

YOGURTS CHOOSE CONSUMERS? ESTIMATION OF RANDOM-UTILITY MODELS VIA TWO-SIDED MATCHING

ODRAN BONNET^b, ALFRED GALICHON[†], YU-WEI HSIEH[◇], KEITH O'HARA[◇],
AND MATT SHUM[§]

ABSTRACT. The problem of *demand inversion* – a crucial step in the estimation of random utility discrete-choice models – is equivalent to the determination of stable outcomes in two-sided matching models. This equivalence applies to random utility models that are not necessarily additive, smooth, nor even invertible. Based on this equivalence, algorithms for the determination of stable matchings provide effective computational methods for estimating these models. For non-invertible models, the identified set of utility vectors is a lattice, and the matching algorithms recover sharp upper and lower bounds on the utilities. Our matching approach facilitates estimation of models that were previously difficult to estimate, such as the pure characteristics model. An empirical application to voting data from the 1999 European Parliament elections illustrates the good performance of our matching-based demand inversion algorithms in practice.

Keywords: random utility models, demand inversion, two-sided matching, discrete-choice demand models, partial identification, pure characteristics model

JEL Classification: C51, C60

Date: February 6, 2022 (First draft: 9/2015). Galichon gratefully acknowledges funding from a grant NSF DMS-1716489, and from a European Research Council (ERC) grant No. 866274. We thank Jeremy Fox, Xavier d'Haultefoeuille, Lars Nesheim, Ariel Pakes, Giovanni Compiani and participants in seminars at Amazon, Caltech, UC-Davis, Emory, Johns Hopkins, Indiana, LSE, Rochester (Simon), Stanford, UNC, USC, Yale, NYU CRATE conference, Banff Applied Microeconomics Conference, SHUFE Econometrics Conference, Toronto Intersections of Econometrics and Applied Micro Conference, WARP (Workshop of Applications of Revealed Preference) webinar, and UCL-Vanderbilt Conference on Econometrics and Models of Strategic Interactions for helpful comments. Alejandro Robinson-Cortes provided excellent research assistance.

^bINSEE/CREST, Paris; odran.bonnet@insee.fr

[†]New York University; ag133@nyu.edu

[◇]Amazon.com; yuweihsieh01@gmail.com and keith.ohara@nyu.edu. Hsieh and O'Hara's contributions to the paper reflect work done prior to their joining Amazon.

[§]California Institute of Technology; mshum@caltech.edu .

1. INTRODUCTION

Discrete choice models play a tremendous role in applied work in economics. In these models, an agent i characterized by a utility shock $\varepsilon_i \in \Omega$ must choose from a finite set of alternatives $j \in \mathcal{J}$ in order to maximize her utility. The random utility framework pioneered by McFadden et al. (1978) assumes that the utility $\mathcal{U}_{\varepsilon_i j}(\delta_j)$ that agent i gets from alternative j depends on δ_j , a systematic utility level associated with alternative j (possibly parameterized as a function of regressors such as characteristics of alternative j) which is identical across all agents, and a realization ε_i of agent i 's random utility shocks. The agent chooses the alternative yielding maximal utility:

$$\max_{j \in \mathcal{J}} \{\mathcal{U}_{\varepsilon_i j}(\delta_j)\}. \quad (1.1)$$

We focus on parametric random utility models, where the function $\mathcal{U} : (\varepsilon, j, \delta) \in \Omega \times \mathcal{J} \times \mathbb{R} \mapsto \mathcal{U}_{\varepsilon j}(\delta) \in \mathbb{R}$ as well as the distribution of ε in the population, denoted P , are known to the researcher. Particular instances of these models are Additive Random Utility Models (hereafter ARUMs), including Logit or Probit models, where $\Omega = \mathbb{R}^{\mathcal{J}}$, $\mathcal{U}_{\varepsilon_i j}(\delta_j) = \delta_j + \varepsilon_{ij}$, and ε follows a Gumbel or Gaussian distribution. However, our results extend to the more general class of Non-Additive Random Utility Models (NARUMs), in which $\mathcal{U}_{\varepsilon_i j}(\delta_j)$ is not (quasi-)linear in δ_j .

Under an assumption guaranteeing that agents are not indifferent between any pair of alternatives (see Assumption 2 below), we can define the vector-valued *demand map* $\sigma(\cdot)$, the j -th component ($j \in \mathcal{J}_0$) of which is defined as the probability that alternative j dominates all the other ones, given the vector of systematic utilities $(\delta_j)_{j \in \mathcal{J}}$:

$$\sigma_j(\delta) = P(\varepsilon : \mathcal{U}_{\varepsilon j}(\delta_j) \geq \mathcal{U}_{\varepsilon j'}(\delta_{j'}), \forall j' \in \mathcal{J}). \quad (1.2)$$

The main focus of the paper pertains to *demand inversion*: given a vector of observed market shares $(s_j)_{j \in \mathcal{J}}$, how can one characterize and compute the full set of utility vectors $(\delta_j)_{j \in \mathcal{J}}$ such that $s = \sigma(\delta)$ – that is, which rationalizes the observed market shares? (In cases where the distribution P of the utility shocks depends on parameters unknown to the researcher, which arises in many applications as well as in our simulations and application below, demand inversion refers to recovering δ_j for a given set of parameters of P .) Additionally, the demand map may also be *non-invertible*, which arises when multiple vectors δ solve the demand inversion problem.

1.1. Contribution. We establish a new equivalence principle between the problem of demand inversion and the problem of stable matchings in two-sided models with Imperfectly

Transferable Utility (ITU). More precisely, we show that a discrete choice model can always be interpreted as a two-sided matching market where consumers and alternatives are viewed as firms and workers; and that the *demand inversion problem*, that is the recovery of utility vectors $(\delta_j)_{j \in \mathcal{J}}$, can be reformulated as the *equilibrium problem* of determining competitive wages in the corresponding matching market. In other words, the identified set of solution vectors δ coincides with the set of equilibrium wages in the matching market. This equivalence implies two important contributions:

- (1) **Characterization of the identified set of δ .** The equivalence to the matching equilibrium implies that the identified set of vectors δ_j is a *lattice*, from which one can construct a very simple data-driven test for point-identification.¹ As such, if the greatest element of the lattice coincides with its smallest element, then the utility index δ is point-identified.² Thus, our approach bypasses the need of verifying *a priori* whether the parameters of a given model are point of partial-identified—a non-trivial exercise in many cases.
- (2) **Computation of the identified set of δ .** Our matching approach has two key features. First, the matching equivalence allows the utilization of several high-performance matching algorithms, for which the convergence properties are well-studied. The use of these matching algorithms for estimating random utility models is new; moreover, they can readily handle situations in which multiple values of δ rationalize the observed market shares. Second, these matching-based algorithms do not require the computation of the demand (market-share) mapping. This is important in specific models, such as the *pure characteristics model* (Berry and Pakes (2007)), which are notorious for their non-smooth market-share mappings. Indeed, using our approach, the demand inversion problem for the pure characteristics model becomes a well-behaved *convex program* (see Section 5 below).

1.2. Related literature. In this paragraph we discuss how our paper relates to the existing literature on (1) demand inversion and (2) two-sided matching.

¹A lattice is a partially ordered set that contains the meet and the join of each pair of its element. For the purposes of this paper, lattices are subsets of vectors in Euclidean space and, for any given pair of vectors, the meet (resp. join) is just the vector containing the componentwise infimum (resp. supremum). For additional discussion of lattices in matching theory, consult (Roth and Sotomayor 1992). Relatedly, Jia-Barwick (2008) exploits the lattice structure of equilibria in supermodular games to estimate a large multi-market entry game between discount retailers.

²Khan, Ouyang, and Tamer (2016) call this an “adaptive” property.

Demand inversion literature. Demand inversion is a crucial intermediate step for estimating aggregate discrete-choice models of product-differentiated markets; see, e.g., Berry (1994) and Berry, Levinsohn, and Pakes (1995) (BLP). It also plays an important role in two-step estimation procedures for dynamic discrete-choice models³ (including Hotz and Miller (1993), Aguirregabiria and Mira (2002), Bajari, Benkard, and Levin (2007), Arcidiacono and Miller (2011), Kristensen, Nesheim, and de Paula (2014)).

Theoretically, there is a large literature that tackles the problem of the invertibility of the demand map utilizing either *differentiability* or *monotonicity* of the demand map. The first approach, based on global univalence theorems, requires the *differentiability* of the mapping $\sigma(\delta)$.⁴ Such results were used by Chiappori and Komunjer (2009) and Kristensen, Nesheim, and de Paula (2014) to study (resp.) multinomial choice and nonadditive random utility models. The results in Gale and Nikaido (1965) focus on uniqueness, leaving existence aside: assuming that the Jacobian of σ is a P-matrix and that the domain is a rectangle, Gale-Nikaido's result guarantees the injectivity of σ —namely, that $\sigma^{-1}(\{s\})$ should have at most one point—but can be empty.

A second approach to demand inversion relies on *gross substitutes*—a form of *monotonicity*—that σ_j is decreasing, or at least nonincreasing, with respect to $\delta_{j'}$. This category of papers includes the literature on ARUMs, where this property holds automatically (eg. McFadden et al. (1978)). Hotz and Miller (1993) study the problem of the nonemptiness of $\sigma^{-1}(s)$, within general ARUMs; Berry (1994) provided a complete argument (which extends to nonadditive random utility models), and also shows uniqueness of the vector of systematic utilities under continuity conditions. Magnac and Thesmar (2002) investigate identification of structural parameters (period utility flows) in dynamic discrete choice models. Norets and Takahashi (2013) focus on the surjectivity of ARUMs, under the assumption of absolute continuity of the distribution of the additive utility shocks. More broadly, Berry, Gandhi, and Haile (2013) show injectivity in general demand systems with gross substitutes (thus going beyond random utility models) under a connected strong substitute assumption.

To date, the existing literature has not provided a general characterization of the identified utility set $\sigma^{-1}(s)$, defined as the set of utility vectors (δ_j) which rationalize a vector

³Specifically, even though we don't pursue the connection here, the demand inversion problem considered in this paper also applies straightforwardly to the problem in dynamic discrete-choice estimation of "inverting" the choice-specific value functions from the conditional choice probabilities.

⁴The Hadamard-Palais univalence theorem (Palais (1959)) asserts that (i) if $\sigma : \mathbb{R}^J \rightarrow \mathbb{R}^J$ is C^1 , (ii) if its Jacobian is invertible at all points, and (iii) if $\|\sigma(\delta)\| \rightarrow \infty$ as $\|\delta\| \rightarrow \infty$, then σ is globally invertible; further results are collected in Parthasarathy (1983) and Radulescu and Radulescu (1980).

of market shares (s_j) . This paper is the first to consider situations when the identified utility set $\sigma^{-1}(\{s\})$ is non-singleton. In addition, most of the papers cited above provide little guidance on computing the identified set.⁵ Besides a handful of models,⁶ there are no well-established procedures for demand-inversion in general (non-additive) random utility models with arbitrary error distributions. Our paper aims to fill this gap. In doing so, it builds on a set of recent papers which have reformulated the problem of demand inversion in ARUMs as an optimal transport problem, using the tools of convex duality. This approach was pioneered by Galichon and Salanié (2015), and was extended to ARUMs with possibly noncontinuous distributions of unobserved heterogeneity by Chiong, Galichon, and Shum (2016), and to continuous choice problems by Chernozhukov, Galichon, Henry, and Pass (2019). However, these papers do not cover nonadditive random utility models, and do not characterize the structure of the identified set.

Two-sided matching literature. We provide a brief and incomplete review of the two-sided matching literature, as one key result in this paper is to show the equivalence between demand inversion and a two-sided matching model. The matching literature is split between models with non-transferable utility (NTU), and transferable utility (TU), with intermediate cases called imperfectly transferable utility (ITU).⁷ A connection was made earlier between TU matching models and ARUMs (see Galichon and Salanié (2015) and Chiong, Galichon, and Shum (2016)), and in the present paper, we are making a novel connection between matching models with ITU and NARUMs.

A large class of algorithms for computing matching models consists of “deferred acceptance algorithms” which revolve around Tarski’s fixed point theorem, and interpreting stable matchings as fixed points of monotone mappings. They apply to NTU and ITU models. These ideas appeared in Adachi (2000), followed by Fleiner (2003), Echenique and Oviedo (2004) and Hatfield and Milgrom (2005). The seminal deferred acceptance algorithms (Gale and Shapley (1962); Crawford and Knoer (1981); Kelso and Crawford (1982)) can be interpreted in this way.

Other methods are “descent methods,” which rely on a reformulation of the problem as a convex optimization problem, and apply in the TU case. Some of the methods involve

⁵As Berry and Haile (2015, p. 10) underline, “(...) the invertibility result of Berry, Gandhi, and Haile (2013) is not a characterization (or computational algorithm) for the inverse”.

⁶These include the logit, nested logit, and random-coefficient logit models (Berry (1994), Berry, Levinsohn, and Pakes (1995), Dubé, Fox, and Su (2012)).

⁷A key reference for the matching literature (TU, NTU and ITU alike) is Roth and Sotomayor (1992), while Galichon (2016) focuses on the TU case, or equivalently, optimal transport methods.

coordinate descent (as the auction algorithm of Bertsekas and Castanon (1989)), while some involve gradient descent (as the semi-discrete approach of Aurenhammer (1987)). The linear programming solution to the optimal assignment problem described in Shapley and Shubik (1971) also belongs in this category. The study of rates of convergence of these algorithms have been the subject of intense study; see the book Peyré and Cuturi (2017) and Burkard, Dell'Amico, and Martello (2009) for results and further references.

1.3. Organization. Section 2 introduces the general random utility framework which is the focus of this paper and provides examples. Section 3 presents our main equivalence result between NARUMs and two-sided matching problems, and discusses the lattice structure of the identified utility set. Based on the equivalence result, in Section 4 we introduce several matching-based algorithms which can solve a wide variety of random utility models. Section 5 contains two simulation investigations of the algorithms, including the pure characteristics model. Section 6 utilizes our matching approach to estimate a spatial voting model using electoral data from the 1999 European Parliament elections. Section 7 concludes. Proofs and several additional theoretical results are collected in the appendix.

2. THE FRAMEWORK

2.1. Basic assumptions. Let $\mathcal{J}_0 = \mathcal{J} \cup \{0\}$ be a finite set of alternatives, where $j = 0$ denotes a special alternative which serves as a benchmark (see section 2.2 below). The agent's program is thus

$$u_{\varepsilon_i} = \max_{j \in \mathcal{J}_0} \{U_{\varepsilon_i j}(\delta_j)\}, \quad (2.1)$$

where u_{ε_i} is the indirect utility of an agent with shock ε_i . The utility agent i derives from alternative j depends on the systematic utility vector δ_j associated with this alternative, and on the realization ε_i of this agent's utility shock. We will work under two assumptions:

Assumption 1 (Regularity of \mathcal{U}). *Assume (Ω, P) is a Borel probability space and for every $\varepsilon \in \Omega$, and for every $j \in \mathcal{J}_0$:*

- (a) *the map $\varepsilon \mapsto (U_{\varepsilon j}(\delta_j))_{j \in \mathcal{J}_0}$ is measurable, and*
- (b) *the map $\delta_j \mapsto U_{\varepsilon j}(\delta_j)$ is increasing from \mathbb{R} to \mathbb{R} and continuous.*

Assumption 2 (No indifference). *For every distinct pair of indices j and j' in \mathcal{J}_0 , and for every pair of scalars δ and δ' ,*

$$P(\varepsilon \in \Omega : U_{\varepsilon j}(\delta) = U_{\varepsilon j'}(\delta')) = 0.$$

These two assumptions are standard in the literature, and are automatically satisfied in ARUMs. Assumption 1 (a) is a standard measurability condition, and (b) is often invoked in the literature on non-separable models (see, e.g., Matzkin (2007)).⁸

Assumption 2 rules out indifference (precisely, an event of measure zero) between two alternatives, and is maintained in practically all the applied discrete choice literature. In earlier versions of the paper (available from the authors upon request), we showed that the results in this paper hold even without Assumption 2, albeit at the greater notational expense of introducing set-valued functions. In the present paper, for simplicity, we maintain Assumption 2 as it suffices for our purposes. Under Assumption 2, the demand of alternative j (defined in (1.2) above) corresponds to the fraction of consumers who prefer weakly *or* strictly alternative j to any other one:

$$\sigma_j(\delta) := P\left(\varepsilon \in \Omega : \mathcal{U}_{\varepsilon j}(\delta_j) \geq \max_{j' \in \mathcal{J}_0} \mathcal{U}_{\varepsilon j'}(\delta_{j'})\right) = P\left(\varepsilon \in \Omega : \mathcal{U}_{\varepsilon j}(\delta_j) > \max_{j' \in \mathcal{J}_0 \setminus \{j\}} \mathcal{U}_{\varepsilon j'}(\delta_{j'})\right). \tag{2.2}$$

Importantly, the uniqueness of the vector of market shares associated with a given utility vector δ does not imply that the demand inversion problem has a unique solution. There may be multiple vectors δ such that $\sigma(\delta) = s$. Under Assumptions 1 and 2, the vector of market shares $s = \sigma(\delta)$ is a probability vector on \mathcal{J}_0 , which prompts us to introduce \mathcal{S}_0 , the set of such probability vectors as

$$\mathcal{S}_0 := \left\{ s \in \mathbb{R}_+^{\mathcal{J}_0} : \sum_{j \in \mathcal{J}_0} s_j = 1 \right\}.$$

We formalize the definition of the demand map.

Definition 1 (Demand map). Under Assumption 1 and 2, the *demand map* is the map $\sigma : \mathbb{R}^{\mathcal{J}_0} \rightarrow \mathcal{S}_0$ defined by expression (2.2).

2.2. Normalization. Any discrete choice model requires some normalization, because the choice probabilities result from the comparison of the relative utility payoffs from each

⁸Assumption 1 also circumscribes the set of models we consider here. Specifically, it rules out purely "horizontal" (Hotelling) choice models, where the utility of consumer ε for store j is $\mathcal{U}_{\varepsilon j}(\delta_j) = -|\delta_j - \varepsilon|$, interpreted as the (minus the) absolute distance between the consumer's home location ε and the store's location δ_j where $\delta_j, \varepsilon \in [0, 1]$. This specification violates the monotonicity condition (Assumption 1(b)). Interestingly, it turns out the *quadratic* version of this model can be reparametrized in a way which satisfies our modelling assumptions; see Section 6 below.

alternative. Throughout the paper we normalize the systematic utility associated to the default alternative to zero:

$$\delta_0 = 0, \tag{2.3}$$

and we use

$$\tilde{\sigma} : \mathbb{R}^{\mathcal{J}} \rightarrow \mathbb{R}^{\mathcal{J}} \tag{2.4}$$

to denote the map induced by this normalization.

In the special ARUM case where $\mathcal{U}_{\varepsilon_{ij}}(\delta_j) = \delta_j + \varepsilon_{ij}$, imposing normalization (2.3) is innocuous because the vector of systematic utilities (δ_j) yields the same choice problem as the vector $(\delta_j + c)$ where c is a constant; hence, for ARUMs, any normalization will yield the same identified utility vectors δ up to an additive constant. However, this is no longer true in nonadditive models, for which the normalization (2.3) entails some loss of generality.⁹ We explore this below in Section 5.

2.3. Examples. Next, we consider several examples of random utility models falling within our framework. Since we assume that market shares are generated by the mapping in Eq. (2.2), we implicitly assume that the random element ε is independently and identically distributed across all consumers in the market. However, we make no restrictions on ε across alternatives: as the examples below show, ε can be individual-specific, choice-specific, or some combination of the two.

Example 2.1 (ARUM). In the additive random utility model (ARUM) one sets, $\Omega = \mathbb{R}^{\mathcal{J}_0}$, so that P is a probability distribution on $\mathbb{R}^{\mathcal{J}_0}$, and

$$\mathcal{U}_{\varepsilon_{ij}}(\delta_j) = \delta_j + \varepsilon_{ij}.$$

There are several well-known instances of ARUMs:

Logit model: if P is the distribution of a vector of size $|\mathcal{J}_0|$ of i.i.d. type 1-Extreme value random variables, then the demand map is given by $\sigma_j(\delta) = \exp(\delta_j) / \left(\sum_{j' \in \mathcal{J}_0} \exp(\delta_{j'}) \right)$.¹⁰ The demand map is analytically invertible, and yields the familiar “log-odds ratio” formula:

⁹While it is possible to reparameterize any NARUM into an ARUM $\mathcal{U}_{\varepsilon_j}(\delta) = \gamma_j + \eta_j$ by defining $\gamma_j \equiv \int \mathcal{U}_{\varepsilon_j}(\delta) dP_{\varepsilon}$ and $\eta_j = \mathcal{U}_{\varepsilon_j}(\delta) - \gamma_j$, it is not clear whether this reparameterization simplifies the demand inversion problem of recovering δ , which is the main point of our paper (in some cases, the reparameterization obfuscates the demand inversion problem).

¹⁰For the logit model, the use of aggregate shares in forming moment conditions has been long recognized in the literature, going back to Berkson (1955); Berry (1994) and Berry, Levinsohn, and Pakes (1995) showed that this connection applies more generally to static discrete-choice models beyond the simple logit model.

$$\delta_j = \log(s_j/s_0). \quad (2.5)$$

Pure characteristics model: In this model, consumers value product j only through its measurable characteristics $x_j \in \mathbb{R}^d$, a vector of dimension d associated to each alternative j , and the utility shock vector ε_i is such that

$$\varepsilon_{ij} = \nu_i^\top x_j = \sum_{k=1}^d \nu_i^k x_j^k \quad (2.6)$$

where ν_i is consumer i 's vector of taste-shifters, drawn from a distribution P_ν on \mathbb{R}^d . In this case, there is no closed-form expression for the demand map¹¹. We shall revisit this model in the simulations and empirical application in sections 5 and 6.

The case where $d = 1$ and there is only one characteristic, the price p_j , is the *vertical differentiation* or *quality ladder* model¹², which has the utility specification

$$\mathcal{U}_{\varepsilon_j}(\delta_j) = \delta_j - \nu_i p_j, \quad \forall j.$$

Here δ_j is interpreted as the quality of brand j , while the nonlinear random utility shock ν_i measures household i 's willingness-to-pay for quality. Below, in Section 5, we consider a numerical example based on this framework which is non-invertible.

Random coefficient logit model: In the random coefficient logit model popularized by BLP and McFadden and Train (2000), the random-utility shock is given by:

$$\varepsilon_{ij} = \nu_i^\top x_j + \zeta_{ij}.$$

This is the sum of two independent terms: one logit term ζ_{ij} and one pure characteristics term $\nu_i^\top x_j$.

Example 2.2 (Risk aversion). Consider a market where consumers are not fully aware of the attributes of a product at the time of purchase. This may characterize consumers' choices in online markets, where they have no opportunity to physically examine the goods under consideration. Let ε_i denote the relative risk aversion parameter (under CRRA utility), and that the price of good j is p_j . Choosing option j yields a consumer surplus of $\delta_j - p_j + \eta_j$ where $\log \eta_j \sim N(0, 1)$ is a quality shock unobservable at the time of the

¹¹See Song (2007) and Nosko (2010) for two empirical applications of the pure characteristics demand model. Pang, Su, and Lee (2015) provide computational algorithms for estimating this model.

¹²See, among others, Prescott and Visscher (1977), Bresnahan (1981), Esteban and Shum (2007).

purchase, and δ_j is the willingness to pay (in dollar terms) associated to alternative j . At the time of the purchase, the consumer's expected utility is

$$\mathcal{U}_{\varepsilon_i j}(\delta_j) = \mathbb{E}_{\eta_j} \left[\frac{(\delta_j - p_j + \eta_j)^{1-\varepsilon_i}}{1 - \varepsilon_i} \right],$$

where the expectation is taken over η_j holding ε_i constant. These kind of models are typically non-additive in ε .¹³

3. EQUIVALENCE OF DISCRETE-CHOICE AND TWO-SIDED MATCHING

In this section, we show a central result of this paper; namely, an equivalence between discrete-choice models and two-sided matching problems. This equivalence is noteworthy as discrete choice problems are traditionally considered to be “one-sided” problems. However, we will demonstrate that they are equivalent to a two-sided “marriage problem” between consumers and yogurts, where both sides of the market must assent to be matched.

Specifically, our main theorem demonstrates an equivalence between the discrete-choice framework described in the previous section and a two-sided matching market with imperfectly transferable utility (see Galichon, Kominers, and Weber (2019)). A consequence of this equivalence is that the demand inversion problem for estimating discrete-choice models can be equivalently formulated as solving for equilibrium utility payoffs from the corresponding two-sided matching problem, a well-understood exercise for which a number of algorithms are available.

We begin by formally defining the object of interest for demand inversion, which is to recover the identified utility set:

Definition 2 (Identified utility set). Given a demand map $\tilde{\sigma}$ defined as in (2.4) where Assumptions 1 and 2 are met, and given a vector of market shares s that satisfies $s_j > 0$ and $\sum_{j \in \mathcal{J}_0} s_j = 1$, the *identified utility set* associated with s is defined by

$$\tilde{\sigma}^{-1}(s) = \{\delta \in \mathbb{R}^{\mathcal{J}} : \tilde{\sigma}(\delta) = s\}. \quad (3.1)$$

Requiring non-zero market shares is a standard assumption in demand inversion of discrete-choice models; see, e.g., Lemma 1 of Berry and Haile (2014).¹⁴

¹³See Cohen and Einav (2007) and Apesteguia and Ballester (2014) for examples.

¹⁴In the ARUM case, a sufficient condition for this is that ε has full-support (a nowhere vanishing density) on $\mathbb{R}^{\mathcal{J}_0}$ (see, e.g., Galichon and Hsieh (2019)), which is satisfied in models such as logit, probit, and mixed-logit. However, it may fail to hold in, e.g., pure characteristic models (see, e.g., Song (2007)), without

3.1. The Equivalence Theorem. Next we introduce a matching game between consumers and yogurts, which is essentially that of Demange and Gale (1985); our presentation of this model is inspired by the presentation in chapter 9 of Roth and Sotomayor (1992)¹⁵. In this matching model, one side of the market consists of a continuum of consumers, distinguished by type ε , while the other side is an equi-massed continuum of jars of yogurt, distinguished by alternative (brand) $j \in \mathcal{J}_0$. We start by introducing some terminology. Let $\mathcal{M}(P, s)$ be the set of probability distributions on $\Omega \times \mathcal{J}_0$ with marginal distributions P and s ; namely, $\pi \in \mathcal{M}(P, s)$ if and only if $\pi(B \times \mathcal{J}_0) = P(B)$ for all B (Borel-measurable subsets of Ω), and $\pi(\Omega \times \{j\}) = s_j$ for all $j \in \mathcal{J}_0$.

Let u_ε (resp. v_j) denote utility payoffs for each consumer ε (resp. yogurt j) from the matching game. These payoffs are endogenously determined in equilibrium, as will be clear below. Let $f_{\varepsilon j}(u)$ be the transfer (positive or negative) needed by a consumer ε in order to reach utility level $u \in \mathbb{R}$ when matched with a yogurt j . Symmetrically, let $g_{\varepsilon j}(v)$ be the transfer needed by a yogurt j in order to reach utility level $v \in \mathbb{R}$ when matched with a consumer ε . (The connection between f and g and the primitives of the discrete choice model will be clarified below, in Eq. (3.2).) The functions $f_{\varepsilon j}(\cdot)$ and $g_{\varepsilon j}(\cdot)$ are assumed increasing for every ε and j .

We describe this game in the context of a dance party where consumers (indexed by ε) seek out jars of yogurt j to dance with. Each consumer ε charges a price of u_ε utils for dancing; similarly, yogurt j charges a price of v_j for dancing. Thus, a dance between consumer ε and yogurt j involves a payment of $g_{\varepsilon j}(v_j)$ from consumer ε to yogurt j , and a payment of $f_{\varepsilon j}(u_\varepsilon)$ from the yogurt to the consumer. These transfers can be negative: for instance, if yogurt j already provides consumer ε with utility exceeding u_ε by dancing with her, then $f_{\varepsilon j}(u_\varepsilon) < 0$ and involves a payment from consumer ε to yogurt j .

Definition 3 (Equilibrium outcome). An equilibrium outcome in the matching problem is an element (π, u, v) , where π is a joint probability measure on $\Omega \times \mathcal{J}_0$, u and v are Borel-measurable functions on (Ω, P) and (\mathcal{J}_0, s) respectively, such that:

- (i) π has marginal distributions P and s : $\pi \in \mathcal{M}(P, s)$.

sufficient variation in the choice-specific characteristics. Moreover, we assume throughout that the demand model is correctly specified, so that the identified utility set in (3.1) is non-empty.

¹⁵Demange and Gale's model is discrete and extends the model of Shapley and Shubik (1971) beyond the transferable utility setting. See also Crawford and Knoer (1981), Kelso and Crawford (1982), Hatfield and Milgrom (2005). We formulate a slight variant here in that we (1) we allow for multiple agents per type and (2) do not allow for unmatched agents. However, this leaves analysis essentially unchanged.

- (ii) there is no blocking pair: $f_{\varepsilon j}(u_\varepsilon) + g_{\varepsilon j}(v_j) \geq 0$ for all $\varepsilon \in \Omega$ and $j \in \mathcal{J}_0$.
- (iii) pairwise feasibility holds: if $(\varepsilon, j) \in \text{Supp}(\pi)$, then $f_{\varepsilon j}(u_\varepsilon) + g_{\varepsilon j}(v_j) = 0$.

Condition (i) implies that if a random vector (ε, j) has distribution $\pi \in \mathcal{M}(P, s)$, then $\varepsilon \sim P$ and $j \sim s$. Hence, π is interpreted as the probability distribution that a consumer with utility shock ε is matched with a yogurt of type j ; in other words, $\pi(j|\varepsilon)$ denotes the conditional probability that an individual with utility shock ε chooses yogurt j , which is degenerate (equaling 0 or 1) under Assumption 2. To understand condition (ii), consider that if there exists a consumer ε and a yogurt of type j for which $f_{\varepsilon j}(u_\varepsilon) + g_{\varepsilon j}(v_j) < 0$, then there exists $u' > u_\varepsilon$ and $v' > v_j$ such that $f_{\varepsilon j}(u') + g_{\varepsilon j}(v') = 0$. In other words, (u', v') are feasible for (ε, j) and strictly improve upon the equilibrium payoffs u_ε and v_j , which is ruled out in equilibrium. The feasibility condition (iii) rules out matchings involving net positive transfers ($f_{\varepsilon j}(u_\varepsilon) + g_{\varepsilon j}(v_j) > 0$) which, intuitively, are those where one (or both) of the agents have equilibrium payoffs which are inachievably high compared to the utility they supply each other from matching.

The next theorem establishes that the demand inversion problem is equivalent to a matching problem. The proofs for this and all subsequent claims are in the appendix.

Theorem 1 (Equivalence theorem). *Under Assumptions 1 and 2, consider a vector of market shares s that satisfies $s_j > 0$ and $\sum_{j \in \mathcal{J}_0} s_j = 1$. Consider a vector $\delta \in \mathbb{R}^{\mathcal{J}}$. Then, the two following statements are equivalent:*

(i) δ belongs to the identified utility set $\tilde{\sigma}^{-1}(s) = \{\delta \in \mathbb{R}^{\mathcal{J}} : \tilde{\sigma}(\delta) = s\}$ associated with the market shares s in the sense of Definition 2 in the discrete choice problem with $\varepsilon \sim P$;

(ii) there exists $\pi \in \mathcal{M}(P, s)$ and $u_\varepsilon = \max_{j \in \mathcal{J}_0} \mathcal{U}_{\varepsilon j}(\delta_j)$ such that $(\pi, u, -\delta)$ is an equilibrium outcome in the sense of Definition 3 in the matching problem with transfer functions

$$f_{\varepsilon j}(u) = u \text{ and } g_{\varepsilon j}(-\delta) = -\mathcal{U}_{\varepsilon j}(\delta). \quad (3.2)$$

We offer some intuition of the equivalence here. A matching equilibrium requires that, given the transfer functions, the equilibrium transfers be set such that *both* sides of the market (consumers and yogurts) are happy with their matched partners. On the consumer side, consumer ε seeks the yogurt j which offers her the highest payoff \mathcal{U}_ε . On the yogurt side, each yogurt j also seeks to maximize its payoff, which is equivalent to *minimizing* the mean utility δ_j that consumers receive from them; for that reason, we have the yogurt's payoff $v_j = -\delta_j$. To understand the intuition for the transfer functions (3.2), note that if

we plug these into Definition 3, the no-blocking pair condition becomes

$$\forall \varepsilon, j : u_\varepsilon = \max_{j' \in \mathcal{J}_0} \mathcal{U}_{\varepsilon j'}(\delta_{j'}) \geq \mathcal{U}_{\varepsilon j}(\delta_j); \tag{3.3}$$

and the feasibility pair condition becomes

$$\text{if } (\varepsilon, j) \in \text{Supp}(\pi) : u_\varepsilon = \max_{j' \in \mathcal{J}_0} \mathcal{U}_{\varepsilon j'}(\delta_{j'}) = \mathcal{U}_{\varepsilon j}(\delta_j) \tag{3.4}$$

which correspond to the conditions characterizing optimal consumer choices in the discrete-choice problem.

Our equivalence result states that the identified set of utilities in the discrete-choice demand problems corresponds to the equilibrium set of some matching problem. This matching equivalence result also raises the possibility of multiplicity of the identified set of utilities. Assumption 2 implies a unique demand map (Eq. (1.2)), and hence a unique allocation in the matching problem. However, just as in Shapley and Shubik (1971), there may be multiple payoffs (corresponding to the δ 's here) which support the equilibrium allocation. We will return to this below.

In the special case when the random utility model is additive (ARUM) as in example 2.1, one has $f_{\varepsilon j}(u) = u$ and $g_{\varepsilon j}(-\delta) = -\mathcal{U}_{\varepsilon j}(\delta) = -\varepsilon_j - \delta_j$, so that the stability conditions become $u_\varepsilon + v_j \geq \varepsilon_j$ with equality for $(\varepsilon, j) \in \text{Supp}(\pi)$. As noted initially in Galichon and Salanié (2015), this problem is now equivalent to a matching problem with transferable utility, where the joint surplus of a match between a consumer ε and a yogurt j is ε_j . We will discuss these in more detail in section 4.2 below.

3.2. Lattice structure of the identified utility set. Next, we show that the set-valued function $s \rightarrow \tilde{\sigma}^{-1}(s)$ is isotone¹⁶ (in a sense to be made precise) and that $\tilde{\sigma}^{-1}(s)$ has a lattice structure.

The lattice structure is very useful for deriving algorithms for computing the identified utility set $\tilde{\sigma}^{-1}(s)$ when demand is non-invertible. The literature on the estimation of discrete choice models has favored an approach based on imposing conditions guaranteeing invertibility of demand, or equivalently situations in which $\tilde{\sigma}^{-1}(s)$ is restricted to a single point. In particular, Berry, Gandhi, and Haile (2013) (hereafter BGH) provide conditions under which $\tilde{\sigma}^{-1}(s)$ should contain at most one point, from which it also follows that the map $s \rightarrow \tilde{\sigma}^{-1}(s)$ is isotone on its domain. In contrast, our approach here imposes minimal assumptions, and one must consider *non-invertible* models, in which the demand map

¹⁶Isotone has the meaning of (monotone) increasing, and is a standard term used in the literature on lattices (see Topkis (1998) for a reference on lattices and isotonicity in economics).

(1.2) is not one-to-one and $\tilde{\sigma}^{-1}(s)$ is a set. In this case, we need to generalize the notion of isotonicity which applies to the identified set $\tilde{\sigma}^{-1}(\cdot)$. The next theorem states that the correct generalization is the notion of isotonicity with respect to *Veinott's strong set order*.¹⁷ For the following, recall the lattice “join” and “meet” operators (\wedge and \vee) are defined by $(\delta \wedge \delta')_j := \min \{ \delta_j, \delta'_j \}$ (componentwise minimum) and $(\delta \vee \delta')_j := \max \{ \delta_j, \delta'_j \}$ (componentwise maximum).

Theorem 2. *The set-valued function $s \rightarrow \tilde{\sigma}^{-1}(s)$ is isotone in Veinott's strong set order, i.e. if $\delta \in \tilde{\sigma}^{-1}(s)$ and $\delta' \in \tilde{\sigma}^{-1}(s')$ with $s \leq s'$, then $\delta \wedge \delta' \in \tilde{\sigma}^{-1}(s)$ and $\delta \vee \delta' \in \tilde{\sigma}^{-1}(s')$.*

In the special case where $s \rightarrow \tilde{\sigma}^{-1}(s)$ is a singleton, we recover the isotonicity of the inverse demand as in BGH. Going further, by taking $s = s'$ in Theorem 2, we obtain that, whenever it is non-empty, the set $\tilde{\sigma}^{-1}(s)$ is a *lattice*¹⁸:

Corollary 1. *Under Assumption 1 and 2, if $\tilde{\sigma}^{-1}(s)$ is non-empty, it is a lattice. That is, if $\delta, \delta' \in \tilde{\sigma}^{-1}(s)$, then both $(\delta \wedge \delta'), (\delta \vee \delta') \in \tilde{\sigma}^{-1}(s)$.*

This result is similar to Demange and Gale (1985), who showed that the set of payoffs which ensures a stable allocation is a lattice whenever it is non-empty. This implies that the set of identified utilities has a “maximal” (resp. “minimal”) element which is composed of the *component-wise* upper- (resp. lower-) bounds among all the utility vectors in $\tilde{\sigma}^{-1}(s)$. The upper bound corresponds to the unanimously most preferred stable allocation for the consumers (“consumer-optimal”) and the unanimously least preferred stable allocation for the yogurts; conversely, the lower bound corresponds to the unanimously most preferred stable allocation for the yogurts (“yogurt-optimal”) and least preferred for the consumers. Formally, define

$$\tilde{\delta}_j^{\min}(s) = \min \{ \delta_j : \delta \in \tilde{\sigma}^{-1}(s) \} \text{ and } \tilde{\delta}_j^{\max}(s) = \max \{ \delta_j : \delta \in \tilde{\sigma}^{-1}(s) \}.$$

Then the lattice property implies:

- (i) The set $\tilde{\sigma}^{-1}(s)$ has a minimal and a maximal element:

$$\tilde{\delta}^{\min}(s) \in \tilde{\sigma}^{-1}(s) \text{ and } \tilde{\delta}^{\max}(s) \in \tilde{\sigma}^{-1}(s).$$

¹⁷See e.g. Veinott (2005). Veinott's strong set order provides an ordering over sets. Let \mathcal{X} and \mathcal{X}' be two subsets in \mathbb{R}^d ; we say that $\mathcal{X} < \mathcal{X}'$ in the Veinott strong set order iff $\forall x \in \mathcal{X}, x' \in \mathcal{X}'$, the “join” (or componentwise minimum) $x \wedge x' \in \mathcal{X}$ and the “meet” (or componentwise maximum) $x \vee x' \in \mathcal{X}'$.

¹⁸Whether the identified set is empty is considered in the Section A.

(ii) Any $\delta \in \tilde{\sigma}^{-1}(s)$ is such that

$$\tilde{\delta}^{\min}(s) \leq \delta \leq \tilde{\delta}^{\max}(s).$$

(iii) $\tilde{\sigma}^{-1}(s)$ is a singleton if and only if

$$\tilde{\delta}^{\min}(s) = \tilde{\delta}^{\max}(s).$$

Practically, most applications of partially identified models focus on computing the component-wise upper and lower bounds of the identified set of parameters; for general partially identified models, the vector of component-wise bounds will typically lie *outside* the (joint) identified set of parameters. In contrast, our lattice result here implies that these component-wise upper and lower bounds constitute *sharp* upper and lower bounds for the parameter vector as a whole, in the sense that they are attainable for selection mechanisms which place all probability on the highest (for upper bound) or lowest (for lower bound) payoffs for consumers.¹⁹

In addition, the matching literature provides algorithms to compute these extremal elements, which can be directly used to assess multiplicity: indeed, $\tilde{\sigma}^{-1}(s)$ is a single element (point-identified) if and only if its minimal and maximal elements coincide. This flexibility in handling models for which the researcher may not know *a priori* whether model parameters are point- or partially-identified is an important contribution of the matching approach developed in this paper. We turn to these algorithms next.

4. MATCHING-BASED ALGORITHMS

The equivalence established in Theorem 1 between matching and discrete-choice models allows us to leverage several matching algorithms for both NARUMs and ARUMs. The use of these algorithms in the empirical discrete-choice literature is new. In addition, matching-based algorithms have several advantages over existing procedures: (i) they can handle non-invertibility of the demand map, as all of these algorithms allow for the case where multiple utility vectors can rationalize the same set of market shares; and (ii) these algorithms do not require smoothness of the demand map and therefore can handle some well-known models—like the pure characteristics model—which have non-smooth demand

¹⁹In addition, the upper- and lower-lattice bounds here are the “widest” possible in the sense that for each j , no utility lower than δ_j^{\min} , or higher than δ_j^{\max} , can rationalize the observed set of market shares. Moreover, the equivalence theorem also implies that the identified utility set $\sigma^{-1}(s)$, while not necessarily convex, is connected, which derives from the properties of the core of matching games (section 9.2 of Roth and Sotomayor (1992)).

maps. In contrast, existing algorithms for demand inversion often rely on directly solving the demand map $s_j = \sigma_j(\delta)$ for δ using fixed-point iterations or nonlinear-equation solvers, which typically requires smoothness of the demand map, and also rules out non-invertibility of the demand map.

We introduce three matching-based algorithms in this section. The first is an algorithm for matching models with imperfectly transferable utility (ITU) which can be used for demand inversion in both ARUMs or NARUMs. This algorithm, called *Market Share Adjustment*, is essentially an “accelerated” version of the classic deferred acceptance algorithm (Gale and Shapley (1962), Crawford and Knoer (1981)), which iteratively adjusts the payoffs of the potential partners to achieve equilibrium. The second and third algorithms, in Section 4.2, are methods for computing stable allocations in two-sided matching models with transferable utility (TU), and apply only to ARUMs. We consider a *linear-programming* approach based on the classic Shapley and Shubik (1971) assignment game, and a version of the *auction algorithm* of Bertsekas (1992), augmented to produce bounds for partially-identified settings.

While the model in this paper assumes a continuum of agents on each side of the market, for computational purposes we approximate this with a finite market populated by an equal (and large but finite) number, denoted N , of consumers and jars of yogurt.²⁰ On the consumer side, each consumer $i \in \{1, \dots, N\}$ is characterized by a value of the utility shock ε_i drawn i.i.d. (across i) from P , the distribution of the utility shocks. For a given vector of market shares (s_0, s_1, \dots, s_J) , the number of jars of each brand j of yogurt are set proportionately to the observed market share; that is, $m_j \approx Ns_j \in \mathbb{N}$ of yogurts of type j and, if needed, m_j has been rounded to an adjacent integer so that $\sum_{j \in \mathcal{J}_0} m_j = N$. Throughout we maintain the utility normalization $\delta_0 = 0$.

4.1. Deferred-acceptance algorithm (for both NARUMs and ARUMs). Theorem 1 establishes an equivalence between the identified utility set and equilibrium payoffs in a two-sided matching game with imperfectly transferable utility. Hence, for computing the identified utilities one could use the deferred-acceptance algorithms developed in Crawford and Knoer (1981) and Kelso and Crawford (1982) which are generalizations of Gale and Shapley’s (1962) deferred-acceptance algorithm. But these algorithms are very slow and

²⁰In Appendix C we present the *semi-discrete* algorithm for ARUMs, which is exact in that it computes the continuum problem directly. However, as we explain there, its use requires the unobserved taste vector to be (jointly) uniformly distributed over a polyhedron, which limits its general application.

inefficient, especially in the common situation where there are fewer products (i.e. “brands of yogurt”) than consumers.

4.1.1. *Market shares adjusting algorithm (MSA)*. As with all deferred-acceptance algorithms, there are two versions of the algorithm – the “consumer-proposing” and “yogurt-proposing” versions – return (resp.) the lattice upper bound or lower bound on the utility parameters. In the case when the model is point identified, the upper bound and lower bound will coincide; hence running both versions of this algorithm yields a data-driven assessment of whether the model is point- or partially-identified.

In the “consumer proposing” version, the utilities (δ 's) start at a high level, and consumers choose, in successive rounds, jars of yogurts which maximize their utilities. Between rounds, the utilities pertaining to the brands of yogurts in excess demand (ie. chosen by more consumers than available jars) are decreased by an adjustment factor. Bidding continues until a round is reached where the reference brand of yogurt ($j = 0$) is in excess demand: that is, when the number of consumers choosing jars of brand 0 is greater or equal to the number of its available jars. In the original Kelso and Crawford (1982) version, the adjustment is done for *each jar* of yogurt separately, leading to very slow convergence with a large number of consumers and jars. The MSA algorithm speeds this up by adjusting the utilities for *all jars of the same brand* of yogurt simultaneously. Since this accelerated process can lead to “overshooting” (in which utilities move below their equilibrium values), the MSA algorithm involves running a deferred-acceptance procedure multiple times with successively smaller adjustment factors.²¹

We present below the pseudo code for the consumer-proposing version of the MSA algorithm, which obtains the lattice upper bound; Appendix B.6 contains a version of the algorithm which yields the lattice lower bound. Define $\bar{\delta}_j = \sup_{i \in \{1, \dots, N\}} \mathcal{U}_{\varepsilon_i j}^{-1}(\mathcal{U}_{i0}(\delta_0))$. Clearly, $\bar{\delta}$ is an upper bound for the stable payoffs and for the lattice upper bound. Let η^{tol} denote a small adjustment factor, which is a design parameter for the algorithm.

²¹This adjustment factor plays a role analogous to the step size in optimization procedures. One typically chooses a larger step size in initial *exploration* phases to move parameters away from regions where the optimum is unlikely to be. Choosing a larger step size speeds up the routine, but if the step size is too large, one might miss (“overshoot”) the optimum. Therefore, in later *exploitation* phases, one decreases the step size to achieve a better accuracy. Similar heuristics are used also to set the temperature parameter in simulated annealing, or the “learning rate” in machine learning procedures.

Algorithm 1 (Consumer-proposing MSA).

Define $\delta_j^{init} = \sup_{i \in \{1, \dots, N\}} \mathcal{U}_{\varepsilon_{ij}}^{-1}(\mathcal{U}_{i0}(\delta_0))$ # Starting values above upper-bound

Begin Adjustment (outer) Loop

Initialize $\eta^{init} \gg \eta^{tol}$ and $\delta_j^{init} = \delta_j^{return}$.

Repeat:

Call **Deferred-Acceptance Loop** with $(\delta_j^{init}, \eta^{init})$ which returns (δ', η')

Set $\delta^{init} = \delta' + 2\eta'$ and $\eta^{init} = \eta'$

Until $\delta_j^{return} < \delta_j^{init}$ for all $j \in \mathcal{J}$. # Algo stops if all δ have decreased

End Adjustment Loop

Begin Deferred-acceptance (inner) Loop

Require $(\delta_j^{init}, \eta^{init})$.

Set $\eta = \eta^{init}$ and $\delta = \delta^{init}$.

While $\eta \geq \eta^{tol}$ # Run as long as tol. factor above threshold η^{tol}

If $j \in \arg \max_j \mathcal{U}_{\varepsilon_j}(\delta_j)$ then $\pi_{ij} = 1$ else $\pi_{ij} = 0$ # i is matched to optimizing j

If $\sum_i \pi_{i0} < m_0$, then for all $j \in \mathcal{J}$ # If brand 0 in excess supply then

if $\sum_i \pi_{ij} > m_j$ then $\delta_j = \delta_j - \eta$. # decrease δ_j , for brands j w/ excess demand

Else # Else if brand 0 in excess demand (overshooting)

$\delta_j = \delta_j + 2\eta$ for all $j \in \mathcal{J}$ # Increase all δ (except δ_0)

$\eta = \eta/4$. # Decrease the tol. factor for next loop

End While

Return $\delta^{return} = \delta$ and $\eta^{return} = \eta$.

End Deferred-acceptance Loop

As shown above, the consumer-proposing MSA consists of a deferred-acceptance loop nested inside of an adjustment loop. In the deferred-acceptance loop, two cases can occur. In the “good case”, the deferred-acceptance loop starts with values $\delta_{j \in \mathcal{J}}$ above the lattice upper bound and an adjusting factor η small enough, then all $\delta_{j \in \mathcal{J}}$ will reach the lattice upper bound without overshooting. In the “bad case”, that is when the deferred-acceptance loop starts with one $\delta_{j \in \mathcal{J}}$ below its upper bound and an adjusting factor η small enough, then this δ_j will not decrease during the loop. The outer adjustment loop repeatedly calls the deferred-acceptance loop for decreasing values of the increment η . It terminates when the utilities outputted by the deferred-acceptance loop have all decreased (indicating that

the approximation loop is in the “good case”); otherwise, the utilities are increased and the deferred-acceptance loop is called again.

While we have not yet formally proven convergence of the MSA algorithm, we find that it runs remarkably quickly in all our simulations, relative to the Crawford and Knoer (1981) algorithm. In practice, one can assess the convergence of the algorithm in the outer loop by comparing the actual market shares with the predicted market shares evaluated using the δ 's returned in each call to the inner loop.

4.2. Transferable Utility Matching algorithms (for ARUMs). For ARUMs, we can use algorithms for matching models with transferable utility (TU). A well-known result in optimal transport theory (see Galichon (2016)) states that the equilibrium matching under transferable utility maximizes the total surplus $\mathbb{E} [\varepsilon_{\tilde{j}}]$ over all the distributions of (ε, \tilde{j}) such that $\varepsilon \sim P$ and $\tilde{j} \sim s$, that is π solves

$$\max_{(\varepsilon, \tilde{j}) \sim \pi \in \mathcal{M}(P, s)} \mathbb{E}_{\pi} [\varepsilon_{\tilde{j}}] \tag{4.1}$$

The equilibrium payoffs u and δ are the parameters which optimize the dual problem:

$$\begin{aligned} \inf_{u, \delta: \delta_0=0} & \left\{ \mathbb{E}_P [u_{\varepsilon}] - \mathbb{E}_s [\delta_{\tilde{j}}] \right\} \\ \text{s.t.} & \quad u_{\varepsilon} - \delta_j \geq \varepsilon_j \quad \forall \varepsilon \in \Omega, \quad j \in \mathcal{J}_0. \end{aligned} \tag{4.2}$$

These are convex programs which can be solved efficiently by linear programming (LP) or auction algorithms. We contribute two novel modifications to the literature. For LP, we introduce a formulation that *simultaneously* inverts multiple demand maps. Moreover, we augment both the LP and auction algorithms to compute lattice bounds for the identified utility set $\tilde{\sigma}^{-1}$.

4.2.1. Linear programming (Shapley-Shubik). For the large but finite discretization described above, the linear program (4.2) becomes

$$\begin{aligned} \inf_{u_i, \delta_j} & \quad \sum_{i=1}^N \frac{1}{N} u_i - \sum_{j \in \mathcal{J}_0} s_j \delta_j \\ \text{s.t.} & \quad u_i - \delta_j \geq \varepsilon_j^i, \quad \forall i, j. \end{aligned} \tag{4.3}$$

which coincides with the dual of Shapley and Shubik’s (1971) classic assignment game. While Shapley and Shubik showed that the set of optimizers in (4.3) is a lattice, with bounds equal to the consumer-optimal and yogurt-optimal payoffs, they did not discuss how

to compute these bounds. In principle, the upper (resp. lower) bounds can be obtained from the following problem:

$$\max_{u, \delta, \pi} \text{ (resp. } \min_{u, \delta, \pi} \text{)} \sum_{j=0}^J \delta_j \quad \text{s.t. } (u, \delta) \in \{\text{arginf (4.3)}\}.$$

This is a “bilevel” program, as the solutions to the LP in (4.3) are used as the inputs into a second LP.²² As is well-known, we can collapse a bilevel LP into a regular LP by replacing the lower-level LP (corresponding to (4.3)) with its optimality conditions. That is, the upper (resp. lower) bounds can be obtained from the following LP:

$$\begin{aligned} \max_{u, \delta, \pi} \text{ (resp. } \min_{u, \delta, \pi} \text{)} \quad & \sum_{j=0}^J \delta_j \\ \text{s.t.} \quad & \sum_{i=1}^N \pi_{ij} = s_j, \quad \forall j \\ & \sum_{j=0}^J \pi_{ij} = \frac{1}{N}, \quad \forall i \\ & \pi_{ij} \geq 0, \quad \forall i, j \\ & u_i - \delta_j \geq \epsilon_{ij}, \quad \forall i, j \\ & \sum_{i=1}^N \sum_{j=0}^J \pi_{ij} \epsilon_{ij} = \frac{1}{N} \sum_{i=1}^N u_i - \sum_{j=0}^J s_j \delta_j, \\ & \delta_0 = 0. \end{aligned} \tag{4.4}$$

Constraints 1-3 in (4.4) are the constraints of the primal assignment game (4.1), whereas constraint 4 appears in the dual problem (4.2). Constraint 5 equates the primal and dual objectives at the optimum. Taken together, these 5 constraints characterize the optimizing (u, δ) from (4.2).²³ Constraint 6 is our maintained normalization.

Combining LP problems. A further benefit of the LP approach is that multiple demand inversion problems for different markets can be *combined* and solved simultaneously. Specifically, suppose there are $m = 1, \dots, M$ markets. Instead of solving Problem (4.3) for each of the M markets separately, we combine these M problems into *one* problem to invert *all* demand maps *simultaneously*:

$$\begin{aligned} \inf_{u_{mi}, \delta_{mj}} \quad & \sum_{m=1}^M \left(\sum_{i=1}^N \frac{1}{N} u_{mi} - \sum_{j \in \mathcal{J}_0} s_{mj} \delta_{mj} \right) \\ \text{s.t.} \quad & u_{mi} - \delta_{mj} \geq \epsilon_{mj}^i, \quad \forall i, j, m \end{aligned} \tag{4.5}$$

²²See Dempe, Kalashnikov, Pérez-Valdés, and Kalashnykova (2015).

²³See, for instance, Mangasarian (1969).

where all of the subscripts are augmented by the market index m . Since the decision variables (u_{mi}, δ_{mj}) and the associated constraints are market-specific, the resulting constraint coefficient matrix has a block-diagonal structure, which enables the use of efficient parallel sparse matrix routines for modern LP solvers. We utilize this simultaneous demand inversion approach in the empirical application below, and confirm how one large but sparse problem (involving a larger number of parameters and constraints) is solved much more quickly than M small problems.

4.2.2. *Auction algorithms.* Auction-type algorithms à la Bertsekas (1992) provide an alternative approach to linear programming methods for solving TU-matching models. In these algorithms, unassigned persons bid simultaneously for objects, decreasing their systematic utilities (or equivalently raising their prices). Once all bids are in, objects are assigned to the highest bidder. The procedure is iterated until no one is unassigned. The description here follows Bertsekas and Castanon (1989). We let $\kappa \in \{1, \dots, N\}$ index jars of yogurt, where $j(\kappa) \in \mathcal{J}_0$ denotes the brand identity for the κ -th jar of yogurt.

We define the prices p_κ as negative systematic utilities: $p_\kappa = -\delta_\kappa$.

Algorithm 2 (Auction). *Start with an empty assignment and a given vector of prices p_κ and set a scale parameter $\eta > 0$.*

Bidding phase

(a) *Each currently unassigned consumer i chooses the jar κ^* to maximize utility:*

$$\mathcal{U}_{\varepsilon_i \kappa^*}(-p_{\kappa^*}) = \max_{\kappa} \mathcal{U}_{\varepsilon_i \kappa}(-p_{\kappa}). \tag{4.6}$$

(b) *Consumer i 's bid is set to:*

$$b_{i \kappa^*} = p_{\kappa^*} + \mathcal{U}_{\varepsilon_i \kappa^*}(-p_{\kappa^*}) - w_i + \eta \tag{4.7}$$

where w_i denotes the utility from consumer i 's second-best choice:

$$w_i = \max_{j(\kappa) \neq j(\kappa^*)} \mathcal{U}_{\varepsilon_i \kappa}(-p_{\kappa}) \tag{4.8}$$

Assignment phase

Jar κ is assigned to its highest bidder i^ and its price is raised to bidder i^* 's bid:*

$$p_\kappa := b_{i^* \kappa} \tag{4.9}$$

Final step

When no one is left unassigned the solution δ_j is recovered as:

$$\delta_j = - \min_{\kappa \in j(\kappa)} p_\kappa \quad (4.10)$$

Intuitively, the algorithm implements a Walrasian-style bidding procedure. In each round, each unassigned consumer bids for his favorite jar of yogurt. His bid (Eq. (4.7)) is equal to the difference between the utilities from his most-preferred and second-most-preferred brands of yogurt (plus an extra $\eta > 0$ factor to ensure that prices are increasing each round). The consumer which makes the highest bid for a jar is assigned to it, and its price is increased by the amount of the bid; a consumer previously assigned to this jar becomes unassigned and bids in the next round. The algorithm stops when all individuals are assigned. The performance of the algorithm is considerably improved by applying it several times, starting with a large value of η and gradually decreasing it.

Computing utility bounds. The auction algorithm as described above works for the case when the demand map is invertible (so that the identified set of utilities is a singleton). In practice, we do not know *a priori* whether the demand map is invertible or not; hence, we must extend the algorithm to allow for multiplicity of solutions. To do this, we note that the auction algorithm described above returns the optimal allocation π_{ij} (for $i = 1, \dots, N$ and $j \in \mathcal{J}_0$) which is equal to one if consumer i matches with a jar of yogurt brand j , and zero otherwise. In principle the utility bounds could be solved from the earlier linear programming problem (4.4), with the matching indicators π_{ij} fixed at the optimal allocation. However, that would be inefficient as it doesn't fully exploit our knowledge of the optimal allocation. Here we present a simpler alternative which is much quicker as it converges monotonically to the bounds of δ .

For the problem (4.4), given the optimal allocation $\{\pi_{ij}\}_{i,j}$, the solutions (u, δ) satisfy the inequalities

$$u_i \geq \mathcal{U}_{\varepsilon_{ij}}(\delta_j) \equiv \delta_j + \varepsilon_{ij} \quad \forall i, j; \quad \text{with equality for } \pi_{ij} > 0.$$

Accordingly, we construct an operator $T : \mathbb{R}^{\mathcal{I} \cup \mathcal{J}} \rightarrow \mathbb{R}^{\mathcal{I} \cup \mathcal{J}}$ given by

$$T_i(u, \delta) = \max(u_i, \max_{j \in \mathcal{J}_0} (\mathcal{U}_{\varepsilon_{ij}}(\delta_j))), \quad T_j(u, \delta) = \max(\delta_j, \max_{i: \pi_{ij} > 0} \mathcal{U}_{\varepsilon_{ij}}^{-1}(u_i)), \quad \delta_0 = 0. \quad (4.11)$$

The set of fixed points of this operator $u_i = T_i(u, \delta), \delta_j = T_j(u, \delta)$ satisfy the inequalities above. To see this, notice $\max(u_i, \max_{j \in \mathcal{J}_0} (\mathcal{U}_{\varepsilon_{ij}}(\delta_j))) = u_i$ if and only if $u_i \geq \mathcal{U}_{\varepsilon_{ij}}(\delta_j)$ for all $j \in \mathcal{J}_0$. Similarly, $\max(\delta_j, \max_{i: \pi_{ij} > 0} \mathcal{U}_{\varepsilon_{ij}}^{-1}(u_i)) = \delta_j$ holds if and only if $\delta_j \geq \max_{i: \pi_{ij} > 0} \mathcal{U}_{\varepsilon_{ij}}^{-1}(u_i) \Leftrightarrow u_i \leq \mathcal{U}_{\varepsilon_{ij}}(\delta_j)$ for all i such that $\pi_{ij} > 0$. Combining these, we obtain

that the fixed point satisfies $u_i = \mathcal{U}_{\varepsilon_{ij}}(\delta_j)$ for all i, j for which $\pi_{ij} > 0$. Hence, by iterating on (4.11) beginning from the initial lowest possible values $\{u_i = -\infty; \delta_j = -\infty, j \neq 0; \delta_0 = 0\}$, we will obtain a non-decreasing sequence of δ 's converging to $\underline{\delta}$.

Analogously, starting with values of $+\infty$ and iterating on the following operator we will obtain a monotonic non-increasing sequence of δ 's converging to the upper bounds $\bar{\delta}$:

$$\delta_j = \min(\delta_j, \min_{i \in \mathcal{I}} \mathcal{U}_{\varepsilon_{ij}}^{-1}(u_i)), \quad u_i = \min(u_i, \min_{j: \pi_{ij} > 0} (\mathcal{U}_{\varepsilon_{ij}}(\delta_j))), \quad \delta_0 = 0. \quad (4.12)$$

Indeed, $\delta_j = \min(\delta_j, \min_{i \in \mathcal{I}} \mathcal{U}_{\varepsilon_{ij}}^{-1}(u_i))$ holds if and only if $\delta_j \leq \min_{i \in \mathcal{I}} \mathcal{U}_{\varepsilon_{ij}}^{-1}(u_i) \Leftrightarrow \mathcal{U}_{\varepsilon_{ij}}(\delta_j) \leq u_i$, for all $i \in \mathcal{I}$. Similarly, $u_i = \min(u_i, \min_{j: \pi_{ij} > 0} (\mathcal{U}_{\varepsilon_{ij}}(\delta_j)))$ holds if and only if $u_i \leq \min_{j: \pi_{ij} > 0} (\mathcal{U}_{\varepsilon_{ij}}(\delta_j)) \Leftrightarrow u_i \leq \mathcal{U}_{\varepsilon_{ij}}(\delta_j)$, for all j such that $\pi_{ij} > 0$.²⁴

4.3. Implementation. We have developed R packages containing fast and efficient implementations of the algorithms described in this section. They are designed to enable user-friendly access for researchers to popular matching methods, and are used in the next section to benchmark the relative performance of each algorithm. Bonnet, Galichon, Hsieh, O'Hara, and Shum (2018a) collects several auction and linear programming algorithms into an R package, utilizing C++ code provided by Walsh and Dieci (2017). Finally, a parallelized implementation of the MSA is contained in a separate package Bonnet, Galichon, Hsieh, O'Hara, and Shum (2018b). Links to the packages and installation instructions are provided in the bibliography references.

5. NUMERICAL EXPERIMENTS

We test our algorithms on two different models: the first one is the additive pure characteristics model, and the second one is a special case of a pure characteristics model which is *not* invertible, so that multiple values of utilities are consistent with a set of market shares.²⁵

5.1. The pure characteristics model. In this section we evaluate the performance of our matching-based algorithms in computing the pure characteristics model. Berry and

²⁴The iterations over the isotone operators (4.11) and (4.12) are instances of the *Bellman-Ford algorithm*, used for solving certain types of network flow problems, cf. Sedgewick and Wayne (2011).

²⁵While it is beyond the scope of this paper to consider formal statistical testing of equivalence between the bounds $\delta_{\min} = \delta_{\max}$, in all our computational and empirical results below we computed the upper and lower bounds for different and increasing values of N , to ensure that any distance between δ_{\min} and δ_{\max} is not driven by approximation error due to small N .

Pakes (2007) (p. 1193) underline that this model is appealing on theoretical grounds,²⁶ but it is infrequently used in empirical work, arguably due to the computational challenges associated with the non-smooth demand map.

We consider the following specification, which is adapted from Dubé, Fox, and Su (2012): The utility of consumer i who chooses product j in market m is generated by:

$$u_{mj} = \beta_0 - \beta_p p_{mj} + \beta_1 x_{mj1} + \beta_2 x_{mj2} + \beta_3 x_{mj3} + \xi_{mj}, \quad (5.1)$$

where x_{mjk} , $k = 1, 2, 3$ are observed, exogenous product attributes, and p_{mj} is the price of product j in market m that is correlated with the unobserved product attribute ξ_{mj} . Instruments are constructed as noisy nonlinear functions of the exogenous x 's. The random coefficients are $(\beta_p, \beta_1, \beta_2, \beta_3)$, distributed independently from normal distributions with means equal to $\bar{\beta}_p, \bar{\beta}_1, \bar{\beta}_2, \bar{\beta}_3$ and unit variances. Following Berry and Pakes (2007), the mean price coefficient $\bar{\beta}_p$ is normalized to -1. Additional details on the simulation setup and implementation are provided in Appendix D.

We compare and contrast the performance of several algorithms. Our algorithm, which we call "Matching-LP", takes the form of a nested iterative procedure a la BLP (Berry, Levinsohn, and Pakes (1995)): that is, the procedure alternates between an outer and inner loop. In the outer loop, a GMM objective function is optimized with respect to the model parameters, while the inner loop performs the demand inversion to recover the mean utilities (δ) at the current candidate values of the model parameters. In the inner loop, we utilize the linear programming (LP) algorithm from Section 4.2.1 to perform the demand inversion.

To benchmark performance, we compare our "Matching-LP" approach to two alternatives from the literature which reflect the current thinking on estimating the pure characteristics model. The first, denoted "BLP-MPEC", comes from Berry and Pakes (2007), who suggest smoothing out the non-smooth pure characteristics demand map by adding small logit errors to each alternative, and then using the BLP estimation procedure.²⁷ The second

²⁶Models with logit errors have properties that may be undesirable for welfare analysis: They restrain substitution patterns and utility grows without bounds as the number of products in the market grows. See Berry and Pakes (2007) and Akerberg and Rysman (2005).

²⁷In implementation, we utilize Dubé, Fox, and Su's (2012) MPEC algorithm, which has demonstrated speed advantages over nested methods. This speed advantage of MPEC arises in part from the use of Newton-type (gradient-based) iterations in the optimization procedure. As Lee and Seo (2016) point out, Newton iterations have speed gains relative to BLP's original algorithm, which utilizes contraction mapping iterations.

alternative, denoted “PSL”, is the algorithm proposed in Pang, Su, and Lee (2015) based on reformulating the model as a linear complementarity problem.

In Table 1, we report the RMSE, bias, proportion of runs converged, and the average runtime across 20 Monte Carlo repetitions. There are 100 markets and 5 products (including the outside goods). In our simulations, we consider discretizations of the market into both $N = 500$ and $N = 1000$ agents on each side of the market. We consider two model specifications: In Model I, we estimate the location parameters and fix all scale parameters. In Model II, we estimate the location parameter and the scale parameter associated with the endogenous price, fixing the rest of the scale parameters.

In Model I, Matching-LP and PSL have similar RMSE, which is roughly half of that of BLP-MPEC. In Model II, Matching-LP clearly dominates the other two alternatives: particularly, it delivers nearly an unbiased estimate for the standard deviation of the random coefficient of price (σ_p), while the other two approaches falter. In terms of computational speed, our method is on average 130 times faster than PSL and 26 times faster than BLP-MPEC in Model I. In model II, our method outperforms PSL by a factor of 7 and outperform BLP-MPEC by a factor of 2. These simulations demonstrate the superior performance of the matching-based approach for estimating the pure characteristics model.

While all three algorithms require random draws by simulating N consumers from the distribution of the random coefficients, there is a fundamental difference in how these algorithms utilize the random draws. Both our matching-based approach and PSL use the random draws to *discretize* the distribution, whereas the traditional BLP approach *averages over* the random draws to approximate the market shares. Since simulation can potentially create bias, it is importantly to investigate the relationship between the simulation errors and the number of draws. By comparing Panel I and Panel II of Table 1, we find that increasing the number of draws noticeably improves the RMSE of both our matching-based algorithm as well as PSL. For BLP-MPEC, however, there is little noticeable improvement.

As a supplementary exercise, we hone in on the performance of our matching algorithm in the demand inversion step itself, and compare the numerical accuracy of our matching-based algorithms to alternative existing approaches. Table 2 summarizes the numerical performance of three matching-based algorithms: (i) LP, (ii) Auction, and (iii) MSA; the (iv) BLP contraction mapping (which adds logit errors to the utilities of each choice to smooth the demand map) is also included as a benchmark.²⁸

²⁸In our simulations we use the most common version of the BLP-contraction mapping algorithm, which utilizes iterations based on successive approximations. Lee and Seo (2016) show that the performance of

TABLE 1. Numerical Performances: Estimating the Demand Parameters in Pure Characteristics Models

| Panel I: Number of Consumers $N = 500$ | | | | | | | | | | | | |
|--|--------------|-----------------|-----------|------------|-----------------|-----------------|-----------------|-----------------|------------|-----------------|-----------------|-----------------|
| algorithm | converge (%) | runtime (secs.) | RMSE | | | | | Bias | | | | |
| | | | β_0 | σ_p | $\bar{\beta}_1$ | $\bar{\beta}_2$ | $\bar{\beta}_3$ | $\bar{\beta}_0$ | σ_p | $\bar{\beta}_1$ | $\bar{\beta}_2$ | $\bar{\beta}_3$ |
| Model I | | | | | | | | | | | | |
| Matching-LP | 100 | 2.31 | 0.08 | | 0.10 | 0.09 | 0.08 | -0.02 | | 0.00 | -0.02 | 0.01 |
| PSL | 85 | 134.42 | 0.08 | | 0.11 | 0.10 | 0.08 | 0.00 | | 0.00 | -0.02 | 0.01 |
| BLP-MPEC | 100 | 32.88 | 0.19 | | 0.17 | 0.18 | 0.14 | 0.09 | | 0.01 | -0.02 | -0.04 |
| ----- | | | | | | | | | | | | |
| Model II | | | | | | | | | | | | |
| Matching-LP | 100 | 76.52 | 0.09 | 0.06 | 0.11 | 0.11 | 0.07 | -0.02 | 0.01 | 0.00 | -0.02 | 0.00 |
| PSL | 90 | 354.72 | 0.11 | 0.26 | 0.10 | 0.09 | 0.12 | -0.06 | -0.18 | 0.03 | 0.01 | 0.06 |
| BLP-MPEC | 100 | 100.83 | 0.19 | 0.11 | 0.19 | 0.19 | 0.13 | 0.07 | -0.07 | 0.02 | 0.00 | -0.01 |

| Panel II: Number of Consumers $N = 1000$ | | | | | | | | | | | | |
|--|--------------|-----------------|-----------|------------|-----------------|-----------------|-----------------|-----------------|------------|-----------------|-----------------|-----------------|
| algorithm | converge (%) | runtime (secs.) | RMSE | | | | | Bias | | | | |
| | | | β_0 | σ_p | $\bar{\beta}_1$ | $\bar{\beta}_2$ | $\bar{\beta}_3$ | $\bar{\beta}_0$ | σ_p | $\bar{\beta}_1$ | $\bar{\beta}_2$ | $\bar{\beta}_3$ |
| Model I | | | | | | | | | | | | |
| Matching-LP | 100 | 3.09 | 0.08 | | 0.07 | 0.07 | 0.06 | -0.03 | | -0.01 | -0.02 | 0.01 |
| PSL | 90 | 409.65 | 0.08 | | 0.07 | 0.07 | 0.06 | 0.01 | | -0.01 | -0.02 | 0.01 |
| BLP-MPEC | 100 | 79.38 | 0.27 | | 0.16 | 0.17 | 0.16 | 0.14 | | 0.02 | 0.01 | -0.09 |
| ----- | | | | | | | | | | | | |
| Model II | | | | | | | | | | | | |
| Matching-LP | 100 | 120.19 | 0.09 | 0.07 | 0.07 | 0.07 | 0.07 | -0.02 | 0.00 | -0.01 | -0.01 | 0.01 |
| PSL | 85 | 834.48 | 0.14 | 0.35 | 0.09 | 0.10 | 0.11 | -0.09 | -0.28 | 0.06 | 0.05 | 0.08 |
| BLP-MPEC | 100 | 260.03 | 0.26 | 0.13 | 0.17 | 0.17 | 0.16 | 0.12 | -0.06 | 0.04 | 0.03 | -0.07 |

Note: The numbers are averages across 20 Monte Carlo repetitions. There are 100 markets and 5 products. The data-generating process is adapted from Dubé, Fox, and Su (2012), and complete details are contained in Appendix D. β_0 denotes the constant in utility, while $\bar{\beta}_1, \bar{\beta}_2, \bar{\beta}_3$ correspond to the means of the random coefficients. σ_p is the scale parameter of price. The true values are: $\beta_0 = 1$, $\bar{\beta}_1 = 0.5$, $\bar{\beta}_2 = 0.5$, $\bar{\beta}_3 = 0.2$, $\sigma = 1$. The inner loop of Matching-LP is based on the combined LP defined in Eq. (4.5).

The numerical accuracy is almost identical across all three matching algorithms (LP, Auction, MSA). In comparison, the RMSE of the BLP contraction mapping is about three times larger. Similar to the finding in Table 1, we find that increasing the number of draws the algorithm can be improved by utilizing Newton iterations instead. However, both approaches require smoothness in the market-share mapping, which is not satisfied in the pure characteristics model considered here.

TABLE 2. Numerical Performances: Demand Inversion in Pure Characteristics Models

| Algorithms | Draws | Brands | RMSE | runtime (secs) |
|--------------------|--------|--------|-------|----------------|
| BLP contract. map. | 1,000 | 5 | 0.155 | 0.032 |
| LP (gurobi) | 1,000 | 5 | 0.072 | 0.082 |
| Auction | 1,000 | 5 | 0.073 | 0.007 |
| MSA | 1,000 | 5 | 0.073 | 16.060 |
| BLP contract. map. | 1,000 | 50 | 0.307 | 0.282 |
| LP (gurobi) | 1,000 | 50 | 0.091 | 0.281 |
| Auction | 1,000 | 50 | 0.094 | 0.037 |
| BLP contract. map. | 1,000 | 500 | 0.564 | 2.754 |
| LP (gurobi) | 1,000 | 500 | 0.100 | 3.364 |
| Auction | 1,000 | 500 | 0.107 | 0.370 |
| BLP contract. map. | 10,000 | 5 | 0.151 | 0.330 |
| LP (gurobi) | 10,000 | 5 | 0.029 | 0.289 |
| Auction | 10,000 | 5 | 0.030 | 0.076 |
| MSA | 10,000 | 5 | 0.029 | 1.126 |
| BLP contract. map. | 10,000 | 50 | 0.306 | 2.876 |
| LP (gurobi) | 10,000 | 50 | 0.039 | 3.895 |
| Auction | 10,000 | 50 | 0.039 | 0.579 |
| BLP contract. map. | 10,000 | 500 | 0.557 | 30.738 |
| LP (gurobi) | 10,000 | 500 | 0.043 | 53.796 |
| Auction | 10,000 | 500 | 0.043 | 4.622 |

Note: The numbers are averaged from 50 Monte-Carlo replications. Demand inversion for the pure characteristics model with 5, 50 and 500 brands of yogurt and 1,000 and 10,000 draws of taste shocks. The column "RMSE" corresponds to the root mean squared errors of the estimated δ_j . The MSA becomes computationally infeasible when the number of brands becomes large, so we report it only for scenarios with 5 brands.

N improves the RMSE of our matching algorithms but not the BLP contraction mapping – this arises from the additional approximation error due to the introduction of additive logit errors to smooth the mapping.

For computational speed, the auction algorithm is the fastest by a wide margin: even under the most demanding scenario (500 brands and 10,000 simulated consumers), it only takes 5 seconds, which far outstrips other approaches.²⁹

On the other hand, while MSA is the slowest algorithm, and does not scale up well with the number of brands, it is a general-purpose algorithm that also applies to NARUMs.³⁰ It is therefore not surprising that MSA is not as fast as the other methods since it does not exploit the additive separability in ARUMs.

5.2. A non-invertible pure characteristics model. For the second simulation exercise, we consider a pure characteristics model which is non-invertible—that is, there are multiple utility vectors which rationalizes the same set of market shares. It underscores an important contribution of our approach, namely its ability to handle models which are not invertible. Moreover, it also suggests that non-invertibility may be a typical feature of the pure characteristics model. Pang, Su, and Lee (2015, pg. 654) likewise note a problem of multiple solutions in results from their estimation procedure for pure characteristics models.

There are three goods $y = 1, 2, 3$, and the unknown parameters are the quality of each good are $\delta_1, \delta_2, \delta_3$, with the normalization $\delta_1 = 0 < \delta_2, \delta_3$. Consumers fall into two segments, which differ in the prices that they face. In segment 1, prices are given by the vector \mathbf{p}_1 , with $p_1^1 < p_2^1 < p_3^1$, while in segment 2, prices are given by the vector \mathbf{p}_2 , with $p_1^2 \leq p_3^2 < p_2^2$. The utility function is given by

$$\mathcal{U}_{\varepsilon_j}(\delta_j) = \delta_j - \frac{\varepsilon^b}{\varepsilon^a} p_j^1 - \frac{(1 - \varepsilon^b)}{\varepsilon^a} p_j^2,$$

which is a pure characteristics specification. Consumer heterogeneity $\varepsilon \equiv (\varepsilon^a, \varepsilon^b)$ consists of both $\varepsilon^a \sim U[0, 1]$, interpreted as willingness-to-pay for quality, and ε^b , an indicator for whether they are in the first segment, which equals 0 or 1 with equal probability.

²⁹The computational speed reported in Table 1 and 2 are not directly comparable for several reasons: First, Table 2 only measures the computational time in the demand-inversion step—the “inner loop” in Table 2. Second, in Table 1, the differences between NFXP and MPEC for searching the structural parameters can contribute to the performance difference. Third, in Table 1, there are 100 markets, whereas in Table 2, there is only one market. The scalability of different demand-inversion methods also affects the overall computational speed.

³⁰Interestingly, with 5 brands, MSA is faster with 10,000 draws than with 1,000: this may appear counter-intuitive, but increasing the number of draws improves accuracy during the loop and therefore reduces total computation times.

Let s_j^d denote the (unobserved) market share of good j among segment $d(\in \{1, 2\})$ customers. The observed market shares are mixtures of market shares across both customer segments:

$$s_j = 0.5(s_j^1 + s_j^2), \quad j = 1, 2, 3. \tag{5.2}$$

This example typifies a common problem in supermarket scanner datasets, that consumers' transactions prices often differ markedly by products' list prices.³¹ In this case, the segment 1 prices \mathbf{p}^1 are analogous to "list" prices (ie., the prices posted on the supermarket shelves) while segment 2 customers have a coupon which gives them a discount on good 3.

In the simulation exercise, we set $p_1^1 = 1$, $p_2^1 = 2$, and $p_3^1 = 3$ be the prices in segment 1, and $p_1^2 = 1$, $p_2^2 = 2$, and $p_3^2 = 1$ in segment 2. The observed market shares are given by $(s_1, s_2, s_3) = (0.25, 0.25, 0.5)$. Given $\delta_1 = 0$, the identified utility set contains multiple values of the quality parameters (δ_2, δ_3) :

$$\tilde{\sigma}^{-1}(s) = \{(\delta_2, \delta_3) : \delta_2 = 2, \delta_3 \in [1, 3]\} \tag{5.3}$$

which is a lattice with minimal element $(0, 2, 1)$ and maximal element $(0, 2, 3)$.

Computational results using the LP algorithm are given in Table 3. The algorithm performs as expected; utilizing the linear programming algorithm (4.3) along with the supplemental program to compute the upper and lower bound of payoffs (4.4) accurately recovers the upper and lower bounds. For comparison, we also performed the demand inversion exercise using the BLP contraction mapping approach, which requires the demand inverse to be unique. Not surprisingly, we found that the contraction mapping algorithm would not converge for this problem.³²

TABLE 3. Non-invertible model: 3 goods

| True δ_j | $N = 100$ | | | $N = 1000$ | | |
|-----------------------|-----------------------|-----------------------------|-----------------|-----------------------|-----------------------------|-----------------|
| | LP - $\bar{\delta}_j$ | LP - $\underline{\delta}_j$ | BLP contr. map. | LP - $\bar{\delta}_j$ | LP - $\underline{\delta}_j$ | BLP contr. map. |
| $\delta_2 = 2$ | 2.110 (0.328) | 2.036 (0.315) | —† | 2.015 (0.092) | 2.005 (0.092) | —† |
| $\delta_3 \in [1, 3]$ | 3.133 (0.334) | 1.017 (0.315) | —† | 3.016 (0.092) | 1.003 (0.092) | —† |

We normalized $\delta_1 = 0$. The aggregate market shares are $(s_1, s_2, s_3) = (0.25, 0.25, 0.5)$, and prices: store 1 $(1, 2, 3)$, store 2 $(1, 2, 1)$. Lower and upper bounds are computed as the solution to the Linear Programming algorithm. For comparison, the BLP contraction mapping, an algorithm which requires the solution to be unique, failed to converge in each trial run. N denotes the number of discretization points used. Standard deviations (across 50 replications) reported in parentheses. †: the BLP contraction mapping algorithm failed to converge.

³¹See Erdem, Keane, and Sun (1998) for detailed discussions and modelling approaches to this problem.

³²On each trial run, the maximum number of function evaluations was reached without convergence.

To continue, we consider an expanded version of this model involving 5 segments, and 8 goods. The prices and market shares for this example are given in Table 4. Since, these market shares were chosen arbitrarily, we do not know *a priori* whether the identified utility set $\tilde{\sigma}^{-1}(s)$ is a singleton or contains multiple values.³³

In the two right-most columns of Table 4, we report the recovered upper and lower bounds for the δ 's. Clearly, for all entries, the lower and upper bounds coincide at the second or third decimal place, suggesting that the identified utility set $\tilde{\sigma}^{-1}$ is a singleton. The results here were computed using $N = 50,000$, a very large value to ensure that the results are not driven by approximation error. In Table 7 in the appendix, we report results for different values of N , demonstrating that these results are robust and that, even with smaller values of N , the conclusions do not change.

TABLE 4. Multisegment price heterogeneity model: 5 segments and 8 goods

| Brand | p_1 | p_2 | p_3 | p_4 | p_5 | MktShare | Lower-bound $\underline{\delta}$ | Upper-bound $\bar{\delta}$ |
|-------|-------|-------|-------|-------|-------|----------|----------------------------------|----------------------------|
| A | 3.32 | 3.36 | 3.45 | 3.37 | 3.35 | 0.071 | 0.000(ref.) | 0.000(ref.) |
| B | 3.88 | 3.60 | 3.53 | 3.39 | 3.07 | 0.057 | -0.405(0.070) | -0.403(0.069) |
| C | 3.70 | 3.30 | 4.16 | 4.31 | 4.25 | 0.201 | 2.319(0.041) | 2.320(0.041) |
| D | 3.98 | 4.12 | 4.06 | 3.11 | 4.09 | 0.389 | 3.186(0.043) | 3.187(0.043) |
| E | 4.20 | 4.34 | 4.21 | 4.29 | 4.35 | 0.155 | 3.547(0.043) | 3.548(0.043) |
| F | 4.49 | 4.82 | 4.25 | 3.73 | 4.86 | 0.076 | 3.614(0.043) | 3.615(0.043) |
| G | 7.13 | 7.92 | 7.95 | 7.99 | 7.71 | 0.006 | 6.762(0.043) | 6.763(0.043) |
| H | 8.34 | 8.37 | 8.59 | 8.62 | 8.67 | 0.045 | 8.048(0.043) | 8.049(0.043) |

Estimates are LP solution computed with Gurobi 8.1. 50 Monte-Carlo estimations were made with different draws of ε^a (standard deviation across 50 replications reported in parentheses). 50,000 discretization points were used for each Monte-Carlo simulation (*i.e.* 50,000 draws of ε^a in the uniform). Brand A is the reference and the systematic utility is normalized to 0.

6. EMPIRICAL APPLICATION: VOTING IN EUROPEAN PARLIAMENT ELECTIONS

Finally, we use our matching-based algorithm to estimate an aggregate spatial voting model using data from the 1999 Parliamentary Elections in the European Union countries, following Merlo and de Paula (2017).

³³That is, unlike the two-segment example used in Table 3, we did not start by computing a market equilibrium for given parameter values. Rather we chose the prices and market shares in Table 4 arbitrarily, and use our approach to determine $\tilde{\sigma}^{-1}$.

6.1. **Model.** We consider a spatial voting framework in which both voters and political parties are characterized by their “location” in the political spectrum, which is the Cartesian plane \mathbb{R}^2 . Voter i is characterized by her ideal point $t_i \in \mathbb{R}^2$ within this space; likewise, candidate (political party) j has an ideological position $C_j \equiv (C_{j1}, C_{j2}) \in \mathbb{R}^2$.

Voters are ideological: voter i from electoral precinct m votes for the party with the platform closest to her ideal point t_{mi} . Specifically, her preferred party is:

$$D_{mi} = \operatorname{argmin}_{j \in J_m} d(t_{mi}, C_{mj}), \tag{6.1}$$

where $d(\cdot, \cdot): \mathbb{R}^2 \rightarrow \mathbb{R}_+$ is the distance function, which is specified as a quadratic function:

$$d(t, C) = (t - C)'W(t - C). \tag{6.2}$$

W is a 2×2 weighting matrix, which for simplicity we assume to be identity: $W = \mathbb{I}_2$. (This is restrictive as it implies that voters weigh both dimensions of the political spectrum equally.) Unlike Merlo and de Paula (2017) who consider the nonparametric identification and estimation of the distribution of voter ideal points t , we assume that they come from a bivariate normal distribution, as follows:

$$t_{mi1} = (X'_m \alpha + \nu_{mi1}), \tag{6.3}$$

$$t_{mi2} = (X'_m \beta + \nu_{mi2}),$$

$$(\nu_{mi1}, \nu_{mi2}) \sim N(\mathbf{0}, \Sigma), \text{ i.i.d. over } i, j, m \tag{6.4}$$

where $X_m = (x_{mk}, k = 1, \dots, K)$, are aggregate demographic and economic variables in precinct m , and α and β denote $K \times 1$ parameter vectors of interest.³⁴

By substituting (6.3) into (6.2), the “disutility” that voter i gets from voting for candidate j becomes:

$$\begin{aligned} U_{mij} &= d(t_{mi}, C_{mj}) = (t_{mi1} - C_{mj1})^2 + (t_{mi2} - C_{mj2})^2 \\ &= (X'_m \alpha)^2 + C_{mj1}^2 - 2(C_{mj1} * (X'_m \alpha)) + (X'_m \beta)^2 + C_{mj2}^2 - 2(C_{mj2} * (X'_m \beta)) \\ &\quad + \nu_{mi1}^2 + 2(X'_m \alpha - C_{mj1}) \nu_{mi1} + \nu_{mi2}^2 + 2(X'_m \beta - C_{mj2}) \nu_{mi2} \end{aligned} \tag{6.5}$$

In the discrete-choice setting, only utility components which vary across candidates j will affect choices. From the above, the terms (ν_{i1}^2, ν_{i2}^2) and $(X'_m \alpha) \nu_{i1}, (X'_m \beta) \nu_{i2}$ are individual-specific, and hence do not affect i 's choice problem. Therefore, our specification can be

³⁴As usual, we normalize the variance of ν_{mi1} to one.

simplified into the following pure characteristics form:³⁵

$$U_{mij} \equiv \delta_{mj} + \varepsilon_{mij} \tag{6.6}$$

where: $\delta_{mj} = C_{mj1}^2 + C_{mj2}^2 - 2(C_{mj1} * (X'_m \alpha)) - 2(C_{mj2} * (X'_m \beta))$;

$$\varepsilon_{mij} = -2C_{mj1}\nu_{mi1} - 2C_{mj2}\nu_{mi2}$$

The parameters in this specification are collectively denoted by $\theta \equiv (\alpha, \beta, \Sigma)$.

As we pointed out above, the computational difficulties of the pure characteristics model has hampered empirical work utilizing it; indeed, despite the theoretical importance of the spatial voting model in political economy (ever since the pioneering work by Downs (1957)), empirical specifications have been relatively sparse. Thus the application here, while primarily illustrative, does have a methodological contribution in showing how the matching-based algorithms introduced in this paper can be used to estimate a work-horse model from political science.

As data we use precinct-level vote shares from the 1999 European Parliament elections. Altogether we have voting data from 822 electoral precincts in 22 regions, which are typically countries but may be sub-national regions distinguished by different sets of political parties.³⁶ Parties' ideological positions are taken from Hix, Noury, and Roland (2006), who computed two ideological positions for each party: $C = (C_{j1}, C_{j2})$, with C_{j1} denoting position on a left-right spectrum, and C_{j2} denoting party's stance on the EU (with larger values denoting, resp. a more right-wing position and more pro-EU stance). We use $K = 3$ precinct-specific socio-economic and demographic variables: the female-to-male ratio, the proportion of the population older than 35 years, and the unemployment rate.³⁷ All the data we use are available from the *Review of Economic Studies* website.³⁸

³⁵Merlo and de Paula (2017) also pointed out the equivalence of the spatial voting and pure characteristics models.

³⁶The 22 regions are: Austria, Finland, France, Germany, Greece, Italy-Center, Italy-Islands, Italy-Northeast, Italy-Northwest, Italy-South, Portugal, Spain, Sweden, Netherlands, UK-East Midlands, UK-Eastern, UK-London, UK-Northwest, UK-Southeast, UK-Southwest, UK-West Midlands, and UK-Yorkshire.

³⁷Merlo and de Paula (2017) include additionally GDP as a regressor, and hence exclude Austria and Italy from their analysis due to missing values of this variable.

³⁸<https://doi.org/10.1093/restud/rdw046>

As the quadratic spatial voting model is a pure characteristics model, we use our “Matching-LP” algorithm for estimation, which we also used for the simulations in Section 5.1.³⁹ Essentially, this resembles the Berry, Levinsohn, and Pakes (1995) estimation algorithm except that the inner loop utilizes a matching-based algorithm in place of the contraction mapping in BLP. We describe the outer and inner loops in turn.

6.1.1. *Estimation: Inner loop.* In each call to the inner loop, for a given parameter vector θ , we solve the demand inversion problems across all precincts simultaneously (see the earlier discussion associated with Eq. (4.5)), by combining all the linear programs across all precincts into the following single large linear program:

$$\inf_{u_{mi}, \delta_{mj}, m \in \mathcal{M}} \sum_{m=1}^M \left[\sum_{i=1}^N \frac{1}{N} u_{mi} - \sum_{j \in \mathcal{J}_m} s_{mj} \delta_{mj} \right] \tag{6.7}$$

$$s.t. \quad u_{mi} - \delta_{mj} \geq -\varepsilon_{mij} \quad \forall m \in \mathcal{M} \tag{6.8}$$

Let $\{\hat{\delta}_{mj}(\theta)\}_{m,j}$ denote the optimized values of the δ_{mj} ’s from this problem.

6.1.2. *Estimation: Outer loop.* For the outer loop, we minimize the least-square differences between the $\hat{\delta}_{mj}(\theta)$ ’s emerging from the inner loop and the functional form for δ_{mj} , as implied in Eq. (6.6), to estimate θ :

$$\min_{\theta} \sum_{m,j} \left[\hat{\delta}_{mj}(\theta) - C_{mj1}^2 - C_{mj2}^2 - 2C_{mj1} * (X'_m \alpha) - 2C_{mj2} * (X'_m \beta) \right]^2 \tag{6.9}$$

6.2. **Results.** Table 5 contains our estimation results for two specifications.⁴⁰ In Model I, we assume that the covariance matrix $\Sigma = \mathbb{I}_2$, thus ruling out correlation between the two ideal point dimensions conditional on precinct characteristics X_m . Model II eliminates

³⁹We focus on the LP algorithm here in order to leverage combining and simultaneously solving the demand inversion problems across different precincts, which saves substantially on computational time. Other algorithms, including the auction algorithm, are not amenable to combining the demand inversion problems.

⁴⁰For this ARUM, without loss of generality, we normalized $\delta_{m1} = 0$, corresponding to the party listed first in alphabetical order in each precinct m . We implemented this normalization by subtracting, in each market, the covariates $C_{j1}^2 + C_{j2}^2 - 2(C_{j1} * (X'_m \alpha)) - 2(C_{j2} * (X'_m \beta))$ for candidate 1 from the utilities for the other candidates $j \neq 1$ (cf. Eq. (6.6)). The pure characteristics model also requires a location normalization (cf. Berry and Pakes (2007, footnote 13)) which is automatically satisfied as the spatial voting model implies that the constant term in the utility is equal to the squared ideological positions $C_{j1}^2 + C_{j2}^2$ which is known by the researcher.

TABLE 5. Estimation Results: 1999 European Parliamentary Elections

| | Model I | Model II |
|--|------------------------|------------------------|
| dim. 1: right-leaning | | |
| female-to-male ratio | 0.81 (0.50,1.12) | 0.44 (0.19,0.68) |
| above 35-year-old(%) | -0.07 (-0.56,0.49) | -0.29 (-0.68,0.12) |
| unemployment rate | -2.62 (-3.02,-2.17) | -1.76 (-2.16,-1.33) |
| dim. 2: pro-EU | | |
| female-to-male ratio | 0.62 (0.29,0.96) | -0.48 (-0.85,-0.10) |
| above 35-year-old(%) | -0.02 (-0.60,0.55) | 0.80 (0.16,1.39) |
| unemployment rate | -2.95 (-3.53,-2.28) | -1.25 (-1.96,-0.62) |
| ρ | 0 (fixed) | -0.66 (-0.70,-0.40) |
| <i>Checking multiplicity:</i> $\max_{m,j} \{\delta_{mj}^{UB} - \delta_{mj}^{LB}\}$ | 1.52×10^{-9} | 9×10^{-4} |

Note: The inner loop utilizes the linear programming algorithm (Section 4.2.1). We report the 95% bootstrap confidence interval constructed from 500 bootstrap samples. We use $N = 100$ voters in each precinct. δ_{mj}^{UB} and δ_{mj}^{LB} refer, respectively, to estimates of the upper and lower bound of the identified utility parameters, computed using the algorithm in Eq. (4.4).

this restriction and allows for correlation between the dimensions, so that $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$. The estimate for ρ in Model II is negative (coef. -0.66) and significantly different from

zero, implying that, after controlling for precinct-level characteristics, voters who are right-leaning (dimension 1) tend to be *less* favorable towards the EU (dimension 2).⁴¹ Given this result, we focus the discussion here to the results from Model II.

The estimated coefficients summarize the contribution of demographic variables on the two dimensions of voters' ideal points. For the first dimension, the coefficient on unemployment rate is negative (-1.76), indicating that left-leaning precincts tend to have higher unemployment rates. Precincts with larger female-to-male ratio tend to be more right-leaning (0.44), while there is no significant relation between age (measured by the proportion of population older than 35 years) on the tendency to be pro-conservative.

For the second dimension, we find that precincts with higher unemployment rate are significantly less supportive of the EU (coef. -1.25). Precincts with higher percentage of older voters are significantly more pro-EU, while those with a higher female-male ratio are less supportive of the EU. Altogether, economic considerations – as exemplified in the unemployment rate – appear to be the strongest and most consistent explainer of voters' preferences across European regions.

We also use our algorithms to check whether the demand map is invertible by solving for the upper and lower bounds on the precinct/party qualities δ_{mj} using the linear program in Eq. (4.4). At the bottom of Table 5, we report the maximal (across all precincts and parties) difference between the estimated upper and lower bounds in the δ_{mj} 's, evaluated at the parameter values reported in Table 5. As is evident, the difference is minuscule, suggesting that multiplicity of these parameters is not an issue for this model.⁴² Finally, the combined linear program in Eq. (6.7) is a big time-saver, as executing the demand inversion problem *simultaneously* across all precincts is ten times faster than performing the demand inversion

⁴¹In comparison, the richer model in Merlo and de Paula (2017) allows the correlation between voter dimensions to differ by country, and they estimate this correlation to be negative in 4 out of the 10 countries, including France and Germany, the two largest EU countries (cf. Table 3 in Merlo and de Paula (2017)).

⁴²With multiple δ 's, the identification and estimation of the structural parameters (the β 's and the parameters in the distribution of random coefficients) is an open question, and we do not consider it here. Estimating these parameters typically relies upon instruments, and the associated moment conditions. The literature on identification and inference in moment condition models with possibly partially-identified parameters is still nascent; see Chen, Christensen, and Tamer (2018) for one recent paper. In ongoing work (Hsieh, Montero, and Shum (2019)) we are tackling this issue in the context of discrete-choice demand models as here, but both the issue regarding identification of β , as well as the computational and inferential issues associated with estimating β in this setting are challenging.

separately for each precinct. This suggests that simultaneous solution of demand inversion problems constitutes a very practical advantage of the linear programming approach.

7. CONCLUDING REMARKS

In this paper we have explored the intimate connection between discrete choice models and two-sided matching models, and used results from the literature on matching under imperfectly transferable utility to derive procedures for demand inversion and estimation of discrete choice models based on the non-additive random utility specification. Although the microeconomics literature distinguishes between “one-sided” and “two-sided” demand problems, our results show that this distinction is immaterial for the purpose of estimating discrete-choice models. Given the matching equivalence, it is as appropriate to consider the discrete choice problem of consumers choosing yogurts as one in which “yogurts choose consumers”.

The connection between discrete choice and two-sided matching is a rich one, and we are exploring additional implications. For instance, the phenomenon of “multiple discrete choice” (consumers who choose more than one brand, or choose bundles of products on a purchase occasion) is challenging and difficult to model in the discrete choice framework⁴³ but is quite natural in the matching context, where “one-to-many” markets are commonplace – perhaps the most prominent and well-studied being the National Residents Matching Program for aspiring doctors in the United States (cf. Roth (1984)). We are exploring this connection in ongoing work.

The data and code underlying this research is available in Zenodo, at <https://doi.org/10.5281/zenodo.5896606>.

REFERENCES

- ACKERBERG, D. A., AND M. RYSMAN (2005): “Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects,” *RAND Journal of Economics*, 36(4), 771–788.
- ADACHI, H. (2000): “On a characterization of stable matchings,” *Economics Letters*, 68, 43–49.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models,” *Econometrica*, 70(4), 1519–1543.
- APESTEGUIA, J., AND M. A. BALLESTER (2014): “Discrete choice estimation of risk aversion,” working paper, Pompeu Fabra.
- ARCIDIACONO, P., AND R. A. MILLER (2011): “Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity,” *Econometrica*, 79(6), 1823–1867.

⁴³See Hendel (1999), Dubé (2004), Fox and Bajari (2013) for some applications.

- AURENHAMMER, F. (1987): “Power diagrams: properties, algorithms and applications,” *SIAM Journal on Computing*, 16(1), 78–96.
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): “Estimating dynamic models of imperfect competition,” *Econometrica*, 75(5), 1331–1370.
- BERKSON, J. (1955): “Maximum likelihood and minimum χ^2 estimates of the logistic function,” *Journal of the American statistical association*, 50(269), 130–162.
- BERRY, S., A. GANDHI, AND P. HAILE (2013): “Connected substitutes and invertibility of demand,” *Econometrica*, 81(5), 2087–2111.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica*, pp. 841–890.
- BERRY, S., AND A. PAKES (2007): “The pure characteristics demand model,” *International Economic Review*, 48(4), 1193–1225.
- BERRY, S. T. (1994): “Estimating discrete-choice models of product differentiation,” *RAND Journal of Economics*, pp. 242–262.
- BERRY, S. T., AND P. HAILE (2015): “Identification in Differentiated Products Markets,” Discussion paper, National Bureau of Economic Research.
- BERRY, S. T., AND P. A. HAILE (2014): “Identification in differentiated products markets using market level data,” *Econometrica*, 82(5), 1749–1797.
- BERTSEKAS, D. P. (1992): “Auction algorithms for network flow problems: A tutorial introduction,” *Computational optimization and applications*, 1(1), 7–66.
- BERTSEKAS, D. P., AND D. A. CASTANON (1989): “The auction algorithm for the transportation problem,” *Annals of Operations Research*, 20(1), 67–96.
- BONNET, O., A. GALICHON, Y.-W. HSIEH, K. O’HARA, AND M. SHUM (2018a): “Auction.Rcpp - R Wrapper to General Auction and Assignment Methods,” <https://github.com/TraME-Project/auction>.
- BONNET, O., A. GALICHON, Y.-W. HSIEH, K. O’HARA, AND M. SHUM (2018b): “Yogurts - Fast IPFP and MSA Methods for R,” <https://github.com/TraME-Project/yogurts>.
- BRESNAHAN, T. F. (1981): “Departures from marginal-cost pricing in the American automobile industry: Estimates for 1977–1978,” *Journal of Econometrics*, 17(2), 201–227.
- BURKARD, R., M. DELL’AMICO, AND S. MARTELLO (2009): *Assignment Problems*. SIAM.
- CHEN, X., T. M. CHRISTENSEN, AND E. TAMER (2018): “Monte Carlo confidence sets for identified sets,” *Econometrica*, 86(6), 1965–2018.
- CHERNOZHUKOV, V., A. GALICHON, M. HENRY, AND B. PASS (2019): “Single market nonparametric identification of multi-attribute hedonic equilibrium models,” *arXiv preprint arXiv:1709.09570*.
- CHIAPPORI, P., AND I. KOMUNJER (2009): “On the Nonparametric Identification of Multiple Choice Models,” working paper, Columbia.
- CHIONG, K. X., A. GALICHON, AND M. SHUM (2016): “Duality in dynamic discrete choice models,” *Quantitative Economics*, 7, 83–115.
- COHEN, A., AND L. EINAIV (2007): “Estimating risk preferences from deductible choice,” *American economic review*, 97(3), 745–788.
- CRAWFORD, V. P., AND E. M. KNOER (1981): “Job matching with heterogeneous firms and workers,” *Econometrica*, pp. 437–450.

- DEMANGE, G., AND D. GALE (1985): "The strategy structure of two-sided matching markets," *Econometrica*, pp. 873–888.
- DEMPE, S., V. KALASSHIKOV, G. A. PÉREZ-VALDÉS, AND N. KALASHNYKOVA (2015): *Bilevel Programming Problems: Theory, Algorithms and Applications to Energy Networks*. Springer-Verlag.
- DOWNS, A. (1957): *An economic theory of democracy*. Harper, New York.
- DUBÉ, J.-P. (2004): "Multiple discreteness and product differentiation: Demand for carbonated soft drinks," *Marketing Science*, 23(1), 66–81.
- DUBÉ, J.-P., J. T. FOX, AND C.-L. SU (2012): "Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation," *Econometrica*, 80(5), 2231–2267.
- ECHENIQUE, F., AND J. OVIEDO (2004): "Core Many-to-one Matchings by Fixed Point Methods," *Journal of Economic Theory*, 115, 358–376.
- ERDEM, T., M. P. KEANE, AND B. SUN (1998): "Missing price and coupon availability data in scanner panels: Correcting for the self-selection bias in choice model parameters," *Journal of Econometrics*, 89(1-2), 177–196.
- ESTEBAN, S., AND M. SHUM (2007): "Durable-goods oligopoly with secondary markets: the case of automobiles," *The RAND Journal of Economics*, 38(2), 332–354.
- FLEINER, T. (2003): "A fixed-point approach to stable matchings and some applications," *Mathematics of Operations Research*, 28.1, 103–126.
- FOX, J. T., AND P. BAJARI (2013): "Measuring the efficiency of an FCC spectrum auction," *American Economic Journal: Microeconomics*, 5(1), 100–146.
- GALE, D., AND H. NIKAIKO (1965): "The Jacobian matrix and global univalence of mappings," *Mathematical Annals*, 159, 81–93.
- GALE, D., AND L. S. SHAPLEY (1962): "College admissions and the stability of marriage," *The American Mathematical Monthly*, 69(1), 9–15.
- GALICHON, A. (2016): *Optimal Transport Methods in Economics*. Princeton.
- GALICHON, A., AND Y.-W. HSIEH (2019): "Aggregate Stable Matching with Money Burning," *working paper*.
- GALICHON, A., S. D. KOMINERS, AND S. WEBER (2019): "An empirical framework for matching with imperfectly transferable utility," *Journal of Political Economy*, 127, 2875–2925.
- GALICHON, A., AND B. SALANIÉ (2015): "Cupid's invisible hand: Social surplus and identification in matching models," *working paper*, NYU.
- HATFIELD, J. W., AND P. R. MILGROM (2005): "Matching with contracts," *American Economic Review*, pp. 913–935.
- HENDEL, I. (1999): "Estimating multiple-discrete choice models: An application to computerization returns," *The Review of Economic Studies*, 66(2), 423–446.
- HIX, S., A. NOURY, AND G. ROLAND (2006): "Dimensions of Politics in the European Parliament," *American Journal of Political Science*, 50, 494–511.
- HOTZ, V. J., AND R. A. MILLER (1993): "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 60(3), 497–529.
- HSIEH, Y.-W., S. MONTERO, AND M. SHUM (2019): "A Semiparametric Discrete-Choice Aggregate Demand Model," *working paper*.

- JIA-BARWICK, P. (2008): "What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry," *Econometrica*, 76(6), 1263–1316.
- KELSO, A. S., AND V. P. CRAWFORD (1982): "Job matching, coalition formation, and gross substitutes," *Econometrica*, pp. 1483–1504.
- KHAN, S., F. OUYANG, AND E. TAMER (2016): "Adaptive Rank Inference in Semiparametric Multinomial Response Models," working paper, Boston College.
- KRISTENSEN, D., L. NESHEIM, AND A. DE PAULA (2014): "CCP and the Estimation of Nonseparable Dynamic Models," Working paper, UCL.
- LEE, J., AND K. SEO (2016): "Revisiting the nested fixed-point algorithm in BLP random coefficients demand estimation," *Economics Letters*, 149, 67–70.
- MAGNAC, T., AND D. THESMAR (2002): "Identifying dynamic discrete decision processes," *Econometrica*, 70(2), 801–816.
- MANGASARIAN, O. (1969): *Nonlinear Programming*. Society for Industrial and Applied Mathematics.
- MATZKIN, R. (2007): "Nonparametric Identification," *Handbook of Econometrics*, 6b.
- McFADDEN, D., ET AL. (1978): *Modelling the choice of residential location*. Institute of Transportation Studies, University of California California.
- McFADDEN, D., AND K. TRAIN (2000): "Mixed MNL models for discrete response," *Journal of Applied Econometrics*, 15(5), 447–470.
- MERLO, A., AND A. DE PAULA (2017): "Identification and Estimation of Preference Distributions When Voters Are Ideological," *Review of Economic Studies*, 84, 1238–1263.
- NORETS, A., AND S. TAKAHASHI (2013): "On the surjectivity of the mapping between utilities and choice probabilities," *Quantitative Economics*, 4(1), 149–155.
- NOSKO, C. (2010): "Competition and quality choice in the cpu market," working paper, Harvard University.
- PALAI, R. S. (1959): "Natural operations on differential forms," *Transactions of the American Mathematical Society*, 92(1), 125–141.
- PANG, J.-S., C.-L. SU, AND Y.-C. LEE (2015): "A constructive approach to estimating pure characteristics demand models with pricing," *Operations Research*, 63(3), 639–659.
- PARTHASARATHY, T. (1983): *On Global Univalence Theorems*. Springer.
- PEYRÉ, G., AND M. CUTURI (2017): *Computational Optimal Transport*. Forthcoming.
- PRESCOTT, E. C., AND M. VISSCHER (1977): "Sequential location among firms with foresight," *The Bell Journal of Economics*, pp. 378–393.
- RADULESCU, M., AND S. RADULESCU (1980): "Global inversion theorems and applications to differential equations," *Nonlinear Analysis: Theory, Methods & Applications*, 92, 951–965.
- ROTH, A. E. (1984): "The evolution of the labor market for medical interns and residents: a case study in game theory," *The Journal of Political Economy*, pp. 991–1016.
- ROTH, A. E., AND M. A. O. SOTOMAYOR (1992): *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press.
- SEDGEWICK, R., AND K. WAYNE (2011): *Algorithms*. Addison-Wesley.
- SHAPLEY, L. S., AND M. SHUBIK (1971): "The assignment game I: The core," *International Journal of Game Theory*, 1(1), 111–130.

- SONG, M. (2007): "Measuring Consumer Welfare in the CPU Market: an Application of the Pure-Characteristics Demand Model," *RAND Journal of Economics*, 38(2), 429–446.
- TOPKIS, D. (1998): *Supermodularity and Complementarity*. Princeton University Press.
- VEINOTT, A. (2005): "Lecture notes in Supply-Chain optimization," Stanford University.
- WALSH, J., AND L. DIECI (2017): "General auction method for real-valued optimal transport," *arXiv preprint arXiv:1705.06379*.