U}(\tilde{x}, q) - P^B(q)$ the utility of a buyer who purchases quantities $(q(z))_{z\in\mathcal{Z}}$ and $\bar{V}(\tilde{y}, Q) + P^S(Q)$ the utility of a seller who produces and sells quantities $(Q(z))_{z\in\mathcal{Z}}$. Each elementary flow now has weight $q(z)$ or $Q(z)$; the balance conditions **C3** are

$$\int_{\tilde{X}} q(z)\mu^D(d\tilde{x}, dq) = \int_{\tilde{Y}} Q(z)\mu^S(d\tilde{y}, dQ)$$

for each variety z. In this reformulation, the scarcity constraints **C2** may be unnecessary if $\bar{U}$ and $\bar{V}$ are concave in q. On the other hand, the link between the functions P^B and P^S will depend on restrictions on pricing. Linear pricing, for instance, would have $P^B(q) = \int_{z\in\mathcal{Z}} q(z)p(z)dz$ and $P^S(Q) = \int_{z\in\mathcal{Z}} Q(z)p(z)dz$ for a common function p.

4.1.1 The Primal Problem

Solving for μ^D and μ^S that minimize the cost of this network flow under the constraints **C2** and **C3** gives us the measures that constitute the hedonic equilibrium. Moreover, the multipliers of the scarcity constraints **C2** are the buyer and seller utilities

$$u(\tilde{x}) = \max\left(0, \max_{z\in\mathcal{Z}}\left(\bar{U}(\tilde{x}, z) - P(z)\right)\right)$$
$$v(\tilde{y}) = \max\left(0, \max_{z\in\mathcal{Z}}\left(\bar{V}(\tilde{y}, z) + P(z)\right)\right),$$

where the 0 accounts for buyers and sellers who may not want to trade. Finally, the equilibrium prices $P(z)$ are the multipliers of the mass balance conditions **C3**.

4.1.2 The Dual Problem

Sometimes the dual form of the problem is easier to solve. The unknowns here are given by the three functions $\tilde{u}$, $\tilde{v}$, and P. We combine them in a *potential* function W over the set $\mathcal{N} \equiv \tilde{\mathcal{Y}} \cup \mathcal{Z} \cup \tilde{\mathcal{X}}$ of all nodes. More precisely, we define $W(\tilde{y}) = v(\tilde{y})$; $W(z) = P(z)$; and $W(\tilde{x}) = u(\tilde{x})$. Clearly, the value of the total joint surplus is the sum of the equilibrium utilities of all buyer and seller types:

$$\int_{\tilde{\mathcal{X}}} W(\tilde{x})\tilde{n}(d\tilde{x}) + \int_{\tilde{\mathcal{Y}}} W(\tilde{y})\tilde{m}(d\tilde{y}).$$

By construction, W is non-negative. The following inequalities hold everywhere:

$$\bar{U}(\tilde{x}, z) - u(\tilde{x}) \leq P(z) \leq v(\tilde{y}) - \bar{V}(\tilde{y}, z) \tag{4.1}$$

with equality on the left (resp. on the right) on the support of $\mu(d\tilde{x}, \cdot)$ (resp. $\mu^S(d\tilde{y}, \cdot)$). These inequalities can be rewritten as

$$W(\tilde{x}) + W(z) \geq \bar{U}(\tilde{x}, z) \text{ and } W(\tilde{y}) - W(z) \geq \bar{V}(\tilde{y}, z)$$

Now define a sign function over the set of all nodes as $\sigma(\tilde{x}) = 1$, $\sigma(z) = -1$ and $\sigma(\tilde{y}) = -1$. We can rewrite the inequalities on W as

$$\begin{aligned} \sigma(\tilde{x})W(\tilde{x}) - \sigma(z)W(z) &\geq \bar{U}(\tilde{x}, z) \\ \sigma(z)W(z) - \sigma(\tilde{y})W(\tilde{y}) &\geq \bar{V}(\tilde{y}, z). \end{aligned}$$

Let f be a function defined over the nodes of the network. In network flow terminology, the difference between the values of f at the two ends of a link $a \to b$ is called its *gradient* $(\nabla f)_{a \to b} = f(b) - f(a)$; it is a function over the links of the network. Since our links go from $\tilde{y}$ to z and from z to $\tilde{x}$, we can finally rewrite the inequalities as

$$\nabla(\sigma W) \geq \psi$$

where the function ψ is defined over the links of the network; it takes the value $\bar{U}(\tilde{x}, z)$ for a $z \to \tilde{x}$ link and $\bar{V}(\tilde{x}, z)$ for a $\tilde{y} \to z$ link.

Putting everything together, we obtain the equilibrium utilities $\tilde{u}, \tilde{v}$ and the equilibrium price function P from the function $W : \mathcal{N} \to \mathbb{R}$ that minimizes

$$\int_{\tilde{\mathcal{X}}} W(\tilde{x})\tilde{n}(d\tilde{x}) + \int_{\tilde{\mathcal{Y}}} W(\tilde{y})\tilde{m}(d\tilde{y})$$

under the constraints $W \geq 0$ and $\nabla(\sigma W) \geq \psi$. Moreover, the measures μ^D and μ^S are obtained as the multipliers of the gradient constraints for the corresponding links at the optimum.

4.2 Solving for equilibrium matchings

Applying convex duality to equation (2.10) gives us

$$\mu_{xy} = \frac{\partial \mathcal{F}}{\partial \Phi_{xy}}(\boldsymbol{\Phi}; \boldsymbol{n}, \boldsymbol{m}),$$

where $\mathcal{F} := (-\mathcal{E})^*$ is the Legendre–Fenchel transform of the opposite of the generalized entropy.

While this formula does give the equilibrium matching patterns for any separable matching problem, there are much better ways to solve for the equilibrium. Consider the logit model of Section 2.2.2 for instance. Let us rewrite (2.12) as

$$\mu_{xy} = \sqrt{\mu_{x0}\mu_{0y}} \exp(\Phi_{xy}/2).$$

Substituting into the adding up equations $\mu_{x0} + \sum_{y=1}^{Y} \mu_{xy} = n_x$ and $\mu_{0y} +$

$\sum_{x=1}^{X} \mu_{xy} = m_y$ gives

$$a_x^2 + a_x \sum_{y=1}^{Y} b_y \exp(\Phi_{xy}/2) = n_x \text{ for } x = 1, \ldots, X \tag{4.2a}$$

$$b_y^2 + b_y \sum_{x=1}^{X} a_x \exp(\Phi_{xy}/2) = m_y \text{ for } y = 1, \ldots, Y \tag{4.2b}$$

where we defined $a_x = \sqrt{\mu_{x0}}$ and $b_y = \sqrt{\mu_{0y}}$. This is a system of $(X+Y)$ quadratic equations in the $(X+Y)$ unknowns $(\boldsymbol{a}, \boldsymbol{b})$. While we know from our general analysis of separable matching that a solution exists and is unique[3], computing it is another matter. We describe here several solutions.

4.2.1 The IPFP Algorithm

Galichon and Salanié (2022) proposed the following algorithm:

1. choose some initial values $\boldsymbol{b}^{(0)}$;
2. taking $\boldsymbol{b} = \boldsymbol{b}^{(0)}$, solve the X quadratic equations (4.2a) for $\boldsymbol{a} = \boldsymbol{a}^{(1)}$;
3. taking $\boldsymbol{a} = \boldsymbol{a}^{(1)}$ as given, solve the Y equations (4.2b) for $\boldsymbol{b} = \boldsymbol{b}^{(1)}$;
4. iterate steps 2 and 3.

They proved that it converges globally to the values $(\boldsymbol{a}^*, \boldsymbol{b}^*)$ that generates the stable matching for the Choo and Siow (2006) model via $\mu_{x0} = (a_x^*)^2, \mu_{0y} = \left(b_y^*\right)^2$, and $\mu_{xy} = a_x^* b_y^* \exp(\Phi_{xy}/2)$.

This is in fact an implementation of the Iterative Projections Fitting Procedure (IPFP), which is also known as the Sinkhorn algorithm. The

[3]See also Decker, Lieb, McCann, and Stephens (2012) for the logit model.

underlying idea of IPFP is that to project a point on the intersection of several convex sets, one can project on the first set, then on the second, etc until convergence. Here the convex sets are represented by the adding up equations in the space of matching patterns. For the Choo and Siow (2006) model, the projection is done according to the Kullback-Leibler divergence. This can be generalized beyond the logit model, by projecting according to the Bregman divergence associated from the generalized entropy $\mathcal{E}$; see Galichon and Salanié (2022) for details[4].

IPFP is especially easy to implement and very fast in the logit model, as each quadratic equation has only one unknown and can be solved in closed form. It is not hard to extend to the heteroskedastic logit and to the nested logit[5].

4.2.2 Alternatives to IPFP

For some specifications, each of the projection steps in IPFP may be costly. In such cases, it can be simpler to minimize the following, convex function of $\boldsymbol{U}$:

$$\sum_{x=1}^{X} n_x G_x(\boldsymbol{U}_{x\cdot}) + \sum_{y=1}^{Y} m_y H_y(\boldsymbol{\Phi}_{\cdot y} - \boldsymbol{U}_{\cdot y}). \tag{4.3}$$

[4]In separable matching models, we derived the generalized entropy to account for the contribution of matching on unobservables to the total joint surplus. The literature on numerical techniques for optimal transport introduces the classical entropy as a regularization device, and also makes extensive use of the Sinkhorn algorithm (see Peyré and Cuturi, 2019).

[5]The Python package `cupid_matching` implements IPFP for these and other variants of the logit model—see Salanié (2023).

At the optimum, we have for each (x, y):

$$n_x \frac{\partial G_x}{\partial U_{xy}}(\boldsymbol{U}_{x\cdot}) = m_y \frac{\partial H_y}{\partial V_{xy}}(\boldsymbol{\Phi}_{\cdot y} - \boldsymbol{U}_{\cdot y}).$$

These are simply the first-order conditions we obtained in (2.8), which define the stable matching patterns[6].

If it is easier to compute the generalized entropy function $\mathcal{E}$ than the Emax functions G_x and H_y, then one can choose to maximize the social welfare directly by using (2.11):

$$\max_{\boldsymbol{\mu}} \left(\sum_{x=1}^{X} \sum_{y=1}^{Y} \mu_{xy} \Phi_{xy} + \mathcal{E}(\boldsymbol{\mu}; \boldsymbol{n}, \boldsymbol{m}) \right). \tag{4.4}$$

Both optimization programs (4.3) and (4.4) are unconstrained, globally convex, and generally well-behaved. Still, they have $X \times Y$ variables; this may be a large number in some applications. Sometimes their dimension can be reduced to $(X + Y)$. The Choo and Siow (2006) model is a case in point. Consider the following function:

$$F(\boldsymbol{u}, \boldsymbol{v}) = \sum_{x \in \mathcal{X}} n_x \left(u_x + e^{-u_x} - 1\right) + \sum_{y \in \mathcal{Y}} m_y \left(v_y + e^{-v_y} - 1\right) + 2 \sum_{\substack{x \in \mathcal{X} \\ y \in \mathcal{Y}}} \sqrt{n_x m_y} e^{\frac{\Phi_{xy} - u_x - v_y}{2}}.$$

The function F is globally convex. Galichon and Salanié (2022, section 4.2)

[6] If the specification allows for zero cells $\mu_{xy} = 0$ at the optimum, then one needs to replace $\boldsymbol{\Phi}_{\cdot y} - U_{\cdot y}$ with $\boldsymbol{\Phi}_{\cdot y} - U_{\cdot y} - d_{\cdot y}$, where the d_{xy} are slack variables that satisfy $d_{xy} \geq 0$ and $\mu_{xy} d_{xy} = 0$.

prove that it is minimized for the equilibrium values of the type-average utilities $\boldsymbol{u} = (u_x)$ and $\boldsymbol{v} = (v_y)$. The equilibrium matching patterns are then obtained from

$$\begin{aligned}
\mu_{x0} &= n_x \exp(-u_x) \\
\mu_{0y} &= m_y \exp(-v_y) \\
\mu_{xy} &= \sqrt{n_x m_y} \exp\left((\Phi_{xy} - u_x - v_y)/2\right).
\end{aligned}$$

Concluding Remarks

The array of methods available in this field has expanded at a rapid pace in the past ten years.

A dominant theme of our survey is that matching (or trade, in hedonic models) is partly driven by characteristics of the partners that are not observed by the econometrician. The separability assumption that underlies much empirical work only allows for matching on such unobservables in a restricted form. We do not yet have broadly applicable methods to deal with markets where the joint surplus from a match, or a trade, involves complementarities between unobservables.

The theoretical literature on matching is well ahead of the applied literature. Theorists have modeled matching on networks, complex supply chains, and two-sided platforms among other things (see Hatfield, Jagadeesan, Kominers, Nichifor, Ostrovsky, Teytelboym, and Westkamp (2023) for a recent survey). As these features often play a crucial role in industrial organization and international trade, we are looking forward to see much more empirical work based on matching these areas.

This trend is already apparent in international trade. Since Eaton and

Kortum (2002), a substantial literature in international trade[7] has embedded discrete choice methods into general equilibrium models. In this approach, the productivities of the various immobile factors determine the probability that a given factor is used in the production of a given good in a given country. The factor productivity shocks are assumed to be independent draws from a Fréchet distribution, which is the multiplicative analog of the logit model of Section 2.2.2. This results in a structure that is formally similar to (while more complex than) that of the Choo and Siow model.

This chapter assumed away two sources of frictions: asymmetric information and search frictions. We refer the reader to the survey by Liu (2025). Search frictions are, of course, at the center of the search-and-matching framework that dominates structural applied work in labor economics. The seminal contribution of Shimer and Smith (2000) emphasized endogenous sorting by heterogeneous agents under matching with transferable utility. Eeckhout (2018) surveys more recent work along these lines.

Finally, we showed in Section 4.1 that hedonic and matching models relate to *flows on networks*. There is a very large mathematically oriented literature; we believe that some of its tools could be used very fruitfully in economics.

[7]See Costinot and Vogel (2015) for a recent survey.

Bibliography

ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High wage workers and high wage firms," *Econometrica*, 67, 251–333.

AGARWAL, N. (2015): "An Empirical Model of the Medical Match," *American Economic Review*, 105, 1939–1978.

AGARWAL, N., AND P. SOMAINI (2023): "Empirical Models of Non-Transferable Utility Matching," in *Online and Matching-Based Market Design*, ed. by F. Echenique, N. Immorlica, and V. V. Vazirani. Cambridge University Press.

AHUJA, R., T. MAGNANTI, AND J. ORLIN (1993): *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall.

AZEVEDO, E., AND J. HATFIELD (2018): "Existence of Equilibrium in Large Matching Markets with Complementarities," Discussion paper, mimeo Wharton.

BAJARI, P., AND C. L. BENKARD (2005): "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy*, 113, 1239–76.

BAJARI, P., AND J. FOX (2013): "Measuring the Efficiency of an FCC Spectrum Auction," *American Economic Journal: Microeconomics*, 5, 100–146.

BECKER, G. (1965): "A Theory of the Allocation of Time," *Economic Journal*, 75, 493–517.

BECKER, G. (1973): "A theory of marriage, part I," *Journal of Political Economy*, 81, 813–846.

——— (1974): "A theory of marriage, part II," *Journal of Political Economy*, 82, S11–S26.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

BERRY, S., AND A. PAKES (2007): "The Pure Characteristics Demand Model," *International Economic Review*, 48, 1193–1225.

BOJILOV, R., AND A. GALICHON (2016): "Matching in closed-form: equilibrium, identification, and comparative statics," *Economic Theory*, 61, 587–609.

BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A distributional framework for matched employer employee data," *Econometrica*, 87, 699–739.

BONNET, O., A. GALICHON, Y.-W. HSIEH, K. O'HARA, AND M. SHUM (2022): "Yogurts choose consumers? Estimation of random-utility models via two-sided matching," *The Review of Economic Studies*, 89(6), 3085–3114.

Brown, J., and H. Rosen (1982): "On the Estimation of Structural Hedonic Price Models," *Econometrica*, 50, 765—-768.

Chernozhukov, V., A. Galichon, M. Henry, and B. Pass (2021): "Identification of hedonic equilibrium and nonseparable simultaneous equations," *Journal of Political Economy*, 129(3), 842–870.

Chiappori, P.-A., A. Galichon, and B. Salanié (2019): "On Human Capital and Team Stability," *Journal of Human Capital*, 13, 236–259.

Chiappori, P.-A., D. L. Nguyen, and B. Salanié (2019): "Matching with Random Components: Simulations," Columbia University mimeo.

Chiappori, P.-A., S. Oreffice, and C. Quintana-Domeque (2012): "Fatter Attraction: Anthropometric and Socioeconomic Characteristics in the Marriage Market," *Journal of Political Economy*, 120, 659–695.

——— (2020): "Erratum: Fatter Attraction: Anthropometric and Socioeconomic Characteristics in the Marriage Market," *Journal of Political Economy*, 128, 4673–5.

Chiappori, P.-A., B. Salanié, and Y. Weiss (2017): "Partner Choice, Investment in Children, and the Marital College Premium," *American Economic Review*, 107, 2109–67.

Choné, P., and F. Kramarz (2022): "Matching Workers' Skills and

Firms' Technologies: From Bundling to Unbundling," Discussion paper, CREST.

Choo, E., and A. Siow (2006): "Who Marries Whom and Why," *Journal of Political Economy*, 114, 175–201.

Ciscato, E., A. Galichon, and M. Goussé (2020): "Like Attract Like: A Structural Comparison of Homogamy Across Same-Sex and Different-Sex Households," *Journal of Political Economy*, 128, 740–781.

Corblet, P. (2022): "Education Expansion, Sorting, and the Decreasing Education Wage Premium," Discussion paper, Sciences Po Paris, mimeo.

Costinot, A., and J. Vogel (2015): "Beyond Ricardo: Assignment Models in International Trade," *Annual Review of Economics*, 7, 31–62.

Dagsvik, J. (1994): "Discrete and Continuous Choice, Max-Stable Processes, and Independence from Irrelevant Attributes," *Econometrica*, 62, 1179—-1205.

Daly, A., and S. Zachary (1978): "Improved Multiple Choice Models," in *Identifying and Measuring the Determinants of Mode Choice*, ed. by D. Henscher, and Q. Dalvi. Teakfields, London.

Decker, C., E. Lieb, R. McCann, and B. Stephens (2012): "Unique Equilibria and Substitution Effects in a Stochastic Model of the Marriage Market," *Journal of Economic Theory*, 148, 778–792.

Diamond, W., and N. Agarwal (2017): "Latent indices in assortative matching models," *Quantitative Economics*, 8, 685–728.

Dupuy, A., and A. Galichon (2014): "Personality traits and the marriage market," *Journal of Political Economy*, 122, 1271–1319.

Dupuy, A., A. Galichon, and Y. Sun (2019): "Estimating matching affinity matrix under low-rank constraints," *Information and Inference*, 8, 677–89.

Eaton, J., and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

Eeckhout, J. (2018): "Sorting in the Labor Market," *Annual Review of Economics*, 10, 1–29.

Ekeland, I., J. Heckman, and L. Nesheim (2004): "Identification and Estimation of Hedonic Models," *Journal of Political Economy*, 112, S60–S109.

Epple, D. (1987): "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products," *Journal of Political Economy*, 95, 59–80.

Epple, D., and H. Sieg (1999): "Estimating Equilibrium Models of Local Jurisdictions," *Journal of Political Economy*, 107, 645–81.

Fox, J. (2010): "Identification in Matching Games," *Quantitative Economics*, 1, 203–254.

——— (2018): "Estimating Matching Games with Transfers," *Quantitative Economics*, 8, 1–38.

FOX, J., C. YANG, AND D. HSU (2018): "Unobserved Heterogeneity in Matching Games with an Appplication to Venture Capital," *Journal of Political Economy*, 126, 1339–1373.

GALE, D., AND L. SHAPLEY (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–14.

GALICHON, A. (2016): *Optimal Transport Methods in Economics*. Princeton University Press.

GALICHON, A., AND M. HENRY (2011): "Set identification in models with multiple equilibria," *The Review of Economic Studies*, 78(4), 1264–1298.

GALICHON, A., S. KOMINERS, AND S. WEBER (2019): "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility," *Journal of Political Economy*, 127, 2875–2925.

GALICHON, A., AND B. SALANIÉ (2017): "The Econometrics and Some Properties of Separable Matching Models," *American Economic Review Papers and Proceedings*, 107, 251–255.

GALICHON, A., AND B. SALANIÉ (2019): "Labeling Dependence in Separable Matching Markets," mimeo Columbia.

GALICHON, A., AND B. SALANIÉ (2022): "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models," *The Review of Economic Studies*, 89, 2600—2629.

——— (2024): "Estimating Separable Matching Models," *Journal of Applied Econometrics*, 39, 1021–1044.

Goussé, M., N. Jacquemet, and J.-M. Robin (2017): "Marriage, Labor Supply, and Home Production," *Econometrica*, 85, 1873–1919.

Gozlan, N., C. Roberto, P.-M. Samson, and P. Tetali (2017): "Kantorovich duality for general transport costs and applications," *Journal of Functional Analysis*, 273, 3327–3405.

Graham, B. (2011): "Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers," in *Handbook of Social Economics*, ed. by J. Benhabib, A. Bisin, and M. Jackson. Elsevier.

——— (2014): "Errata on "Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers"," mimeo Berkeley.

Graham, B., G. Imbens, and G. Ridder (2014): "Complementarity and Aggregate Implications of Assortative Matching: A Nonparametric Analysis," *Quantitative Economics*, 5, 29–66.

Griliches, Z. (1961): "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change," in *The price statistics of the federal government*, pp. 173–196. National Bureau of Economic Research.

Guadalupe, M., V. Rappoport, B. Salanié, and C. Thomas (2024): "The Perfect Match: Assortative Matching in Mergers and Acquisitions," Columbia mimeo.

Hatfield, J. W., R. Jagadeesan, S. D. Kominers, A. Nichifor,

M. Ostrovsky, A. Teytelboym, and A. Westkamp (2023): "Generalized Matching: Contracts and Networks," in *Online and Matching-Based Market Design*, ed. by F. Echenique, N. Immorlica, and V. Vazirani. Cambridge University Press.

Heckman, J., and J. Scheinkman (1987): "The Importance of Bundling in a Gorman-Lancaster Model of Earnings," *Review of Economic Studies*, 54, 243–255.

Heckman, J.-J., R. Matzkin, and L. Nesheim (2010): "Nonparametric Identification and Estimation of Nonadditive Hedonic Models," *Econometrica*, 78, 1569–1591.

Kelso, A. S., and V. P. Crawford (1982): "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50, 1483—-1504.

Lancaster, K. (1966): "A New Approach to Consumer Theory," *Journal of Political Economy*, 74, 132–157.

Lindenlaub, I. (2017): "Sorting multidimensional types: theory and application," *Review of Economic Studies*, 84, 718—-789.

Liu, Q. (2025): "Matching under Asymmetric Information," in *Handbook of the Economics of Matching, volume I*, ed. by Y.-K. Che, P.-A. Chiappori, and B. Salanié. Elsevier.

Mailath, G., A. Postlewaite, and L. Samuelson (2013): "Pricing and investments in matching markets," *Theoretical Economics*, 8, 535–590.

Manski, C. F. (1975): "Maximum score estimation of the stochastic utility model of choice," *Journal of Econometrics*, 3, 205—228.

Matzkin, R. L. (2003): "Nonparametric estimation of nonadditive random functions," *Econometrica*, 71(5), 1339–1375.

Menzel, K. (2015): "Large Matching Markets as Two-Sided Demand Systems," *Econometrica*, 83, 897–941.

Monge, G. (1781): "Mémoire sur la théorie des déblais et des remblais," *Histoire de l'Académie Royale des Sciences de Paris*, pp. 6666–704.

Pakes, A. (2003): "A Reconsideration of Hedonic Price Indexes with an Application to PC's," *American Economic Review*, 93, 1578–1596.

Peyré, G., and M. Cuturi (2019): "Computational Optimal Transport: With Applications to Data Science," *Foundations and Trends in Machine Learning*, 11(5-6), 355–607.

Queyranne, M. (2011): personal communication.

Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 82, 34–55.

Roth, A. (1984): "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, 92, 991–1016.

Roth, A. (2015): *Who Gets What—and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt.

SALANIÉ, B. (2015): "Identification in Separable Matching with Observed Transfers," Columbia University mimeo.

——— (2023): "`cupid_matching`: a Python package to solve and estimate separable matching models," `https://pypi.org/project/cupid_matching/`.

SHAPLEY, L., AND M. SHUBIK (1972): "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–130.

SHIMER, R., AND L. SMITH (2000): "Assortative matching and Search," *Econometrica*, 68, 343–369.

SKANS, O. N., P. CHONÉ, AND F. KRAMARZ (2021): "When Workers' Skills Become Unbundled: Some Empirical Consequences for Sorting and Wages," Discussion paper, Uppsala University.

STRASSEN, V. (1965): "The existence of probability measures with given marginals," *The Annals of Mathematical Statistics*, 36(2), 423–439.

TINBERGEN, J. (1956): "On the Theory of Income Distribution," *Weltwirtschaftliches Archiv*, 77, 155–73.

VILLANI, C. (2003): *Topics in Optimal Transportation*. American Mathematical Society.

WAUGH, F. (1928): "Quality Factors Influencing Vegetable Prices," *Journal of Farm Economics*, 10, 185–196.

Williams, H. (1977): "On the Formulation of Travel Demand Models and Economic Measures of User Benefit," *Environment and Planning A*, 9, 285–344.

Williamson, D. (2019): *Network Flow Algorithms*. Cambridge University Press.