A MECHANISM DESIGN APPROACH TO IDENTIFICATION AND ESTIMATION

BRADLEY LARSEN AND ANTHONY LEE ZHANG

Abstract. This paper provides a new, nonparametric identification and estimation approach for a variety of incomplete information games, both static and dynamic. The approach relies on the Revelation Principle, exploiting the incentive compatibility of the direct revelation mechanism corresponding to the underlying and unspecified game, rather than attempting to solve for or specify the extensive form or equilibrium strategies of the game directly. We illustrate the approach using simulated and actual data from bargaining settings.

Keywords: Incomplete information game estimation, mechanism design, revelation principle, nonparametric identification and estimation, incentive compatibility

---

Date: January 17, 2017.
First version: July 2016. We thank Susan Athey, Lanier Benkard, Evgeni Drynkin, Han Hong, Mike Ostrovsky, Karl Schurter, Paulo Somaini, Caio Waisman, and Ali Yurukoglu, as well as participants as the California Econometrics Conference, the Empirics and Methods in Economics Conference, and Paul Milgrom's workshop for helpful comments.
Larsen: Stanford University, Department of Economics and NBER; bjlarsen@stanford.edu.
Zhang: Stanford University, Graduate School of Business; anthonyz@stanford.edu.

1. Introduction

Incomplete information games have posed difficult challenges for empirical work in economics. The empirical literature has largely proceeded by designing identification strategies for specific *extensive forms:* for any given extensive form, the analyst solves for a Bayes-Nash equilibrium, and uses this equilibrium to determine the mapping between observed equilibrium strategies and players' unobserved types. In some settings, such as commonly studied auction games, clean models of equilibrium behavior enable empirical researchers to identify and estimate underlying primitives from observed auction outcomes, yielding a rich methodological and applied literature. However, for a large set of extensive-form games that are important in practice, equilibrium characterization is difficult; multiple equilibria often exist, with different equilibria yielding qualitatively different outcomes, and often no complete characterization of these equilibria exists. This class encompasses, for example, certain types of bargaining games, non-standard auctions, signaling games, games with persistent private information, nonstationary games, and oligopoly pricing games with incomplete information. Relative to the rich empirical literature on standard auctions, proposals for identification and estimation under many of these extensive forms have been scarce.

The theoretical literature on mechanism design, pioneered by Myerson (1981) and others, has proposed a different approach. The *revelation principle* allows the analyst to study incomplete information games, independently of specific extensive forms, by studying *revelation mechanisms* – the mappings between agent types and physical outcomes induced by the Bayes-Nash equilibria of extensive form games. Any such mapping from types to outcomes must be *incentive compatible* for all types; conversely, any incentive compatible mapping can be supported as an equilibrium of some incomplete-information game. Hence, studying incentive compatible revelation mechanisms is equivalent to studying the full class of outcomes that can be supported as equilibria of incomplete information games. This abstract notion of revelation mechanisms appears unamenable to empirical settings. A contribution of our paper is to make a precise connection between the mechanism design revelation principle approach and empirical work, and then to propose an extension of this concept which allows for identification and estimation in settings, such as sequential games, where the econometrician often does not observe an agent's full vector of contingent actions.

In this paper, we propose a class of techniques for empirical analysis of incomplete information games, independently of specific extensive forms. We suppose that the econometrician observes multiple repetitions of a trading game, and that in each repetition she observes, for each agent, the final allocation (i.e. whether the good traded hands and, if so, to whom), any transfers paid by each agent, and one additional piece of information, analogous to the report made by the agent to the mechanism designer in a theoretical mechanism design framework. This additional piece of information may either be an action taken by the agent, or some *proxy*—a variable such as an agent's initial action in a sequential game or some characteristic of the agent known to the econometrican but not to the opposing parties. The identification arguments and estimation procedure are analogous in the observed action and proxy cases, but the proxy case is much more general. Rather than analyzing strategies in the context of specific extensive form equilibria, we think of agents' actions/proxies as choices from a menu of feasible *expected physical outcomes* induced by the Bayes-Nash equilibrium of the game. We show that this expected physical outcome menu is sufficient to summarize agents' choices in equilibrium; moreover, this menu can be estimated by the econometrician. If an agent is observed choosing a given point on the menu of feasible physical outcomes, the marginal costs of other feasible options on the equilibrium menu allow the econometrician to derive bounds on the agent's unobserved type. In the case that all types play distinct actions in equilibrium, we show that these bounds collapse to single points, so that the mapping between observed actions and agent types is point-identified independently of the extensive form of the game. If types do not play distinct actions, our approach gives bounds on the values of agents playing any given action; these bounds are the best possible, in the sense that no information about the extensive form of the game can allow more precise identification of values.

For estimation we propose procedures for both discrete and continuous actions cases. In the discrete case, we propose an *empirical ironing* procedure reminiscent of Myerson (1981) to enforce convexity of the outcome menu in the process of estimation. For the continuous case, we propose a nonparametric *local polynomial regression* estimator for values, which estimates pointwise the mapping from actions to values as the ratio of two nonparametric derivative estimates. Additionally, we propose a nonparametric continuous analog of the discrete ironing procedure, using *constrained cubic splines* to estimate continuous menus with convexity constraints.

We illustrate this approach using Monte Carlo simulations of the bilateral bargaining game studied in Satterthwaite and Williams (1989). This game, referred to as a k double auction, has a continuum of qualitatively different equilibria. Without prior knowledge of the precise equilibrium played and the bargaining power weights, the traditional structural approach of inverting player's best-response function (as is done in first price auctions in Guerre, Perrigne, and Vuong (2000), for example), fails. The identification and estimation approach we propose, on the other hand, does not require this prior knowledge, and we demonstrate that it performs well in practice in estimating player's valuations.

We show that our method can be extended in a number of directions, using tools developed in the empirical auctions literature for particular extensive forms. We show that the assumption of independent values can be relaxed; in a setting with non-independent private values, one can proceed by estimating type-contingent menus and using their subgradients for value estimates. We show that we can combine our approach with tools used to analyze settings with unobserved heterogeneity, pioneered by Krasnokutskaya (2011), in a large class of *generalized bidding games*, using deconvolution arguments to recover unobserved-heterogeneity-corrected menus from observed probability/transfer outcomes.

We apply our estimation approach to data from a secret reserve price auction followed by dynamic, two-sided bargaining. This mechanism is used in business-to-business transactions between used-car dealers as well as other settings (Elyakime, Laffont, Loisel, and Vuong (1997)). This game has multiple equilibria and no complete theoretical characterization. In the data, we observe final transaction price, an indicator for whether or not trade occurred (the allocation), and the secret reserve price of the seller, which we assume is strictly increasing in the seller's underlying valuation. We combine our proxy variable approach with our methods to correct for unobserved heterogeneity to estimate the mapping from observed reserve price offers from sellers to seller's values. We find that the estimated distribution of values is close to bounds on seller valuations estimated in Larsen (2014), which used the same data but exploited outcomes of the bargaining game. An advantage of the approach we propose is that it is applicable even in settings where only the final outcome is observed, and not the intermediate actions, such as intermediate offers in a bargaining game.

We use our estimates to compute counterfactual revenue in a setting where, rather than participating in an auction—which increases competition among buyers and thus decreases market power of buyers—the high bidder and seller face each other in a single, take-it-or-leave-it offer bargaining game, with the offer made by the buyer. We find that sellers' gains from trade would decrease by approximately $390–960 per car, suggestive that, in the current mechanism, seller's benefit from a substantial degree of market power.

Several previous papers in the structural estimation literature propose methods that rely on similar ideas to the revelation-principle identification we present here. In particular, the past two decades have seen a number of innovations that yield identification of primitives of interest by plugging in directly observable agent actions, choice probabilities, or outcome probabilities rather than fully solving for equilibria of games. For example, Guerre, Perrigne, and Vuong (2000) demonstrated that valuations in a first price auction can be identified directly from distributions and densities of observed bids. The approach of Guerre, Perrigne, and Vuong (2000) can be thought of as a special case of ours, where ours generalizes the idea to arbitrary incomplete information trading games. Our approach is also related to Tamer (2003), which derived identification results in static discrete games relying on plugging in empirical measures of probabilities that cannot be pinned down uniquely by a model (due to multiplicity of equilibria). In dynamic games, Bajari, Benkard, and Levin (2007) and others proposed two-step methods in which the first step involves estimating policy functions directly from observed choice probabilities rather than from solving the model. Similiar procedures are also adopted in Athey and Nekipelov (2010) applied to search position auctions, in Nekipelov, Syrgkanis, and Tardos (2015) applied to ad auctions, and in Hortaçsu and McAdams (2010) applied to treasury auctions. Agarwal and Somaini (2014) present a method for estimating preferences from reported rankings in a matching game; as in our setting, the authors treat agents as choosing an expected outcome from a menu that can be estimated in the data. In contemporaneous work, Kline (2016) provides an identification argument that is closely related to ours for trading games with observed actions, deriving stronger results than our paper in the non-independent private values case under an additional assumption about equilibrium monotonicity. Our work differs in focusing on both estimation and identification, deriving results relating to menu convexity, and applying our approach to cases where actions are not fully observed.

A number of other papers have built on earlier insights in Guerre, Perrigne, and Vuong (2000), Tamer (2003), and others to achieve feasible estimation approaches in particular settings. For the most part these settings have been cases in which the equilibrium of the game can be characterized and the extensive form is known, and the advantage of plugging in empirical objects in these cases is that it avoids the need to solve for the equilibria. A contribution of our approach is that it yields identification, and a corresponding estimation approach, in arbitrary incomplete information trading settings in which the full characterization of equilibria and the extensive form may be unknown.

2. Model

Throughout, agents will be indexed by i. Uppercase X_i will denote random variables or vectors, lowercase x_i will denote realizations, and bold $\mathbf{x}_i(\cdot)$ will denote functions. For all such objects, we will use a $-i$ subscript to denote the vector of objects for all agents other than i. For example, $X_{-i} \equiv (X_1 \ldots X_{i-1}, X_{i+1} \ldots X_m)$, where m is the number of agents.

We consider an *incomplete information trading game.* In this section, we consider an information environment with asymmetric independent private values; in Subsection 5.1 we extend these results to allow for arbitrarily correlated values. Agents $i \in \{1, 2, \ldots m\}$ have values V_i for a single indivisible good, where each V_i is drawn independently from a continuous bounded distribution F_i, supported on $[\underline{v}_i, \bar{v}_i]$.[1] Agent i's value is observed only by i. All agents are risk-neutral. Let x_i be an indicator representing i attaining the good, and $t_i \in \mathbb{R}$ any net payment made by i. If agent i has value $V_i = v_i$, her utility for the pair (x_i, t_i) is linear in her value, as is standard in the theoretical mechanism design literature:

$$v_i x_i - t_i.$$

Agents play trading game $\mathscr{G}$. The solution concept is Bayes-Nash equilibrium.[2] We will analyze $\mathscr{G}$ in normal form (thus, we do not require refinements such as perfection). First, values V_i are drawn $F_i(\cdot)$ and observed by each agent i. Having observed their types V_i, agents choose (potentially mixed) strategies: $\mathbf{s}_i : \mathbb{R} \to \Delta\mathcal{A}_i$, mapping values $v_i \in \mathbb{R}$ into actions $a_i \in \Delta\mathcal{A}_i$, where $\mathcal{A}_i$ is the space of actions available to i. The outcome allocation and transfers for all agents

$$(x_1, t_1), (x_2, t_2) \ldots (x_m, t_m)$$

are calculated as a function of actions $a_1 \ldots a_m$. We will refer to the individual allocation and transfer functions as $\mathbf{x}_i(a_1 \ldots a_m), \mathbf{t}_i(a_1 \ldots a_m)$. We assume nothing about the structure of $\mathscr{G}$, except that each agent i has some *outside option* $\bar{a}_i$ which leads to some outcome $\bar{x}_i$, and transfer normalized to $\bar{t}_i = 0$.

[1]Our results apply to discrete distributions as well, but assuming continuous types throughout simplifies exposition.

[2]A variety of processes can lead agents to play Bayes-Nash equilibria, from assuming agents have full common knowledge of game rules and prior distributions, to assuming that agents are naive and converge to playing best responses by various learning processes. We do not take a stance on any particular set of assumptions underlying the Bayes-Nash equilibrium concept.

Fixing some strategy $\mathbf{s}_i$, we define $\Sigma_i(v_i)$ as the set of all actions $a_i \in \mathcal{A}_i$ played by type v_i with positive probability under strategy $\mathbf{s}_i(\cdot)$. Let

$$\mathbf{s}_i^{-1}(a_i) = \{v_i : \ a_i \in \Sigma_i(v_i)\},$$

that is, $\mathbf{s}_i^{-1}(a_i)$ is the set of types v_i which play a_i with positive probability under strategy $\mathbf{s}_i$.

Two examples of such incomplete information trading games are the following:

Example 1. *Auction*: Agents $\{1 \ldots m\}$ participate in an auction. Actions a_i belong to a space that depends on the rules of the auction. For example, in a sealed-bid auction, the actions are sealed bids in $\mathbb{R}$. In an ascending or multi-round auction, actions are history-contingent bidding strategies. Agents' outside options are to leave without participating in the auction, leading to $\bar{x}_i = 0$.

Example 2. *Bargaining game*: Seller (player 1) and buyer (player 2) bargain over an indivisible good. The seller's outside option is $\bar{x}_1 = 1$, and the buyer's outside option is $\bar{x}_2 = 0$. Once again, the form of the actions a_i depends on the specific rules of the bargaining game; the game could be a take-it-or-leave-it offer by one party or an alternating-offer bargaining game, or could follow any other bargaining protocol.

Assuming all other agents play according to their equilibrium strategies, if type v_i of i plays action a_i, she attains some expected *physical outcome* $(P_i(a_i), T_i(a_i))$, defined as:

$$P_i(a_i) \equiv E[\mathbf{x}_i(a_i, A_{-i})], \ T_i(a_i) \equiv E[\mathbf{t}_i(a_i, A_{-i})]$$

that is, the expectation of the allocation $x_i(a_i, A_{-i})$ and transfer $t_i(a_i, A_{-i})$ over the actions A_{-i} of players $-i$ (which, from i's perspective, is a random vector). The expected utility that type v_i of agent i attains from playing action a_i, relative to her outside option, is:

$$v_i P_i(a_i) - T_i(a_i) - v_i \bar{x}_i$$

In Bayes-Nash equilibrium, all types v_i of each agent i must be optimally choosing actions with respect to the distributions of opponents' actions A_{-i}. This implies that, for all i, ν_i, the following incentive compatibility condition must hold:

$$a_i \in \Sigma_i(\nu_i) \implies a_i \in \arg\max_{a_i'} \nu_i P_i(a_i') - T_i(a_i') - \nu_i \bar{x}_i \tag{1}$$

Note that, in addition to incentive compatibility, we require *individual rationality:* a_i must do better than the outside option, so the total utility $\max_{a_i'} \nu_i P_i(a_i') - T_i(a_i') - \nu_i \bar{x}_i$ must be nonnegative. However, this condition will not play a major role in our primary identification and estimation arguments, with the exception of the unobserved heterogeneity correction in Subsection 5.2.

Equivalently, we can write:

$$\nu_i \in \mathbf{s}_i^{-1}(a_i) \implies \nu_i P_i(a_i) - T_i(a_i) \geqslant \nu_i P_i(a_i') - T_i(a_i') \ \forall a_i' \tag{2}$$

(1) is a necessary and sufficient condition for strategies $\mathbf{s}_i(\nu_i)$ to constitute a Bayes-Nash equilibrium. Importantly, (1) does not directly reference either the extensive form of the game — that is, the functions $\mathbf{x}_i(a_1 \ldots a_m), \mathbf{t}_i(a_1 \ldots a_m)$ — or the distribution of opponents' actions A_{-i}, avoiding specifying beliefs or any particular equilibrium refinement. This is because neither of the objects $\mathbf{x}_i(a_1 \ldots a_m)$ and $\mathbf{t}_i(a_1 \ldots a_m)$ enter directly into the utility function of type ν_i of agent i. From the perspective of agent i, the equilibrium of $\mathscr{G}$ defines a *menu* of feasible expected physical outcomes $(P_i(a_i), T_i(a_i))$, indexed by action choices a_i. This menu is a *sufficient statistic* for i's choice in equilibrium – each type ν_i of agent i chooses the item $(P_i(a_i), T_i(a_i))$ from the equilibrium menu which affords her the highest utility. We will exploit this menu in a variety of ways below to obtain identification and estimation results.

3. Identification

In this section we derive identification results for the model described above. In Subsection 3.1, we discuss identification in the case in which the econometrician observes agents' actions directly, as in many simultaneous-move trading games such as a sealed-bid auctions or a double auction. Here we prove that the slope of the

(P, T) menu, evaluated at the action chosen by the agent, provides identification of the agent's valuation. The arguments we present in this case are related to a variety of arguments already known in the literature to some extent, although, to our knowledge, this paper, along with Kline (2016), is the first general exposition of how these arguments can be used for identification. We demonstrate that this identification argument holds regardless of whether the structure of the equilibria of the game is known. We focus on these results first because they provide the necessary backdrop for our main result.

Our main result, which is new to the literature, is found in Subsection 3.2. Here we generalize our approach to trading games—such as alternating-offer bargaining or other sequential bargaining or multistage auction games—in which the econometrician cannot observe all contingent actions of agents. Here we require that the econometrician observe a proxy for agents' types. This proxy may be an initial action of a dynamic game or a first-stage bid, or some other feature of the data, such as demographic characteristics about an agent that the econometrican observes but other agents do not.

In all cases, we assume the econometrician observes multiple independent instances of the trading game $\mathscr{G}$, where instances of $\mathscr{G}$ are indexed by j. Thus, in each instance of the game, values V_{ij} are independently drawn from F_i. We assume that in each instance j of the game, the econometrician observes outcomes x_{ij} (the allocation) and t_{ij} (the transfer). If players are asymmetric (i.e. not exchangeable in their indices, i), we assume the econometrican also observes the identity i of any player whose value is to be identified. If players are symmetric, $F_i = F$ for all i.

3.1. **Fully Observed Actions Case.** In this section, we assume that in each instance j of the game, in addition to observing x_{ij} and t_{ij}, the econometrician observes agents' actions a_{ij}. Examples of cases in which the econometrican may observe agents' actions are any sealed-bid trading game or any simultaneous-move trading game. This includes not only first price or second price auctions, where the structure of equilibria is well-known in the theoretical and empirical literature, but also any arbitrary sealed-bid trading game where such properties may be less well-known or less well-behaved, such as the Medicare median-price auction discussed in Cramton, Ellermeyer, and Katzman (2015). This also includes sealed-bid trading games with multiple equilibria, such as a k double auction (Satterthwaite and Williams 1989). We assume nothing else about the structure of the game, or the particular equilibrium being played (except that the same equilibrium is played in each instance j). In particular, the equilibrium need not be increasing or in pure strategies; the median-price auction is such an example. However, we will obtain stronger results when such an increasing, pure-strategy equilibrium does exist.

It is also worth noting that for our identification argument to hold, it need not be the case that all agents are behaving according to an equilibrium. In particular, for identification of the valuation of a particular agent i, it need only be the case that i is best-responding to other agents' actions, regardless of whether those other agents' actions themselves represent best-responses.

In Section 2, we argued that the expected outcome functions $(P_i(a_i), T_i(a_i))$ if sufficient to summarize agents' choices in equilibrium. The basis of our identification approach is that these expected outcome functions can also be estimated by the econometrician. While estimation will be discussed in more detail in Section 4, we simply note here that we from observing n instance of the game, we can consistently estimate $(P_i(a_i), T_i(a_i))$ by taking the empirical averages of x_{ij}, t_{ij} conditional on agent i choosing action a_i:

$$\hat{P}_i(a_i) = \frac{\sum_{j=1}^{n} x_{ij} 1_{a_{ij}=a_i}}{\sum_{j=1}^{n} 1_{s_{ij}=a_i}}, \ \ \hat{T}_i(a_i) = \frac{\sum_{j=1}^{n} t_{ij} 1_{a_{ij}=a_i}}{\sum_{j=1}^{n} 1_{a_{ij}=a_i}}$$

For any given action value a_i, the econometrician can then use (2) to bound the values of any type $v_i \in \mathbf{s}_i^{-1}(a_i)$, that is, any type v_i that plays a_i with positive probability in equilibrium. We state this identification result as the following theorem:

Theorem 1. *For any* a_i*, all* $v_i \in \mathbf{s}_i^{-1}(a_i)$ *satisfy:*

$$v_i \geqslant \frac{T_i(a_i) - T_i(a_i')}{P_i(a_i) - P_i(a_i')} \; \forall a_i': \; P_i(a_i') < P_i(a_i)$$

$$v_i \leqslant \frac{T_i(a_i') - T_i(a_i)}{P_i(a_i') - P_i(a_i)} \; \forall a_i': \; P_i(a_i') > P_i(a_i)$$

Proof. Follows immediately from (2). □

The intuition behind this identification result is as follows. The econometrician observes multiple instances of equilibrium play; hence, the econometrician can take sample averages conditional on any observed action value a_i to estimate the expected physical outcome $(P_i(a_i), T_i(a_i))$ indexed in equilibrium by action a_i, that is, the *menu* of expected physical outcomes $\{(P_i(a_i), T_i(a_i))\}$ available to agent i in equilibrium. Ranking these actions according to their expected allocation $P(\cdot)$, agent i's chosen action reflects how the agent traded off $P(\cdot)$ and the expected transfer $T(\cdot)$, yielding bounds on the agent's value.

In Figure 1, we illustrate a hypothetical equilibrium menu in a setting where agents possible actions are $a_i' \in \{a_i^1, ..., a_i^5\}$. If we observe an agent choosing point a_i^3, it must be the case that the agent's value $v_i \in \mathbf{s}_i^{-1}(a_i^3)$ is lower than the "marginal cost"

$$\frac{T_i(a_i') - T_i(a_i^3)}{P_i(a_i') - P_i(a_i^3)}$$

of items $a_i' \in \{a_i^4, a_i^5\}$ with $P_i(a_i') > P_i(a_i^3)$. Likewise, the agent's value must be higher than the marginal cost

$$\frac{T_i(a_i^3) - T_i(a_i')}{P_i(a_i^3) - P_i(a_i')}$$

from items $a_i' \in \{a_i^1, a_i^2\}$ with $P_i(a_i') < P_i(a_i^3)$. Thus, the value of any agent type choosing point a_i^3 lies between the upper and lower marginal costs from point $(P(a_i^3), T(a_i^3))$, represented by the slopes of the green lines labeled $\underline{v}(a_i^3), \bar{v}(a_i^3)$ respectively.

Since any action played in equilibrium must be optimal for some type, the inequalities in Theorem 1 must have nonempty intersection; in particular, this implies that

FIGURE 1. Hypothetical menu

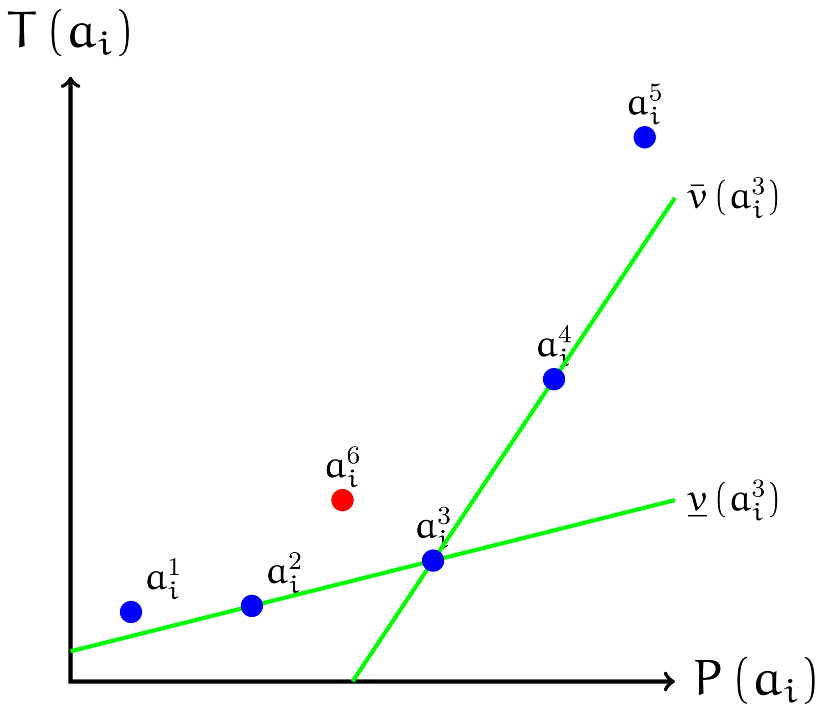

Notes: Hypothetical menu. The slopes of the green lines are upper and lower bounds for the value of an agent choosing action a_i^3.

the menu $\{(P_i(a_i), T_i(a_i))\}$ of actions played with positive probability in equilibrium must be convex, ruling out the existence of points such as a_i^6 in Figure 1.

3.1.1. *Pure Strategies.* We can derive further results if we assume that the equilibrium of $\mathscr{G}$ is in pure strategies; that is, each v plays a single action with positive probability in equilibrium, so that $\Sigma_i(v_i)$ contains only a single value a_i for any v_i. Then we can think of the strategy $\mathbf{s}_i(v_i)$ as a function mapping values to actions $a_i \in \mathcal{A}_i$. The informativeness of the bounds in Theorem 1 depends on the degree to which different types play different actions in game $\mathscr{G}$.[3] Specifically, suppose agents with types δ apart play strictly different actions; that is, $\mathbf{s}_i(v_i + \delta) \neq \mathbf{s}_i(v_i)\ \forall v_i$. Then, we have

[3]We will think of actions a, a' which induce the same expected physical allocation and transfer $(P(a), T(a))$ as identical. Hence, without loss of generality, distinct actions a, a' lead to distinct physical outcomes.

for any ν:

$$\nu_i \leqslant \frac{T_i(\mathbf{s}_i(\nu_i+\delta)) - T_i(\mathbf{s}_i(\nu_i))}{P_i(\mathbf{s}_i(\nu_i+\delta)) - P_i(\mathbf{s}_i(\nu_i))} \leqslant \nu_i + \delta \tag{3}$$

$$\nu_i - \delta \leqslant \frac{T_i(\mathbf{s}_i(\nu_i)) - T_i(\mathbf{s}_i(\nu_i-\delta))}{P_i(\mathbf{s}_i(\nu_i)) - P_i(\mathbf{s}_i(\nu_i-\delta))} \leqslant \nu_i \tag{4}$$

Hence, for any a_i, $\mathbf{s}_i^{-1}(a_i)$ is an interval with length at most 2δ. In particular, if $\mathbf{s}_i(\cdot)$ fully separates types, the intervals $\mathbf{s}_i^{-1}(a_i)$ all collapse onto single points, leading to the following result:

Corollary 1. *If, in game $\mathscr{G}$, each type ν_i has a distinct best response action $\mathbf{s}_i(\nu_i)$, the inverse mapping $\mathbf{s}_i^{-1}(a_i)$ from actions to types is pointwise identified.*

Proof. Follows immediately from Equations 3 and 4. □

As we demonstrate below, the menu $\{(P_i(a_i), T_i(a_i))\}$ is convex, which implies that the inequalities in Theorem 1 are tightest for those values of $P_i(a_i')$ which are closest to $P_i(a_i)$. Supposing we have ordered actions a_i s.t. $P_i(a_i)$ is strictly increasing, then in the case where $\mathbf{s}_i(\cdot)$ is strictly increasing as well, we have the following:

$$\nu_i = \lim_{\delta \to 0} \frac{T_i(\mathbf{s}_i(\nu_i)) - T_i(\mathbf{s}_i(\nu_i-\delta))}{P_i(\mathbf{s}_i(\nu_i)) - P_i(\mathbf{s}_i(\nu_i-\delta))}$$

3.1.2. *Differentiable, Increasing Actions.* In many examples the functions T_i and P_i are smoothly increasing, in addition to $\mathbf{s}_i(\cdot)$, and all three objects are differentiable, in which case this expression simplifies further.

Corollary 2. *If $a_i \in \mathbb{R}$ and the functions $T_i, P_i, \mathbf{s}_i$ are monotonically increasing and differentiable, we have:*

$$\nu_i = \mathbf{s}_i^{-1}(a_i) = \frac{\frac{dT_i}{da_i}}{\frac{dP_i}{da_i}} = \frac{T_i'(a_i)}{P_i'(a_i)}$$

In Subsection 4.2, we will describe an estimation strategy based on Corollary 2. We note in the following example that existing identification arguments for some easily solveable trading games, such as first price auctions, are special cases of our argument.

Example 3. Consider a first price auction in a symmetric IPV environment. Bidder i chooses a bid b_i, and the maximum opposing bid is given by $M_i \sim G(\cdot)$. In this setting, $P_i(b_i) = G(b_i)$ and $T_i(b_i) = b_i G(b_i)$. Player i's value is then given by

$$v_i = \frac{\frac{dT_i}{db_i}}{\frac{dP_i}{db_i}} = \frac{b_i g(b_i) + G(b_i)}{g(b_i)} = b_i + \frac{G(b_i)}{g(b_i)}$$

This expression is equivalent to that derived in the identification argument presented in Guerre, Perrigne, and Vuong (2000). Note, however, that this explicit solution requires knowing the rules/extensive form of the game, whereas our approach does not.

3.1.3. *Theoretical Properties of the* (P, T) *Menu.* We now provide several observations about the structure of the equilibrium menu $\{(P_i(a_i), T_i(a_i))\}$. We first define a number of terms from convex analysis; see, for example, Rockafellar (1997) for more details regarding these and related objects.[4]

Let $\{(P_i(a_i), T_i(a_i))\}$ denote the set of all $(P_i(a_i), T_i(a_i))$ pairs. We will define a *subgradient* of a set $\{(P_i(a_i), T_i(a_i))\}$ at point a_i as any value v such that

$$T_i(a_i') \geqslant T_i(a_i) + v(P_i(a_i') - P_i(a_i)) \ \forall a_i',$$

that is, a line in p, t space of slope v passing through $(P_i(a_i), T_i(a_i))$ which lies weakly below all points in $\{(P_i(a_i), T_i(a_i))\}$. We define the *graph* of $\{(P_i(a_i), T_i(a_i))\}$ as the function obtained by joining the points in order of increasing $P_i(a_i)$ values.

Proposition 1. (1) *The graph of* $\{(P_i(a_i), T_i(a_i))\}$ *is convex.*

(2) *For any* a_i, $\mathbf{s}_i^{-1}(a_i)$ *for any* a_i *is the collection of subgradients of* $\{(P_i(a_i), T_i(a_i))\}$ *at* $P_i(a_i)$. *Each* $\mathbf{s}_i^{-1}(a_i)$ *is a closed interval, and the union of all* $\mathbf{s}_i^{-1}(a_i)$ *contains the interval of values* $[\underline{v}_i, \bar{v}_i]$.

(3) *If we order actions* a_i *by the values of* $P_i(a_i)$, $\mathbf{s}_i^{-1}(a_i)$ *is setwise increasing in* a_i. *For any* a_i, a_i', *the intervals* $\mathbf{s}_i^{-1}(a_i), \mathbf{s}_i^{-1}(a_i')$ *intersect at at most one point.*

[4]Note that our notation is adapted to our setting, and does not correspond exactly to Rockafellar (1997)

Proof. See Appendix A.1. □

Remark. Proposition 1 can be interpreted as follows. Part 1 formalizes the sense in which we refer to the menu $\{(P_i(a_i), T_i(a_i))\}$ as convex. Part 3 states that $\mathbf{s}_i^{-1}(a_i)$ is higher for values of a_i with higher probabilities $P_i(a_i)$. This is related to the classic fact in single-crossing mechanism design that implementable allocation rules must be monotone, assigning higher bundles to higher types. Intuitively, under our "convex menu" interpretation of equilibria, convex menus have monotonically increasing marginal costs, hence agents that choose bundles with higher $P_i(a_i)$ pay higher marginal costs, and thus must have higher values.

Together, parts 2 and 3 also state that each $\mathbf{s}_i^{-1}(a_i)$ is an interval, and distinct intervals $\mathbf{s}_i^{-1}(a_i), \mathbf{s}_i^{-1}(a_i')$ intersect at no more than a single point. This implies that the bounds of Theorem 1 effectively partitions the interval of values $[\underline{v}_i, \bar{v}_i]$. While in general this does not allow us to identify the exact types of each agent, this identification result is the best possible, in the sense that different types in the same interval $v_i \in \mathbf{s}_i^{-1}(a_i)$ are observationally equivalent from the perspective of the econometrician observing x_{ij}, t_{ij}, a_{ij}, regardless of the extensive form of the game played. Thus, the bounds in Theorem 1 capture the full empirical content of the incomplete information games model. This allows us to draw a parallel to mechanism design: if the econometrician observes x_{ij}, t_{ij}, a_{ij}, the extensive-form structure of the game is largely irrelevant for the question of identification of $\mathbf{s}_i^{-1}(a_i)$; the extensive form matters only insofar as it affects the equilibrium mapping of types v_i to expected physical outcomes $\{(P_i(a_i), T_i(a_i))\}$.

3.2. **Main Result: Actions Not Fully Observed.** In many contexts, it is impossible for the econometrician to observe the entire action vector a_{ij} in any instance j of the game. For example, in a multiple-offer bargaining game, observing a_i would entail observing all actions *contingent on all possible sequences* of back-and-forth offers from other agents, or in an ascending auction, all bid strategies over all sequences of bids, within a single instance j of the game. Thus, observing a_i would not simply mean observing what actions i *took*, but what actions i *would have taken* in every history of the game, including those not reached; this would be equivalent to observing a set of instructions player i would have given to another agent to play on her behalf.

However, in many cases the econometrician will be able to observe a *proxy* Z_i which is correlated with A_i, but uncorrelated with A_{-i}. In this section, we show that such proxies Z_i allow the econometrician to derive lower bounds on the utility of any given type in equilibrium. Moreover, if the proxy fully separates types of the agent, we can once again recover the entire mapping from types to actions.[5]

Definition 1. Z_i is a proxy for A_i if Z_i and A_i are not independent, and Z_i is independent of A_{-i}.

Two examples in which this condition is satisfied are:

Example 4. Suppose that A_i specify strategies in a multiple-round bargaining game. If i always makes the first offer in the game, the first offer cannot depend on other actions A_{-i}. So the first offer satisfies the conditions of Definition 1.

Example 5. Suppose we observe characteristics Z_i of agent i, such as demographic information or information about the agent's behavior in other settings, which are correlated with her value V_i (and hence her action A_i), but are unobserved by other players $-i$. Then Z_i must be independent of A_{-i}, since other agents can't condition their actions on i's private information. So these characteristics satisfy the conditions of Definition 1. For example, in the setting of Ambrus, Chaney, and Salitsky (2014)—that of Spanish rescue parties negotiating with North African pirates—the amount of earmarked money raised by the captive's family back home is known to the econometrician and to the buyer (the rescue party) but is unobserved to the seller (the pirates). This earmarked money can serve as an proxy for the rescue party's action.

[5]The "proxy variable" terminology is used similarly in Levinsohn and Petrin (2003), who use flexible functions of proxy variables such as investments or material inputs to control for unobserved productivity.

Supposing that Z_i satisfies our Definition 1, we know that:

$$\mathsf{E}\left[\left(x_i\left(A_i, A_{-i}\right), t_i\left(A_i, A_{-i}\right)\right) \mid Z_i = z_i\right] =$$
$$\mathsf{E}\left[\mathsf{E}\left[\left(x_i\left(A_i, A_{-i}\right), t_i\left(A_i, A_{-i}\right)\right) \mid A_i, Z_i = z_i\right] \mid Z_i = z_i\right]$$

Since we have assumed Z_i is independent of A_{-i}, we can ignore the inner conditioning on Z_i:

$$= \mathsf{E}\left[\mathsf{E}\left[\left(x_i\left(A_i, A_{-i}\right), t_i\left(A_i, A_{-i}\right)\right) \mid A_i\right] \mid Z_i = z_i\right]$$

$$= \mathsf{E}\left[\left(P_i\left(A_i\right), T_i\left(A_i\right)\right) \mid Z_i = z_i\right]$$

Hence, conditional expectations of x_i, t_i with respect to z_i recover convex combinations of $\{(P_i(a_i), T_i(a_i))\}$.

In some cases, there is some proxy Z_i which can be shown to be a one-to-one function of type V_i, that is, $V_i = \boldsymbol{z}_i^{-1}(Z_i)$.[6] For example, suppose the game $\mathscr{G}$ is a multiple-round bargaining game, with a first sealed-bid stage in which the optimal bid is a strictly increasing function of type v_i. In this case, the mapping $\boldsymbol{z}_i(\cdot)$ is fully identified from the data.

If $\boldsymbol{z}_i$ is a one-to-one function of type, then $\left(P_i(\boldsymbol{z}_i), T_i(\boldsymbol{z}_i)\right)$ is exactly the physical outcome attained by the unique type $\boldsymbol{z}_i^{-1}(z)$. Moreover, for any other z_i', the physical outcome $\left(P_i(z_i'), T_i(z_i')\right)$ is attainable by type $\boldsymbol{z}_i^{-1}(z_i)$. Also, there are types $v_i + \delta, v_i - \delta$ playing different actions $\boldsymbol{z}_i(v_i+\delta), \boldsymbol{z}_i(v_i-\delta)$. As in Subsection 3.1, this implies the following bounds for any δ:

$$v_i \leqslant \frac{T_i(\boldsymbol{z}_i(v_i+\delta)) - T_i(\boldsymbol{z}_i(v_i))}{P_i(\boldsymbol{z}_i(v_i+\delta)) - P_i(\boldsymbol{z}_i(v_i))} \leqslant v_i + \delta \tag{5}$$

$$v_i - \delta \leqslant \frac{T_i(\boldsymbol{z}_i(v_i)) - T_i(\boldsymbol{z}_i(v_i-\delta))}{P_i(\boldsymbol{z}_i(v_i)) - P_i(\boldsymbol{z}_i(v_i-\delta))} \leqslant v_i \tag{6}$$

[6]Once again, we will treat different values of z as identical if they induce the same physical outcome $(P(z), T(z))$.

Thus, as in Subsection 3.1, the bounds collapse to a single point, and the entire mapping $z_i(\cdot)$ is identified, which we state as the following extension of Corollary 1:

Corollary 3. *If, in game $\mathscr{G}$, each type v_i is one to one with z_i, the inverse mapping $v_i = z_i^{-1}(z_i)$ from proxies to types is pointwise identified.*

Proof. Follows immediately from Equations 5 and 6. □

We also obtain the immediate extension of this result, that if the $(P_i(z), T_i(z))$ menu, as a function of the proxy, is differentiable, it's slope directly corresponds to the valuation of agent i:

Corollary 4. *If $z_i \in \mathbb{R}$ and the functions T_i, P_i, z_i are monotonically increasing and differentiable, we have:*

$$v_i = z_i^{-1}(z_i) = \frac{\frac{dT_i}{dz_i}}{\frac{dP_i}{dz_i}} = \frac{T_i'(z_i)}{P_i'(z_i)}$$

More generally, using any proxy which satisfies our Definition 1 — even if that proxy is not strictly increasing in the agent's value—we can derive a lower bound for the utility of any given type of the agent in equilibrium. Since the graph of $\{(P_i(a_i), T_i(a_i))\}$ is convex, any conditional expectations with respect to z_i fall within its convex hull. That is to say, for any z_i that satisfies Definition 1, the graph $\{(P_i(z_i), T_i(z_i))\}$ lies strictly above the graph $\{(P_i(a_i), T_i(a_i))\}$. This allows us to lower-bound the utility of any given type v_i in equilibrium:

Corollary 5. *For any proxy satisfying Definition 1, a lower bound on the equilibrium utility of v_i is given by* $\max_{z_i} v_i P_i(z_i) + T(z_i)$

Intuitively, bundles $\{(P_i(z_i), T_i(z_i))\}$ are probably distributions over outcomes $\{(P_i(a_i), T_i(a_i))\}$ from different actions a_i. Hence, agents can achieve the physical outcome associated with any value of z_i by using a mixed strategy corresponding to the distribution over a_i induced by z_i.

4. Estimation

We now present an approach for estimating valuations in incomplete information trading games. The approach follows the identification argument above closely. We focus this section on the case where the agent's actions are either fully observable and are increasing in the agent's type, or where a proxy is observed that is increasing in the agent's type; in either of these two cases, the estimation strategy will be the same. We first discuss estimation in the case of discrete actions/proxies, and then in the case of continuous actions/proxies.

Our goal is to estimate a player's valuation using observations of the game, where in each instance of the game we observe the outcomes (allocation and transfer) and either an action or proxy. Given that the estimation strategy is the same in the action and proxy cases, we will, without loss of generality, refer to "actions" in this section, rather than actions/proxies. Throughout the estimation section, we will focus on estimation for a single agent, thus we will omit subscripts i, writing for example $a, v, P(\cdot), T(\cdot)$ to mean $a_i, v_i, P_i(\cdot), T_i(\cdot)$.

4.1. **Discrete Actions.**

Suppose that there are a finite number of actions, so that $s \in \{a_1 \dots a_K\}$, with generic element a_k. As above, we order the values of a_k in terms of increasing probability $P(a_k)$ of attaining the asset. We wish to identify the set of types $\mathbf{s}^{-1}(a_k)$ choosing each action value a_k. Again, we suppose that the econometrician observes multiple instances of the trading game, and that in each instance she observes the action a_j, the trade outcome x_j and the transfer t_j. We can construct a family of two-step estimators as follows. First, we construct estimates $\hat{P}(a_k), \hat{T}(a_k)$ as the averages of x_j, t_j respectively conditional on actions a_k. Utilizing the convex structure of the set of pairs $\{(P(a_k), T(a_k))\}$, we can then choose, as in Theorem 1:

$$\max_{k'<k}\left[\frac{\hat{T}(a_k)-\hat{T}(a_{k'})}{\hat{P}(a_k)-\hat{P}(a_{k'})}\right] \leqslant \hat{\mathbf{s}}^{-1}(a_k) \leqslant \min_{k'>k}\left[\frac{\hat{T}(a_{k'})-\hat{T}(a_k)}{\hat{P}(a_{k'})-\hat{P}(a_k)}\right]$$

Asymptotically, all ratios $\frac{\hat{T}(a_k)-\hat{T}(a_{k'})}{\hat{P}(a_k)-\hat{P}(a_{k'})}$ converge to their population equivalents, hence $\hat{\mathbf{s}}^{-1}(a_k)$ consistently estimates the bounds of the set $\mathbf{s}^{-1}(a_k)$.

A disadvantage of this estimator is that, in finite samples, the set of $\{(P(a_k), T(a_k))\}$ pairs may not be convex, in which case the lower and upper bounds may cross for some values of a. An alternative strategy is to adopt an "empirical ironing" procedure: rather than using the $\{(P(a_k), T(a_k))\}$ graph directly, we take its convex hull, and use the subgradients of the convex hull to estimate values.

For a given collection of $\left[\hat{P}(\cdot), \hat{T}(\cdot)\right]$ pairs, we define the supporting hyperplane $H(p;\ \nu)$ of slope ν, as the highest line of slope ν which lies below all $\{(P(a_k), T(a_k))\}$ pairs:

$$b(\nu) \equiv \max\{b:\ T(a_k) \geqslant b + \nu P(a_k)\ \ \forall a_k\}$$

$$H(p;\nu) \equiv b(\nu) + \nu p$$

We construct the convex hull of $[P(\cdot), T(\cdot)]$ at any point p by taking the supremum over all supporting hyperplanes:

$$F(p) = \sup_{\nu} H(p;\nu)$$

We will estimate $\hat{\mathbf{s}}^{-1}(a_k)$ using the set of *subgradients* of $F(p)$ at point $P(a_k)$; that is, the set of slopes ν such that $H(p;\nu)$ attains the supremum at point $P(a_k)$:

$$\hat{\mathbf{s}}^{-1}(a_k) = \{\nu:\ H(P(a_k);\nu) = F(P(a_k))\}$$

$F(p)$ is an upper envelope of linear functions $H(p;\nu)$, so it is convex. Thus, it admits subgradients at any point p, and the collection of subgradients is setwise increasing in p. Asymptotically, since the true graph $\{(P(a_k), T(a_k))\}$ is convex, the inferred $\hat{\mathbf{s}}^{-1}(a_k)$ has the same limit as the first estimator. However, using the convex hull of $\{(P(a_k), T(a_k))\}$ ensures that the estimator produces attainable bounds in finite samples.

In the discrete case, estimating the sets of values $\hat{\mathbf{s}}^{-1}(a_k)$ is equivalent to estimating the subgradients of the convex graph $\{(P(a_k), T(a_k))\}$. We have described a simple two-step procedure which accomplish this estimation by estimating the $\hat{P}(a_k), \hat{T}(a_k)$ values, then calculating subgradients based on these. However, it is conceivable that one could construct a more efficient estimator by estimating the subgradients directly from the observed data, rather than estimating the $\hat{P}(a_k), \hat{T}(a_k)$ functions as an intermediate step.

4.2. **Continuous Actions.** In many cases of interest the equilibrium strategies (or the transformations between proxies and values) are smooth functions of values, and the equilibrium $P(a), T(a)$ are also smooth. In this case, we can estimate the mapping from actions to values using nonparametric regression. In particular, assume that the mappings $P(a), T(a)$ are differentiable, and the function $v = \mathbf{s}^{-1}(a)$ is continuous. Corollary 2 implies that:

$$\mathbf{s}^{-1}(a) = \frac{\frac{dT}{ds}}{\frac{dP}{ds}}$$

If we can nonparametrically estimate the derivatives $\hat{T}'(a), \hat{P}'(a)$ as functions of actions a, their ratio is a consistent estimator for $\mathbf{s}^{-1}(a)$. Nonparametric derivative estimation of smooth functions can be done using local polynomial regression (Fan and Gijbels, 1996). The local polynomial regression estimator for $T(a)$ at a given point a with degree p, bandwidth h, kernel K_h is:

$$\left[\hat{\beta}_0(a), \hat{\beta}_1(a) \ldots \hat{\beta}_p(a)\right] = \arg\min_{\beta} \sum_j \left[\left[t_j - \sum_{k=0}^{p} \beta_k (a_j - a)^k\right]^2 K_h(a_j - a)\right] \tag{7}$$

In this expression, p represents the degree of the local polynomial fit; Fan and Gijbels suggest using even polynomial orders $p = k + 2m + 1$ for estimating first derivatives, hence local quadratic regression with $p = 2$ is appropriate for our case. $K_h(\cdot)$ is a kernel function of bandwidth h; common kernel functions include Gaussian or Epanechnikov kernels. The coefficient β_k estimates the kth derivative of T. Therefore, an estimate of the first derivative $\hat{T}'(a)$ is given by performing a local polynomial regression of the observed transfer, t_j, on the observed action, a_j, and taking the coefficient on the linear term in (7), $\hat{\beta}_1$. Similarly, an estimate of the first derivative $\hat{P}'(a)$ is given by performing a local polynomial regression of the observed allocation, x_j (i.e., an indicator for whether the player won), on the observed action, a_j, and taking the coefficient on the linear term in the regression.

We note here that Fan and Gijbels, chap. 4.2, describe the following rule-of-thumb bandwidth selection procedure for local quadratic regression. First, one fits a global quintic polynomial by standard OLS:

$$\hat{\hat{T}}(a) = \alpha_0 + \alpha_1 a \ldots + \alpha_5 s^5$$

Let the residual variance from the regression be $\tilde{s}^2$. The rule of thumb bandwidth is then equal to the following "variance components"-like formula:

$$\hat{h} = C_{\nu,p}(K) \left[\frac{\tilde{s}^2}{\sum_{i=1}^{n} \left(\hat{\hat{T}}''(a_i) \right)^2} \right]^{\frac{1}{7}}$$

Where $C_{\nu,p}(K)$ is a kernel-specific constant, which is approximately 1 for the Gaussian kernel and 2 for the Epanechnikov kernel. Intuitively, this procedure chooses smaller bandwidths for functions that can be fitted better by polynomials. A similar approach applies to estimation of $\hat{P}(\cdot)$.

As in Subsection 4.1, this estimation procedure may result in a nonconvex $\{(P(a), T(a))\}$ menu, and it may be desirable to "iron" the empirical menu function, constraining it to be convex during estimation. In addition, it is often desirable to enforce monotonicity of the $P(a)$ function. In a manner similar to Judd (1998) and Schumaker (1983), we propose a spline-based procedure to nonparametrically estimate the $P(\cdot), T(\cdot)$ functions while imposing convexity of the $[P(\cdot), T(\cdot)]$ menu. In Appendix B, we describe the construction of the quadratic and cubic spline bases shown in Figure 2. Constraining the quadratic (cubic) spline coefficients to be nonnegative ensures that the target function is nondecreasing (convex). By construction, the quadratic splines have two continuous derivatives, and the cubic splines three continuous derivatives.

FIGURE 2. Quadratic/cubic spline basis

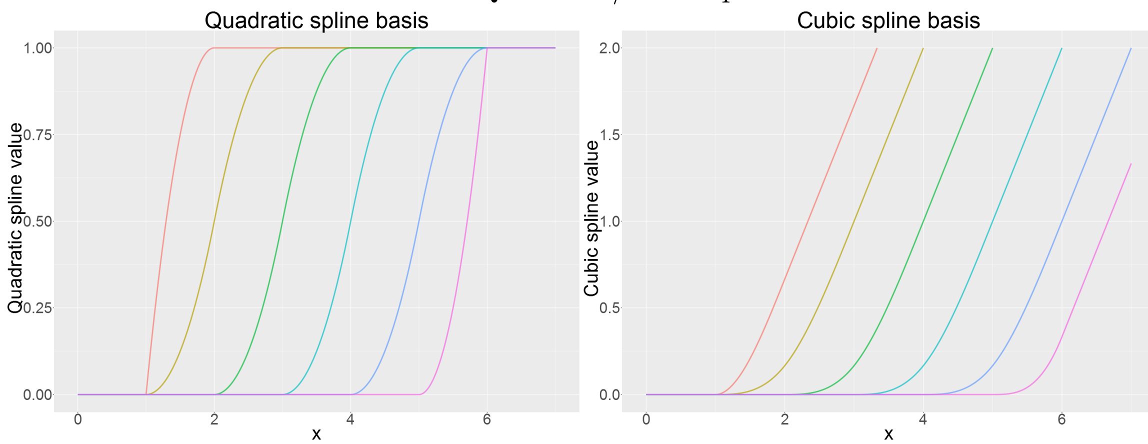

Notes: Quadratic and cubic spline basis functions, with knots at $x = (1, 2, 3, 4, 5, 6)$. In addition, the quadratic spline basis includes an intercept term, and the cubic spline basis includes slope and intercept terms.

Estimation then proceeds in two stages: first, $P(a)$ is nonparametrically estimated as a smooth function of a, possibly constrained to be monotonic using quadratic splines. Then, $T(\cdot)$ is estimated as the composite function $\hat{T}(P(a))$, where $\hat{T}(\cdot)$ is constrained to be a convex cubic spline. Since $\hat{T}(p)$ is a cubic spline, the estimated mapping $\mathbf{s}^{-1}(a) = \frac{dT}{dP}$ is guaranteed to be continuous and differentiable.

4.3. **Simulations.** To illustrate our method, we choose a setting that previously existing methods are incapable of handling: a k double auction. A k double auction is a bilateral bargaining game of incomplete information in which both parties simultaneously submit sealed offers. If the buyer's offer (p^B) exceeds that of the seller (p^S), trade occurs at price $p = kp^S + (1-k)p^B$, where $k \in [0, 1]$. The parameter k can be considered a bargaining power weight. A k double auction with $k = 1$ corresponds to the seller-optimal mechanism (a take-it-or-leave-it offer by the seller) and a k double auction with $k = 0$ corresponds to the buyer-optimal mechanism (a take-it-or-leave-it offer by the buyer).

As demonstrated in Satterthwaite and Williams (1989), this game has infinitely many equilibria that can be qualitatively quite different. Therefore, it is impossible to back out buyer and seller valuations from observed offer data using equilibrium first order conditions, as is done in first price auctions in Guerre, Perrigne, and Vuong (2000) and the follow-on literature, for example, where the equilibrium is unique.[7] Also, even if the model were to have a unique equilibrium, solving for equilibria in k double auctions is somewhat more involved, as described below. The mechanism design approach we propose herein solves both of these issues by identifying/estimating valuations through exploiting the incentive compatibility constraints that must be satisfied by players' observed choices, rather than actually specifying or solving for the equilibrium.

Satterthwaite and Williams (1989) demonstrated that, for any $k = [0, 1]$, a continuum of increasing, differentiable equilibria exist satisfying the following linked differential equations:

$$p^{B(-1)}(p^S(s)) = p^S(s) + kp^{S\prime}(s)\left(s + \frac{F_s(s)}{f_s(s)}\right) \tag{8}$$

$$p^{S(-1)}(p^B(b)) = p^B(b) + (1-k)p^{B\prime}(b)\left(b - \frac{1 - F_b(b)}{f_b(b)}\right) \tag{9}$$

where $s \sim F_s$ is the seller's value, $b \sim F_b$ is the buyer's value, and $p^{B(-1)}(\cdot)$ and $p^{S(-1)}(\cdot)$ are the inverses of the buyer's and seller's strategies. Satterthwaite and Williams (1989) provided an approach for solving for equilibria numerically given knowledge of the distributions of player types. A point (s, b, p) is chosen in the set $\mathcal{P} = \{(s, b, p) : \underline{s} \leqslant s \leqslant p \leqslant b \leqslant \overline{b}, s \leqslant \overline{s}, b \geqslant \underline{b}\}$, and then a one-dimensional manifold passing through this point is traced out using differential equations defined by (8) and (9). This path traces out an equilibrium. An example of a solution path crossing through a point in $\mathcal{P}$ is shown in Figure 3. This approach does not allow for identification of players' value distributions, only for solving for equilibria given knowledge of the distributions. We use their approach to solve for an equilibrium and simulate data from equilibrium play, then apply our method for estimating the underlying valuations to illustrate the estimator's performance.

[7]It is important to note that, as in much of the structural literature, we require that the same equilibrium be selected at all observations in a given sample. This assumption does not imply that the researcher knows *which* equilibrium is selected, only that it be the same in each realization of the game observed in the data.

FIGURE 3. An Equilibrium in the k Double Auction

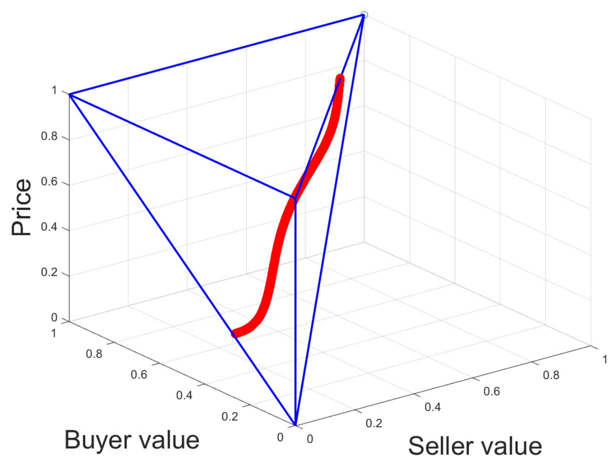

Notes: A solution to the $k = 1/2$ double auction, lying with the tetrahedron $\mathcal{P} = \{(s, b, p) : \underline{s} \leqslant s \leqslant p \leqslant b \leqslant \overline{b}, s \leqslant \overline{s}, b \geqslant \underline{b}\}$. This solution passes through the point $(s, b, p) = (0.375, 0.625, 0.5)$. Buyer valuations are drawn from a $N(0.6, 0.3)$ and seller valuations from a $N(0.5, 0.2)$, with each distribution truncated to lie between $[0, 1]$.

We draw 5,000 realizations of buyer valuations from a $N(0.6, 0.3)$ and seller valuations from a $N(0.5, 0.2)$, with each distribution truncated to lie between $[0, 1]$. We set $k = 1/2$. We choose an equilibrium passing through the point $(s, b, p) = (0.375, 0.625, 0.5)$, which is the equilibrium path illustrated in Figure 3. We solve for this equilibrium using the Satterthwaite and Williams (1989) approach, and then use the simulated draws of buyer and seller valuations to simulate offers and outcomes (the allocation and transfer). We treat these 5,000 realizations of the buyer offer, allocation, and transfer as our data and estimate the (P, T) menu and infer valuations. We focus on estimating buyer valuations for this exercise. For estimation, we apply the local polynomial approach described in Section 4.2.

The estimated menu for the buyer is displayed in Figure 4. As with the illustrative, hypothetical menu displayed in Figure 1, the horizontal axis is the expected probability of trade corresponding to different offers and the vertical axis is the expected transfer at these offers. The expected probability of trade and expected transfer are estimated in separate local polynomial regressions. The estimates are displayed in red dots. Dashed lines display pointwise 95% confidence bands computed from a nonparametric bootstrap constructed by resampling from the data 200 times and performing estimation on each bootstrap sample. The menu is estimated quite precisely.

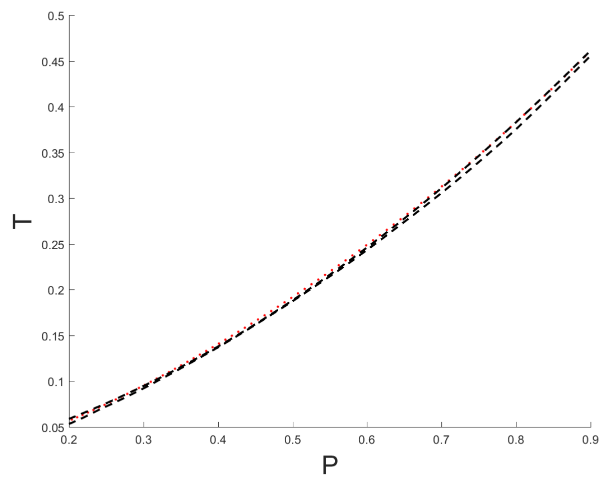

FIGURE 4. Estimated (P, T) Menu from k Double Auction

Notes: (P, T) menu estimated from local polynomial regressions. Red dots indicate estimates, dashed lines indicate 95% confidence bands from 200 bootstrap replications.

It is important to note that Figure 4—the menu—is only a display of data; the menu by itself does not yet impose the structure of our method. That structure comes into play when we interpret the slope of the menu as providing information about buyer valuations. Estimates of this slope are given by the linear terms in the local polynomial regressions. In Figure 5 we plot the observed buyer offers on the horizontal axis and the estimated buyer values (in red dots) on the vertical axis. Dashed lines correspond to pointwise 95% confidence bands. The solid blue line represents the true valuations. The estimated values reflect the true values quite well, with the 95% confidence bands containing the truth over most of the range of offers.

We also remark here that this estimation exercise did not exploit any information about the value of k (the bargaining power), the offers made by the seller, or the particular equilibrium being played. Recall that *any* point in the tetrahedron displayed

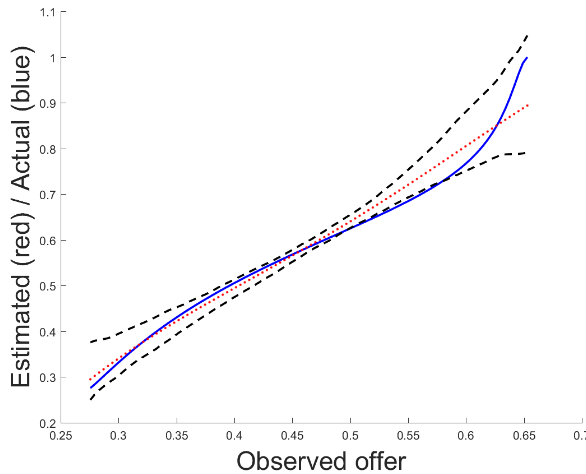

FIGURE 5. Estimated Buyer Values from k Double Auction

Notes: Estimated buyer values from derivatives of (P, T) menu estimated from local polynomial regressions. Red dots indicate estimates, dashed lines indicate 95% confidence bands from 200 bootstrap replications, and solid blue line indicates true values.

in Figure 3 has an equilibrium passing through it, and these equilibria will vary depending on the value of k. Indeed, the generated data in our simulation exercise could have come from any fixed value of k and any fixed equilibrium, and the mechanism design approach would still have returned reasonable estimates of valuations based solely on the observed buyer offer, the allocation, and the transfer.

We conjecture that, in some settings, modeling a bargaining game as k double auction in this fashion may serve as a reasonable alternative to imposing a Nash bargaining structure on the game. Such a framework would allow the presence of incomplete information, unlike Nash bargaining, and would allow for bargaining power (k in this case) to be flexible. The method would require repeated realizations of the game, with observations of the final transfer, the allocation (both from cases where trade occurred and did not occur), and some initial offer (such as a list price or indicative bid) or proxy for the player's valuation. We leave further exploration of this possibility for future work.

5. Extensions: Correlated Values and Unobserved Heterogeneity

We now provide two important extensions to our identification arguments provided above. In Subsection 5.1, we show that our approach generalizes to non-independent private value settings. In Subsection 5.2, we extend our approach to allow for unobserved game-level heterogeneity.

5.1. Non-independent private values.

In this subsection, we relax the assumption that values of different agents are independent. Suppose that agents' values $V_1 \ldots V_m$ are drawn from some joint distribution $F(v_1 \ldots v_m)$, which is common knowledge to all agents. This incorporates and generalizes, for example, the affiliated private value model of first-price auctions analyzed by Li, Perrigne, and Vuong (2002). As above, we suppose that the agents play trading game $\mathscr{G}$. We assume that the equilibrium of the game is separating: equilibrium strategies are described by the $\mathbf{s}_i(v_i)$, where each $\mathbf{s}_i$ is invertible. We show that, as in Subsection 3.1, we can derive bounds on the inverse functions $\mathbf{s}_i^{-1}(\cdot)$ for each a_i.

Let $\mathbf{s}_i(\cdot)$ denote the equilibrium strategy of agent i. Since values are not independent, equilibrium actions will be given by some joint distribution $G(a_1 \ldots a_n)$, derived from $F(v_1 \ldots v_n)$ and the equilibrium strategy $\mathbf{s}_i(\cdot)$. Fix any given value v_i of player i; conditional on v_i, the distribution over values of agents $-i$ is some $F(v_{-i} \mid v_i)$. This conditional distribution of values, combined with the equilibrium strategies of other players $\mathbf{s}_{-i}$, induces a conditional distribution over opponents' actions $G(a_{-i} \mid v_i)$. Thus, in equilibrium, if type v_i of agent i plays action a_i', she attains the physical outcome $[P_i^{v_i}(a_i'), T_i^{v_i}(a_i')]$, defined as the expectation of the physical outcomes $\mathbf{x}_i(a_i', A_{-i}), \mathbf{t}_i(a_i', A_{-i})$ when $A_{-i} \sim G(a_{-i} \mid v_i)$. That is,

$$P_i^{v_i}(a_i') = E[\mathbf{x}_i(a_i', A_{-i}) \mid A_{-i} \sim G(a_{-i} \mid v_i)]$$

$$T_i^{v_i}(a_i') = E[\mathbf{t}_i(a_i', A_{-i}) \mid A_{-i} \sim G(a_{-i} \mid v_i)]$$

In order for type ν_i to play action $a_i = \mathbf{s}_i(\nu_i)$ in equilibrium, as in Subsection 3.1, $\mathbf{s}_i(\nu_i)$ must then satisfy:

$$\mathbf{s}_i(\nu_i) = \arg\max_{a_i} \nu_i P_i^{\nu_i}(a_i) - T_i^{\nu_i}(a_i) - \nu_i \bar{x}_i \tag{10}$$

As in Subsection 3.1, this allows us to bound $\mathbf{s}_i^{-1}(a_i)$, the unique type that plays a_i in equilibrium.

Proposition 2. *for each a_i value, the unique $\nu = \mathbf{s}_i^{-1}(a_i)$ satisfies:*

$$\nu \geqslant \frac{T_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i) - T_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i')}{P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i') - P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i')} \; \forall \left\{a_i' : \; P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i') < P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i)\right\} \tag{11}$$

$$\nu \leqslant \frac{T_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i') - T_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i)}{P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i') - P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i)} \; \forall \left\{a_i' : \; P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i') > P_i^{\nu_i=\mathbf{s}_i^{-1}(a_i)}(a_i)\right\} \tag{12}$$

In the case that the distribution $F(\nu_1 \ldots \nu_n)$ has full support, these bounds collapse to a single point.

Proof. Follows from (10). ☐

If the distribution $F(\nu_1 \ldots \nu_n)$ has full support on the rectangle $[\min \nu_1, \max \nu_1] \times [\min \nu_2, \max \nu_2] \times \ldots$, then the equilibrium probability distribution over action tuples $G(a_1 \ldots a_n)$ will likewise have full support on the product rectangle of actions played; thus, by observing multiple independent repetitions of $\mathscr{G}$, the econometrician can consistently estimate both the equilibrium action distribution $G(a_1 \ldots a_n)$, and the outcomes conditional on all action tuples:

$$P_i(a_1 \ldots a_n) = E[x_i(a_1, \ldots a_n) \mid a_1, \ldots a_n]$$

$$T_i(a_1 \ldots a_n) = E[t_i(a_1, \ldots a_n) \mid a_1, \ldots a_n]$$

Note that $G\left(a_{-i} \mid \nu_i = \mathbf{s}_i^{-1}(a_i)\right)$, the equilibrium action distribution conditional on $\nu = \mathbf{s}_i^{-1}(a_i)$, involves the unknown quantity $\mathbf{s}_i^{-1}(a_i)$. However, this is equivalent to the conditional distribution $G(a_{-i} \mid a_i)$, which can be derived from $G(a_1 \ldots a_n)$.

Thus, for any a_i, the econometrician can estimate the functions:

$$P_i^{v_i = \mathbf{s}_i^{-1}(a_i)}(a_i') = E\left[\mathbf{x}_i(a_i', A_{-i}) \mid A_{-i} \sim G(a_{-i} \mid a_i)\right]$$

$$T_i^{v_i = \mathbf{s}_i^{-1}(a_i)}(a_i') = E\left[\mathbf{t}_i(a_i', A_{-i}) \mid A_{-i} \sim G(a_{-i} \mid a_i)\right]$$

These functions, plugged into the equations in Proposition 2, allow us to identify the unique $\mathbf{s}_i^{-1}(a_i)$.

This approach is related to that of Li, Perrigne, and Vuong (2002), although it is more general, as it applies to incomplete information trading games more broadly, rather than just auctions, showing that our identification strategy applies to private value settings with correlated values. The argument utilizes the fact that any type v_i must play an equilibrium action that is a best response to the distribution of opponents' actions conditional on her type. These conditional distributions can be estimated by the econometrician, allowing us to identify types essentially as in Subsection 3.1.

Our approach in this section requires that the equilibrium strategy $\mathbf{s}_i(v_i)$ is strictly separating. This assumption is necessary because it allows us to estimate the distribution $G(a_{-i} \mid v_i)$ for the unique $v_i = \mathbf{s}_i^{-1}(a_i)$ using the observed $G(a_{-i} \mid a_i)$. If $\mathbf{s}_i$ is not invertible, in general $\mathbf{s}_i^{-1}(a_i)$ is a set of v_i values; thus, the observed $G(a_{-i} \mid a_i)$ is a mixture over distributions $G(a_{-i} \mid v_i)$ for different values $v_i \in \mathbf{s}_i^{-1}(a_i)$. We thus cannot use $G(a_{-i} \mid a_i)$ to consistently estimate $P_i^{v_i}(a_i'), T_i^{v_i}(a_i')$ for any given type v_i when $\mathbf{s}_i$ is not invertible.

5.2. **Unobserved heterogeneity.** We now consider an extension of our identification arguments to a setting with unobserved game-level heterogeneity, similar to the unobserved auction-level heterogeneity in the model of Krasnokutskaya (2011), but applied to the general incomplete information trading games we consider here, rather than only static first price auctions. We refer to the class of games we study here as *generalized bidding games*, although, as before, these games need not be auctions; many bargaining games would also fit into this class. The important feature of games in this class is that actions consist of a price offer. Throughout this section, we refer to the observed actions case, although these results apply to the proxy case as well.

We define the class of *generalized bidding games* as follows. In the first stage, common component W is drawn from $H(\cdot)$ and commonly observed by all agents $1 \ldots m$, but not the econometrician. In the second stage, agents' private values V_i are drawn independently from distributions $F_i(\cdot)$; agents' values are then

$$\tilde{V}_i = V_i + W$$

In the third stage, agents take actions $\tilde{a}_i \in \mathbb{R}$, which are observed by the econometrician. We require the game to satisfy the following property:

Definition 2. We say that game $\mathcal{G}$ satisfies the *generalized bidding game property* for all $\tilde{a}_1 \ldots \tilde{a}_m, \Delta$, and for all i:

$$\mathbf{x}_i(\tilde{a}_1 + \Delta, \tilde{a}_2 + \Delta, \ldots \tilde{a}_m + \Delta) = \mathbf{x}_i(\tilde{a}_1, \tilde{a}_2, \ldots \tilde{a}_m) \tag{13}$$

$$\mathbf{t}_i(\tilde{a}_1 + \Delta, \tilde{a}_2 + \Delta, \ldots \tilde{a}_m + \Delta) = \mathbf{t}_i(\tilde{a}_1, \tilde{a}_2, \ldots \tilde{a}_m) + \Delta(x_i(\tilde{a}_1, \tilde{a}_2, \ldots \tilde{a}_m) - \bar{x}_i) \tag{14}$$

In the case of an auction, (13) implies that if all agents increase bids by a constant amount Δ, the price paid by the winning bidder increases by Δ.

Since the common component W is observed by all agents prior to agents' action choices, agents can condition their strategies on the common component W; thus, we can think of agents' strategies in generalized bidding games as functions $\mathbf{s}_i(v_i, w)$ mapping common components and private values into actions. Bayes-Nash equilibrium in the full game requires that agents' strategies constitute Bayes-Nash equilibria conditional on any value of w. Fixing a given value of w, the game is identical to that of Subsection 3.1. Let A_i^w denote the random variable representing i's equilibrium action when the common component is w. As in Subsection 3.1, we define the expected probability and transfer that i achieves when playing a_i in equilibrium as:

$$P_i^w(a_i) \equiv E\left(\mathbf{x}_i\left(a_i, A_{-i}^w\right)\right), \quad T_i^w(a_i) \equiv E\left(\mathbf{t}_i\left(a_i, A_{-i}^w\right)\right)$$

Type v_i of i's expected utility from playing a_i when the common component is w is:

$$(v_i + w) P_i^w (a_i) - T_i^w (a_i) - (v_i + w) \bar{x}_i$$

Analogous to Subsection 3.1, equilibrium in a generalized bidding game with common component w requires that i's strategy $s_i (v_i, w)$ maximizes her utility, in expectation over the distributions of other agents' actions A_{-i}^w. That is, fixing w, for all i, v_i, we require

$$s_i (v_i, w) \in \arg\max_{a_i} (v_i + w) P_i^w (a_i) - T_i^w (a_i) - (v_i + w) \bar{x}_i$$

In the following proposition, we show that the equilibria of generalized bidding games have a "translation invariance" property with respect to w_j – if actions $a_1 \ldots a_n$ are equilibrium actions conditional on w, actions $a_i + w' - w$ are equilibrium actions under w'.

Proposition 3. *Fix some value of* w*, and suppose that bidding strategies:*

$$s_1 (v_1, w) \ldots s_m (v_m, w)$$

constitute an equilibrium. Then, for any common component w'*, bidding strategies:*

$$s_1 (v_1, w') = s_1 (v_1, w) + w' - w$$

$$\vdots$$

$$s_m (v_m, w') = s_m (v_m, w) + w' - w$$

constitute an equilibrium.

Proof. See Appendix A.2. □

Motivated by this theorem, we will define *markup equilibria* by requiring that agents play the same equilibrium for any common component w:

Definition 3. A *markup equilibrium* is a set of *markup strategies* $\mathbf{s}_i(\nu_i)$, such that:

$$\mathbf{s}_i(\nu_i, w) = \mathbf{s}_i(\nu_i) + w$$

and $\mathbf{s}_i(\nu_i)$ constitute equilibrium strategies for $w = 0$.

Suppose $w_j = 0$. We will define the *markup function* as the expected transfer for $w = 0$:

$$M_i(a_i) = E\left[T_i^0(a_i)\right] = E\left[t_i\left(a_i, A_{-i}^0\right)\right]$$

In the equilibrium conditional on $w = 0$, $T_i(a_i) = M_i(a_i)$, hence we have as before in Subsection 3.1:

$$\mathbf{s}_i^{-1}(a_i) \geqslant \frac{M_i(a_i) - M_i(a_i')}{P_i(a_i) - P_i(a_i')} \; \forall a_i' : P_i(a_i) > P_i(a_i')$$

$$\mathbf{s}_i^{-1}(a_i) \leqslant \frac{M_i(a_i') - M_i(a_i)}{P_i(a_i') - P_i(a_i)} \; \forall a_i' : P_i(a_i) < P_i(a_i')$$

Thus, if we can recover the function $P_i(a_i), M_i(a_i)$, we can bound values as in Subsection 3.1.

Again, we suppose that the econometrician observes multiple independent observations of a generalized bidding game. The econometrician can estimate probability and transfer functions conditional on observed actions which we will refer to as:

$$\tilde{P}_i(\tilde{a}_i) \equiv E\left[\mathbf{x}_i\left(\tilde{a}_i, \tilde{A}_{-i}\right)\right], \; \tilde{T}_i(a_i) \equiv E\left[\mathbf{t}_i\left(\tilde{a}_i, \tilde{A}_{-i}\right)\right]$$

Proposition 4. $P_i(\cdot), M_i(\cdot)$ *are uniquely determined by* $\tilde{P}_i(\tilde{a}_i), \tilde{T}_i(\tilde{a}_i), f_W, f_{a_i}$. *All of these objects are identified from observing multiple independent repetitions of a generalized bidding game. Thus,* $P_i(\cdot), M_i(\cdot)$ *are identified.*

Proof. See Appendix A.3. □

The intuition behind our identification result is as follows. Since the unobserved private value components ν_i are independent by assumption, any correlation in observed actions $\tilde{a}_i$ must be caused by to the unobserved heterogeneity W. Using a method similar to Krasnokutskaya (2011), we can thus separately recover the distribution f_W

of the unobserved heterogeneity term W, and the distribution of actions f_{a_i} generated by the markup strategies, from the observed distribution of actions $f_{\tilde{a}_i}$, through a deconvolution argument. We then show that the distributions f_W, f_{a_i} allow us to recover $P_i(a_i), M_i(a_i)$ from the functions $\tilde{P}_i(\tilde{a}_i), \tilde{T}_i(\tilde{a}_i)$ through a series of deconvolutions and convolutions against $H(\cdot)$. The functions $\tilde{P}_i(\tilde{a}_i), \tilde{T}_i(\tilde{a}_i)$ are expected values of observables, hence they are identified; thus, the functions $P_i(a_i), M_i(a_i)$ are identified.

Once we have recovered $P_i(a_i), M_i(a_i)$, we can bound the value $\mathbf{s}_i^{-1}(a_i)$ of any agent playing the unobserved action a_i. We can then calculate the distribution of values conditional on any observed action $\tilde{a}_i$ by integrating against the distribution $H(\cdot)$.

6. Application to Secret Reserve Auction with Bargaining

In this section, we apply our approach to estimate the valuations of used-car sellers in wholesale used-car markets. In wholesale used car markets, used-car dealers sell cars to other used-car dealers at auction houses. The mechanism employed by the auction houses consists of a secret reserve price, set by the seller, followed by an ascending price auction between multiple potential buyers. If the secret reserve price is not met, the highest bidder and the seller enter into an alternating-offer bargaining game. While the full equilibrium of this game is difficult to characterize, Larsen (2014) proves that the seller's optimal secret reserve price is a strictly increasing function of her value. Hence, the secret reserve price satisfies our conditions in Definition 1 for a proxy which is a one-to-one function of type. In addition, while the game does not exactly satisfy our definition of generalized bidding games in Subsection 5.2, Larsen (2014) shows that equilibria of this bargaining game satisfy the equilibrium translation property of Subsection 3. Thus, we can combine the approaches described in Subsections 5.2 and 3.2 to estimate the equilibrium mapping from sellers' unobserved markup offers to sellers' values.

The data consists of 135,000 realizations of the auction/bargaining game. For each game, the primary variables we observe are the seller's reported secret reserve price, the final transaction price, the final allocation (i.e. an indicator for whether the car sold), as well as the high bid from the auction.

6.1. **Estimation.** We handle observed and unobserved heterogeneity as in Larsen (2014); we summarize these aspects only briefly here. First, we control for observed heterogeneity following the homogenization approach of Haile, Hong, and Shum (2003) by estimating a linear regression of reserve prices and auction high bids on a large set of observable characteristics and treating the residuals from this regression as homogenized reserve prices/auction bids. Let $\tilde{r} = r + w$ represent the residualized reserve price, where w is an additively separable, game-level, *unobserved* heterogeneity scalar term as in Subsection 5.2 and r is the reserve price net of any observed/unobserved heterogeneity. We estimate the densities of w and r, f_w and f_r, using a likelihood approach, modeling each as normal distributions.

With these densities in hand, our main estimation steps are then the following:

(1) Nonparametrically estimate the functions $\tilde{P}(\tilde{r}), \tilde{T}(\tilde{r})$
(2) Using the estimated densities f_r, f_w and the estimated $\tilde{P}(\tilde{r}), \tilde{T}(\tilde{r})$, correct for unobserved heterogeneity to estimate the underlying menu functions $P(r), M(r)$
(3) Take derivatives of the menu $P(r), M(r)$ to get value estimates $v(r)$

We describe each step in turn.

6.1.1. *Nonparametric estimation of* $\tilde{P}(\tilde{r}), \tilde{T}(\tilde{r})$. For step 1, we use local linear regressions of x_j, t_j on $\tilde{r}_j$ to estimate the functions $\tilde{P}(\tilde{r}),\ \tilde{T}(\tilde{r})$. We use normal kernels, and we choose a bandwidth this is larger than the statistically optimal bandwidth, as later steps of the estimation benefit from smoothness of the $\tilde{P}(\tilde{r}),\ \tilde{T}(\tilde{r})$ functions in this stage.

6.1.2. *Unobserved heterogeneity.* In Proposition 4, we showed that $M(r), P(r)$ are identified from the functions

$$\tilde{T}(\tilde{r}),\ \tilde{P}(\tilde{r}),\ f_r(\cdot),\ f_w(\cdot)$$

In particular, $P(r)$ solves:

$$\tilde{P}(\tilde{r}) = \frac{\int_{r=-\infty}^{\infty} P(r) f_r(r) f_w(\tilde{r}-r)\, dr}{\int f_r(r) f_w(\tilde{r}-r)\, dr}$$

and, $M(r)$ solves

$$\tilde{T}(\tilde{r}) - E(w\Delta P \mid \tilde{r}) = \frac{\int_{r=-\infty}^{\infty} M(r) f_r(r) f_w(\tilde{r}-r)\, dr}{\int f_r(r) f_w(\tilde{r}-r)\, dr}$$

Given the estimates of $f_r(\cdot), f_w(\cdot), \tilde{P}(\tilde{r}), \tilde{T}(\tilde{r})$, we solve for $P(r), M(r)$ using minimum weighted distance, using our shape-constrained spline basis functions. We model $P(r)$ as a quadratic spline with 7 knots, constrained to be nondecreasing. We choose spline coefficients to minimize the following objective function:

$$\min_{\hat{P}(\cdot)} \left[\left(\int \hat{P}(r) f_r(r) f_w(\tilde{r}-r)\, dr\right) - \left(\tilde{P}(\tilde{r}) \int f_r(r) f_w(\tilde{r}-r)\, dr\right)\right]^2$$

With the estimated $P(r)$ function, we can then plug this in to estimate the term $E(w\Delta P \mid \tilde{r})$ as:

$$\int P(r)(\tilde{r}-r) f_r(r) f_w(\tilde{r}-r)\, dr$$

We model $M(r)$ indirectly as the function $M(P(r))$, where $M(\cdot)$ is a cubic spline with 7 knots, constrained to be convex. We will choose $M(\cdot)$ to minimize:

$$\min_{\hat{M}(\cdot)} \left[\left(\int \hat{M}(P(r)) f_r(r) f_w(\tilde{r}-r)\, dr\right) - \left(\tilde{T}(\tilde{r}) - E(w\Delta P \mid \tilde{r})\right)\left(\int f_r(r) f_w(\tilde{r}-r)\, dr\right)\right]^2$$

We use standard gradient descent methods to perform spline optimization.

In Figure 6, we show the local linear estimates of $\tilde{P}(\tilde{r}),\ \tilde{T}(\tilde{r})$, as well as the unobserved heterogeneity corrected estimates $P(r),\ T(r)$. Intuitively, the unobserved heterogeneity corrections work as follows. For probabilities, the $\tilde{P}(\tilde{r})$ function is essentially a noisy version of the $\mathsf{P}(r)$ function; thus, correcting for unobserved heterogeneity will imply that $\mathsf{P}(r)$ is steeper than $\tilde{\mathsf{P}}(\tilde{r})$. For transfers, unobserved heterogeneity necessitates two corrections to the $\tilde{\mathsf{T}}(\tilde{r})$ function. First, we subtract from $\tilde{\mathsf{T}}(\tilde{r})$ the term $\mathsf{E}(w\Delta\mathsf{P} \mid \tilde{r})$, which represents the expected value of the unobserved heterogeneity conditional on $\tilde{r}$. Intuitively, for higher values of $\tilde{r}$, we will observe that trades tend to happen at higher prices, but much of this is due to the unobserved heterogeneity term w being higher on average, rather than the markup $\mathsf{M}(r)$ being higher. Comparing the "locpoly" line to the "debiased" line, correcting for $\mathsf{E}(w\Delta\mathsf{P} \mid \tilde{r})$ makes the slope of $\tilde{\mathsf{T}}(\tilde{r})$ significantly less negative. Secondly, $\mathsf{M}(r)$ is essentially a de-noised version of $\tilde{\mathsf{T}}(\tilde{r}) - \mathsf{E}\left(w\Delta\tilde{\mathsf{P}} \mid \tilde{r}\right)$, and thus the slope and concavity of $\mathsf{M}(r)$ are both larger in absolute value than that of $\tilde{\mathsf{T}}(\tilde{r}) - \mathsf{E}(w\Delta\mathsf{P} \mid \tilde{r})$. The net effect is that $\mathsf{M}(r)$ is much less negatively sloped—and somewhat more concave—than the original nonparametric estimate $\tilde{\mathsf{T}}(\tilde{r})$.

6.1.3. *Value estimation.* Since our menu $\mathsf{M}(\mathsf{P}(r))$ is represented as a convex sum of splines, we can analytically take its derivatives, giving us the final estimated mapping $v(r)$ from reserve prices to values.

6.2. **Results and counterfactual.** In the left panel of Figure 7, we show the estimated $[\mathsf{P}(r), \mathsf{M}(r)]$ menu. In the right panel, we show the estimated mapping $v(r)$ between the reserve price r and the inferred value $v(r)$. The estimated reserve-value mapping $v(r)$, combined with the distribution f_r of reserve prices, gives us an estimated distribution F_v of sellers' values, and we plot this in Figure 8.

We use the value estimate to compute a simple counterfactual measuring the how the seller's expected gains from trade would decrease, relative to the current mechanism, if all market power were given to buyers. The current mechanism, with a first-stage auction followed by a second stage of alternating-offer bargaining, may award the lion's share of market power to the seller, as competition between buyers in the auction reduces market power on the buyer side. We simulate a counterfactual mechanism where, instead, the high-bidder from the auction and the seller meet in a one-time, take-it-or-leave-it offer bargaining game, with the offer made by the

FIGURE 6. Deconvolution graphs

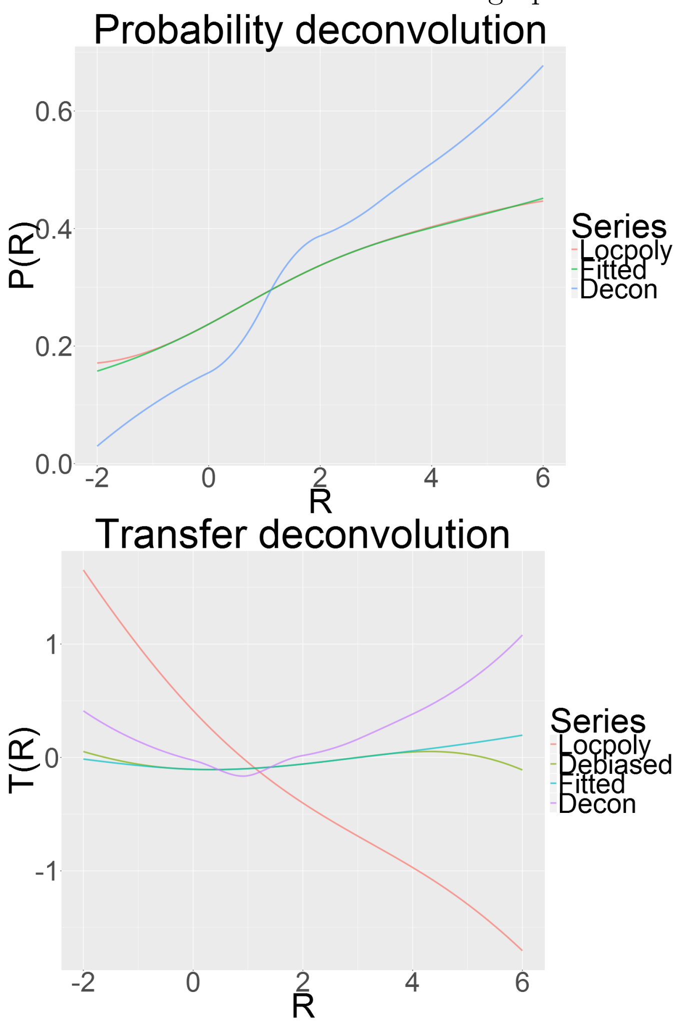

Notes: "Locpoly" lines represent the local polynomial regression estimates of $\tilde{P}(\tilde{r}), \tilde{T}(\tilde{r})$. "Decon" lines are the spline estimates of $P(R), T(r)$. The "debiased" line represents $\tilde{T}(\tilde{r}) - E\left(w\Delta\hat{P} \mid \tilde{r}\right)$. "Fitted" lines are the minimum-distance fits of $P(r)$, $M(r)$ to target functions.

buyer. For simplicity, we compute this counterfactual with the buyer's value set to the mean high bid from the auction (which, after controlling for observed and unobserved heterogeneity, is simply zero).

Before discussing the results, it is necessary to comment on individual rationality (IR) constraints. Throughout our identification arguments and estimation process, we have only used the incentive compatibility conditions of sellers—that is, the condition that outcomes under the reserve prices chosen are preferred to the outcomes

Figure 7. Menu and value estimates

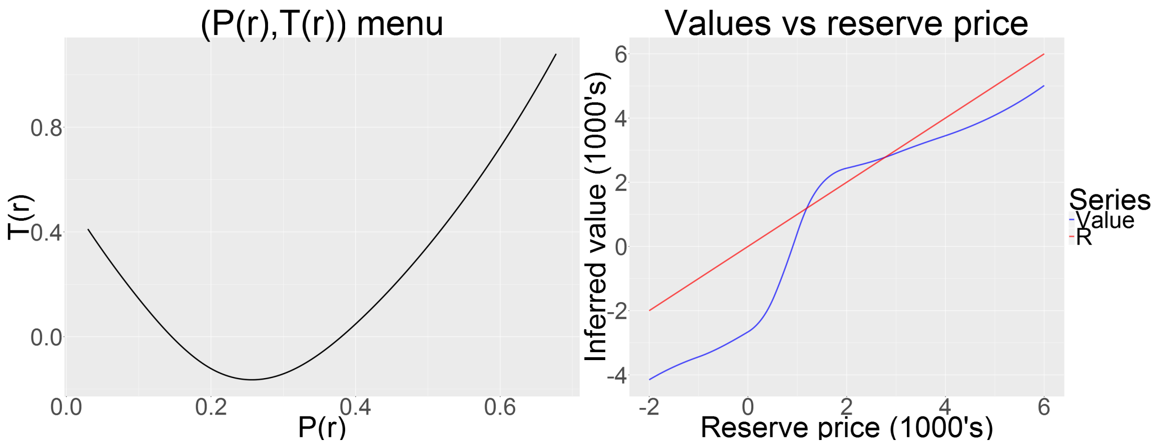

Notes: The left panel shows the final estimated menu. The right panel shows the estimated mapping from reserve prices to values. The reserve price itself is shown in red for comparison.

Figure 8. Estimated value CDF/PDF

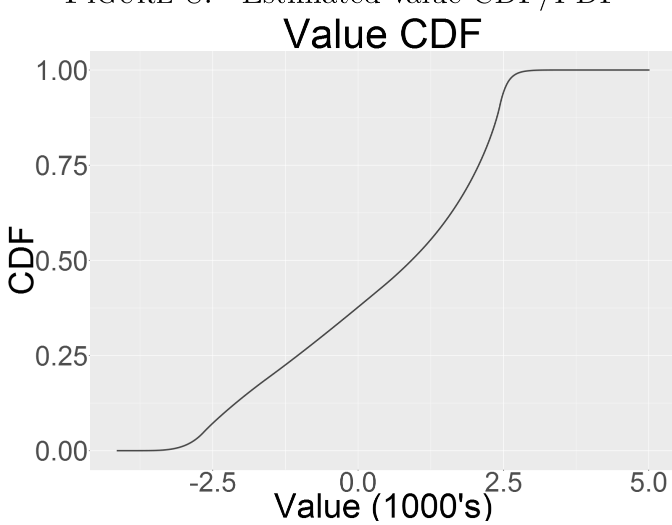

Notes: Value CDF, estimated from the CDF of reserve prices combined with the mapping shown in Figure 7.

from any other possible choice of reserve price. Our results do not rely on or impose IR/participation constraints. This feature can be considered a strength of our approach, in that it relies on weaker conditions than IR constraints would require, but for the purposes of counterfactual analysis, it is necessary to know which seller's would participate in the counterfactual mechanism. For our counterfactual exercise, we enforce the IR constraint in the current mechanism, assuming that sellers whose

Figure 9. Menu vs counterfactual

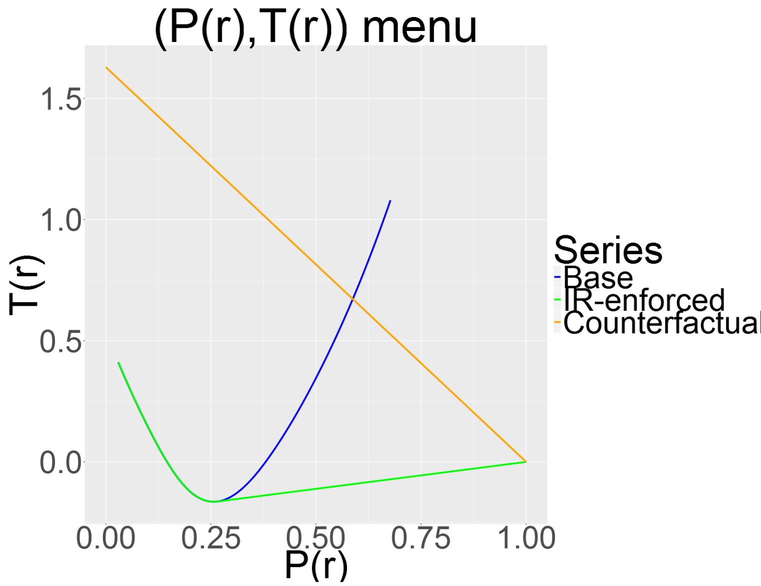

Notes: "Base" is the original menu, from the left panel of Figure 7. "IR-enforced" is the convex hull of the menu and the IR point $(\mathsf{P}, \mathsf{T}) = (1, 0)$. "Counterfactual" is the menu available to the seller under the counterfactual of a single take-it-or-leave-it offer from a buyer.

IR constraint is violated do not trade. An IR-enforced menu can be computed as the convex hull of the original menu and the individual rationality point $(\mathsf{P}, \mathsf{T}) = (1, 0)$.

In Figure 9, the original menu for the current mechanism is shown in blue. The modified, IR-enforced menu is shown green. The orange line in Figure 9 represents the "menu" faced by sellers under this buyer-offer counterfactual bargaining process; it lies strictly above, and thus is dominated by, the IR-constrained menu. We find that in the counterfactual mechanism, giving all the bargaining power to the buyer would reduce the average seller's gains from trade by \$385. As stated above, this welfare change assigns zero change to seller types for whom the IR constraint was binding. If we instead consider only seller types in the range where the IR constraint was non-binding, the average decrease in the sellers' gains from trade is \$961.

7. Conclusion

This paper provided a new, nonparametric identification and estimation approach for trading games of incomplete information. The approach relied on exploiting the incentive compatibility of the direct revelation mechanism corresponding to the actual, underlying (and unknown) extensive form game, rather than attempting to solve for or exploit the equilibrium of this game directly. The main result demonstrated how this approach can be applied in settings where players' actions may not be observable. We believe the approach has the potential to be a useful identification and estimation tool in a number of incomplete-information settings where closed-form equilibrium solutions may not exist, or where players' actions may be difficult to fully observe, such as incomplete information sequential bargaining games.

References

Agarwal, N., and P. Somaini (2014): "Demand analysis using strategic reports: An application to a school choice mechanism," NBER Working paper 20775.

Ambrus, A., E. Chaney, and I. Salitsky (2014): "Pirates of the Mediterranean: An empirical investigation of bargaining with transaction costs," Working paper, Harvard University.

Athey, S., and D. Nekipelov (2010): "A structural model of sponsored search advertising auctions," in *Sixth Ad Auctions Workshop*, vol. 15.

Bajari, P., C. L. Benkard, and J. Levin (2007): "Estimating dynamic models of imperfect competition," *Econometrica*, 75(5), 1331–1370.

Cramton, P., S. Ellermeyer, and B. Katzman (2015): "Designed to fail: The Medicare auction for durable medical equipment," *Economic Inquiry*, 53(1), 469–485.

Elyakime, B., J.-J. Laffont, P. Loisel, and Q. Vuong (1997): "Auctioning and bargaining: An econometric study of timber auctions with secret reservation prices," *Journal of Business & Economic Statistics*, 15(2), 209–220.

Fan, J., and I. Gijbels (1996): *Local Polynomial Modelling and its Applications: Monographs on Statistics and Applied Probability*, vol. 66. CRC Press.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal nonparametric estimation of first price auctions," *Econometrica*, 68(3), 525–574.

Haile, P., H. Hong, and M. Shum (2003): "Nonparametric tests for common values at first-price sealed-bid auctions," NBER Working Paper 10105.

Hortaçsu, A., and D. McAdams (2010): "Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the Turkish treasury auction market," *Journal of Political Economy*, 118(5), 833–865.

Judd, K. (1998): *Numerical Methods in Economics*. Cambridge: The MIT Press.

Kline, B. (2016): "Robust identification in mechanisms," Working Paper, UT Austin.

Krasnokutskaya, E. (2011): "Identification and estimation of auction models with unobserved heterogeneity," *Review of Economic Studies*, 78(1), 293–327.

Larsen, B. (2014): "The efficieny of real-world bargaining: Evidence from wholesale used-auto auctions," NBER Working Paper 20431.

Levinsohn, J., and A. Petrin (2003): "Estimating production functions using inputs to control for unobservables," *Review of Economic Studies*, 70(2), 317–341.

Li, T., I. Perrigne, and Q. Vuong (2002): "Structural estimation of the affiliated private value auction model," *RAND Journal of Economics*, pp. 171–193.

Myerson, R. B. (1981): "Optimal auction design," *Mathematics of Operations Research*, 6(1), 58–73.

Nekipelov, D., V. Syrgkanis, and E. Tardos (2015): "Econometrics for learning agents," in *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 1–18. ACM.

Rockafellar, R. T. (1997): *Convex Analysis*. Princeton University Press.

Satterthwaite, M., and S. Williams (1989): "Bilateral trade with the sealed bid k-double auction: Existence and efficiency," *Journal of Economic Theory*, 48(1), 107–133.

Schumaker, L. I. (1983): "On shape-preserving quadratic spline interpolation," *SIAM Journal on Numerical Analysis*, 20(4), 854–864.

Tamer, E. (2003): "Incomplete simultaneous discrete response model with multiple equilibria," *Review of Economic Studies*, 70(1), 147–165.

Appendix

Appendix A. Proofs

A.1. **Proof of Proposition 1.**

Proof. Part 2. We can rearrange the inequalities in Theorem 1 to:

$$\left(P_i\left(a_i'\right) - P_i\left(a_i\right)\right) v_i \leqslant T_i\left(a_i'\right) - T_i\left(a_i\right) \ \ \forall a_i' \tag{15}$$

Hence, any type $v_i \in \mathbf{s}_i^{-1}\left(a_i\right)$ is a subgradient of $\{\left(P_i\left(a_i\right), T_i\left(a_i\right)\right)\}$ at a_i. Each individual inequality in (15) describes a closed convex set, hence their intersection is a closed convex set, which is a closed interval in $\mathbb{R}$. Since each type v_i must play some action $\mathbf{s}_i\left(v_i\right)$, the union of $\mathbf{s}_i^{-1}\left(a_i\right)$ for all actions a_i is the support of values v_i.

Part 3. Fix some a_i, a_i' and suppose that $P_i\left(a_i'\right) > P_i\left(a_i\right)$. For any $v_i \in \mathbf{s}_i^{-1}\left(a_i\right)$, by (15), we must have:

$$T_i\left(a_i'\right) - T_i\left(a_i\right) \geqslant v_i\left(P_i\left(a_i'\right) - P_i\left(a_i\right)\right)$$

For any $v_i' \in \mathbf{s}_i^{-1}\left(a_i'\right)$, we must have:

$$T_i\left(a_i'\right) - T_i\left(a_i\right) \leqslant v_i'\left(P_i\left(a_i'\right) - P_i\left(a_i\right)\right)$$

Since by assumption $P_i\left(a_i'\right) > P_i\left(a_i\right)$, we have $v_i' \geqslant v_i$. Moreover, if $v_i = v_i'$, we must have

$$v_i = \frac{T_i\left(a_i'\right) - T_i\left(a_i\right)}{P_i\left(a_i'\right) - P_i\left(a_i\right)}$$

Thus, the intersection of $\mathbf{s}_i^{-1}\left(a_i\right), \mathbf{s}_i^{-1}\left(a_i'\right)$ contains at most a single point.

Part 1. Given a set of points $\left(p_i, t_i\right)$, one can verify that the graph of $\left(p_i, t_i\right)$ is convex if and only if every secant line lies above every point $\left(p_i, t_i\right)$. Formally, for any p_1, p_2, p_3, if $tp_1 + (1-t)\,p_2 = p_3$, and $0 \leqslant t \leqslant 1$, then

$$t_3 \leqslant tp_1 + (1-t)\,p_2$$

Consider any a_i. Suppose we have a_i', a_i'' s.t. $tP_i(a_i') + (1-t)P_i(a_i'') = P_i(a_i)$, and $0 \leqslant t \leqslant 1$. We know that the $\{(P_i(a_i), T_i(a_i))\}$ graph admits some subgradient v at a_i. Thus,

$$T_i(a_i) + (P_i(a_i') - P_i(a_i))v \leqslant T_i(a_i')$$

$$T_i(a_i) + (P_i(a_i'') - P_i(a_i))v \leqslant T_i(a_i'')$$

Hence,

$$T_i(a_i) + (tP_i(a_i') + (1-t)P_i(a_i'') - P_i(a_i))v(a_i) \leqslant tT_i(a_i') + (1-t)T_i(a_i'')$$

$$T_i(a_i) \leqslant tT_i(a_i') + (1-t)T_i(a_i'')$$

as desired.

□

A.2. **Proof of Proposition 3.**

Proof. Fixing common component w, we have

$$P_i^w(a_i) \equiv E\left[x_i\left(a_i, A_{-i}^w\right)\right]$$

$$T_i^w(a_i) \equiv E\left[t_i\left(a_i, A_{-i}^w\right)\right]$$

If strategies $\mathbf{s}_i(v_i, w)$ constitute an equilibrium under w, it must be that, for all i, v_i:

$$\mathbf{s}_i(v_i, w) \in \arg\max_{a_i}(v_i + w)P_i^w(a_i) - T_i^w(a_i) - (v_i + w)\bar{x}_i$$

We wish to show that conjectured equilibrium strategies $\mathbf{s}_i'(v_i, w') = \mathbf{s}_i(v_i, w) + (w' - w)$ constitute an equilibrium under w'. Let $A_i^{w'} = \mathbf{s}_i'(V_i, w')$ denote the random variable representing i's action under w', assuming that i plays according to the conjectured equilibrium strategies $\mathbf{s}_i$. We define $\tilde{P}_i^{w'}(a_i), \tilde{T}_i^{w'}(a_i)$ as the expected allocation and transfer i achieves under w', assuming opponents' actions are distributed as $A_{-i}^{w'}$. That is,

$$P_i^{w'}(a_i) = E\left[x_i\left(a_i, A_{-i}^{w'}\right)\right]$$
$$T_i^{w'}(a_i) = E\left[t_i\left(a_i, A_{-i}^{w'}\right)\right]$$

In order to prove our theorem, we need to show that:

$$\mathbf{s}_i'(v_i, w) + (w' - w) \in \arg\max_{a_i} (v_i + w') P_i^{w'}(a_i) - T_i^{w'}(a_i) - (v_i + w')\bar{x}_i$$

We will show a slightly stronger result: for any a_i, the expected utility from playing $a_i + (w' - w)$ under $w', A_{-i}^{w'}$ (net of the outside option) is the same as the expected utility from playing a_i under w, A_{-i}^{w}. That is,

$$(v_i + w') P_i^{w'}(a_i + (w' - w)) - T_i^{w'}(a_i + (w' - w)) - (v_i + w')\bar{x}_i = \\ (v_i + w) P_i^{w}(a_i) - T_i^{w}(a_i) - (v_i + w)\bar{x}_i \quad (16)$$

Thus, if a_i maximizes the RHS, $a_i + (w' - w)$ maximizes the LHS, and we are done. □

A.2.1. *Proof of (16).*

Proof. We have:

$$P_i^{w'}(a_i + (w' - w)) = E\left[x_i\left(a_i + (w' - w),\ A_{-i}^{w'}\right)\right]$$

By construction of $\mathbf{s}_i'(v_i, w') = \mathbf{s}_i(v_i, w) + (w' - w)$, the random variable $A_i^{w'}$ has the same distribution as $A_i^{w} + (w' - w)$. Thus,

$$E\left[\mathbf{x}_i\left(a_i + (w' - w), A_{-i}^{w'}\right)\right] = E\left[\mathbf{x}_i\left(a_i + (w' - w), A_{-i}^{w} + (w' - w)\right)\right]$$

By the generalized bidding game property in Definition 2,

$$\mathbf{x}_i\left(a_i + (w' - w), A_{-i}^{w} + (w' - w)\right) = \mathbf{x}_i\left(a_i, A_{-i}^{w}\right),$$

hence,

$$\mathsf{E}\left[\mathbf{x}_i\left(a_i+(w'-w),A^{w}_{-i}+(w'-w)\right)\right]=\mathsf{E}\left[\mathbf{x}_i\left(a_i,A^{w}_{-i}\right)\right]$$

Thus we have shown that

$$P^{w'}_i\left(a_i+(w'-w)\right)=P^{w}_i\left(a_i\right) \tag{17}$$

For transfers, we have:

$$T^{w'}_i\left(a_i+(w'-w)\right)=\mathsf{E}\left[\mathbf{t}_i\left(a_i+(w'-w),A^{w}_{-i}\right)\right]$$

$$=\mathsf{E}\left[\mathbf{t}_i\left(a_i+(w'-w),A^{w}_{-i}+(w'-w)\right)\right]$$

Again using Definition 2, we have:

$$\mathbf{t}_i\left(a_i+(w'-w),A^{w}_{-i}+(w'-w)\right)=\mathbf{t}_i\left(a_i,A^{w}_{-i}\right)+(w'-w)\left(\mathbf{x}_i\left(a_i,A^{w}_{-i}\right)-\bar{x}_i\right)$$

Hence,

$$\mathsf{E}\left[\mathbf{t}_i\left(a_i+(w'-w),A^{w}_{-i}+(w'-w)\right)\right]= \\ \mathsf{E}\left[\mathbf{t}_i\left(a_i,A^{w}_{-i}\right)\right]+\mathsf{E}\left[(w'-w)\left(\mathbf{x}_i\left(a_i,A^{w}_{-i}\right)-\bar{x}_i\right)\right] \tag{18}$$

The term $\mathsf{E}\left[t_i\left(a_i,A^{w}_{-i}\right)\right]=T^{w}_i\left(a_i\right)$. Using linearity of expectations, the right term simplifies to:

$$\mathsf{E}\left[(w'-w)\left(\mathbf{x}_i\left(a_i,A^{w}_{-i}\right)-\bar{x}_i\right)\right]=(w'-w)\left(\mathsf{E}\left[\mathbf{x}_i\left(a_i,A^{w}_{-i}\right)\right]-\bar{x}_i\right)$$

$$=(w'-w)\left(P^{w}_i\left(a_i\right)-\bar{x}_i\right)$$

Hence, we have shown that:

$$T^{w'}_i\left(a_i+(w'-w)\right)=T^{w}_i\left(a_i\right)+(w'-w)\left(P^{w}_i\left(a_i\right)-\bar{x}_i\right) \tag{19}$$

Once again, i's expected utility from playing $a_i + (w' - w)$ when common component is w' and opponents' actions are distributed as $A_{-i}^{w'}$ is:

$$(v_i + w')\, P_i^{w'}(a_i + (w' - w)) - T_i^{w'}(a_i + (w' - w)) - (v_i + w')\, \bar{x}_i$$

Using the expressions in (17) and (19), this is:

$$= [(v_i + w')\, P_i^{w}(a_i) - T_i^{w}(a_i) - (w' - w)\,(P_i^{w}(a_i) - \bar{x}_i)] - (v_i + w')\, \bar{x}_i$$

$$= [(v_i + w)\, P_i^{w}(a_i) - T_i^{w}(a_i)] - (v_i + w)\, \bar{x}_i$$

Hence we have proved equality in (16). □

A.3. **Proof of Proposition 4.** As before, we observe actions $\tilde{a}_{ij}$, allocations x_{ij}, transfers t_{ij} for a number of repetitions j of the game. We suppose that:

$$\tilde{v}_{ij} = v_{ij} + w_j$$

Following our definition of markup equilibria,

$$\tilde{a}_{ij} = a_{ij} + w_j$$

Following Krasnokutskaya (2011), we can identify the distributions f_W, f_{a_i} using correlation in actions a_i across players. We can also empirically estimate the functions:

$$\tilde{P}_i(\tilde{a}_i) \equiv E\left[\mathbf{x}_i\left(\tilde{a}_i, \tilde{A}_{-i}\right)\right],\ \tilde{T}_i(a_i) \equiv E\left[\mathbf{t}_i\left(\tilde{a}_i, \tilde{A}_{-i}\right)\right]$$

The functions involved in the markup equilibrium are probability of trade and "markup" $M(a_i)$ as a function of the markup action a_i:

$$P_i(a_i) = E\left[\mathbf{x}_i(a_i, A_{-i})\right],\ M_i(a_i) = E\left[\mathbf{t}_i(a_i, A_{-i})\right]$$

Below, we show that the functions $P_i(a_i), M_i(a_i)$ are identified from the functions $P_i(\tilde{a}_i), T_i(\tilde{a}_i), f_W, f_{a_i}$.

A.3.1. *Probabilities.* First, note that:

$$\tilde{P}_i(\tilde{a}_i) = E\left[\mathbf{x}_i\left(\tilde{a}_i, \tilde{A}_{-i}\right)\right]$$

$$= E\left[E\left[\mathbf{x}_i(W + A_i, W + A_{-i}) \mid A_i, W + A_i = \tilde{a}_i\right] \mid W + A_i = \tilde{a}_i\right]$$

By Definition 2, we have

$$\mathbf{x}_i(w + a_i, w + A_{-i}) = \mathbf{x}_i(a_i, A_{-i}) \;\; \forall w$$

$$\implies E\left[\mathbf{x}_i(W + A_i, W + A_{-i}) \mid A_i, \tilde{a}_i = W + A_i\right] = E\left[\mathbf{x}_i(A_i, A_{-i}) \mid A_i\right] = P_i(A_i)$$

Hence,

$$\tilde{P}_i(\tilde{a}_i) = E\left[P_i(A_i) \mid W + A_i = \tilde{a}_i\right]$$
$$\tilde{P}_i(\tilde{a}_i) = \frac{\int_{a_i=-\infty}^{\infty} P_i(a_i) f_{a_i}(a_i) f_w(\tilde{a}_i - a_i)\, da_i}{\int_{a_i=-\infty}^{\infty} f_{a_i}(a_i) f_w(\tilde{a}_i - a_i)\, da_i}$$

This shows that $\tilde{P}_i(\tilde{a}_i)$ is equal to $P_i(a_i)$ convolved against the function

$$\frac{f_{a_i}(a_i) f_w(\tilde{a}_i - a_i)}{\int_{a_i=-\infty}^{\infty} f_{a_i}(a_i) f_w(\tilde{a}_i - a_i)\, da_i}$$

Convolution mappings are invertible, and $f_{a_i}(\cdot), f_w(\cdot)$ are identified, hence $P_i(a_i)$ is identified from the data.

A.3.2. *Markup transfers.* First, note that

$$\tilde{T}_i(\tilde{a}_i) = E\left(\mathbf{t}_i\left(\tilde{a}_i, \tilde{A}_{-i}\right)\right)$$

$$= E\left[E\left[\mathbf{t}_i(W + A_i, W + A_{-i}) \mid A_i, W + A_i = \tilde{a}_i\right] \mid W + A_i = \tilde{a}_i\right]$$

By Definition 2, we have

$$\mathbf{t}_i(W + A_i, W + A_{-i}) = \mathbf{t}_i(A_i, A_{-i}) + W\left(\mathbf{x}_i(A_i, A_{-i}) - \bar{x}_i\right)$$

Taking expectations,

$$E\left[\mathbf{t}_i(W + A_i, W + A_{-i}) \mid W, A_i\right]$$

$$= E\left[\mathbf{t}_i\left(A_i, A_{-i}\right) \mid W, A_i\right] + E\left[W\left(P\left(A_i\right) - \bar{x}_i\right) \mid W, A_i\right]$$

$$= M_i\left(A_i\right) + W\left(P_i\left(A_i\right) - \bar{x}_i\right)$$

Hence,

$$\tilde{T}_i\left(\tilde{a}_i\right) = E\left[\left(M_i\left(A_i\right) + W\left(P_i\left(A_i\right) - \bar{x}_i\right)\right) \mid W + A_i = \tilde{a}_i\right]$$

Since we are conditioning on the event $W + A_i = \tilde{a}_i$, we can substitute out for W:

$$= E\left[\left(M_i\left(A_i\right) + \left(\tilde{a}_i - A_i\right)\left(P_i\left(A_i\right) - \bar{x}_i\right)\right) \mid W + A_i = \tilde{a}_i\right]$$

In integral form, this equation is:

$$\tilde{T}_i\left(\tilde{a}_i\right) = \frac{\int_{a_i=-\infty}^{\infty} M_i\left(a_i\right) f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i}{\int_{a_i=-\infty}^{\infty} f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i} + \frac{\int_{a_i=-\infty}^{\infty} \left(\tilde{a}_i - a_i\right)\left(P_i\left(a_i\right) - \bar{x}_i\right) f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i}{\int_{a_i=-\infty}^{\infty} f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i} \tag{20}$$

The rightmost term represents the average common component of payment. We can define this as:

$$E\left(w\Delta P \mid \tilde{a}_i\right) \equiv \frac{\int_{a_i=-\infty}^{\infty} \left(\tilde{a}_i - a_i\right)\left(P_i\left(a_i\right) - \bar{x}_i\right) f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i}{\int_{a_i=-\infty}^{\infty} f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i}$$

Since we have shown that $P_i\left(a_i\right)$ is identified, $E\left(w\Delta P \mid \tilde{a}_i\right)$ can be calculated for any $\tilde{a}_i$. We can rearrange (20) to:

$$\tilde{T}_i\left(\tilde{a}_i\right) - E\left(w\Delta P \mid \tilde{a}_i\right) = \frac{\int_{a_i=-\infty}^{\infty} M_i\left(a_i\right) f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i}{\int_{a_i=-\infty}^{\infty} f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i}$$

The term $\tilde{T}_i\left(\tilde{a}_i\right)$ can be estimated from the data, so the entire LHS is known. The RHS is a convolution of $M_i\left(a_i\right)$ against

$$\frac{f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right)}{\int_{a_i=-\infty}^{\infty} f_{a_i}\left(a_i\right) f_w\left(\tilde{a}_i - a_i\right) da_i}$$

hence, it is invertible, and thus $M\left(a_i\right)$ is identified.

Appendix B. Spline construction

Let the dependent variable be x, and suppose we wish to estimate $\hat{F}(x)$ using splines. Given a knot sequence $x_1 \ldots x_n$, we define the quadratic splines $S_1^q \ldots S_{n+1}^q$:

$$S_1^q = \frac{2}{x_2 - x_1}\left[|x - x_1|_+ - \frac{|x - x_1|_+^2 - |x - x_2|_+^2}{2(x_2 - x_1)}\right]$$

$$\vdots$$

$$S_k^q = \frac{2}{x_{k+1} - x_{k-1}}\left[\frac{|x - x_{k-1}|_+^2 - |x - x_k|_+^2}{2(x_k - x_{k-1})} - \frac{|x - x_k|_+^2 - |x - x_{k+1}|_+^2}{2(x_{k+1} - x_k)}\right]$$

$$\vdots$$

$$S_n^q = \frac{2}{x_n - x_{n-1}}\left[\frac{|x - x_{n-1}|_+^2 - |x - x_n|_+^2}{2(x_2 - x_1)} - |x - x_n|_+\right]$$

$$S_{n+1} = 1$$

As shown in Figure 2, each S_k^q behaves like a smoothened step function, increasing on the interval $[x_{k-1}, x_{k+1}]$. By constraining coefficients $\beta_1 \ldots \beta_n$ to be nonnegative, we can constrain the target function to be nondecreasing. Moreover, each S_k^q has the property that $\lim_{x\to\infty} S_k^q(x) = 1$; hence, to constrain the target function to always lie below some bound M, we need only constrain $\sum_{k=1}^{n+1} \beta_k \leqslant M$. We impose the constraint that $\sum_k \beta_k \leqslant 1$ in estimating $\hat{P}(\sigma)$.

The family of cubic splines S_k^c we use are integrals of the S_k^q functions, hence they are quadratic splines in first derivative space; that is,

$$S_1^c = \int_{-\infty}^{x} S_1^q(x)\,dx$$

$$\vdots$$

$$S_n^c = \int_{-\infty}^{x} S_n^q(x)\,dx$$

$$S_{n+1}^c = x$$

$$S_{n+2}^c = 1$$

Constraining the coefficients $\beta_1 \ldots \beta_n$ to be nonnegative ensures that the target function is convex.