

Estimating Matching Games with Transfers

Jeremy T. Fox*

Rice University and NBER

January 2017

Abstract

I explore the estimation of transferable utility matching games, encompassing many-to-many matching, marriage and matching with trading networks (trades). Computational issues are paramount. I introduce a matching maximum score estimator that does not suffer from a computational curse of dimensionality in the number of agents in a matching market. I apply the estimator to data on the car parts supplied by automotive suppliers to estimate the valuations from different portfolios of parts to suppliers and automotive assemblers.

*Thanks to SupplierBusiness as well as Thomas Klier for help with the work on automotive supplier specialization. I thank the National Science Foundation, the NET Institute, the Olin Foundation, and the Stigler Center for generous funding. Thanks to helpful comments from colleagues and workshop participants at many universities and conferences. Chenchuan Li, David Santiago, Louis Serranito and Chenyu Yang provided excellent research assistance. Email: jeremyfox@gmail.com.

1 Introduction

There are many situations in which economists have data on relationships, including marriages between men and women and partnerships between upstream and downstream firms. Economists wish to use the data on the set of realized relationships to estimate the valuations of agents over the characteristics of potential partners and other measured aspects of the relationships. This is a challenging task compared to estimating valuations using more traditional data because we observe only the equilibrium relationships and not each agent’s equilibrium choice set: the identity of the other agents who would be willing to match with a particular agent. We must infer utility parameters from the sorting seen in the data.

This paper presents an estimator for transferable utility matching games. Transferable utility matching games feature prices (or transfers) for relationships but this paper’s method does not use data on the prices. I model the formation of relationships as a competitive equilibrium to the matching with trading networks model of Azevedo and Hatfield (2015, Section 6), which uses a continuum of agents. This model includes many special cases of empirical interest.¹ Equilibrium existence and uniqueness are generically satisfied. In this model, a generalization of a match is called a trade. A trade can include other aspects in addition to the identity of the match partners. In a labor market, a trade could specify the number of hours an employee is to work per week and the number of vacation days per year.

Using this structure, the paper explores the estimation of valuation functions, which represent the structural preferences of agents for matches or trades. The estimator does not require data on the equilibrium prices of trades despite those prices being present in the underlying matching model. Computational challenges are key in matching and a computationally simple matching maximum score objective function is introduced to address the computational challenges (Manski, 1975). In an application, the paper uses the matching maximum score estimator to empirically answer questions related to the car parts industry. I first describe the methodological contribution and then the empirical application.

Computational issues in matching games are paramount and, in my opinion, have

¹Some initial papers on one-to-one, two-sided matching with transferable utility are Koopmans and Beckmann (1957); Gale (1960); Shapley and Shubik (1972); Becker (1973). This paper uses the term “matching game” to encompass a broad class of transferable utility models, including some games where the original theoretical analyses used different names. Hatfield, Kominers, Nichifor, Ostrovsky and Westkamp (2013) is a related trading networks model to Azevedo and Hatfield (2015, Section 6).

limited the use of matching games in empirical work. Matching markets often have hundreds of agents in them, compared to, say, the two to five agents often modeled as potential entrants in applications of Nash entry games in industrial organization. In the car parts data, there are 2627 car parts in one so-called car component category. There is a lot of information on agent characteristics and unknown parameters that can be learned from the observed sorting of car part suppliers to car assemblers. To take advantage of rich data sets, a researcher must use an estimator that allows for the dimensionality of typical problems.

The solution in maximum score is to introduce inequalities that are computationally simple to evaluate. The objective function is proportional to the number of inequalities that are true at a value for the unknown parameters in the valuation function. The inequalities involve only observable characteristics of trades and unknown parameters. While there are unobservables in the true matching model, the maximum score inequalities do not require numerically integrating out unobservables, as in simulation estimators (e.g., McFadden, 1989; Pakes and Pollard, 1989). Further, the integrand in a simulation estimator for a game theoretic model often involves a nested fixed point procedure to compute an equilibrium to the game (e.g., Ciliberto and Tamer, 2009). No nested fixed point procedure is used in maximum score. In many applications the number of possible inequalities will be intractable; only a subset of the valid inequalities in maximum score can be included without losing the estimator's (point or set) consistency.² Unfortunately, the maximum score objective function is discontinuous and requires global optimization methods to construct a point estimator unless smoothed alternatives are considered. Numerical optimization is less of an issue under set identification.

The use of a model with a continuum of agents as the true or limiting matching model dates to the pioneering work on estimating matching games by Choo and Siow (2006). Choo and Siow studied the case of one-to-one, two-sided matching or marriage. They assume that the unobservables have the type I extreme value distribution, resulting in a logit choice model at the agent level and closed form formulas for matching

²Simulation estimators (sometimes based on moment conditions chosen for tractability) can be used for matching when required at least when markets are small and computational resources are large, as in Boyd, Lankford, Loeb and Wyckoff (2013), Sørensen (2007), Agarwal and Diamond (2016) and Fox, Yang and Hsu (forthcoming). Also, Galichon and Salanié (2015) introduce an estimator for transferable utility, one-to-one matching games that takes a parametric approach that might use numerical integration to compute say choice probabilities or expected utilities but avoids computing equilibria for different realizations of simulation draws.

patterns.³ Chiappori, Salanié and Weiss (2016) and Galichon and Salanié (2015) in part highlight a key assumption in Choo and Siow that restricts how agents' unobservable components of valuations vary across matches or trades. The identification results in this paper also rely on this type of assumption, which is discussed below.

Previous versions of this paper introduced the matching maximum score estimator for many-to-many, two-sided matching and used the closed form logit formulas from Choo and Siow to show the estimator's consistency for the simpler marriage model. However, it was Graham (2011, Theorem 4.1) who, in a survey article discussing previous drafts of this paper in his Section 4.3, first proved set identification for the marriage model under semiparametric conditions nearly identical to those used for proving set identification for single agent maximum score models (Manski, 1975; Matzkin, 1993; Briesch, Chintagunta and Matzkin, 2002; Fox, 2007). The current version of this paper introduces the matching maximum score estimator, uses the setup of the matching with trading networks model and extends the argument of Graham for marriage to the more general setting. The main methodological contribution of the current paper is not the identification result itself but the idea of using a computationally simple objective function to estimate a complicated matching game with many agents.

This paper was originally part of a larger project including Fox (2007) on maximum score methods for single agent multinomial choice, Fox (2010) on nonparametric identification in matching games, and Fox and Bajari (2013) on an empirical application of the matching maximum score estimator to an FCC spectrum auction. The identification paper does not discuss estimation and the auctions paper only states the estimator for the auction application and does not demonstrate identification or discuss econometric properties.

Graham (2011) and Chiappori and Salanié (2016), both mostly for marriage, as well as Mindruta, Moeen and Agarwal (2016), for the academic field of strategy, are three published surveys that discuss the matching maximum score estimator. The Mindruta et al. survey is particularly useful for those wishing to compare a matching game theoretic approach to working with relationship data to other empirical approaches. I have made code for the matching maximum score estimator available (Santiago and

³Dagsvik (2000) provides logit-based methods for studying matching games where other aspects of a relationship than money are part of the pairwise stable matching. Although he does not emphasize it, one-to-one matching games with transferable utility are a special case of his analysis. Matching games with transfers are also related to models of hedonic equilibria, where estimators typically use data on the prices of trades (Rosen, 1974; Ekeland, Heckman and Nesheim, 2004; Heckman, Matzkin and Nesheim, 2010).

Fox, 2009).

Menzel (2015) studies marriage and shows that a class of semiparametric non-transferable utility matching models (not considered here) converge, as the market grows large, to a parametric matching model with matching formulas quite similar to the matching formulas in the logit transferable utility marriage matching model with a continuum of agents in Choo and Siow (2006). The current paper makes semiparametric assumptions in the continuum. I cite some other methodological papers on matching in the rest of the text.

1.1 Empirical Application to Car Parts

A car is one of the most complex goods that an individual consumer will purchase. Cars are made up of hundreds of parts and the performance of the supply chain is critical to the performance of automobile assemblers and the entire industry. This paper investigates two related questions that are relevant to policy debates on the car parts industry. The first question relates to the productivity loss to suppliers from breaking up large assemblers of cars. During the previous large recession, North American-based automobile assemblers went through financial distress. As a consequence, North American-based assemblers divested or closed both North American brands (General Motor's Saturn) and European brands (Ford's Volvo) and seriously considered the divestment of other brands (GM's large European subsidiary Opel). One loss from divesting a brand is that future product development will no longer be coordinated across as many brands under one parent company. If GM were to divest Opel, which was a serious policy debate in Germany in 2009, then any benefit from coordinating new products across Opel and GM's North American operations would be lost. This is a loss to GM, but also to the suppliers of GM, who will no longer be able to gain as much from specializing in supplying GM. I estimate the valuations to suppliers like Johnson Controls and to assemblers for different portfolios of car parts.

The second question this paper investigates is the extent to which the presence of foreign and in particular Japanese and Korean (Asian) assemblers in North America improves the North American supplier base. There is a general perception that Asian automobile assemblers produce cars of higher quality (e.g., Kamath and Liker, 1994; Langfield-Smith and Greenwood, 1998; Liker and Wu, 2000). Part of producing a car of higher quality is sourcing car parts of higher quality. Therefore, Asian assemblers located in North America might improve North American suppliers' qualities. Under-

standing the role of foreign entrants on the North American supplier base is important for debates about trade barriers that encourage Asian assemblers to locate plants in North America in order to avoid those barriers. Trade barriers might indirectly benefit North American assemblers by encouraging higher quality North American suppliers to operate in order to supply Asian-owned assembly plants in North America.

I answer both of the above questions using the identities of the companies that supply each car part. The data list each car model, each car part on that model and, importantly, the supplier of each car part. The intuition is that the portfolio of car parts that each supplier manufactures informs us about the factors that make a successful supplier. If each supplier sells car parts to only two assemblers, it may be that suppliers benefit from specialization at the assembly firm level. If North American suppliers to Asian-owned assemblers are also likely to supply parts to North American-owned assemblers, it may be because of a quality advantage that those suppliers have.

This paper takes the stand that the sorting pattern of sellers (suppliers like Bosch and Delphi) to buyers (assemblers like General Motors and Toyota) informs us about so-called valuation functions, key components of total profits, generating the payoffs of particular portfolios of car part trades to assemblers and to suppliers. In turn, the valuation functions for assemblers and suppliers help us answer the policy questions about government-induced divestitures and foreign assembler plants in North America. The loss to a supplier from GM divesting Opel occurs when supplying two car parts to a large parent company generates more valuation than supplying one car part each to two different assemblers. Thus, the valuation of a portfolio of trades is not necessarily the sum of the valuations from individual trades. The portfolio of trades of each firm is critical for valuations. Therefore, valuation functions are not additively separable across multiple trades, as they are in most prior empirical work on matching not employing the maximum score estimator introduced in this paper.

I model the market for car parts as a two-sided, many-to-many matching game, with the two sides being assemblers and suppliers. In a competitive equilibrium, each firm will form the trades, car part transactions, that maximize its profits at the market-clearing prices. However, those prices are confidential contractual details not released to researchers. This motivates the use of the matching maximum score estimator without price data.

2 Matching Game

This paper discusses estimation of the transferable utility matching game with a continuum of agents in the matching with trading networks model of Azevedo and Hatfield (2015, Section 6), or AH (I refer to their Section 6 because they consider several other models). Valuations in the matching with trading networks model encompass valuations for a great many applications of empirical interest, including marriage and many-to-many, two-sided matching with valuations defined over sets of matches. The trading networks model has desirable properties: a unique competitive equilibrium exists under fairly innocuous technical conditions.

The arguments in this section lead to an estimation approach that uses data on one large matching market, as considered previously for one-to-one matching by Choo and Siow (2006) and the related work cited in the introduction. The asymptotic argument increases the sample size as the number of agents observed in the data grow large. In an early application, Fox and Bajari (2013) use the matching maximum score estimator and the large market asymptotic argument to study a large spectrum auction.

In the large market asymptotic argument, the limiting matching game has a continuum of agents and there exists a unique equilibrium to this matching game. The equilibrium is deterministic in the aggregate. As a researcher collects more data, the asymptotic fiction is that the researcher is observing more agents from this limiting game. Therefore, the asymptotic fiction of collecting more data does not alter the outcome of the matching game in question; the researcher is merely learning more about an existing market with a continuum of agents.

2.1 The Matching Game

I borrow much of the terminology and the notation from AH. I first lay out the general model and then discuss examples below. Let there be a set of **full agent types** I and a *finite* set of **trades** Ω . An agent of type $i \in I$ has a **valuation function** $v^i(\Phi, \Psi)$, where $\Phi \subseteq \Omega$ is the set of trades for which agent i is a **buyer** and $\Psi \subseteq \Omega$ is the set of trades for which agent i is a **seller**. The valuation function $v^i(\Phi, \Psi)$ takes on values in $[-\infty, \infty)$. The empty set \emptyset refers to making no trades as, say, a buyer; $v^i(\emptyset, \emptyset)$ is normalized to 0.

Consider a **price** p_ω for each trade $\omega \in \Omega$. Let $p^\Omega = (p_\omega)_{\omega \in \Omega}$ be the **price vector** for all trades $\omega \in \Omega$. Under transferable utility, the **profit** of an agent i who buys

trades Φ and sells trades Ψ at the prices p^Ω is

$$v^i(\Phi, \Psi) - \sum_{\omega \in \Phi} p_\omega + \sum_{\omega \in \Psi} p_\omega. \quad (1)$$

There is a measure $\eta(i)$ over the set of agent types $i \in I$.⁴ An **allocation** A is a map from the set of agent types I to the space of distributions over the product space formed by two power sets of Ω ,

$$\mathcal{P}(\Omega) \times \mathcal{P}(\Omega).$$

For each type $i \in I$, the allocation A specifies a distribution $A^i(\Phi, \Psi)$ over sets of trades as a buyer Φ and as a seller Ψ ; each $A^i(\Phi, \Psi)$ is the fraction of agents of type i that conduct the trades Φ and Ψ .

An **arrangement** (A, p^Ω) is comprised of an allocation A and a price vector p^Ω . The allocation A is **incentive compatible** given the price vector p^Ω if each agent maximizes its profits (1) in the sense that $A^i(\Phi, \Psi) > 0$ only if

$$(\Phi, \Psi) \in \arg \max_{\tilde{\Phi} \subseteq \Omega, \tilde{\Psi} \subseteq \Omega} \left(v^i(\tilde{\Phi}, \tilde{\Psi}) - \sum_{\omega \in \tilde{\Phi}} p_\omega + \sum_{\omega \in \tilde{\Psi}} p_\omega \right).$$

The allocation A is **feasible** if the **excess demand** for each trade $\omega \in \Omega$,

$$\int_I \left(\sum_{\Phi \supseteq \{\omega\}} \sum_{\Psi} A^i(\Phi, \Psi) - \sum_{\Psi \supseteq \{\omega\}} \sum_{\Phi} A^i(\Phi, \Psi) \right) d\eta(i), \quad (2)$$

equals 0. In the definition of excess demand, the sums are over subsets of the finite set of trades Ω and $\Phi \supseteq \{\omega\}$ means sum over sets where the trade ω is an element. The arrangement (A, p) is a **competitive equilibrium** if the allocation A is incentive compatible given the price vector p^Ω and is feasible.⁵

AH prove that a competitive equilibrium exists and is efficient in the sense of the

⁴I place the technical conditions from AH in footnotes. The measure η is defined with respect to some σ -algebra and satisfies $\eta(I) < \infty$. Further, the valuation function v^i must be a measurable function of i .

⁵AH require two further technical conditions. Using their words and skipping their notation: 1) The integral of absolute values of utility is finite as long as agents are not given bundles for which they have utility of $-\infty$ and 2) agents can supply any sufficiently small net demand for trades.

allocation component A being the supremum of the social welfare function

$$\int_I \left(\sum_{\Phi} \sum_{\Psi} v^i(\Phi, \Psi) A^i(\Phi, \Psi) \right) d\eta(i), \quad (3)$$

where the supremum is taken over feasible allocations A . If the distribution of full types $\eta(i)$ ensures uniqueness of the maximizer of the social welfare function, then there is a unique allocation A corresponding with a competitive equilibrium. Further, AH prove that if $\eta(i)$ has full support in a precise sense then the price vector p^Ω in the competitive equilibrium is unique. Given these relatively weak conditions for equilibrium existence and uniqueness, I maintain existence and uniqueness in what follows. What actually matters for the empirical approach is that the same competitive equilibrium is being played by all agents in the continuum market.

For reasons of technical complexity within AH's proofs, the set of trades Ω is finite. This finiteness of Ω is a key reason for a technical distinction between set and point identification that will be referred to several times.

2.2 Observables and Unobservables

The matching with trading networks model has complete information in the sense that no attributes of a trade are privately observed. Still, econometricians wish to distinguish between attributes of agents measured in the data and attributes not measured in the data. Recall i indexes a full agent type. Let j index an **observable agent type** in a finite set of observable agent types J and let k index an **unobservable agent type** in some, likely infinite, set K , so that each full agent type i uniquely corresponds to a pair of agent types (j, k) . In the logit marriage model of Choo and Siow (2006), j is the demographics of an agent and k is the realizations of the agent's logit errors.

Assumption 1. *Each trade $\omega \in \Omega$ encodes the observable agent type $b(\omega) \in J$ for the buyer and the observable agent type $s(\omega) \in J$ for the seller on trade ω but trades do not encode the unobservable agent types k .*

Agents have preferences defined over trades. All aspects of a trade ω are in the data, including the observable agent types of the buyer $b(\omega) \in J$ and seller $s(\omega) \in J$. An agent $i \in I$ is allowed to have preferences over the observable agent types $j \in J$ of the counterparties to trades ω . In some examples, a trade ω encodes *only* the observable agent types of the buyer $b(\omega)$ and seller $s(\omega)$.

Importantly, Assumption 1 does not allow agents valuations $v^i(\Phi, \Psi)$ to be a function of the unobservable types of partners in trades. Therefore, valuations are defined over the observable characteristics of trades only. In the marriage case, Assumption 1 as implemented in papers following Choo and Siow (2006), such as Galichon and Salanié (2015) and Chiappori, Salanié and Weiss (2016), states that agents have valuations over the observable demographics of partners and not the unobserved characteristics of partners. For example, a potential partner might be good looking but, if appearance is not in the data, Assumption 1 rules out agents having valuations defined over the appearances of partners. The advantage of Assumption 1 here is that it leads to a maximum score objective function that is computationally simple to evaluate. Fox, Yang and Hsu (forthcoming) explore an alternative assumption that allows valuations to depend on the unobservable types of partners.

Let the valuation of an agent of full type $i \in I$ corresponding to observable and unobservable types $(j, k) \in J \times K$ be

$$v^i(\Phi, \Psi) = \pi^j(\Phi, \Psi) + \epsilon_{\Phi, \Psi}^k,$$

where π^j is the **valuation function for observable agent type j** over trades (Φ, Ψ) and $\epsilon_{\Phi, \Psi}^k$ is the **unobservable valuation component** for unobservable agent type k and the set of trades (Φ, Ψ) . The valuation function over trades π^j is a function of only the observable agent type j and the sets of trades Φ and Ψ .⁶ An agent of full type i also has an unobservable valuation component over the trades Φ and Ψ , given by $\epsilon_{\Phi, \Psi}^k$. The unobservable valuation components are separate for each set (Φ, Ψ) and so vary for k based on the observable agent types $b(\omega)$ and $s(\omega)$ of the agents on the other side of trades in (Φ, Ψ) . As explained before, $\epsilon_{\Phi, \Psi}^k$ does not depend on the unobservable agent types for the counterparties on trades in (Φ, Ψ) .

Recall the distribution $\eta(i)$ over agent full agent types $i \in I$. Each full agent type i is also the realization of an observable type j and the realization of an unobservable type k that itself indexes a realization of the vector $\epsilon^k = (\epsilon_{\Phi, \Psi}^k)_{\Phi \subseteq \Omega, \Psi \subseteq \Omega}$. The vector ϵ^k is of finite length because the set of trades Ω is finite.⁷ Therefore, $\eta(i)$ induces a joint distribution (CDF) $F(\epsilon^k | j)$ for each observable agent type j .⁸ The vector ϵ^k is independently distributed across agents, conditional on the observable types $j \in J$.

⁶Let $\pi^j(\Phi, \Psi) = -\infty$ if j is not $b(\omega)$ for any $\omega \in \Phi$ or not $s(\omega)$ for any $\omega \in \Psi$.

⁷The notation ϵ^k means the unique realization of the vector of unobservables corresponding to the unobservable agent type k . When k is a random variable, then so is ϵ^k .

⁸One can safely drop the $\epsilon_{\Phi, \Psi}^k$ for a particular j and (Φ, Ψ) when $\pi^j(\Phi, \Psi) = -\infty$ for that j .

I will focus on the semiparametric case, where the valuation function

$$\pi^j(\Phi, \Psi) = \pi_\theta(j, \Phi, \Psi)$$

is known up to a finite vector of parameters θ (the “parametric” in “semiparametric”) and $F(\epsilon^k | j)$ is not known up to a finite number of parameters for each j (so is nonparametrically specified). If the parameters in π^j vary across j , collect them all in θ . I do not assume that $F(\epsilon^k | j)$ is common across j , so heteroskedasticity is allowed. To illustrate practical implementation, I further restrict the valuation function to be linear in the parameters θ ,

$$\pi_\theta(j, \Phi, \Psi) = X(j, \Phi, \Psi)' \theta,$$

where $X(j, \Phi, \Psi)$ is a vector of observables chosen by the researcher.

The parameterization of $\pi_\theta(j, \Phi, \Psi)$ is for empirical convenience; Matzkin (1993) studies single agent, multinomial choice maximum score estimation when each $\pi^j(\Phi, \Psi)$ is nonparametrically specified and Fox (2010) studies nonparametric identification of aspects of $\pi^j(\Phi, \Psi)$ in a matching model without price data. Ideally, the elements of $X(j, \Phi, \Psi)$ should be chosen based on intuition arising from formal results on nonparametric identification of $\pi^j(\Phi, \Psi)$ in the particular matching game being estimated. In what follows, I formally establish only set identification of θ but am motivated by models where the elements of $X(j, \Phi, \Psi)$ are specified in such a way that the coming matching maximum score inequalities do not always difference out an element of the vector θ . Without data on prices, a scale normalization on the vector θ is needed. I pick the normalization that one element of θ is either +1 or -1. The sign of θ will typically be identifiable from the data.

2.3 Examples

Example 1. Consider the monogamous, heterosexual marriage setting in Choo and Siow (2006), Galichon and Salanié (2015) and Chiappori, Salanié and Weiss (2016), which uses a continuum of agents. This is an example of one-to-one, two-sided matching.⁹ Divide agents into males and females. Each observable agent type j corresponds

⁹The two-sidedness (heterosexual marriage) defines Example 1 and is used in the cited empirical literature. The trading networks model can also be specialized to one-sided models of marriage: homosexual marriage or a model with both heterosexual and homosexual marriage.

to a sex (male or female) and other observable demographic characteristics, such as age and race. Age is measured in integer years to have finite support. A trade ω corresponds to a marriage: a male observable type $b(\omega)$ and a female observable type $s(\omega)$. The price p_ω of a trade is exchanged between males and females. Define $\pi^j(\Phi, \Psi)$ to be $-\infty$ if any agent is engaged in more than one marriage or married to an agent of the same sex. Therefore, valuations for a male type i or (j, k) from matching with a female observable type $j_f = s(\omega)$ specialize to

$$\pi^j(j_f) + \epsilon_{j_f}^k,$$

where $\pi^j(j_f)$ is the valuation function for males with observable demographics in j matching to females with observable demographics j_f and $\epsilon_{j_f}^k$ is the preference of a male of unobservable type k for females with demographic characteristics j_f . A symmetric valuation exists for females of type i . The key assumption used in the marriage papers cited above means that males have preferences over female demographics, not the unobservable type of the female they match with. While not considered in the cited empirical literature, it is straightforward to include aspects other than demographics into a trade ω . For example, a trade ω could specify the number of children or the hours of work of each spouse in a marriage. Then the valuation for a male of type i or (j, k) for trade ω would be

$$\pi^j(\omega) + \epsilon_\omega^k.$$

A similar valuation exists for females. Empirical implementation of the more general notion of a trade requires the extra elements to be observable in the data, as in data on labor supply and the number of children for each marriage that occurs in the data.

Example 2. Say an agent is defined to be either a buyer or a seller ex ante, as in the empirical work on the car parts industry later in this paper. Then this is an example of two-sided, many-to-many matching. Define $\pi^j(\Phi, \Psi)$ to be $-\infty$ if an agent whose observable type $j \in J$ corresponds to a buyer conducts trades as a seller, and similarly for a seller type. A trade ω specifies the buyer observable type $b(\omega)$ and the seller observable type $s(\omega)$ in addition to possible other attributes, such as the quantity and quality of goods to deliver (if quantity and quality are specified on a finite grid and observable in the data for actual matches). A buyer of full type i or (j, k) then has

profits of

$$\pi^j(\Phi) + \epsilon_{\Phi}^k - \sum_{\omega \in \Phi} p_{\omega}.$$

As in marriage, the buyer's unobservable valuation component ϵ_{Φ}^k depends on the trades and hence on the observable types $s(\omega) \in J$ of the seller partners. Similarly, a seller full type i or (j, k) has profits of

$$\pi^j(\Psi) + \epsilon_{\Psi}^k + \sum_{\omega \in \Psi} p_{\omega}.$$

Recall that a competitive equilibrium exists without ruling out empirically relevant cases, such as a function $\pi^j(\Phi)$ exhibiting complementarities across multiple trades involving the same agent (e.g., Hatfield and Milgrom, 2005). Complementarities across multiple trades involving the same agent are vital to the empirical application to the car parts industry.

An agent's valuation is directly a function of only the trades where that particular agent is a buyer or a seller. The model assumes away **externalities**: valuations defined over trades to which the agent does not participate. Competition for trades certainly affects the price vector for trades, p^{Ω} , although such competition for trades is not a valuation defined over trades to which the agent does not participate. True externalities could be important in applications; for example if buyers are retailers and sellers are wholesalers and buyers compete with each other for retail customers (outside of the matching game) after matching to sellers. Baccara, Imrohroglu, Wilson and Yariv (2012) use the matching maximum score estimator introduced in this paper to estimate a matching game with externalities.

Example 3. Consider mergers between agents. An agent is not restricted to be a buyer or a seller ex ante. If an agent acquires other agents in equilibrium, it ends up conducting only trades as a buyer although this is not specified ex ante. Likewise, an agent acquired by another agent ends up conducting only a single trade as a seller (if partial acquisitions are not modeled). Therefore, mergers are an example of one-sided matching, also called coalition formation. If desired, one can define $\pi^j(\Phi, \Psi)$ to be $-\infty$ if an agent of type j conducts trades as both a buyer and a seller or if an agent conducts two or more trades as a seller (target). The price p_{ω} of a trade ω captures the price the buyer (acquirer) pays the seller (target). A trade ω specifies the buyer observable agent

type $b(\omega) \in J$ and the seller (target) observable agent type $s(\omega) \in J$. As in Uetake and Watanabe (2016), a trade may also specify other observable attributes, such as the equity split of the post-merger firm or the awarding of board seats to representatives of the acquirer and target. Akkus, Cookson and Hortacsu (2016) use a variant of the matching maximum score estimator with data on the prices of trades p_ω to estimate a matching game of mergers.

As mentioned above, the model does not incorporate externalities such as changes in the post-merger competition for customers between firms in the same industry. Therefore, the model is a better model of mergers when focusing on, say, across-industry conglomerate mergers.

Example 4. Hatfield et al. (2013) mention the example of trading between dealers of used cars. There is a lively secondhand market in used cars. Dealers may both buy and sell used cars to other dealers. Here a trade ω specifies the observable attributes of the used car in question and buyer and seller observable characteristics $b(\omega), s(\omega) \in J$, including the dealer locations. A buyer might have valuations defined over the location of a seller in order to minimize transportation costs. Dealers might have complex preferences over the set of used cars on their lot. For example, valuations might be higher from ending up with cars of only a certain brand or from having a diverse set of cars. The model does not restrict the valuations of dealers over the set of observable trades they undertake. The set of cars that a dealer is endowed with (and possibly does not trade) can be included in the observable agent type $j \in J$.

3 Identification and Estimation

The main purpose of this paper is to propose a tractable estimator for the case where data on the prices of trades p^Ω are *not* available. The majority of this section discusses the matching maximum score estimator that does not use data on the prices of trades.

3.1 Econometric Assumptions and Background

Manski (1975) introduces maximum score estimators for single agent multinomial choice. To discuss some background from single agent multinomial choice maximum

score estimators, say for this subsection only that the researcher has data on the bundles of trades (Φ_i, Ψ_i) for $i = 1, \dots, N$ agents as well as the prices p_ω for all trades $\omega \in \Omega$. As described in Fox (2007), the following conditions are sufficient for the (set) identification of θ and the (set) consistency of a single agent maximum score estimator for θ .

Assumption 2.

1. *The observations (Φ_i, Ψ_i) are i.i.d.*
2. *ϵ^k has full support in $\mathbb{R}^{\dim(\epsilon^k)}$.*
3. *ϵ^k has an exchangeable distribution for each j .*
4. *The parameter space of θ is compact.*

Let ρ be a permutation of the elements of ϵ^k . An **exchangeable distribution** satisfies $F(\epsilon^k | j) = F(\rho(\epsilon^k) | j)$ for all such permutations ρ . Exchangeable distributions allow certain types of equicorrelation across the elements of ϵ^k but rule out some common empirical specifications, such as the random coefficients logit where θ is interpreted as the mean of the random coefficients.

Define the **choice probability** for observable type j (in equilibrium) to be

$$\Pr_j(\Phi, \Psi) = \int_{\epsilon^k} 1 \left[(\Phi, \Psi) \in \arg \max_{\tilde{\Phi} \subseteq \Omega, \tilde{\Psi} \subseteq \Omega} \left(X(j, \tilde{\Phi}, \tilde{\Psi})' \theta + \epsilon_{\tilde{\Phi}, \tilde{\Psi}}^k - \sum_{\omega \in \tilde{\Phi}} p_\omega + \sum_{\omega \in \tilde{\Psi}} p_\omega \right) \right] dF(\epsilon^k | j). \quad (4)$$

This is the same choice probability (or market share equation) from the literature on estimating single agent multinomial choice models (McFadden, 1973). Note that while prices for trades p^Ω are determined in equilibrium, under Assumption 1 prices are not statistically endogenous in the sense of being statistically dependent with k (or ϵ^k).

Fix the observable agent type j . Under Assumption 2, Goeree et al. (2005) and Fox (2007) show, for the matching notation used here, that a **single agent rank order property** holds: $\Pr_j(\Phi_1, \Psi_1) \geq \Pr_j(\Phi_2, \Psi_2)$ if and only if

$$X(j, \Phi_1, \Psi_1)' \theta - \sum_{\omega \in \Phi_1} p_\omega + \sum_{\omega \in \Psi_1} p_\omega \geq X(j, \Phi_2, \Psi_2)' \theta - \sum_{\omega \in \Phi_2} p_\omega + \sum_{\omega \in \Psi_2} p_\omega.$$

Roughly speaking, one can interpret $X(j, \Phi, \Psi)' \theta$ as the mean valuation of the observable type j and the rank order property says that choices with higher mean valuation

plus prices are made more often. The single agent rank order property is a statement about the properties of the exchangeable distribution $F(\epsilon^k | j)$ and how the unobservables enter the choice model. The single agent rank order property is an intermediate result that leads to the (set) identification of θ and the (set) consistency of single agent maximum score.

Further, θ is point identified if one element of the vector $X(j, \Phi, \Psi)$ has **full support** (equal to \mathbb{R}) conditional on the other elements of $X(j, \Phi, \Psi)$, on the price vector p^Ω (which are not random variables in a competitive equilibrium), and also on the vectors $X(j, \Phi, \Psi)$ for other sets of trades (Φ, Ψ) . For technical reasons in equilibrium existence proofs in AH, the set of trades Ω in each market is finite so that there is no story in the model where any element of $X(j, \Phi, \Psi)$ could have support on an interval in \mathbb{R} if only a single matching market is modeled. Formally speaking, if Ω is indeed finite then θ will be set identified.

3.2 Matching Maximum Score Inequalities

To use a maximum score estimator without data on the prices of trades p^Ω , in this subsection I define matching maximum score inequalities and the matching maximum score objective function. In the next subsection, I prove a rank order property for matching without data on prices that a further result uses to prove that the matching maximum score inequalities lead to set identification of the parameter vector θ . The main contribution is the computationally simple objective function for a complex game, not the proofs of theorems.

For intuition, a convenient property of matching with trading networks model is that the allocation portion A of any competitive equilibrium (A, p^Ω) is efficient in the sense of maximizing the social welfare function (3). Therefore, any other allocation should weakly lower social welfare. Note that the competitive equilibrium being efficient is used for intuition for motivating matching maximum score inequalities in the main text. However, the formal proofs in the appendix do not rely on a competitive equilibrium being efficient.

First consider a version of the model where the same pair of two agents can undertake two different trades. In this case, we can base an inequality around two specific trades, ω_1 and ω_2 . The multi-set (allowing duplicates in the set) of the buyer and seller observable types for trade ω_1 , $b(\omega_1), s(\omega_1) \in J$, must equal the multi-set of the buyer and seller observable types for trade ω_2 . The buyer observable type $b(\omega_1)$ on

trade ω_1 could be either the buyer or the seller observable type on trade ω_2 , although a particular maximum score inequality fixes the role of $b(\omega_1)$ on trade ω_2

The deviation from trade ω_1 to trade ω_2 for the two observable types $b(\omega_1), s(\omega_1) \in J$ is feasible because there is both a buyer and a seller for each trade under both circumstances. On the left side of the inequality, agent $b(\omega_1)$ conducts the total trades $(\Phi_{b(\omega_1)}, \Psi_{b(\omega_1)})$ and agent $s(\omega_1)$ conducts the total trades $(\Phi_{s(\omega_1)}, \Psi_{s(\omega_1)})$. Further, let $(\bar{\Phi}_{s(\omega_1)}, \bar{\Psi}_{b(\omega_1)})$ and $(\bar{\Phi}_{s(\omega_1)}, \bar{\Psi}_{s(\omega_1)})$ be the respective trades when the agents $b(\omega_1), s(\omega_1) \in J$ switch from trade ω_1 to ω_2 . Then a **matching maximum score inequality** based on trades ω_1 and ω_2 (and on $(\Phi_{b(\omega_1)}, \Psi_{b(\omega_1)})$ and $(\Phi_{s(\omega_1)}, \Psi_{s(\omega_1)})$) is

$$X(b(\omega_1), \Phi_{b(\omega_1)}, \Psi_{b(\omega_1)})' \theta + X(s(\omega_1), \Phi_{s(\omega_1)}, \Psi_{s(\omega_1)})' \theta \geq X(b(\omega_1), \bar{\Phi}_{b(\omega_1)}, \bar{\Psi}_{b(\omega_1)})' \theta + X(s(\omega_1), \bar{\Phi}_{s(\omega_1)}, \bar{\Psi}_{s(\omega_1)})' \theta. \quad (5)$$

The intuition behind the inequality is that the social welfare for trade ω_1 must be greater than the social welfare when the agent observable types instead engage in trade ω_2 . This motivation is only intuition as the inequality drops the unobservable agent types k (the unobservables in each ϵ^k) and so we must prove a rank order property to show that a maximum score estimator based on this inequality will set identify the true θ .

In some examples of the model, the inequality (5) will not be informative. Returning to Example 1, consider one-to-one, two-sided matching (marriage) where trades ω encode *only* the observable agent types of the buyer and the seller. As trades encode no other features than observable types $j \in J$, a male observable type conducting a marriage trade $\omega_1 \in \Omega$ with a female observable type cannot instead conduct a distinct marriage trade $\omega_2 \neq \omega_1$ with that same female observable type. I now introduce notation for a matching maximum score inequality that generalizes both inequalities for the marriage example as well as the inequalities (5) just introduced.

Let the more general matching maximum score inequality be indexed by g out of some finite set G of possible inequalities. The set G is finite as the set of trades Ω is finite. The set of trades Ω might be infinite in some other matching model like Dupuy and Galichon (2014); this is not a challenge for maximum score. The set G is specified in part by the researcher and will not need to include all feasible inequalities; the set of all feasible inequalities has a combinatorial structure that will often make it computationally intractable to itemize over all feasible inequalities. Each inequality

$g \in G$ will involve either a strict $>$ relation or a weak \geq relation. I will explain the weak \geq relation inequality first.

An inequality $g \in G$ will focus on the two trades ω_1 and ω_2 in the multi-set (allowing duplicates) $\Omega_g = \{\omega_1, \omega_2\}$ on the inequality's left side and the two other trades $\bar{\Omega}_g = \{\omega_3, \omega_4\}$ on the inequality's right side. The trades can include the option of not making a trade in order to explore agents dropping or adding trades. The set of observable types of agents is the same for the left and right sides: the multi-set $H_g = \{b(\omega_1), s(\omega_1), b(\omega_2), s(\omega_2)\}$ must equal the multi-set $\bar{H}_g = \{b(\omega_3), s(\omega_3), b(\omega_4), s(\omega_4)\}$. Further, there is some unique mapping between the agents in H_g and \bar{H}_g in the case of identical observable agent types. For each observable type $j \in H_g$, let (Φ_j^g, Ψ_j^g) be $j \in H_g$'s total trades on the left side of the inequality; the corresponding trade $\omega \in \Omega_g$ where j is a buyer or seller must be in (Φ_j^g, Ψ_j^g) . Likewise, each $(\bar{\Phi}_j^g, \bar{\Psi}_j^g)$ is $j \in H_g$'s total trades on the right side of the inequality, where the corresponding trade $\omega \in \bar{\Omega}_g$ where j is a buyer or seller must be in $(\bar{\Phi}_j^g, \bar{\Psi}_j^g)$ and $(\bar{\Phi}_j^g, \bar{\Psi}_j^g)$ must be equal to (Φ_j^g, Ψ_j^g) for each $j \in H_g$ except for the trade $\omega \in \bar{\Omega}_g$ replacing the corresponding trade $\omega \in \Omega_g$.

The **matching maximum score inequality** g based on the trades $\omega_1-\omega_4$ and the corresponding sets (Φ_j^g, Ψ_j^g) and $(\bar{\Phi}_j^g, \bar{\Psi}_j^g)$ for $j \in H_g$ is defined to be

$$\begin{aligned} X \left(b(\omega_1), \Phi_{b(\omega_1)}^g, \Psi_{b(\omega_1)}^g \right)' \theta + X \left(s(\omega_1), \Phi_{s(\omega_1)}^g, \Psi_{s(\omega_1)}^g \right)' \theta + \\ X \left(b(\omega_2), \Phi_{b(\omega_2)}^g, \Psi_{b(\omega_2)}^g \right)' \theta + X \left(s(\omega_2), \Phi_{s(\omega_2)}^g, \Psi_{s(\omega_2)}^g \right)' \theta \geq \\ X \left(b(\omega_1), \bar{\Phi}_{b(\omega_1)}^g, \bar{\Psi}_{b(\omega_1)}^g \right)' \theta + X \left(s(\omega_1), \bar{\Phi}_{s(\omega_1)}^g, \bar{\Psi}_{s(\omega_1)}^g \right)' \theta + \\ X \left(b(\omega_2), \bar{\Phi}_{b(\omega_2)}^g, \bar{\Psi}_{b(\omega_2)}^g \right)' \theta + X \left(s(\omega_2), \bar{\Phi}_{s(\omega_2)}^g, \bar{\Psi}_{s(\omega_2)}^g \right)' \theta. \quad (6) \end{aligned}$$

The inequality states that the sum of the valuation functions from the two trades in $\Omega_g = \{\omega_1, \omega_2\}$ is greater than the sum of the valuation functions from the two trades in $\bar{\Omega}_g = \{\omega_3, \omega_4\}$.

If $\omega_1 = \omega_2$ and $\omega_3 = \omega_4$, this definition encompasses (5); the inequality (5) would be the appropriate (6) divided by 2 on both sides. For the marriage Example 1, only the sum of the valuations of the male and female is identifiable, so let

$$\tilde{X}(j_m, j_f) = X \left(b(\omega_1), \Phi_{b(\omega_1)}, \Psi_{b(\omega_1)} \right) + X \left(s(\omega_1), \Phi_{s(\omega_1)}, \Psi_{s(\omega_1)} \right)$$

for trade ω_1 where $b(\omega_1) = j_m$ and $s(\omega_1) = j_f$ and plug this definition into (6) when

constructing inequalities.

The matching maximum score inequality g in (6) can be notationally simplified. The parameter vector θ multiplies all four X vectors in the inequality. Therefore we can collect terms by defining the vector

$$Z_g = \sum_{\omega \in \Omega_g} \left(X \left(b(\omega), \Phi_{b(\omega)}^g, \Psi_{b(\omega)}^g \right) + X \left(s(\omega), \Phi_{s(\omega)}^g, \Psi_{s(\omega)}^g \right) \right) - \sum_{\omega \in \bar{\Omega}_g} \left(X \left(b(\omega), \bar{\Phi}_{b(\omega)}^g, \bar{\Psi}_{b(\omega)}^g \right) + X \left(s(\omega), \bar{\Phi}_{s(\omega)}^g, \bar{\Psi}_{s(\omega)}^g \right) \right).$$

Then the matching maximum score inequality g in (6) can be written as $Z'_g \theta \geq 0$.

The researcher chooses the set G of possible matching maximum score inequalities to use in estimation. A *possible* inequality $g \in G$ becomes an *actual* inequality in estimation whenever the configuration of observable agent types and sets of trades on the left side of the inequality (6) is sampled in the data.

The researcher has a lot of freedom to choose the set G of possible inequalities. However, for set identification, inequalities need to be included with both weak \geq and strict $>$ relations. The strict inequality \tilde{g} corresponding to the weak inequality g can be written as either $Z'_g \theta < 0$ or $Z'_{\tilde{g}} \theta > 0$, where $Z_{\tilde{g}} = -Z_g$. Let \succ_g be equal to \geq for inequalities g with weak relations and $>$ for inequalities g with strict relations.

Assumption 3. *The choice of inequalities G must be a union of pairs $\{Z'_g \theta \geq 0, Z'_{\tilde{g}} \theta > 0\}$ where $Z_{\tilde{g}} = -Z_g$.*

The assumption rules out two otherwise mutually exclusive inequalities being true for the same parameter value θ and neither being true. The choice of G likely affects the size of confidence regions just like the distribution of regressors does in regression. The limiting distributions of maximum score estimators are complex enough to prevent a formal yet intuitive analysis of how choices of G affect the size of confidence regions.

3.3 Matching Maximum Score Objective Function

The main contribution of the paper is to introduce a computationally simple objective function for matching games with many agents and large sets of trades Ω . The matching maximum score objective function for a sample of data on the trades (Φ_i, Ψ_i) of $i =$

$1, \dots, N$ agents, but not the prices of trades p^Ω , is

$$\sum_{g \in G_N} 1 [Z'_g \theta \succ_g 0], \quad (7)$$

where G_N are the inequalities to use for this sample. G_N may be a multi-set as the same inequality could appear multiple times if agents of the same observable type $j \in J$ are observed. The matching maximum score, or maximum rank correlation as explained below, objective function checks whether each matching maximum score inequality is true. If an inequality is true for a guess of the parameter vector θ , the objective function increases by 1. Not all inequalities will be true even at the true value of the parameter vector θ because of the unobservables ϵ^k .

Computationally, this linear-in-parameters form is the same form as the inequalities in single agent maximum score and maximum rank correlation estimators (Manski, 1975, 1985; Han, 1987; Fox, 2007). Evaluating the maximum score objective function is computationally simple. The matching maximum score estimator avoids nonparametric estimates of choice probabilities and distributions of unobservables, numerical integration, and algorithms to compute competitive equilibria. The set G of possible inequalities can be chosen for computational convenience. Therefore, evaluating the matching maximum score objective function could be computationally practical when certain alternatives are not. As the introduction indicates, the main methodological contribution of this paper is to introduce a computationally simple matching objective function that facilitates otherwise intractable estimation problems.

Under set identification, one typically evaluates (7) on a grid of θ values as part of a procedure to construct a confidence set. This can be computationally straightforward. Under point identification, one maximizes (7); any maximizer of the step function will provide a consistent estimator. Such a numerical optimization problem may be difficult as the objective function is a step function. Some sort of global optimizer should be used, such as simulated annealing or differential evolution. Horowitz (1992) proposes smoothing the indicator function in the objective function in order to improve the statistical rate of convergence of the estimator, which is not as relevant here as the point identified version of the estimator will converge at the maximum \sqrt{N} rate without smoothing (Sherman, 1993). In terms of optimization, smoothing allows derivative-based Newton solvers to search for local maxima although many starting values will be needed to hope to find a global optimum. To my knowledge finding a global optimum

of a non-convex objective function cannot be ensured for any optimization algorithm.

The matching maximum score objective function (7) can be rewritten in a way that facilitates calculating its expectation and its probability limit under i.i.d. sampling of the trades of agents $i \in I$. With an appropriate normalizing constant, the matching maximum score objective function is also, for $N \geq 4$,

$$\binom{N}{4}^{-1} \sum_{i_1=1}^{N-3} \sum_{i_2=i_1+1}^{N-2} \sum_{i_3=i_2+1}^{N-1} \sum_{i_4=i_3+1}^N \sum_{g \in G} 1 \left[\{(\Phi_i, \Psi_i)\}_{i=i_1, i_2, i_3, i_4} = \{(\Phi_j^g, \Psi_j^g)\}_{j \in H_g} \right] 1 [Z'_g \theta \succ_g 0]. \quad (8)$$

Here the outer four summations form all sets of four agents i in the sample. The inner sum is over all matching maximum score inequalities $g \in G$. For each inequality g , the objective function checks with the multi-set of trades of the four agents is equal to the multi-set of trades on the left side of the inequality g . A trade ω contains the observable types of the buyer and seller, and so the objective function also checks if the observable agent types match the inequality g . If the inequality is satisfied, then the corresponding matching maximum score inequality is included in the maximum score objective function as the indicator

$$1 \left[\{(\Phi_i, \Psi_i)\}_{i=i_1, i_2, i_3, i_4} = \{(\Phi_j^g, \Psi_j^g)\}_{j \in H_g} \right] \quad (9)$$

equals 1 and so the inequality itself, $Z'_g \theta \succ_g 0$, enters the objective function. The model's dependent variable, the trades (Φ_i, Ψ_i) that are undertaken by the N agents in the data, enters the matching maximum score objective function (8) through (9).

3.4 Rank Order Property and Set Identification

The parameter vector θ is set identified using data on trades and not the prices of trades p^Ω when a rank order property holds. The following proposition defines the **rank order property** for matching without data on the prices of trades and states that it holds. Recall that $\Pr_j(\Phi, \Psi)$ is the fraction or choice probability of agents of observable type $j \in J$ that conduct the trades Φ and Ψ . To define a choice probability for one out of two options, focus on the trades (Φ_1, Ψ_1) and the trades (Φ_2, Ψ_2) and let

$$\Pr_j(1 \mid \Phi_1, \Psi_1, \Phi_2, \Psi_2) = \frac{\Pr_j(\Phi_1, \Psi_1)}{\Pr_j(\Phi_1, \Psi_1) + \Pr_j(\Phi_2, \Psi_2)}$$

be the probability that observable type $j \in J$ picks the trades (Φ_1, Ψ_1) conditional on the event that j picks either (Φ_1, Ψ_1) or (Φ_2, Ψ_2) . Define $\Pr_j(2 | \Phi_1, \Psi_1, \Phi_2, \Psi_2)$ similarly. Make the following technical assumption.

Assumption 4. *Assume that $F(\epsilon^k | j)$ has bounded, continuous derivatives.*

Then we can prove a rank order property for data on matches but not prices.

Proposition 1. *A matching maximum score inequality g in (6) for the relation \succ_g holds if and only if the following inequality holds:*

$$\Pr_{b(\omega_1)} \left(1 | \Phi_{b(\omega_1)}^g, \Psi_{b(\omega_1)}^g, \bar{\Phi}_{b(\omega_1)}^g, \bar{\Psi}_{b(\omega_1)}^g \right) \cdot \Pr_{b(\omega_2)} \left(1 | \Phi_{b(\omega_2)}^g, \Psi_{b(\omega_2)}^g, \bar{\Phi}_{b(\omega_2)}^g, \bar{\Psi}_{b(\omega_2)}^g \right) \succ_g \Pr_{b(\omega_1)} \left(2 | \Phi_{b(\omega_1)}^g, \Psi_{b(\omega_1)}^g, \bar{\Phi}_{b(\omega_1)}^g, \bar{\Psi}_{b(\omega_1)}^g \right) \cdot \Pr_{b(\omega_2)} \left(2 | \Phi_{b(\omega_2)}^g, \Psi_{b(\omega_2)}^g, \bar{\Phi}_{b(\omega_2)}^g, \bar{\Psi}_{b(\omega_2)}^g \right). \quad (10)$$

The probability statement involves only buyer probabilities as buyer probabilities are related to seller probabilities by feasibility, (2).¹⁰ It is important that the same two observable agent types' choice probabilities are on the left and right sides of the probability statement (10). Recall that the buyer on, for example, trade $\omega_1 \in \Omega_g$ might be a seller on trade $\omega_3 \in \bar{\Omega}_g$.

In words, the rank order property for matching without price data states that the conditional probability of observing the configuration of trades on the left side of (6) is greater than the conditional probability of observing the configuration of trades on the right side of (6) whenever the sum of valuations involving observable types j and trades ω on the left side of (6) exceed those on the right side of (6). The rank order property allows an estimator based on maximizing matching maximum score inequalities involving only measured observable agent types j and trades ω to be (set) consistent in the presence of unobservables ϵ^k .

The simple matching maximum score inequality for two trades without price data is (5). Using that inequality's notation, Proposition 1 plus algebraic simplification states that the inequality (5) holds if and only if

$$\Pr_{b(\omega_1)} \left(1 | \Phi_{b(\omega_1)}, \Psi_{b(\omega_1)}, \bar{\Phi}_{b(\omega_1)}, \bar{\Psi}_{b(\omega_1)} \right) \geq \Pr_{b(\omega_1)} \left(2 | \Phi_{b(\omega_1)}, \Psi_{b(\omega_1)}, \bar{\Phi}_{b(\omega_1)}, \bar{\Psi}_{b(\omega_1)} \right),$$

as there is only one trade on each side of the inequality (5).

¹⁰The trades of sellers in the inequality g other than ω_1 – ω_4 contribute to the sampling of inequalities in the data. This does not affect the statement of the rank order property, as the proof indicates.

The proof of Proposition 1 is in an appendix. A full proof of the proposition uses Assumption 2 to apply the single agent rank order property for multinomial choice, mentioned above (Manski, 1975; Fox, 2007). In addition to single agent maximum score results, the full proof of the theorem also uses properties that only hold in competitive equilibrium; thus the proof uses matching theory in addition to manipulating the integrals in the definition of a choice probability (4). The survey of Graham (2011, Theorem 4.1) was the first to prove that the rank order property holds in a semiparametric model of marriage, namely the assumption that $F(\epsilon^k | j)$ is an exchangeable distribution for all j . Proposition 1 extends Graham’s result from marriage to the full generality of matching with trading networks. To clarify the intellectual contribution, the proof of Proposition 1 in the appendix cites Graham’s result for marriage and shows how the inequalities in Graham can be modified to derive the probability statement (10).¹¹ Note that previous circulating working paper versions of the current paper pointed out that the rank order property was implied by the closed form allocation (matching) probabilities in the parametric Choo and Siow (2006) model of marriage, which uses the type I extreme value (logit) distribution for each element of the vector ϵ^k .

Using the rank order property in Proposition 1, we can prove that the model is **set identified**, meaning that the set of maximizers of the expectation of the maximum score objective function (8) contains the true parameter. Under Assumption 2, the expectation is equal to the probability limit based on i.i.d. sampling of agents i and their trades (Φ_i, Ψ_i) .

Theorem 1. *The set of maximizers θ of the expectation of the maximum score objective function (8) contains the true parameter vector.*

Say one further imposes that an element of the vector Z_g has support on \mathbb{R} conditional on other elements of Z_g , contradicting the finite set of trades Ω in the matching with trades model, but not the marriage model of Dupuy and Galichon (2014). Say also that the elements of Z_g are linearly independent. The proof of the (point) consistency of maximum score would then verify that set identification reduces to point identification and that other conditions in a general consistency theorem for extremum

¹¹Graham’s proof works by inverting choice probabilities. See the erratum Graham (2013). Note that Graham (2011, Theorem 4.1) is stated for the marriage equivalent of independent and identical $\epsilon_{\Phi, \Psi}^k$ conditional on j instead of an exchangeable $F(\epsilon^k | j)$. The first two steps of Graham’s proof reproduce Manski (1975) and Fox (2007), so the assumption of an exchangeable $F(\epsilon^k | j)$ can be used with little change, as in Fox (2007). Graham (2011, Theorem 4.1) allows heteroskedasticity as well.

estimators are satisfied (Newey and McFadden, 1994). Fox (2007) provides such a proof for single-agent, multinomial choice maximum score and all the steps go through for the matching case as well.

Set inference can use a method such as the subsampling approach of Romano and Shaikh (2008), which was used in single agent maximum score under set identification by Bajari, Fox and Ryan (2008). The key regularity condition to apply subsampling is that an appropriately normalized version of the objective function has a limiting distribution. Indeed, an input into a procedure such as Romano and Shaikh is the rate of convergence of the objective function. One can recognize that the objective function has the same double (or more) summations as the maximum rank correlation objective function of Han (1987). Sherman (1993) shows that the maximum rank correlation estimator is \sqrt{N} consistent (under point identification) and asymptotically normal and Subbotin (2007) shows that the bootstrap is valid for inference. Under set identification, the code of Santiago and Fox (2009) in part implements the method of Romano and Shaikh (2008) to construct valid 95% confidence intervals using the rate of convergence \sqrt{N} as an input.

3.5 Multiple Markets

Let there now be data on D markets, indexed by d . Then the matching maximum score objective function (7) has an extra summation over markets,

$$\sum_{d=1}^D \sum_{g \in G_{N_d}^d} 1 [Z'_{g,d} \theta \succ_{g,d} 0], \quad (11)$$

where now the set of inequalities to include $G_{N_d}^d$ is specific to market d .

Maintain that each market is truly a continuum of agents and we merely have data on a subset of N_d agents from market d . There are two asymptotic arguments. The first asymptotic argument fixed the number of markets D and increases the number of agents with recorded data N_d for each d by some common amount \bar{N} such that $N_d = \nu_d \cdot \bar{N}$, for the fixed-with- \bar{N} and market-specific proportionality constants ν_d . The estimator will have the maximum rank correlation asymptotics in some notion of the number of agents in each market \bar{N} , meaning the rate of convergence of the point identified case will be $\sqrt{\bar{N}}$.¹²

¹²When proving that the true parameter maximizes the probability limit of the objective function,

Now consider fixing the number of agents with measured data N_d in each market. Each true matching market is still a continuum. A second asymptotic argument makes the number of markets D increase to infinity. Comparing the expectation of the multi-market objective function

$$\frac{1}{D} \sum_{d=1}^D \sum_{i_1=1}^{N_d-3} \sum_{i_2=i_1+1}^{N_d-2} \sum_{i_3=i_2+1}^{N_d-1} \sum_{i_4=i_3+1}^{N_d} \sum_{g \in G} 1 \left[\{(\Phi_{i,d}, \Psi_{i,d})\}_{i=i_1, i_2, i_3, i_4} = \left\{ \left(\Phi_j^{g,d}, \Psi_j^{g,d} \right) \right\}_{j \in H_{g,d}} \right] 1 [Z'_{g,d} \theta \succ_{g,d} 0] \quad (12)$$

to the expectation of the single-market matching maximum score objective function in the proof of Theorem 1, the expectation of (12) involves an outer expectation over the distribution of full agent types $\eta_d(i)$ for market d and hence the resulting competitive equilibrium. This extra outer expectation over $\eta_d(i)$ does not change the conclusion of Theorem 1: the true parameter value is a global maximizer of the expectation of the objective function. Indeed, the parameter θ could be point identified if the vector $Z_{g,d}$ includes one element with full support across markets, even if the set of trades Ω_d is finite for each market d separately.

For the matching maximum score objective function in (12), the number of terms in each quadruple summation does not increase with D and so the estimator will have the maximum score asymptotics in Kim and Pollard (1990). For point identification, the estimator converges at the rate of $D^{1/3}$. Delgado, Rodriguez-Poo and Wolf (2001) show that subsampling is valid for inference. Under set identification, the code of Santiago and Fox (2009) in part implements the method of Romano and Shaikh (2008) to construct valid 95% confidence intervals using the rate of convergence $D^{1/3}$ as an input.¹³

Maximum score estimators allow heteroskedasticity. In matching, even if the observable type $j \in J$ has the same meaning across markets, the distribution of unobservables for j need not be the same across markets. One can notate the distribution

as in Theorem 1, one applies the rank order property in Proposition 1 to inequalities from each market separately.

¹³For the example of single agent binary choice, sufficient conditions for set identification in the maximum score model of Manski (1975) allow for heteroskedasticity (here $F(\epsilon^k | j)$ varies with j) while known sufficient conditions for set identification in the maximum rank correlation model of Han (1987), applied to binary choice, require homoskedasticity (here $F(\epsilon^k | j)$ does not vary with j). In matching, the “maximum score” and “maximum rank correlation” asymptotic arguments both allow for heteroskedasticity, based on the rank order property in Proposition 1. The proof of Proposition 1 relies on the properties of competitive equilibrium and so the theorem is not an analog to the conditions for set identification for binary choice in Han (1987).

of unobservables as $F(\epsilon^k | j, d)$ for observable type $j \in J$ in market $d \in D$.

4 Car Parts Industry

I now present an empirical application about the matching of assemblers to car parts suppliers in the automobile industry. Automobile assemblers are well-known, large manufacturers, such as BMW, Ford or Honda. Automotive suppliers are less well-known to the public, and range from large companies such as Bosch to smaller firms that specialize in one type of car part. A car is one of the most complicated manufacturing goods sold to individual consumers. Making a car be both high quality and inexpensive is a technical challenge. Developing the supply chain is an important part of that challenge. More so than in many other manufacturing industries, suppliers in the automobile industry receive a large amount of coverage in the industry press because of their economic importance.

A matching opportunity in the automotive industry is an individual car part that is needed for a car model. A particular trade $\omega \in \Omega$ encodes an individual car part that is needed for a named car model as well as buyer and seller observable types. Each car model itself has a brand. For the Chevrolet Impala, Chevrolet is the brand and Impala is the model. There are multiple consummated trades ω for the Chevrolet Impala because each model uses multiple parts. Finally, each brand is owned by an assembler, in Chevrolet's case General Motors. General Motors is the buyer observable type $b(\omega) \in J$ on all the trades for car parts used on the Chevrolet Impala. The seller observable type $s(\omega) \in J$ on each trade is a particular car parts supplier, like Bosch. Therefore, each named firm represents a separate observable type in the matching model.¹⁴ In the matching game, there is one trade $\omega \in \Omega$ for the windshield on the Chevrolet Impala for each supplier $s(\omega)$ that supplies at least one windshield for any of the models in the data and so could (according to the model) counterfactually supply the windshield for the Impala.

I refer to Example 2 and model the car parts industry as an explicitly two-sided market, where each supplier conducts trades only as a seller and each assembler conducts

¹⁴The economic questions considered here focus on supplier and assembler specialization and so I need to allow each firm to be its own observable type $j \in J$ to properly measure specialization. The AH model uses a finite number of trades $\omega \in \Omega$ and a continuum of full agent types $i \in I$. In my empirical version of the AH model, the number of observable types $j \in J$ is also finite. Here, the fiction mapping the continuum AH model to the finite data is that there is a continuum of firms of the General Motors observable type but only one such firm is sampled in the data.

trades only as a buyer. From an automotive engineering perspective, an assembler needs a specific set of car parts to make a particular model. For example, each car model needs a single windshield.¹⁵ For an assembler i , define $v^i(\Phi) = -\infty$ for any set of trades Φ that do not contain exactly one trade for every car part opportunity in the data. On the right side of a matching maximum score inequality (6), I drop inequalities where $v^i(\bar{\Phi}) = -\infty$. Therefore, a counterfactual trade ω_3 corresponds to, for example, General Motors using a different supplier for the windshield for the Chevrolet Impala, not General Motors installing two different windshields on the Chevrolet Impala or replacing the windshield with a tire. By the market definition discussed below, only suppliers that make at least one windshield in the data can be the seller observable type on the counterfactual trade ω_3 in a matching maximum score inequality (6) based on swapping windshield suppliers.

The car parts data come from SupplierBusiness, an analyst firm. I merge them with car sales data collected from several sources for the United States and several large countries in Western Europe. I focus on 30 large component categories, such as air conditioning parts, body parts and transmission parts. In the merged and cleaned data, there are 941 suppliers, 11 assemblers (parent companies), 46 car brands, 260 car models, and 34,836 car parts. While the data cover different model years, for simplicity I ignore the time dimension and treat each market as clearing simultaneously.¹⁶ I treat each component category as a statistically independent matching market.¹⁷ Therefore, I use the matching maximum score objective function for multiple markets in (11).¹⁸

One of the empirical applications focuses on General Motors divesting Opel, a brand it owns in Europe. In order to model the interdependence of the European and North American operations of General Motors and suppliers to General Motors, the definition of a matching market is car parts in a particular component category used in cars assembled in Europe and North America. Most of the assemblers and many

¹⁵The data do not report back up or secondary suppliers for a part on a particular car model.

¹⁶Car models are refreshed around once every five years.

¹⁷The same supplier may appear in multiple component categories, and so a researcher might want to model spillovers across component categories. Pooling component categories into one large market creates no new issues with the AH model or the matching maximum score estimator. The history of the industry shows that many US suppliers were formed in the 1910's and 1920's around Detroit (Klier and Rubenstein, 2008). Some firms chose to specialize in one or a few component categories and others specialized in more component categories. The particular historical pattern of what component categories each supplier produces lies outside of the scope of this investigation.

¹⁸The parameter estimates in this paper would presumably change if SupplierBusiness aggregated or disaggregated car parts into component categories in different ways.

of the larger suppliers operate on multiple continents.¹⁹ However, the point estimates found when splitting Europe and North America into separate matching markets are similar to those presented here, suggesting that geographic market definitions do not play a large role in identifying the parameters. Note that many of the estimated gains to specialization to a supplier likely come from plant co-location: using one supplier plant to supply the same type of car part to multiple car models assembled in the same plant or in nearby plants. Thus, an empirical regularity of certain suppliers being more prevalent in one continent than another is consistent with the gains to specialization that I seek to estimate.²⁰ The data have poor coverage for car models assembled in Asia, so I cannot include the corresponding car parts in the empirical work. I do focus heavily on car parts used on cars assembled in Europe and North America by assemblers with headquarters in Asia.

The automotive supplier empirical application is a good showcase for the strengths of the matching maximum score estimator. The matching markets modeled here contain many more agents than the markets modeled in many non-marriage papers on estimating matching games. The computational simplicity of maximum score, or some other approach that avoids repeated computations of model outcomes, is needed here. I focus on specialization in the portfolio of matches for suppliers and assemblers. Along with our related use of the estimator introduced here in Fox and Bajari (2013), an earlier draft of the current paper was the first empirical application to a many-to-many matching market where the valuation from a set of matches (or trades) is not additively separable across the individual matches. Finally, the prices of the car parts are not in publicly available data. The matching estimator does not require data on the prices of trades, even though prices are present in the economic model being estimated.

¹⁹Nissan and Renault are treated as one assembler because of their deep integration. Chrysler and Daimler were part of the same assembler during the period of the data.

²⁰A few suppliers are owned by assemblers. I ignore the vertical integration decision in my analysis, in part because I lack data on supplier ownership and in part because vertical integration is just an extreme version of specialization, the focus of my investigation. If a supplier sends car parts to only one assembler, that data are recorded and used as endogenous matching outcomes. Vertical integration in automobile manufacturing has been studied previously (Monteverde and Teece, 1982; Novak and Eppinger, 2001; Novak and Stern, 2008, 2009).

5 Costs of Assemblers Divesting Brands

5.1 General Motors and Opel

In 2009, General Motors (GM), the world's largest automobile assembler for most of the twentieth century, declared bankruptcy. As part of the bankruptcy process, GM divested or eliminated several of its brands, including Pontiac and Saturn in North America and SAAB in Europe. Economists know little about the benefits and costs of large assemblers in the globally integrated automobile industry divesting brands. This paper seeks to use the matching patterns in the car parts industry to estimate one aspect of the costs of divestment.

A major public policy issue during 2009 was whether General Motors should also divest its largest European brands, Opel and Vauxhall.²¹ Opel is based in Germany and Vauxhall is based in the United Kingdom. Consistent with the close link between Opel and Vauxhall, they will be grouped together as one brand, Opel, in the empirical work. Over the period of the data, Opel also had assembly plants in Belgium, Hungary, Poland, and Russia.

A major advocate of GM divesting Opel was the German government, which desired to protect jobs at Opel assembly plants, at Opel dealers and at suppliers to Opel, but was reluctant to subsidize a bankrupt North American firm. During most of 2009, the presumption by GM was that Opel would be divested. Indeed, GM held an auction and agreed to sell Opel to a consortium from Canada and Russia. In November 2009, GM canceled the sale and kept Opel as an integrated subsidiary of GM. Opel and the North American operations of GM share many common platforms for basing individual models on. One reason for keeping Opel integrated is that a larger, global assembler will have gains from specialization in its own assembly plants and in the plants of suppliers. Increasing the gains to suppliers from specializing in producing car parts for GM may indirectly benefit GM through lower prices for car parts.

5.2 Valuation Functions of Observable Types

This section estimates the parameters in the valuation functions over observable types for assemblers and suppliers for the portfolio of car part trades each firm buys or sells. Let the notation for the observable type j^s emphasize that the firm in question is a

²¹GM has owned Opel since 1929, although its control temporarily lapsed during the second World War.

supplier or seller and the notation j^b emphasize that the firm is a buyer or assembler, as the car parts industry is an explicitly two-sided market. I use the functional forms $\pi_\theta(j^b, \Phi) = X(j^b, \Phi)' \theta^b$ for buyers and $\pi_\theta(j^s, \Psi) = X(j^s, \Psi)' \theta^s$ for sellers, with $\theta = (\theta^b, \theta^s)$. The elements of the vectors $X(j^b, \Phi)$ and $X(j^s, \Psi)$ are measures of how specialized each portfolio of car part trades is at several levels.

5.2.1 Valuation Functions for Suppliers

For suppliers, $X(j^s, \Psi)$ tracks specialization in four areas: parts (in the same component category) for an individual car, parts for cars from a particular brand (Chevrolet, Audi), parts for cars from a particular parent company or assembler (General Motors, Volkswagen) and parts for cars for brands with headquarters on a particular continent (Asia, Europe, North America).

The choice of a measure of specialization is somewhat arbitrary. I use the Herfindahl-Hirschman Index (HHI) because economists are familiar with its units, which range between 0 and 1. For example, say the North American firms of Chrysler, General Motors and Ford are the only three assemblers. Then the corresponding parent group scalar element $X_{\text{PG}}(j^s, \Psi)$ of the vector $X(j^s, \Psi)$ is

$$X_{\text{PG}}(j^s, \Psi) = \left(\frac{\# \text{ Chrysler parts in } \Psi}{\# \text{ total parts in } \Psi} \right)^2 + \left(\frac{\# \text{ Ford parts in } \Psi}{\# \text{ total parts in } \Psi} \right)^2 + \left(\frac{\# \text{ GM parts in } \Psi}{\# \text{ total parts in } \Psi} \right)^2. \quad (13)$$

As this specialization measure enters the valuation function for a supplier, $X_{\text{PG}}(j^s, \Psi)$ is 1 if the supplier sells parts only to, say, GM and 1/3 if it sells an equal number of parts to each assembler. The use of the HHI differs from antitrust; here the HHI is a measure of specialization for a portfolio Ψ of car part trades for a particular supplier j^s and is not a measure of concentration in the overall industry for car parts. The specialization measure $X_{\text{PG}}(j^s, \Psi)$ can be computed both for the trades Ψ for a supplier in the data and in the counterfactual trades in a matching maximum score inequality (6).²²

When I consider the counterfactual of GM divesting Opel and making it an independent assembler or parent company, the changes in total valuation will be generated

²²Many other upstream firm characteristics would be endogenous at the level of the competitive equilibrium considered here. For example, many of the benefits of specialization occur through plant co-location and so suppliers and assembler plant locations should be considered endogenous matching outcomes rather than exogenous firm characteristics. With just-in-time production at many assembly sites, supplier factories are built short distances away so parts can be produced and shipped to the assembly site within hours, in many cases. Plant location could be added as an extra element to a trade ω in other work.

by the estimated parameter on the importance of specialization at the parent company level, relative to the values of the other parameters.

The pattern of sorting across trades in the car parts market is used to measure the relative importance of specializing at different levels of aggregation. The management literature has suggested that supplier specialization may be a key driver of assembler performance (Dyer, 1996, 1997; Novak and Wernerfelt, 2012).

By construction, two parts for the same car model also have the same brand, parent group and continent. Two car parts for cars from the same brand are automatically in the same parent group and the brand only has one headquarters, so the parts are from a brand with a headquarters in the same continent as well. Two cars from the same parent group are not necessarily from the same continent, as Opel is a European brand of GM and Chevrolet is a North American brand of GM.

The four specialization measures in $X(j^s, \Psi)$ are highly correlated. Just as univariate linear least squares applied to each covariate separately produces different slope coefficients than multivariate linear least squares when the covariates are correlated, a univariate matching theoretic analysis (such as Becker (1973)) on each measure in $X(j^s, \Psi)$ separately will be inadequate here. A univariate analysis of say $X_{PG}(j^s, \Psi)$ would just amount to saying that the corresponding element of θ is positive when each supplier does more business with certain parent groups than others. In principle, even this conclusion about the sign of the parameter could be wrong if the correlation with the other three characteristics is not considered in estimation. Here I measure the relative importance of each of the four types of specialization: at which level do the returns to specialization occur?

5.2.2 Valuation Functions for Assemblers

The valuation function of assemblers has a similar functional form, focusing on specializing in a small number of suppliers. Let Φ be a portfolio of car part trades for buyer or assembler j^b . I consider parent group, brand and model specialization in the vector $X(j^b, \Phi)$. For conciseness, I do not include a term for specialization at the continent of brand headquarters level.

Consider the Herfindahl index for the concentration of suppliers selling parts to an assembler. Given a portfolio Φ , let $s(\Phi)$ be the set of distinct suppliers who sell at

least one car part trade in Φ . Then define the scalar

$$X_{\text{PG}}(j^b, \Phi) = \sum_{i \in s(\Phi)} \left(\frac{\# \text{ trades sold by supplier } i \text{ in } \Phi}{\# \text{ total trades in } \Phi} \right)^2.$$

Next, $X_{\text{Br}}(j^b, \Phi)$ is the mean of such a Herfindahl index computed for each brand separately. Consider GM and say that the only two brands of GM are Chevrolet (Chevy) and Opel and let $s(\Phi, \text{Opel})$ be the set of suppliers selling parts to Opel in Φ . Then, for GM,

$$\begin{aligned} X_{\text{Brand}}(j^b, \Phi) &= \frac{1}{2} \sum_{i \in s(\Phi, \text{Opel})} \left(\frac{\# \text{ trades sold by supplier } i \text{ to Opel in } \Phi}{\# \text{ total trades for Opel in } \Phi} \right)^2 \\ &\quad + \frac{1}{2} \sum_{i \in s(\Phi, \text{Chevy})} \left(\frac{\# \text{ trades sold by supplier } i \text{ to Chevy in } \Phi}{\# \text{ total trades for Chevy in } \Phi} \right)^2. \end{aligned} \quad (14)$$

Likewise, $X_{\text{Model}}(j^b, \Phi)$ is the mean across car models sold by GM of the Herfindahl index calculated for the sellers of parts to each car model separately. As with suppliers, the elements of $X(j^b, \Phi)$ can be computed for the counterfactual trades in the matching maximum score inequalities (6).

The matching maximum score inequalities used in estimation keep the number of car part trades sold by each supplier (and, more obviously, the set of car parts needed on each car model) the same. With strong returns to specialization, it may be more efficient to have fewer but individually larger suppliers. The optimality of supplier size is not imposed as part of the estimator. Nor can the gains from assembler scale be identified from matching maximum score inequalities where each car part and each car model are weighted equally. This paper models the car parts industry, not the market for corporate control of car brands and car models. Not imposing the optimality of supplier and assembler sizes might be an advantage, as other concerns such as capacity constraints and antitrust rules could limit firm size. On the other hand, one of the benefits of GM not divesting Opel is keeping a larger scale, and the matching maximum score inequalities used in estimation do not identify a pure scale economy for GM owning Opel. Instead, the matching maximum score inequalities focus on the gains to assemblers and particularly to suppliers from specialization, for a fixed number of car part trades.

5.3 Estimates for Valuation Functions

Table 1 presents the point estimates and confidence intervals for the parameter vector θ in the valuation functions for observable types, for both assemblers and suppliers. I randomly sample a maximum of 10,000 matching maximum score inequalities (6) per component category. All theoretically valid inequalities with two different suppliers are sampled with an equal probability. I use the set identified subsampling procedure of Romano and Shaikh (2008) to construct confidence regions. See Appendix B for details on estimation and inference.

The parameter on assembler (parent group) specialization for suppliers is normalized to be ± 1 . The other parameters in Table 1 are interpreted relative to the parameter on parent group specialization. One finding is that the point estimates of the assembler parameters have a much lower order of magnitude than the supplier parameters and the assembler parameters have wide confidence bands, always including 0. For assemblers, the upper bounds of the confidence bands for parent group, brand and model specialization are lower than the lower bounds for the analogous specialization measures for suppliers. Therefore, for these specialization measures one can at least statistically conclude that supplier specialization measures are more important. This difference between the point estimates for assemblers and suppliers is not because of a difference in the units of $X(j^b, \Phi)$ and $X(j^s, \Psi)$; the rightmost columns of Table 1 report the means and standard deviations of the specialization measures for realized matches for both suppliers and assemblers. The specialization (HHI) measures are about the same magnitudes for both suppliers and assemblers. What is possibly explaining the small magnitude effects is that two economic forces may offset each other: assemblers prefer to have a diverse supplier base to avoid placing their success in the hands of one supplier (hold up) while there may be some manufacturing benefits from having a fewer number of suppliers. Regardless, the point estimates show that assembler specialization is much less important than supplier specialization in the valuation functions. One caveat is that the confidence interval for assembler specialization at the model level does contain larger, in absolute value, coefficient magnitudes.

For suppliers, Table 1 shows that all four coefficients on supplier specialization are positive, meaning as expected specialization on these dimensions increases the valuation of suppliers. The point estimates show that a given level of specialization at the parent-group level is about as important in valuation as the same level of specialization at the continent-of-brand-headquarters level. At the same time, the standard

deviation of parent-group-specialization HHI, across realized matches, is 0.18, meaning the variation in parent-group specialization across suppliers is lower than for some other specialization measures. A naive researcher might be inclined to interpret this level of dispersion as evidence parent-group specialization is unimportant. This would be wrong: the matching maximum score estimator accounts for the fact that more available matching opportunities occur across firm boundaries than within them. An estimate of a structural parameter such as the coefficient on parent group tells us the importance of parent group in the valuation from a set of trades.

Table 1 also shows that supplier specialization at the brand and model levels is even more important than specialization at the parent-group level, as the brand and model confidence intervals do not contain +1. The high point estimate of 376 for model specialization possibly comes from supplier and assembler plant co-location: car models of even the same brand may be built in separate plants and some benefits from specialization may occur from saving on the need to have multiple supplier plants for each model. Also, the technological compatibility of car parts occurs mainly at the model level. Notice how the standard deviation of the HHI-specialization measure is about the same (0.26–0.28) for the continent, brand and model measures, with parent group specialization being a little lower at 0.18. Again, naive researchers might use the HHI means to conclude that specialization at the model level is less important or use the standard deviations to conclude that specialization at the continent, brand and model levels are equally important. The estimates of the valuation functions give statistically consistent estimates of the relative importance of the types of specialization in the valuation functions for supplier relationships.

Table 1 also shows that there are 298,272 inequalities used in estimation. Of those, 82% are satisfied at the reported point estimates. The fraction of satisfied inequalities is a measure of statistical fit.

Appendix C presents estimates where the HHI specialization measures use different weighting schemes, including weighting schemes using data on car model sales in Europe and North America. The specifications in Appendix C result in lower numbers of inequalities being satisfied at the parameter estimates. Therefore, these alternatives result in statistically worse fit and so are not presented in the main text. However, a common finding in Appendix C is that the assembler parameters become more important, particularly the parameter on assembler model specialization.

5.4 Supplier Valuation Loss From GM Divesting Opel

Encouraging General Motors to divest Opel was a major policy issue in Germany during 2009. The revealed preference of GM to back away from selling Opel to outside investors suggests that GM felt that Opel was important to its performance. One possibility is that GM feared a loss of economies of scale (total size) or scope (strength in fuel efficient cars that could be transferred from Europe to North America, say) from such a divestiture. Matching in the car parts industry is not necessarily informative about assembler economies of scale and scope.

Using information from the car parts industry, and in particular in light of the minuscule point estimates on assembler specialization above, the major estimated effect of GM divesting Opel will come from suppliers to GM being less specialized as GM's and Opel's models technologically diverge. This will hurt GM through equilibrium prices of trades: suppliers will charge higher prices to GM for car parts. In each component category, I construct the counterfactual sum of valuations from observable types to suppliers if Opel and the rest of GM are now treated as separate assemblers, or parent groups. The same sellers supply the same car part trades to the same car models, but now the Opel models are produced by an independent parent group. In (13), some car parts are transferred to a new parent group and so the measure of parent group specialization weakly decreases for any supplier that sells any parts to Opel. The decrease in the parent group specialization measure $X_{PG}(j^s, \Psi)$ times its estimated parameter θ_{PG}^s gives the decrease in the valuation for each supplier who sells at least one part to Opel. I focus on a percentage decrease measure

$$\frac{\theta_{PG}^s \Delta X_{PG}(j^s, \Psi)}{X(j^s, \Psi)' \theta},$$

for a particular supplier with the matches Ψ in the data. Note that this measure imposes a cardinal (up to scale) interpretation of a supplier's valuation function, as opposed to identifying a supplier's valuation function only up to a positive monotonic transformation. Fox (2010) proves that the cardinal aspects of a related function are identified nonparametrically in a related matching game with transfers.

Table 2 reports statistics for the distribution of percentage changes in valuation for suppliers. A supplier in the table is a real-life supplier in a particular component category. Only suppliers who sell at least one part to Opel and one car part to another GM brand are affected and so included in the table. The mean loss is tiny, at 0.04%.

The other quantiles are tiny as well. This partly reflects suppliers where either Opel is a small fraction of car parts or a very large fraction of parts, so GM divesting Opel makes little difference in how specialized the supplier is. This result also follows from the parameter estimates in Table 1, where the point estimates for the coefficients on brand and especially model specialization are many times larger than the coefficient on parent group specialization.²³

6 Benefits to Domestic Suppliers From Foreign Assemblers

European and North American countries have imposed formal and political-pressure based trade barriers to imports of automobiles from Asia. Consequently, most Asian assemblers who sell cars in Europe and North America also assemble cars in Europe and North America. While some car parts are imported from Asia, Asian assembly plants in Europe and North America use many parts produced locally as well (perhaps because of more political pressure). As Klier and Rubenstein (2008) document for Asian assemblers in North America, a key part of operating an assembly plant is developing a network of high-quality suppliers.

Despite some occasional quality setbacks, the magazine *Consumer Reports* and other sources routinely record that brands with headquarters in Asia (Japan, Korea) have higher quality automobiles than brands with headquarters in Europe or North America. The parts supplied to higher-quality cars must typically also be of higher quality. Liker and Wu (2000) document that suppliers to Japanese-owned brands in the US produce fewer parts requiring reworking or scrapping, for example. Because of this emphasis on quality, the suppliers to, say, Toyota undergo a rigorous screening and training program, the Supplier Development Program, before producing a large volume of car parts for Toyota (Langfield-Smith and Greenwood, 1998). Indeed, there is a hierarchy of suppliers, with more trusted Toyota suppliers being allowed to supply more car parts (Kamath and Liker, 1994; Liker and Wu, 2000).

It is possible that the need by Asian assemblers for higher-quality suppliers benefits the entire domestic supplier bases in Europe and North America. If a supplier is of high-

²³I compute but do not report the small changes in GM's and Opel's valuations from divesting Opel. Because the coefficient estimates on assembler specialization in Table 1 are small in magnitude, the overwhelming effect in profit levels is estimated to be on suppliers.

enough quality to deal with an Asian assembler, non-Asian assemblers that also source parts from that supplier may also benefit. If this potential effect is causal (the suppliers were not of sufficiently high quality before the Asian assemblers' entry), it is evidence that trade barriers that promote Asian-owned assembly plants in Europe and North America may indirectly aid non-Asian (domestic) assemblers, as those producers now have access to higher-quality suppliers. This is an underexplored channel by which foreign-direct investment in assembly plants may raise the quality of producers in upstream markets. Indeed, there is evidence in the management literature that Asian assemblers do causally upgrade the quality of their suppliers: the Supplier Development Program mentioned above, for example (Langfield-Smith and Greenwood, 1998).

This section complements the management literature by providing evidence from sorting in the market for car parts that might be consistent with suppliers to Asian assemblers being higher quality than other suppliers. Measures of car part quality by individual suppliers are presumably observed by assemblers, but are not publicly available. In this section, a measure of quality will be a supplier's share of the market for supplying parts to Asian assemblers. If Asian assemblers together demand 100 parts in a particular component category, and one supplier sells 30 of them, its quality measure will be 0.30. In notation, one aspect of an observable firm type j^s for a supplier is

$$j_{\text{Asia}}^s = \frac{\# \text{ Asian assembler parts supplied}}{\text{total } \# \text{ Asian assembler parts all suppliers}}.$$

This is not a specialization measure, as a firm could sell many parts to Asian assemblers and many parts to non-Asian assemblers. This quality measure j_{Asia}^s is treated as an aspect of observable firm type j^s of a supplier. If it were recomputed for new portfolios Φ without interactions in a valuation function, it would difference out of the matching maximum score inequalities (6). Instead, the vector $X(j^s, \Psi)$ contains a new element that is the interaction of the above Asian quality measure with specialization by the continent headquarters of the brand, discussed earlier:

$$X_{\text{AsiaCont}}(j^s, \Psi) = j_{\text{Asia}}^s \cdot X_{\text{Cont}}(j^s, \Psi).$$

The interpretation of the corresponding supplier parameter in θ , if it is estimated to be negative, is that suppliers with higher j_{Asia}^s (greater shares of the industry for supplying Asian assemblers) gain less benefit from selling car parts to only one continent of assembler than suppliers with lower j_{Asia}^s . Thus, suppliers with higher Asian shares

can go out and win business from non-Asian assemblers, which is consistent with those firms have a competitive edge (possibly from higher quality parts) over other suppliers. The empirical pattern in the data might be that suppliers with high j_{Asia}^s have diverse (across continents of assembler origin) portfolios of car parts that they supply. This diversity might be interpreted as a sign of quality.

Even if the parameter on $X_{\text{AsiaCont}}(j^s, \Psi)$ is negative and economically large in magnitude, it does not prove that the presence of Asian assemblers causally upgrades the quality of suppliers in Europe and North America. It could have been that the suppliers with high j_{Asia}^s were of high quality before the creation of plants outside Asia by Asian assemblers. However, when combined with the evidence from the management literature about supplier development programs, it does seem as if some portion of supplier quality differences are due to the presence of the Asian assemblers.

A separate concern is that this approach treats j_{Asia}^s as an economically exogenous characteristic, rather than recomputing the Asian market share for counterfactual sets of trades Φ in the right sides of matching maximum score inequalities. I have explored the specification where notationally j_{Asia}^s is replaced by $X_{\text{Asia}}(j^s, \Psi)$, which is recomputed for counterfactual sets of trades Ψ . The corresponding point estimate on the interaction is ≈ 0 , with a wide confidence interval in terms of economic magnitudes. An explanation for the point estimate close to 0 is that a new effect is introduced to the model: the inequalities ask why more suppliers do not choose to supply parts to Asian assemblers if there is some quality upgrade from doing so? A reason outside of the model why this does not happen is the fixed cost of having an additional supplier participate in a supplier development program. Having explored an alternative, I return to the preferred specification, where a supplier's competitive advantage is an economically exogenous supplier characteristic j_{Asia}^s .

Table 3 presents the point estimates from the preferred specification. The other covariates are the assembler and supplier specialization measures in Table 1, which have similar point estimates. The scale normalization is still on parent group specialization. With the interaction term $X_{\text{AsiaCont}}(j^s, \Psi)$ involving continent specialization, the normalization can only be understood by substituting a typical value for j_{Asia}^s into the interaction term $X_{\text{AsiaCont}}(j^s, \Psi)$ and comparing also the coefficient on continent specialization without an interaction.

The new addition to Table 3 is the estimate on the interaction term $X_{\text{Asia}}(j^s, \Psi)$, which would use an estimated decrease in the importance of specialization at the

continent-of-brand level for suppliers to Asian brands' assembly plants in Europe and North America as evidence that suppliers to Asian assemblers have higher quality. These suppliers possibly can win business from non-Asian assemblers. The estimate of the parameter on $X_{\text{Asia}}(j^s, \Psi)$ is -0.261 and the mean and standard deviation of j_{Asia}^s , not listed in the table, are 0.069 and 0.102, respectively. Therefore, a one-standard deviation change in j_{Asia}^s creates a change of $-0.261 \cdot 0.102 = -0.0266$ in the coefficient on the degree of specialization at the continent-of-brand level. A car parts supplier with a market share among Asian assemblers that is one standard deviation higher than the mean, a share of 0.171, will have a total coefficient on continent-of-brand specialization of $+1.03 - 0.261 \cdot 0.171 = 0.985$, or approximately 1. This is a small magnitude change. The confidence region for the interaction parameter is -30 to 32. The data do not closely pin down this effect.

With a bigger in absolute value and more precisely estimated effect, the interpretation would have been that suppliers to Asian assemblers can go out and win business from non-Asian assemblers as well, but suppliers to European and North American assemblers cannot win as much business from assemblers from other continents. Thus, the evidence from sorting in the market for car parts would have suggested that domestic suppliers to assemblers with headquarters in Asia are in a unique competitive position, consistent with them having a quality advantage.

7 Conclusions

This paper discusses the estimation of valuation functions in matching games with transferable utility. A matching maximum score estimator is introduced for a model of matching with trading networks, which has many special cases of empirical relevance. The matching maximum score objective function is computationally simple and the econometric model is semiparametric.

The empirical work answers two policy questions surrounding the automotive industry. First, the paper estimates the relative loss in valuation to suppliers from decreased specialization when General Motors divests Opel. A divestiture ends up hurting most suppliers only a little as the point estimates to the gains to specialization at the model level, which is not affected by the divestment, are higher than the gains to specialization at the parent group level. Second, the paper estimates the gain to, say, North American suppliers from the presence of Asian-based assemblers in North America.

Both estimates are inferred from a new type of data, the portfolios of car part trades from each supplier.

A For Online Publication: Proofs

A.1 Proposition 1: Rank Order Property

The proof cites Graham (2011, Theorem 4.1) and the erratum Graham (2013). His theorem is stated for one-to-one, two-sided matching or marriage and, in his notation, considers two observable types of men k and m and two observable types of women l and n . Graham's proof works by considering the so-called suballocation of the two observable types of men and the two observable types of women, the suballocation $klmn$ in his notation. In the suballocation $klmn$, Graham considers the probability within the suballocation of k and l matching, which is called r^{klmn} . The probability of k matching with n is $p^{klmn} - r^{klmn}$, the probability of m matching with l is $q^{klmn} - r^{klmn}$, and the probability of m matching with n is $1 - p^{klmn} - q^{klmn} + r^{klmn}$.

The statement that drawing two matches between k and l and m and n are more likely than drawing two matches between k and n and m and l is

$$(1 - p^{klmn} - q^{klmn} + r^{klmn}) r^{klmn} > (p^{klmn} - r^{klmn}) (q^{klmn} - r^{klmn}). \quad (15)$$

Algebra at the end of Section 4.3.2 in Graham, after correcting typos, shows that this inequality is equivalent to $r^{klmn} > p^{klmn} q^{klmn}$. The conclusion of Theorem 4.1 on the same page of Graham can be rewritten to state that

$$\delta_{mn} + \delta_{kl} \geq \delta_{ml} + \delta_{kn} \quad (16)$$

if and only if $r^{klmn} > p^{klmn} q^{klmn}$, which as just said is equivalent to the inequality (15). The inequality (16) can be seen as a matching maximum score inequality (6). Therefore, Graham (2011, Theorem 4.1) is a rank order property for one-to-one, two-sided matching without data on prices.

The setup in the matching with trading networks model is more general than the model of one-to-one, two-sided matching or marriage in Graham. However, the extra generality can be handled by conditioning. In a matching maximum score inequality $g \in G$ in my notation, there are the two observable agent types of buyers and the

two observable agent types of sellers in $H_g = \{b(\omega_1), s(\omega_1), b(\omega_2), s(\omega_2)\}$. While the matching with trading networks mode model does not necessarily assign roles of buyers and sellers ex ante, a matching maximum score inequality (6) does condition on these roles on the left side of the inequality and the right side of the inequality. A complication is the observable type $b(\omega_1)$ that is the buyer on ω_1 on the left side of the inequality could be a seller on, say, trade ω_3 on the right side of the inequality. This switch of roles of buyer and seller does not change the proof in Graham.

Likewise, agents in the matching with trading networks mode model make sets of trades (Φ, Ψ) . This can be handled by conditioning on the set of trades (Φ_j, Ψ_j) other than ω_1 – ω_4 for all four observable agent types in $j \in H_g$. In other words, condition on the joint event C that $b(\omega_1)$ picks either $(\Phi_{b(\omega_1)}^g, \Psi_{b(\omega_1)}^g)$ or $(\bar{\Phi}_{b(\omega_1)}^g, \bar{\Psi}_{b(\omega_1)}^g)$, $s(\omega_1)$ picks either $(\Phi_{s(\omega_1)}^g, \Psi_{s(\omega_1)}^g)$ or $(\bar{\Phi}_{s(\omega_1)}^g, \bar{\Psi}_{s(\omega_1)}^g)$, and similarly for $b(\omega_2)$ and $s(\omega_2)$.

In what follows, abbreviate $\Pr_{b(\omega_1)}(1 \mid \Phi_{b(\omega_1)}, \Psi_{b(\omega_1)}, \bar{\Phi}_{b(\omega_1)}, \bar{\Psi}_{b(\omega_1)})$ with $\Pr_{b(\omega_1)}(1 \mid g)$. When forming choice probabilities conditional on the joint event C mentioned just above, the conditional choice probabilities for say $b(\omega_1)$ will multiplicatively factor into $\Pr_{b(\omega_1)}(1 \mid g)$ and three choice probabilities for the other three agents. In an inequality such as the probability statement (10) in the statement of the proposition, the choice probabilities for the other three agents are the same multiplicative correction for $\Pr_{b(\omega_1)}(1 \mid g)$ and $\Pr_{b(\omega_1)}(2 \mid g)$ and cancel out on either side of the inequality. The choice probabilities multiplicatively factor as the choice probabilities of agents are mutually independent conditional on the observable type $j \in J$ in the matching game.

Divide, on the left and right sides of (15), $(1 - p^{klmn} - q^{klmn} + r^{klmn})$ and $(q^{klmn} - r^{klmn})$ by the constant

$$(1 - p^{klmn} - q^{klmn} + r^{klmn}) + (q^{klmn} - r^{klmn}).$$

This changes those terms into, in my notation, the conditional choice probabilities $\Pr_{b(\omega_1)}(1 \mid g)$ and $\Pr_{b(\omega_1)}(2 \mid g)$, respectively, after canceling out the multiplicatively factorable choice probabilities for the three agents other than $b(\omega_1)$ for conditioning on the joint event C , as just discussed. A similar argument applies to give $\Pr_{b(\omega_2)}(1 \mid g)$ and $\Pr_{b(\omega_2)}(2 \mid g)$. Therefore, the probability statement (15) becomes the probability statement (10) in the statement of the proposition. Hence, we have proved the rank order property for the matching with trading networks game without data on the prices of trades.

A.2 Theorem 1: Set Identification

In what follows, abbreviate $\Pr_{b(\omega_1)}(1 \mid \Phi_{b(\omega_1)}, \Psi_{b(\omega_1)}, \bar{\Phi}_{b(\omega_1)}, \bar{\Psi}_{b(\omega_1)})$ with $\Pr_{b(\omega_1)}(1 \mid g)$. By the law of iterated expectations, some algebra and Assumption 3, the expectation of the matching maximum score objective function (8) can be written

$$\frac{1}{2} \sum_{g \in G} \Pr[g \text{ included}] \cdot \{ \Pr_{b(\omega_1)}(1 \mid g) \cdot \Pr_{b(\omega_2)}(1 \mid g) \cdot 1 [Z'_g \theta \succ_g 0] + \Pr_{b(\omega_1)}(2 \mid g) \cdot \Pr_{b(\omega_2)}(2 \mid g) \cdot 1 [Z'_g \theta \prec_g 0] \}.$$

The relation \prec_g refers to the opposite of \succ_g . If \succ_g is the strict $>$ greater than relation, then \prec_g refers to the weak \leq less than relation. If \succ_g is the weak \geq greater than relation, then \prec_g refers to the strict $<$ less than relation. The $1/2$ is to remove double counting: counting an inequality g once in the summation where $g \in G$ is indeed the index and once in the summation where g enters the summation index for some \tilde{g} where $Z'_g \theta \succ_g 0$ equals $Z'_{\tilde{g}} \theta \prec_{\tilde{g}} 0$. The calculation $\Pr[g \text{ included}]$ is over the four observable types of i_1-i_4 in (7) and all the trades except $\omega_1-\omega_4$ in inequality g and, hence, the reverse direction inequality.

The forward direction inequality $Z'_g \theta \succ_g 0$ is mutually exclusive with $Z'_g \theta \prec_g 0$. Only one of the two inequalities can enter the objective function with nonzero weight for a given parameter θ . By the rank order property in Proposition 1, the maximum of the two products of conditional choice probabilities of the form

$$\Pr_{b(\omega_1)}(1 \mid g) \cdot \Pr_{b(\omega_2)}(1 \mid g)$$

will be included at the true parameter value for θ . Any other parameter value θ results in either the same objective function value or a lower objective function value where some smaller weight contributes to the objective function value. Therefore, the matching maximum score objective is (perhaps not uniquely) globally maximized at the true parameter value.

B For Online Publication: Estimation and Inference for the Car Parts Empirical Work

I first describe point estimation and then inference. I sample 10,000 inequalities for each of 29 matching markets given by component categories. For the 30th component category, I use the maximum number of inequalities, 8,272, for that market. To create one of the 10,000 inequalities for a market, I randomly sample two car parts from different suppliers and include the resulting matching maximum score inequality where the two car parts are exchanged between the suppliers.

For point estimation, I use the numerical optimization routine differential evolution, in Mathematica. For differential evolution, I use a population of 250 points and a maximum of 15,000 iterations. The numerical optimization is run 20 times with different initial populations of 250 points. For the specification in Table 1, all 20 runs found the same objective function value and the same point estimates, up to two significant digits. Therefore, it might be reasonable to consider the model to be point identified in the limit.

In the maximum score objective function, an inequality is satisfied if the left side plus the constant $+0.0000000001$ exceeds the right side. This small perturbation to the sum of profits on the left side ensures that inequalities such as $0 > 0$ are counted as being satisfied consistently, rather than inconsistently because of some numerical-approximation error resulting in, say, $2.0 \times 10^{-15} > 1.0 \times 10^{-15}$.

I use the inference procedure of Romano and Shaikh (2008), which is valid under both set identification and point identification. Let Θ_0 be the identified set for the parameter θ . The authors provide a subsampling procedure to construct a confidence region \mathcal{C} called, by those authors, a “confidence region for the identifiable parameters that is uniformly consistent in level”, under the conditions of their Theorem 3.3. This definition of the properties of a confidence region is equation (3) of their paper. Under easier to verify conditions (namely that a limiting distribution exist for the objective function), their procedure produces “confidence regions for identifiable parameters that are pointwise consistent in level”, or their equation (2). Computationally, I implement equations (12) and (10) in their paper.

I have 30 component categories. I treat these component categories as separate markets for the industry-specific reasons discussed in the main text. I perform asymptotics in the number of suppliers in the market, as this is an explicitly two-sided market

where a supplier cannot also be an assembler, as in Example 2. Subsampling requires a choice of subsample size. Unfortunately, the literature has not provided a data driven method to pick this tuning parameter. I use 33% of the suppliers from each of the 30 component categories. Given each subsampled set of suppliers, that subsample uses only matching maximum score inequalities where both suppliers whose valuation functions are in the inequality are in the subsample. I use 500 subsamples; in experiments results are robust when using more subsamples.

In the car parts empirical work, the maximum score or maximum rank correlation objective function is, for $D = 30$ component categories,

$$Q(\theta) = \binom{N}{2}^{-1} \sum_{d=1}^D \sum_{i_1=1}^{N_d-1} \sum_{i_2=i_1+1}^{N_d} \sum_{g \in G_{i_1, i_2}^d} 1 [Z'_{g,d} \theta \geq 0],$$

where G_{i_1, i_2}^d is the set of included matching maximum score inequalities where i_1 and i_2 exchange one car part each. Romano and Shaikh (2008) is written where one minimizes a function and where the population objective function's minimum value is 0. In maximum score, one maximizes the objective function value and the value of the population objective function can only be estimated using the finite sample objective function. Following the suggestion at the beginning of their Section 3.2, I work with the objective function $-Q(\theta) - (-\max_{\theta} Q(\theta))$.

I define the “sample size” for the entire procedure to be $N = \max_{d \in D} N_d$. This choice is arbitrary and does not impact the reported confidence regions. I implement the procedure of Romano and Shaikh (2008) using the rate of convergence for the objective function of \sqrt{N} , following the results of Sherman (1993) for the case of point identified maximum rank correlation estimators.

The confidence regions reported in tables are projections onto the axes of the confidence region \mathcal{C} from Romano and Shaikh (2008). Computationally, the lower bound for the confidence region for a scalar parameter θ_1 is found by minimizing the parameter θ_1 subject to the constraint that the entire parameter vector θ is in \mathcal{C} . Likewise, the upper bound for the confidence region θ_1 is found by maximizing the parameter θ_1 subject to the entire parameter vector being in \mathcal{C} .

C For Online Publication: Alternative Weighting Schemes for Specialization Measures

This appendix discusses versions of the estimates in Table 1 where the included HHI specialization measures use different weighting schemes, many involving data on car sales for car models in Europe and the United States. These alternative weighting schemes result in worse statistical fit than the estimates in Table 1, in the sense that fewer inequalities are satisfied. For this reason, I report the estimates in this appendix instead of the main text.

In terms of sales, car models primarily sold in Europe are matched to European sales from Western Europe and car models primarily from North America are matched to sales from the United States. Note that all weighting schemes affect the HHI specialization measures. I do not explore weighting the maximum score inequalities in the objective function, although that could be pursued.

Let w_c be a weight for a particular car part c . The example parent group supplier HHI measure in (13) becomes

$$X_{\text{PG}}(j^s, \Psi) = \left(\frac{\text{sum Chrysler weights in } \Psi}{\text{sum all weights in } \Psi} \right)^2 + \left(\frac{\text{sum Ford weights in } \Psi}{\text{sum all weights in } \Psi} \right)^2 + \left(\frac{\text{sum GM weights in } \Psi}{\text{sum all parts in } \Psi} \right)^2.$$

Analogous schemes are used for other supplier HHI specialization measures and for assembler specialization at the parent group level. The example brand assembler HHI measure in (14) becomes

$$X_{\text{Brand}}(j^b, \Phi) = \bar{w}_{\text{Opel}} \sum_{i \in s(\Phi, \text{Opel})} \left(\frac{\text{sum weights sold by supplier } i \text{ to Opel in } \Phi}{\text{sum weights for Opel in } \Phi} \right)^2 + \bar{w}_{\text{Chevy}} \sum_{i \in s(\Phi, \text{Chevy})} \left(\frac{\text{sum weights sold by supplier } i \text{ to Chevy in } \Phi}{\text{sum weights for Chevy in } \Phi} \right)^2,$$

where, for example,

$$\bar{w}_{\text{Chevy}} = \frac{\text{sum weights for Chevy in GM}}{\text{sum weights in GM}}.$$

Note that the brand weights like \bar{w}_{Chevy} are not recomputed for counterfactual trades Φ , as indeed Chevrolet and other brands still produce the same car models with the

same need for car parts in the counterfactuals in matching maximum score inequalities.

Table 4 reports the estimates with weights. There are three specifications, each with different weights. Note that the parameters multiply different explanatory variables across the three specifications, so the parameters have slightly different interpretations. However, the means and standard deviations of the weighted HHI measures are qualitatively similar to those reported in Table 1 for the baseline specification.

The first weighting scheme in Table 4 does not use sales data; it sets

$$w_c = \frac{1}{\# \text{ car parts for model in market } d}.$$

This scheme weights a car part on a car model with more parts in a component category less, in order to equalize the contribution of car models to HHI calculations. The second weighting scheme uses the weights

$$w_c = \text{sales of car model}.$$

The third weighting scheme combines the previous two, as in

$$w_c = \frac{\text{sales of car model}}{\# \text{ car parts for model in market } d}.$$

Compared to Table 1, Table 4 shows less of a role for supplier model specialization and much more of a role for assembler model HHI specialization. In all three of the specifications, the lower bound for the confidence region for supplier HHI model specification is negative and large in magnitude relative to the +1 normalization for supplier parent group specialization. The point estimate is negative in two of the three specifications. In all three specifications, the upper bound of the confidence region for assembler model specialization is unbounded. The point estimate for assembler model specialization is unbounded in two specifications and still quite large, at 202, for the other specification.

Also compared to Table 1, there is now statistical uncertainty about the sign of supplier brand specialization, although the point estimate is about the same as for parent group specialization, which is normalized to +1. Similarly to Table 1, 0 is still in the confidence regions for assembler parent group and brand specialization.

Overall, I emphasize Table 1 in the main text because of its higher statistical fit.

References

- Agarwal, Nikhil and William Diamond**, “Latent Indices in Assortative Matching Models,” 2016. MIT working paper.
- Akkus, Oktay, J. Anthony Cookson, and Ali Hortacsu**, “The Determinants of Bank Mergers: A Revealed Preference Analysis,” *Management Science*, 2016, 62 (8), 2241–2258.
- Azevedo, Eduardo M. and John William Hatfield**, “Existence of Equilibrium in Large Matching Markets with Complementarities,” 2015. University of Pennsylvania working paper.
- Baccara, Mariagiovanna, Ayse Imrohoroglu, Alistair Wilson, and Leeat Yariv**, “A Field Study on Matching with Network Externalities,” *American Economic Review*, 2012, 102 (5).
- Bajari, Patrick, Jeremy T. Fox, and Stephen Ryan**, “Evaluating Wireless Carrier Consolidation Using Semiparametric Demand Estimation,” *Quantitative Marketing and Economics*, 2008, 6 (4), 299–338.
- Becker, Gary S.**, “A Theory of Marriage: Part I,” *Journal of Political Economy*, July-August 1973, 81 (4), 813–846.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff**, “Analyzing the Determinants of the Matching Public School Teachers to Jobs: Estimating Compensating Differentials in Imperfect Labor Markets,” *Journal of Labor Economics*, 2013, 31 (1), 83–117.
- Briesch, Richard A., Pradeep C. Chintagunta, and Rosa L. Matzkin**, “Semiparametric estimation of brand choice behavior,” *Journal of the American Statistical Association*, December 2002, 97 (460), 973–982.
- Chiappori, Pierre-André and Bernard Salanié**, “The Econometrics of Matching Models,” *Journal of Economic Literature*, 2016, 54 (3), 832–861.
- , – , and **Yoram Weiss**, “Partner Choice and the Marital College Premium: Analyzing Marital Patterns Over Several Decades,” 2016. Columbia University working paper.

- Choo, Eugene and Aloysius Siow**, “Who Marries Whom and Why,” *The Journal of Political Economy*, November 2006, *114* (1), 175–201.
- Ciliberto, Federico and Elie Tamer**, “Market Structure and Multiple Equilibria in Airline Markets,” *Econometrica*, November 2009, *77* (6), 1791–1828.
- Dagsvik, John K.**, “Aggregation in Matching Markets,” *International Economic Review*, February 2000, *41* (1), 27–57.
- Delgado, M.A., J.M. Rodriguez-Poo, and M. Wolf**, “Subsampling inference in cube root asymptotics with an application to Manski’s maximum score estimator,” *Economics Letters*, 2001, *73* (2), 241–250.
- Dupuy, Arnaud and Alfred Galichon**, “Personality traits and the marriage market,” *Journal of Political Economy*, 2014, *122* (6), 1271–1319.
- Dyer, Jeffrey H.**, “Specialized Supplier Networks as a Source of Competitive Advantage: Evidence from the Auto Industry,” *Strategic Management Journal*, 1996, *17* (4), 271–291.
- , “Effective interim collaboration: how firms minimize transaction costs and maximize transaction value,” *Strategic Management Journal*, 1997, *18* (7), 535–556.
- Ekeland, Ivar, James J. Heckman, and Lars Nesheim**, “Identification and Estimation of Hedonic Models,” *Journal of Political Economy*, 2004, *112* (S1), 60–109.
- Fox, Jeremy T.**, “Semiparametric estimation of multinomial discrete-choice models using a subset of choices,” *RAND Journal of Economics*, Winter 2007, *38* (4), 1002–1019.
- , “Identification in Matching Games,” *Quantitative Economics*, 2010, *1* (2), 203–254.
- **and Patrick Bajari**, “Measuring the Efficiency of an FCC Spectrum Auction,” *American Economic Journal: Microeconomics*, 2013, *5* (1), 100–146.
- , **Chenyu Yang, and David Hsu**, “Unobserved Heterogeneity in Matching Games,” *Journal of Political Economy*, forthcoming.
- Gale, David**, *The Theory of Linear Economic Models*, McGraw-Hill, New York, 1960.

- Galichon, Alfred and Bernard Salanié**, “Cupid’s Invisible Hand: Cupid’s Invisible Hand: Social Surplus and Identification in Matching Models,” 2015. New York University working paper.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey**, “Regular Quantal Response Equilibrium,” *Experimental Economics*, December 2005, 8, 347–367.
- Graham, Bryan S.**, “Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers,” in J. Benhabib, A. Bisin, and M. Jackson, eds., *Handbook of Social Economics*, Elsevier, 2011, chapter 19.
- , “Errata in “Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers”,” 2013. Online erratum.
- Han, A.K.**, “Non-parametric analysis of a generalized regression model,” *Journal of Econometrics*, 1987, 35, 303–316.
- Hatfield, John William and Paul R. Milgrom**, “Matching with Contracts,” *American Economic Review*, September 2005, 95 (4), 913–935.
- , **Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp**, “Stability and Competitive Equilibrium in Trading Networks,” *Journal of Political Economy*, 2013, 121 (5), 966–1005.
- Heckman, James J., Rosa L. Matzkin, and Lars Nesheim**, “Nonparametric Identification and Estimation of Nonadditive Hedonic Models,” *Econometrica*, September 2010, 78 (5), 1569–1591.
- Horowitz, Joel L.**, “A Smoothed Maximum Score Estimator for the Binary Response Model,” *Econometrica*, May 1992, 60 (3), 505–551.
- Kamath, Rajan R. and Jeffrey K. Liker**, “A Second Look at Japanese Product Development,” *Harvard Business Review*, November-December 1994.
- Kim, J. and D. Pollard**, “Cube root asymptotics,” *The Annals of Statistics*, 1990, pp. 191–219.
- Klier, Thomas and James Rubenstein**, *Who Really Made Your Car? Restructuring and Geographic Change in the Auto Industry*, Upjohn Institute, 2008.

- Koopmans, Tjalling C. and Martin Beckmann**, “Assignment Problems and the Location of Economic Activities,” *Econometrica*, January 1957, 25 (1), 53–76.
- Langfield-Smith, K. and M.R. Greenwood**, “Developing Co-operative Buyer-Supplier Relationships: A Case Study of Toyota,” *Journal of Management Studies*, 1998, 35 (3), 331–353.
- Liker, Jeffrey K. and Yen-Chun Wu**, “Japanese Automakers, U.S. Suppliers and Supply-Chain Superiority,” *Sloan Management Review*, Fall 2000.
- Manski, C.F.**, “Semiparametric analysis of discrete response: Asymptotic properties of the maximum score estimator,” *Journal of Econometrics*, 1985, 27 (3), 313–333.
- Manski, Charles F.**, “Maximum Score Estimation of the Stochastic Utility Model of Choice,” *Journal of Econometrics*, 1975, 3 (3), 205–228.
- Matzkin, Rosa L.**, “Nonparametric identification and estimation of polychotomous choice models,” *Journal of Econometrics*, 1993, 58, 137–168.
- McFadden, D.**, “Conditional Logit Analysis of Qualitative Choice Behavior,” *Frontiers in Econometrics*, 1973, pp. 105–142.
- McFadden, Daniel L.**, “A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration,” *Econometrica*, September 1989, 57 (5), 995–1026.
- Menzel, Konrad**, “Large matching markets as two-sided demand systems,” *Econometrica*, 2015, 83 (3), 987–941.
- Mindruta, Denisa, Mahka Moeen, and Rajshree Agarwal**, “A Two-Sided Matching Approach for Partner Selection and Assessing Complementarities in Partners’ Attributes in Inter-Firm Alliances,” *Strategic Management Journal*, 2016, 37 (1), 206–231.
- Monteverde, Kirk and David J. Teece**, “Supplier Switching Costs and Vertical Integration in the Automobile Industry,” *The Bell Journal of Economics*, Spring 1982, 13 (1), 206–213.
- Newey, W.K. and D. McFadden**, “Large Sample Estimation and Hypothesis Testing,” in “Handbook of Econometrics,” Vol. 4, Elsevier, 1994, pp. 2111–2245.

- Novak, Sharon and Birger Wernerfelt**, “On the Grouping of Tasks into Firms: Make-or-Buy with Interdependent Parts,” *Journal of Economics and Management Strategy*, 2012, *21* (1), 53–77.
- **and Scott Stern**, “How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry,” *Management Science*, December 2008, *54* (12), 1963–1979.
- **and –**, “Complementarity Among Vertical Integration Decisions: Evidence from Automobile Product Development,” *Management Science*, 2009, *55* (2), 311–332.
- **and Steven D. Eppinger**, “Sourcing by Design: Product Complexity and the Supply Chain,” *Management Science*, 2001, *47* (1), 189–204.
- Pakes, A. and D. Pollard**, “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 1989, *57* (5), 1027–1057.
- Romano, Joseph P and Azeem M Shaikh**, “Inference for identifiable parameters in partially identified econometric models,” *Journal of Statistical Planning and Inference*, 2008, *138* (9), 2786–2807.
- Romano, J.P. and A.M. Shaikh**, “Inference for the identified set in partially identified econometric models,” *Econometrica*, 2010, *78* (1), 169–211.
- Rosen, Sherwin**, “Hedonic prices and implicit markets: product differentiation in pure competition,” *Journal of Political Economy*, 1974, *82* (1), 34–55.
- Santiago, David and Jeremy T. Fox**, “A Toolkit for Matching Maximum Score Estimation and Point and Set Identified Subsampling Inference,” 2009. Rice University working paper.
- Shapley, Lloyd S. and Martin Shubik**, “The assignment game I: the core,” *International Journal of Game Theory*, 1972, *1*, 111–130.
- Sherman, Robert P.**, “The Limiting Distribution of the Maximum Rank Correlation Estimation,” *Econometrica*, January 1993, *61* (1), 123–137.
- Sørensen, Morten**, “How Smart is Smart Money? A Two-Sided Matching Model of Venture Capital,” *Journal of Finance*, December 2007, *LXII* (6), 2725–2762.

Subbotin, Viktor, “Asymptotic and bootstrap properties of rank regressions,”
November 2007. Northwestern University working paper.

Uetake, Kosuke and Yasutora Watanabe, “Entry by Merger: Estimates from
a Two-Sided Matching Model with Externalities,” 2016. Yale University working
paper.

Table 1: Specialization By Suppliers and Assemblers

HHI Measure	Valuation Function Estimates		Sample Statistics for HHI Measures	
	Point Estimate	95% CI Set Identified	Mean	Standard Deviation
Suppliers				
Parent Group	+1	Superconsistent	0.35	0.28
Continent	1.04	(0.0482,9.45)	0.76	0.18
Brand	23.9	(1.29,121)	0.25	0.27
Model	376	(278,933)	0.17	0.26
Assemblers				
Parent Group	-0.007	(-1.30,0.202)	0.14	0.11
Brand	-0.005	(-1.99,0.705)	0.35	0.33
Model	-0.003	(-3.36,33.5)	0.58	0.60
# Inequalities	298,272			
% Satisfied	82.3%			

The parameter on parent group specialization is fixed at +1. Estimating it with a smaller number of inequalities always finds the point estimate of +1, instead of -1. The estimate of a parameter that can take only two values is superconsistent, so I do not report a confidence interval. See Appendix B for details on estimation and inference.

Table 2: Percentage Valuation Change By Suppliers From GM Divesting Opel

Quantile	
0	-0.0032
0.10	-0.0014
0.25	-0.0008
0.50 (median)	-0.0004
0.75	-0.0002
0.90	-0.00008
1	~0

This table uses the point estimates from Table 1 to calculate the valuations from observable types of suppliers before and after GM divests Opel. In the model, Opel becomes a separate parent group. For each supplier, selling one or more parts to Opel and one or more cars to another GM brand, I calculate $\frac{\theta_{PG}^s \Delta X_{PG}(j^s, \Psi)}{X(j^s, \Psi)^\theta}$. Each supplier that operates in multiple component categories (markets) is treated separately in each component category.

Table 3: Supplier Competitive Advantages From Asian Assemblers

	Valuation Function Estimates	
HHI Measure	Point Estimate	95% CI
Suppliers		
Parent Group	+1	Superconsistent
Continent	1.03	(0.045,13.7)
Brand	24.2	(1.09,235)
Model	388	(363,898)
Competitive Advantage	-0.261	(-30.0,32.2)
Assemblers		
Parent Group	-0.0101	(-1.50,0.224)
Brand	-0.00789	(-2.07,0.831)
Model	-0.00437	(-3.64,34.2)
# Inequalities	298,272	
% Satisfied	82.3%	

The parameter on parent group specialization is fixed at +1. Estimating it with a smaller number of inequalities always finds the point estimate of +1, instead of -1 . The estimate of a parameter that can take only two values is superconsistent, so I do not report a confidence interval. See Appendix B for details on estimation and inference.

Table 4: Specialization By Suppliers and Assemblers with Different Weighting Schemes

Weights→	1 Divided by #Parts		Sales		Sales Divided by #Parts	
HHI Measure	Point Estimate	95% CI	Point Estimate	95% CI	Point Estimate	95% CI
Suppliers						
Parent Group	+1	Supercon.	+1	Supercon.	+1	Supercon.
Continent	0.982	(0.405,4.66)	0.733	(0.582,1.29)	0.717	(0.298,3.12)
Brand	0.988	(-0.154,8.52)	0.967	(0.158,3.19)	0.894	(-0.52,6.06)
Model	0.0806	(-8.79,2.52)	-1.97	(-4.14,-1.31)	-1.88	(-6.18,-0.374)
Assemblers						
Parent Group	0.474	(-0.0902,4.70)	≈ 0	(-0.0318,0.0654)	-0.0448	(-0.36,0.0333)
Brand	-0.380	(-3.77,1.79)	≈ 0	(-0.0435,0.556)	0.0462	(-0.212,1.2)
Model	202	(100,+ ∞)	+ ∞	(+ ∞ , + ∞)	+ ∞	(+ ∞ , + ∞)
# Inequalities	298,272		298,272		298,272	
% Satisfied	76.3%		74.4%		74.3%	

The parameter on parent group specialization is fixed at +1. Estimating it with a smaller number of inequalities always finds the point estimate of +1, instead of -1. The estimate of a parameter that can take only two values is superconsistent, so I do not report a confidence interval. See Appendix B for details on estimation and inference.

