

A Mechanism Design Approach to Identification and Estimation*

Bradley J. Larsen[†] and Anthony Lee Zhang[‡]

January 31, 2026

Abstract

This paper presents a two-step identification argument for a large class of quasilinear utility trading games, imputing agents' values using revealed preference based on their choices from a convex menu of expected outcomes available in equilibrium. This generalizes many existing two-step approaches in the auctions literature and applies to many cases for which there are no existing tools and where the econometrician may not know the precise rules of the game, such as incomplete-information bargaining settings. We also derive a methodology for settings in which agents' actions are not perfectly observed, bounding menus and agents' utilities based on features of the data that shift agents' imperfectly observed actions. We propose nonparametric value estimation procedures based on our identification results for general trading games. Our procedures can be combined with previously existing tools for handling unobserved heterogeneity and non-independent types. We illustrate the method using simulated data from a first price auction and an incomplete-information, simultaneous-move bargaining game (a k double auction).

Keywords: Incomplete information games, nonparametric identification and estimation, Revelation Principle

*First version: July 2016.

[†]Washington University in St. Louis, Olin Business School and NBER; blarsen@wustl.edu.

[‡]University of Chicago Booth School of Business; anthony.zhang@chicagobooth.edu

1 Introduction

A frequent goal of the growing empirical literature studying mechanism and market design is to compare outcomes of a mechanism that is used in practice to outcomes under a theoretical counterfactual benchmark, such as some benchmark notion of efficiency, optimality, or equity. In order to make such counterfactual comparisons, the econometrician must first estimate primitives (principally, distributions of players' valuations) under the mechanism used in practice. Given advances in the methodological literature over the past twenty years, such an exercise is straightforward for certain types of real-world mechanisms, such as first price auctions. A number of games observed in the real world, however, are quite complex and do not have clear-cut equilibrium characterizations from which the researcher can infer valuations with off-the-shelf tools. Examples of such complex incomplete-information trading games include alternating-offer bargaining games, simultaneous-move bargaining games, non-standard auctions (such as median-price or average-price auctions), auctions preceded by a knockout auction (as in cases of collusion), and auctions followed by bargaining (or the reverse). In such settings, analytically solving for equilibria is difficult, and general proposals for identification and estimation have been scarce.

The theoretical mechanism design literature tackles analysis of such games by abstracting away from the particular rules of the game or extensive form of the game and instead focusing on the direct mechanism corresponding to an equilibrium of the original game (following the celebrated Revelation Principle, Myerson 1979). In this paper, we adopt a similar approach for empirical analysis, abstracting away from the rules or extensive form of the game and instead focusing on revealed preference arguments to develop a unified approach to identification and estimation of players' valuations in such trading games.

We focus on a class of single-unit, transferable utility trading games in which players have independent private values and their utilities are linear in these values. This nests a number of games, such as various kinds of auctions and bargaining games. In such a setting, a player with valuation v who plays action a (such an action could be a bid in an auction or an offer in bargaining) receives an expected payoff of $vP(a) - T(a)$, where $P(a)$ is the probability of winning the good and $T(a)$ is the transfer paid by the player in expectation when choosing action a . Our identification argument is

simple: An agent’s action can be considered to be a choice from a convex menu of possible (P, T) pairs, and subgradients of this menu provide bounds on the agent’s value. When the action space is continuous, the derivative of this menu, evaluated at the action chosen by the agent, corresponds precisely to the agent’s value. The data requirements for identifying an agent’s value are thus that the econometrician observe the final allocation, the final transfer, and the agent’s action from many independent instances of the same underlying game; the objects $P(\cdot)$ and $T(\cdot)$ are then essentially observed in the data, and subgradients of the menu identify agents’ values.

This is a simple revealed preference argument, and it is by no means new to the theoretical mechanism design literature. The novelty of our paper is to exploit this argument to obtain a general identification—and accompanying estimation—approach that applies to a broad class of games regardless of the specific rules of the game, which is applicable even in cases where the econometrician may be unaware of the actual rules of the game, such as when prices arise from some unknown negotiation process. Furthermore, our approach can apply to games in which, even if the rules are known, equilibrium outcomes may not be uniquely predicted by the model (due to the possibility of multiple equilibria) or may simply be difficult to characterize analytically. Each of these features of our approach are not particularly advantageous in a well-understood setting such as a first price auctions, where the theoretical model has a unique, simple-to-characterize equilibrium and where the rules of the game are known. In a number of other settings, however, our approach can be useful.¹

A particular novel feature of this paper is the application of these identification and estimation concepts to cases where actions are imperfectly observed by the econometrician. This can include settings where only a portion of an agent’s actions are observed, such as a multistage auction mechanism where the econometrician observes an indicative (first-stage) bid and the final allocation and price, but not intermediate-stage actions; or a bilateral bargaining setting where the econometrician may observe a seller’s list price and an indicator for whether or not a sale occurs and (if a sale does occur) the final negotiated price, but not the full sequence of intermediate bargaining actions. The imperfectly observed action case can also include cases where the action

¹Our approach can be thought of as a generalization of the first price auction approach of Guerre, Perrigne, and Vuong (2000) to settings where the rules of the game or equilibrium of the game are unknown.

is not observed at all but the econometrician observes a variable in the data that is correlated with that agent’s action and not with other agents’ actions, similar to an instrumental variable. We provide conditions under which our identification and estimation approach can be used to approximate equilibrium menus and players’ values, and we provide conditions under which this approximation will work well.

We derive nonparametric estimation procedures to accompany these identification arguments. We consider both continuous- or discrete-action settings. For the continuous action case, we propose a local polynomial regression procedure. The general trading game model imposes the restriction that equilibrium menus must be convex, and therefore we also propose *empirical ironing* procedures that enforce menu convexity during estimation through flexible shape-constrained splines.

We analyze the performance of our estimation procedure in simulations of first price auctions relative to the approach of Guerre, Perrigne, and Vuong (2000). While our estimator performs well in terms of finite-sample mean squared error, the approach of Guerre, Perrigne, and Vuong (2000) clearly outperforms our estimator, as it exploits the knowledge that the game is indeed a first price auction. We then simulate the bilateral bargaining game studied in Satterthwaite and Williams (1989) and tie our hands, treating the simulated data as though it came from a bargaining game with an unknown extensive form, as is this case with bargaining data in many real-world settings.² We demonstrate that our method again performs well. We emphasize that if a researcher is analyzing a game with known rules and an existing, well-established estimation approach, as in the case of first price auctions, existing tools should be used. The advantage of our approach is in analyzing settings which, in spite of being quite common forms of trade—such as bargaining—do not have well-established tools or well-understood equilibrium characterizations.

The tools we propose can be extended in a number of ways. First, we demonstrate that in certain games our identification and estimation approach can accommodate observed as well as unobserved game-level heterogeneity; that is, a component of player valuations that is common to all players. Second, we show that independence of agents’ values can be relaxed when the equilibrium is assumed to be separating in

²Fudenberg and Tirole (1991) refer to the choice of a particular extensive form as a “thorny issue” when modeling bargaining, as the actual extensive form in many real-world negotiations is unknown to the researcher.

agents' types.

We see the results presented in this paper as having potential to be used to identify and estimate agents' preferences from complex games or mechanisms in a variety of settings, such as wage bargaining (Cramton and Tracy 1992), pre-trial settlement (Silveira 2017; Prescott, Spier, and Yoon 2014), bargaining in developing countries (Keniston 2011), state-run procurement mechanisms (Best, Hjort, and Szakonyi (2023)), international trade negotiations (Bagwell, Staiger, and Yurukoglu 2020), bargaining between hospitals and suppliers (Grennan and Swanson 2020), public housing allocation mechanisms (Waldinger 2021), and others.

Related Literature

Our identification and estimation approach is related to existing approaches for specific types of games, such as Guerre, Perrigne, and Vuong (2000) for first price auctions. Similar approaches have since been applied to many other specific settings. For example, the ad auctions literature (Varian, 2009; Athey and Nekipelov, 2010; Nekipelov, Syrgkanis, and Tardos, 2015) discusses value identification in a manner that is close to ours. Related two-step approaches have also been applied to treasury auctions (Hortaçsu and McAdams, 2010), auctions for bundles of goods (Gentry, Komarova, and Schiraldi, 2023), assignment mechanisms (Agarwal and Somaini, 2018), demand curve estimation in e-commerce (Einav, Kuchler, Levin, and Sundaresan, 2015), conservation auctions (Aspelund and Russo, 2025), and general dynamic games (Bajari, Benkard, and Levin, 2007).

In a related, contemporaneous study, Kline (2026) focuses on identification, but not estimation, in a class of games that overlaps with the class we study: trading games with monotone equilibria and perfectly observable actions. We see his study as complementary to ours; whereas his results apply to settings with correlated values, we only discuss correlated values briefly in Section 7. Our study instead focuses primarily on the case of independent values, deriving additional identification results regarding menu convexity, imperfectly observable actions, and unobserved game-level heterogeneity, and presenting estimation approaches for each of these cases. In another complementary study, Agarwal, Li, and Somaini (2023) derived results that nest a number of related identification arguments, including some of ours. The au-

thors focus on extensions to multidimensional private information but do not consider differences in fully vs. partially observable actions or game-level unobserved heterogeneity.

Our results for handling heterogeneity relate to arguments in Krasnokutskaya (2011), which have previously been applied primarily to particular auction settings and not incomplete information trading games more broadly. Our extension to correlated private values settings relates to arguments from Li, Perrigne, and Vuong (2002).

The most closely related paper is our own, Larsen and Zhang (2021). In that paper, we apply a special case of the identification and estimation approaches we derive herein: the continuous, smooth case with partially observable actions. The emphasis of Larsen and Zhang (2021) is on an empirical application of the approach to used-car markets and the estimation of bargaining power under incomplete information. The current paper instead offers more general identification results, including for fully observed actions, discrete actions, and correlated values, and evaluates the method's performance in simulated data. We clarify when we are restating results already shown in Larsen and Zhang (2021),

Our exposition aims to emphasize several key points. First, while our identification arguments in the case of *perfectly observable* are related to ideas that have been applied to other specific settings in the past, we illustrate here that these identification results largely only require taking a stance on the structure of agents' utility functions, not the specific rules of the game being played. These approaches can be applied in any kind of trading procedure with observed actions, so long as agents' utility functions are assumed to be affine in expected trade probabilities and transfers. We hope that our general exposition can help to conceptually unify the many independent identification results derived in specific settings in the literature. We also highlight that our approach in the observed actions case can apply in games without unique equilibria and that are currently understudied empirically, such as median- or average-price auctions (Cramton, Ellermeyer, and Katzman 2015; Chang, Chen, and Salmon 2014; Decarolis 2018) or double auctions (Chatterjee and Samuelson 1983; Satterthwaite and Williams 1989; Li and Liu 2015).

Our approach to identification and estimation with *imperfectly observed* actions

in Section 4 is new and unrelated to previous approaches. It can allow the econometrician to avoid modeling a complex game and instead use a partially observed action vector or a characteristic of an agent not observed by other agents to identify payoff menus and infer underlying valuations. This can be particularly useful in many complex games that do not immediately fit the mold of existing empirical frameworks and tools, such as sequential bilateral bargaining (Keniston 2011; Ambrus, Chaney, and Salitskiy 2018), auctions followed by bargaining (Elyakime, Laffont, Loisel, and Vuong 1997; Larsen 2021; An and Tang 2019), auctions with collusion (Asker 2010; Kawai and Nakabayashi 2022), or other multi-stage games of incomplete information.

Furthermore, as we show in Section 3, the general trading game model imposes a simple empirical restriction, independently of the extensive-form game being played: equilibrium menus $\{(P_i(a_i), T_i(a_i))\}$ must be convex. This unifies a number of known restrictions on equilibrium outcomes in specific settings: menu convexity is exactly equivalent to the restrictions on first-price auction equilibrium outcomes in Guerre, Perrigne, and Vuong (2000), and the condition that marginal revenues are monotone in the ad auctions (Varian, 2009; Athey and Nekipelov, 2010; Nekipelov, Syrgkanis, and Tardos, 2015) and e-commerce (Einav, Kuchler, Levin, and Sundaresan, 2015) literatures. Our paper demonstrates how to exploit this restriction. We propose estimation procedures in Section 5 below that enforce convexity of equilibrium menus, and our methodology for handling imperfectly observed actions in Section 4 also relies on menu convexity. In principle, menu convexity could be used as a test of the utility and equilibrium assumptions of the trading game model, independently of the particular rules of the game.

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. We will use a $-i$ subscript to denote the vector of objects for all agents other than i . For example, $X_{-i} \equiv (X_1 \dots X_{i-1}, X_{i+1} \dots X_m)$, where m is the number of agents.

We consider an *incomplete information trading game* with asymmetric independent private values. Each agent $i \in \{1, 2, \dots, m\}$ has a value V_i for a single indivisible

good, where each V_i is drawn independently from a continuous bounded distribution $F_i(\cdot)$, supported on $[\underline{v}_i, \bar{v}_i]$. 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:

$$v_i x_i - t_i.$$

Agents play trading game \mathcal{G} . We assume that agents play according to a single Bayes-Nash equilibrium; if there are multiple equilibria of the game, we assume the same equilibrium is played across all observations. We will analyze \mathcal{G} in normal form (thus, we do not require refinements such as perfection).

First, values V_i are drawn from $F_i(\cdot)$ and observed by each agent i . Having observed their types V_i , agents choose (potentially mixed) strategies: $\mathbf{s}_i : \mathbb{R} \rightarrow \Delta \mathcal{A}_i$, mapping values $v_i \in [\underline{v}_i, \bar{v}_i]$ into probability distributions over actions $a_i \in \mathcal{A}_i$, where \mathcal{A}_i is the space of actions available to i . The outcome allocation and transfers for all agents,

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

are calculated as a function of all agents' actions $a_1 \dots a_m$. We will denote the individual allocation and transfer functions as $\mathbf{x}_i(a_1 \dots a_m), \mathbf{t}_i(a_1 \dots a_m)$. These functions are what constitute the *rules of the game*; they determine the outcome for each player given the full set of agent actions. We assume that each agent i has some *outside option* \bar{a}_i which leads to some outcome \bar{x}_i , and transfer normalized to $\bar{t}_i = 0$, independently of the actions of other agents a_{-i} . We assume nothing else about the structure of \mathcal{G} .

For a given 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

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

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

Example 1. Auction: Agents $\{1 \dots m\}$ 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:* A seller (player 1) and a 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; for example, the game could consist of a take-it-or-leave-it offer by one party or an alternating-offer bargaining game, or could follow any other bargaining protocol.

If player i plays action a_i when her value is v_i , she attains some expected outcome $(P_i(a_i), T_i(a_i))$, defined as

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

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 agent i attains where her value is v_i and she plays action a_i , relative to her outside option, is

$$v_i P_i(a_i) - T_i(a_i) - v_i \bar{x}_i. \tag{1}$$

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

$$a_i \in \Sigma_i(v_i) \implies a_i \in \arg \max_{a'_i} v_i P_i(a'_i) - T_i(a'_i) - v_i \bar{x}_i, \tag{2}$$

Or, equivalently,

$$v_i \in \mathbf{v}_i(a_i) \implies v_i P_i(a_i) - T_i(a_i) \geq v_i P_i(a'_i) - T_i(a'_i) \quad \forall a'_i. \tag{3}$$

In addition to incentive compatibility, we require *individual rationality*: if the outside option is not chosen, the action that is chosen must yield a better outcome

than the outside option, so the total utility $\max_{a'_i} v_i P_i(a'_i) - T_i(a'_i) - v_i \bar{x}_i$ must be nonnegative. This condition will not play a major role in our identification and estimation arguments, with the exception of the unobserved heterogeneity correction in Section 7.2.

Equation (3) is a necessary and sufficient condition for strategies $\mathbf{s}_i(v_i)$ to constitute a Bayes-Nash equilibrium. Importantly, (3) does not directly reference either the rules of the game—that is, the functions $\mathbf{x}_i(a_1 \dots a_m), \mathbf{t}_i(a_1 \dots a_m)$ —or the distribution of opponents’ actions A_{-i} . This is because neither of the objects $\mathbf{x}_i(a_1 \dots a_m)$ and $\mathbf{t}_i(a_1 \dots a_m)$ enter directly into the expected utility function of type v_i of agent i . From the perspective of agent i , the equilibrium of \mathcal{G} defines a *menu* of feasible expected 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: given her type v_i , each agent i chooses the item $(P_i(a_i), T_i(a_i))$ from the equilibrium that affords her the highest utility. The equilibrium menu is also sufficient for characterizing agents’ welfare gains from participating in trade; the expected welfare gain of type v_i is the maximized value of (2).

In what follows, we will assume that the econometrician observes data from many independent realizations of the same underlying trading game. As mentioned above, we will further assume that all realizations of the game represent play from the same Bayes-Nash Equilibrium. While this assumption is common in the structural empirical literature, it is certainly strong. We emphasize here, however, that this assumption does not imply that the underlying game need only have a single equilibrium, only that the data represents play from a single equilibrium. This therefore allows the data to speak as to the equilibrium selection from a (potentially infinite) set of equilibria rather than requiring the researcher to impose an equilibrium refinement. This is particularly useful in the settings we seek to analyze, as many of these are cases with multiple equilibria that are qualitatively very different and where theoretical equilibrium refinements do not always yield a unique equilibrium.³ Moreover, all of our results in this paper still hold if the assumption of a single equilibrium in the data is relaxed to instead assume that, if the equilibrium does indeed differ from one

³For example, in the case of bilateral bargaining, proposed equilibrium refinements often fail to exist or result in equilibria that predict immediate agreement. See discussion in Ausubel, Cramton, and Deneckere (2002) or results in Perry (1986).

realization of the game to another, then which equilibrium is selected depends only on features in the data that are observable to the econometrician and can thus be controlled for. We discuss this further in Section 8.4.

By allowing all value distributions F_i and action sets \mathcal{A}_i to be distinct, we allow agents to be completely asymmetric. This is important for games in which players play different roles, such as in bilateral bargaining settings, as in Example 2. In some cases below, we will use the subscripts i to denote a particular role played by the agent, such as using $i \in \{S, B\}$ indexing whether the player is a seller or buyer in bargaining. In practice, each instance of the trading game may involve classes of agents who are indistinguishable. For example, different bidders in an auction may be indistinguishable to the econometrician, or only distinguishable based on certain observables; an example is the case of loggers vs. mills bidding in timber auctions (Athey, Coey, and Levin 2013). In such settings, the econometrician may wish to model all players from a certain class as drawing valuations from a single distribution. All of our identification and estimation results can be thought of as applying to such cases as well.

3 Identification with Fully Observed Actions

In this section, we derive identification results for the model described above, assuming that the econometrician fully observes agents' actions from many instances of equilibrium play of the trading game \mathcal{G} . We first analyze the case where the equilibrium is not necessarily increasing or in pure strategies. We then demonstrate stronger results that can be obtained when such an increasing, pure-strategy equilibrium does exist. In each instance of the game, values V_i , for each agent $i \in \{1, \dots, m\}$, are independently drawn from F_i (and these actions are independent across instances of the game). Agents then take actions corresponding to a set of equilibrium strategies $(\mathbf{s}_1(\cdot), \dots, \mathbf{s}_m(\cdot))$. We assume in this section that in each instance of the game, the econometrician observes x_i (the allocation for agent i), t_i (the transfer for agent i), and agent i 's action a_i .

Examples of cases in which the econometrician might 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 strategies less well-behaved, such as the median-price auction used in Medicare durable good equipment auctions, which is known to potentially yield equilibria that are non-increasing (Cramton, Ellermeyer, and Katzman 2015); or sealed-bid bilateral bargaining games, such as the k double auction, which are known to have multiple, qualitatively different equilibria (Chatterjee and Samuelson 1983; Satterthwaite and Williams 1989; Li and Liu 2015). Another example is the average-bid auction, which Chang, Chen, and Salmon (2014) and Decarolis (2018) demonstrate can have multiple equilibria and, in some cases, no equilibria in increasing strategies. Decarolis (2018) highlights that this mechanism is used for public procurement by state agencies in Florida and New York, as well as in Chile, China, Colombia, Italy, Japan, Peru, Malaysia, Switzerland, and Taiwan.

3.1 Fully Observed Actions: General Case

In Section 2, we argued that the expected outcome functions $(P_i(a_i), T_i(a_i))$ are sufficient to summarize agents' choices in equilibrium, without requiring knowledge of the full functions $\mathbf{x}_i(a_1 \dots a_m)$ and $\mathbf{t}_i(a_1 \dots a_m)$. These expected outcome functions are simply conditional expectations of outcomes x_i, t_i with respect to actions a_i . Hence value identification follows from a revealed preference argument: for any given action value a_i , the econometrician can use the incentive compatibility conditions in (3) to bound the values of any type $v_i \in \mathbf{v}_i(a_i)$, that is, any type v_i that plays a_i with positive probability in equilibrium. We state this result as the following theorem, which corresponds to Theorem 1 of Larsen and Zhang (2021):

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

$$v_i \geq \frac{T_i(a_i) - T_i(a'_i)}{P_i(a_i) - P_i(a'_i)} \quad \forall a'_i : P_i(a'_i) < P_i(a_i)$$

$$v_i \leq \frac{T_i(a'_i) - T_i(a_i)}{P_i(a'_i) - P_i(a_i)} \quad \forall a'_i : P_i(a'_i) > P_i(a_i).$$

Proof. Follows immediately from (3). □

Agents' values are inferred from the points they choose on the *menu* of expected outcomes $\{(P_i(a_i), T_i(a_i))\}$. In Figure 1, we illustrate a hypothetical equilibrium menu in a setting where agent i 's possible actions are $a'_i \in \{a_i^1, \dots, a_i^5\}$. Indifference curves in this figure correspond to straight lines, with higher utility achieved by the agent on curves lying more to the southwest of the figure.

Suppose we observe an agent choosing point a_i^3 . Compared to a_i^3 , points a_i^4 and a_i^5 have higher probability $P_i(a_i)$ of receiving the good, in exchange for higher transfers $T_i(a_i)$. If the agent prefers point a_i^3 to a_i^4 and a_i^5 , her value must be lower than the average cost of purchasing the additional probability when moving from a_i^3 to a_i^4 and a_i^5 ; that is,

$$v_i \leq \frac{T_i(a'_i) - T_i(a_i^3)}{P_i(a'_i) - P_i(a_i^3)},$$

for items $a'_i \in \{a_i^4, a_i^5\}$ with $P_i(a'_i) > P_i(a_i^3)$. Similarly, compared to point a_i^3 , points a_i^1 and a_i^2 have lower probability of receiving the good in exchange for lower transfers. If the agent prefers point a_i^3 , her value must be higher than the average cost of purchasing the additional probability of getting the good when moving from a_i^1 or a_i^2 to a_i^3 ; that is,

$$v_i \geq \frac{T_i(a_i^3) - T_i(a'_i)}{P_i(a_i^3) - P_i(a'_i)},$$

for 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 must lie between 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 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. This is a known restriction on the structure of incentive-compatible revelation mechanisms from the theoretical mechanism design literature; we formalize this observation and its implications in Proposition 1 below.

Let $\{(P_i(a_i), T_i(a_i))\}$ denote the set of all $(P_i(a_i), T_i(a_i))$ pairs. A *subgradient* of a set $\{(P_i(a_i), T_i(a_i))\}$ at point a_i is any value ν such that

$$T_i(a'_i) \geq T_i(a_i) + \nu(P_i(a'_i) - P_i(a_i)) \quad \forall a'_i,$$

that is, a value ν such that a line in \mathbb{R}^2 of slope ν passing through $(P_i(a_i), T_i(a_i))$ lies weakly below all points in $\{(P_i(a_i), T_i(a_i))\}$. The *graph* of $\{(P_i(a_i), T_i(a_i))\}$ is 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 , $\mathbf{v}_i(a_i)$ for any a_i is the collection of subgradients of $\{(P_i(a_i), T_i(a_i))\}$ at $P_i(a_i)$. Each $\mathbf{v}_i(a_i)$ is a closed interval, and the union of all $\mathbf{v}_i(a_i)$ contains the interval of values $[\underline{v}_i, \bar{v}_i]$.*

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

The proof of Proposition 1 and all other involved proofs are found in the Appendix. Part 1 of Proposition 1 corresponds to Theorem 2 of Larsen and Zhang (2021), and states that the menu is convex. Part 3 states that $\mathbf{v}_i(a_i)$ is higher for values of a_i with higher probabilities $P_i(a_i)$. This is related to the classic fact in 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 slopes, where the slope represents the average cost of additional probability. Agents who choose bundles with higher $P_i(a_i)$ pay higher average costs per unit probability, and thus must have higher values.

Together, parts 2 and 3 also state that each $\mathbf{v}_i(a_i)$ is an interval, and distinct intervals $\mathbf{v}_i(a_i), \mathbf{v}_i(a'_i)$ intersect at no more than a single point. This implies that the bounds of Theorem 1 partition 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 $v_i \in \mathbf{v}_i(a_i)$ are observationally equivalent if the econometrician only observes x_i, t_i, a_i for different agents i in each instance of the game (and does not observe, for example, actions that are never chosen in equilibrium or the probabilities with which agents mix across actions in a mixed-strategy equilibrium). Thus, the bounds in Theorem 1 capture the full empirical content of the incomplete information trading game model; in other words, the structure of the game only matters for value identification insofar as it affects the equilibrium menu $\{(P_i(a_i), T_i(a_i))\}$.

3.2 Pure Strategy or Increasing Equilibrium

We can derive stronger results than those in Theorem 1 if we assume that the equilibrium of \mathcal{G} is in pure strategies, so that the strategy $\mathbf{s}_i(v_i)$ is 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 \mathcal{G} . Specifically, suppose agents with types δ apart play strictly different actions, so that $\mathbf{s}_i(v_i + \delta) \neq \mathbf{s}_i(v_i) \forall v_i$, and that distinct actions lead to distinct expected outcomes $(P_i(a_i), T_i(a_i))$. Then, we have, for any v_i ,

$$v_i \leq \frac{T_i(\mathbf{s}_i(v_i + \delta)) - T_i(\mathbf{s}_i(v_i))}{P_i(\mathbf{s}_i(v_i + \delta)) - P_i(\mathbf{s}_i(v_i))} \leq v_i + \delta, \quad (4)$$

$$v_i - \delta \leq \frac{T_i(\mathbf{s}_i(v_i)) - T_i(\mathbf{s}_i(v_i - \delta))}{P_i(\mathbf{s}_i(v_i)) - P_i(\mathbf{s}_i(v_i - \delta))} \leq v_i. \quad (5)$$

Hence, for any a_i , $\mathbf{v}_i(a_i)$ is an interval with length at most 2δ . In particular, if $\mathbf{s}_i(\cdot)$ fully separates types, the interval $\mathbf{v}_i(a_i)$ for any a_i collapses to a single point, leading to the following result:

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

Proof. Follows immediately from (4) and (5). □

Under corollary 1, we also have:

$$v_i = \lim_{\delta \rightarrow 0} \frac{T_i(\mathbf{s}_i(v_i)) - T_i(\mathbf{s}_i(v_i - \delta))}{P_i(\mathbf{s}_i(v_i)) - P_i(\mathbf{s}_i(v_i - \delta))}. \quad (6)$$

If the problem is sufficiently smooth, as described in the following corollary, this expression simplifies further.

Corollary 2. *If $a_i \in \mathbb{R}$, P_i and $\mathbf{s}_i(\cdot)$ are both increasing and continuous, and P_i, T_i are both differentiable, then*

$$v_i(a_i) = \frac{T'_i(a_i)}{P'_i(a_i)}.$$

In Section 5, we will describe an estimation strategy based on Corollary 2.

As highlighted above, one example of an incomplete information trading game that fits into our framework is a first-price auction. We now demonstrate that the standard identification argument for first-price auctions, derived in Guerre, Perrigne, and Vuong (2000), is analytically equivalent to our approach; however, theirs relies explicitly on knowing that the game is indeed a first-price auction whereas ours does not.

Example 3. Consider an m -bidder first-price auction in an independent private values environment. Assume for simplicity that bidders are symmetric. Bidder i 's action a_i is a bid. Let the distribution of bids be written $G(\cdot)$, with density $g(\cdot)$. In a first-price auction, the expected probability of winning, (P), and expected transfer (T) are given by

$$P(a_i) = G(a_i)^{m-1}, \quad T(a_i) = a_i G(a_i)^{m-1},$$

where i subscripts are omitted from P and T because of symmetry. Player i 's value is then given by

$$\begin{aligned} \frac{dT(a_i)}{da_i} &= \frac{a_i(m-1)G(a_i)^{m-2}g(a_i) + G(a_i)^{m-1}}{(m-1)G(a_i)^{m-2}g(a_i)} \\ &= a_i + \frac{G(a_i)}{(m-1)g(a_i)}. \end{aligned}$$

This expression is equivalent to that derived in the identification argument of Guerre, Perrigne, and Vuong (2000).

4 Identification with Imperfectly Observed Actions

In many contexts, it is impossible for the econometrician to observe all elements of a realization of the action a_i . For example, in a multiple-offer bargaining game, observing a_i would entail observing a vector of all of i 's actions *contingent on all possible sequences* of offers from other agents, or in an ascending auction, all of i 's bid strategies over all sequences of opponent bids, within a single instance 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 a third-party agent to play on her behalf.

However, in many of these cases where the action is a multidimensional vector, the econometrician will be able to observe some portion of an agent’s action vector or will be able to observe some other feature in the dataset that is correlated with an agent’s unobserved action vector. We refer to such an object—either a partially observed action or a separate feature in the dataset correlated with an agent’s action—as an *action shifter*, defined as follows:

Definition 1. Z_i is an *action shifter* if Z_i and A_i are not independent, but Z_i and A_{-i} are independent.

We now provide several examples in which this condition is satisfied. The first three pertain to a partially observed action.

Example 4. *First Offer in Bargaining:* 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 will generally depend on i ’s value V_i in equilibrium, but it cannot depend on other agents’ values V_{-i} . Since actions A_i are functions of values V_i in equilibrium, this first offer is a Z_i satisfying the conditions of Definition 1.

Example 5. *Knockout Bid in Bidding Ring Auction Followed by Target Auction:* In the stamp auction case of Asker (2010), members of a bidding ring participate in a sealed-bid knockout auction. The winning bid from the knockout auction then becomes the cartel’s bid in the target (English) auction. Asker (2010) provides a full model of this game. The knockout bid of an agent, which comes from a first, sealed-bid stage of the game, would be a Z_i satisfying Definition 1.

Example 6. *Indicative Bid:* In many procurement settings with auctions or bargaining, interested bidders are required to first submit an *indicative bid* (Ye 2007; Quint and Hendricks 2018). This indicative bid only represents a portion of an agent’s full action vector to be played in the full game, but can satisfy Definition 1 as a shifter Z_i for that unobserved action vector.

The arguments provided in this section can allow the researcher to identify objects of interest using the above action shifters Z_i without a full model of the (potentially

complex) game agents play. The next two examples pertain to other observables in the data that are arguably correlated with one agent’s action but not with other agents’ actions.

Example 7. *Agent Characteristics in Bargaining:* In the setting of Ambrus, Chaney, and Salitskiy (2018), Spanish ransom parties haggle with North African pirates to buy back Spanish captives. In the data, the authors observe the amount of earmarked money raised by a given captive’s family back home. This amount is known to the econometrician and to the buyer (the rescue party) but is unobserved to the seller (the pirates). This earmarked money is arguably correlated with the buyer’s valuation and thus with his action and can serve as a shifter Z_i for the rescue party’s action.⁴

Example 8. *Assets in Auctions:* In many auction datasets, such as auto auctions or real estate auctions, the econometrician may observe information about an agent’s assets that is only known to that agent and not to opponents, and that may be correlated with that agent’s willingness to pay. For example, in auto auctions, the data may contain information on previous cars purchased by that agent or credit constraints of that agent. In real estate auctions, the data may contain information on the size of the loan the potential buyer has been approved for, known to the buyer but not the seller. In each case, the level of these assets may serve as a Z_i satisfying Definition 1.

We now prove several properties of our menu approach in imperfectly observed action cases. These properties can be used to identify objects of interest without a complete model of the full underlying (potentially multi-stage) game.

4.1 Bounds on Utility with Imperfectly Observed Actions

Suppose that Z_i satisfies our Definition 1. Let $H(a_i, z_i)$ be the joint distribution of actions and action shifters, which is unobserved by the econometrician, and let $\check{P}_i(z_i)$ and $\check{T}_i(z_i)$ represent conditional expectations of x_i, t_i with respect to z_i , given by

⁴Ambrus, Chaney, and Salitskiy (2018) estimated buyer’s private valuations by instead imposing the structure of a one-sided-offers bargaining game from Fudenberg, Levine, and Tirole (1985). Through our approach, such structure could potentially be tested, as it would allow the researcher to identify certain objects without imposing a particular bargaining model.

$$\begin{aligned}
(\check{P}_i(z_i), \check{T}_i(z_i)) &\equiv E[(x_i(A_i, A_{-i}), t_i(A_i, A_{-i})) \mid Z_i = z_i] \\
&= E[E[(x_i(A_i, A_{-i}), t_i(A_i, A_{-i})) \mid A_i, Z_i = z_i] \mid Z_i = z_i] \\
&= E[E[(x_i(A_i, A_{-i}), t_i(A_i, A_{-i})) \mid A_i] \mid Z_i = z_i] \\
&= E[(P_i(A_i), T_i(A_i)) \mid Z_i = z_i].
\end{aligned}$$

The equality on the second line follows by the law of iterated expectations and the equality on third line follows by the fact that Z_i is independent of A_{-i} , so the inner conditioning on Z_i can be ignored. The final equality follows by the definition of P_i and T_i , the expected outcomes defining the menu. The above argument shows that conditional expectations of x_i, t_i with respect to z_i recover convex combinations of $(P_i(a_i), T_i(a_i))$ pairs. Since the graph of $\{(P_i(a_i), T_i(a_i))\}$ is convex, any such convex combinations lie above the graph of $\{(P_i(a_i), T_i(a_i))\}$. In other words, for any z_i that satisfies Definition 1, the graph of $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$ lies strictly above the graph of $\{(P_i(a_i), T_i(a_i))\}$. This allows us to obtain a lower bound on the utility of any given type v_i in equilibrium, or, equivalently, an upper bound on the graph $\{(P_i(a_i), T_i(a_i))\}$:

Corollary 3. *For any Z_i satisfying Definition 1, $\max_{z_i} v_i \check{P}_i(z_i) + \check{T}_i(z_i)$ is a lower bound on the equilibrium utility of type v_i .*

While we can thus derive an upper bound for the menu, in general, neither the true menu $\{(P_i(a_i), T_i(a_i))\}$ nor the joint distribution of action shifters and values $G(z, v)$ are identified. In Figure 2, we show the imperfectly observed menu $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$ generated from a true menu $\{(P_i(a_i), T_i(a_i))\}$, where we generate action shifters such that $V_i = Z_i + \epsilon_i$, where $\epsilon_i > 0$ is a uniformly distributed error term. The imperfectly observed menu is convex; thus, if we were to only observe $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$, we would be unable to distinguish the true data generating process from an alternative process in which $Z_i = V_i$, and the menu $\{(P_i(a_i), T_i(a_i))\}$ were exactly the imperfectly observed menu $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$.

4.2 One-to-one Shifters

In some cases, the econometrician may be willing to assume that Z_i is a one-to-one function of type V_i , that is, $V_i = \mathbf{z}_i^{-1}(Z_i)$.⁵ For example, suppose the game \mathcal{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 $\mathbf{z}_i(\cdot)$ is fully identified from the data.

If \mathbf{z}_i is a one-to-one function of the type v_i , then $(\check{P}_i(z_i), \check{T}_i(z_i))$ is exactly the expected outcome attained by the unique type $\mathbf{z}_i^{-1}(z_i)$. Moreover, for any other z'_i , the expected outcome $(\check{P}_i(z'_i), \check{T}_i(z'_i))$ is attainable by type $\mathbf{z}_i^{-1}(z'_i)$. Also, there are types $v_i + \delta, v_i - \delta$ playing different actions $\mathbf{z}_i(v_i + \delta), \mathbf{z}_i(v_i - \delta)$. As in Section 3, this implies the following bounds for any δ :

$$v_i \leq \frac{\check{T}_i(\mathbf{z}_i(v_i + \delta)) - \check{T}_i(\mathbf{z}_i(v_i))}{\check{P}_i(\mathbf{z}_i(v_i + \delta)) - \check{P}_i(\mathbf{z}_i(v_i))} \leq v_i + \delta, \quad (7)$$

$$v_i - \delta \leq \frac{\check{T}_i(\mathbf{z}_i(v_i)) - \check{T}_i(\mathbf{z}_i(v_i - \delta))}{\check{P}_i(\mathbf{z}_i(v_i)) - \check{P}_i(\mathbf{z}_i(v_i - \delta))} \leq v_i. \quad (8)$$

One-to-oneness implies that these bounds will collapse to a single point, and the entire mapping $\mathbf{z}_i(\cdot)$ is identified, which we state as the following extension of Corollary 1:

Corollary 4. *If, in game \mathcal{G} , each type v_i is one-to-one with z_i , the inverse mapping $v_i = \mathbf{z}_i^{-1}(z_i)$ from action shifters to types is pointwise identified.*

Proof. Follows immediately from (7) and (8). □

Corollary 2 also immediately extends to the imperfectly observed actions case, as follows. This corresponds to Corollary 2 of Larsen and Zhang (2021):

Corollary 5. *If each type v_i is one-to-one with z_i , $z_i \in \mathbb{R}$, the functions $\check{P}_i, \mathbf{z}_i(\cdot)$ are increasing and continuous, and the functions \check{T}_i, \check{P}_i are differentiable, then*

$$v_i = \frac{\check{T}'_i(z_i)}{\check{P}'_i(z_i)}.$$

⁵Once again, we will treat different values of z as identical if they induce the same expected outcomes $(\check{P}_i(z), \check{T}_i(z))$.

As with Corollary 2 in the fully observed action case, we see Corollary 5 as particularly useful for empirical work in the imperfectly observed action case. Section 5 will discuss estimation for these cases.

4.3 Approximating $E[V | z]$

For general noisy action shifters, the joint distribution of action shifters and values $G(z, v)$ is not identified. However, under certain conditions, it is possible to approximate the function $E[V | Z = z]$ (which we will write in short as $E[V | z]$), the conditional expectation of an agent's valuation given the observed Z . This object may be of direct interest as a best guess of an agent's valuation given the data, and it will also be useful in allowing us to construct an approximate lower bound for gains from trade for the agent in the game.

Throughout this and the next subsection, we suppress subscripts i , writing for example z to mean z_i . We will also write $P(v)$ and $T(v)$ as shorthand for $P(\mathbf{s}(v))$ and $T(\mathbf{s}(v))$, rather than writing $P(a)$ and $T(a)$, because we are primarily interested here in the relationship between the action shifter Z and values V . We also assume Z is real-valued and supported on a continuous bounded interval; $P(\cdot), T(\cdot)$ are everywhere strictly increasing and differentiable; $G(v | z)$ is differentiable in v for any z ; and $G(v | z)$ is continuous in z for almost all v . These assumptions make our results easier to state and prove, but they are stronger than we need; all of these assumptions can be relaxed without significantly affecting the results.

We impose one final restriction that cannot easily be relaxed: we assume that $G(v | z)$ is ordered by stochastic dominance, that is

$$z_2 > z_1 \implies G(v | z_2) >_{FOSD} G(v | z_1).$$

Under these assumptions, we aim to find conditions under which we can approximate the conditional expectation of an agent's value given the agent's observed Z , that is, $\frac{\check{T}'(z)}{\check{P}'(z)} \approx E[V | z]$. To begin, Proposition 2 characterizes the ratio $\frac{\check{T}'(z)}{\check{P}'(z)}$ in terms of the conditional distributions $G(v | z)$ and the menu $\{(\check{P}(\cdot), \check{T}(\cdot))\}$:

Proposition 2. *Under the conditions stated above for Z , $\check{P}(\cdot)$, $\check{T}(\cdot)$, and $G(v | z)$,*

$$\frac{\check{T}'(z)}{\check{P}'(z)} = \frac{\int_0^{\bar{v}} v \left(-\frac{dG(v|z)}{dz}\right) P'(v) dv}{\int_0^{\bar{v}} \left(-\frac{dG(v|z)}{dz}\right) P'(v) dv}. \quad (9)$$

Since both $\left(-\frac{dG(v|z)}{dz}\right)$ and $P'(v)$ are positive and finite by assumption, their product $\left(-\frac{dG(v|z)}{dz}\right) P'(v)$ can be thought of as a weighting function; hence the expression on the right hand side of (9) can be thought of as a weighted average of values v with respect to $\left(-\frac{dG(v|z)}{dz}\right) P'(v)$.

Proposition 2 implies that if the weighting function $\left(-\frac{dG(v|z)}{dz}\right) P'(v)$ behaves similarly to the probability density function $g(v | z)$ then the ratio $\frac{\check{T}'(z)}{\check{P}'(z)}$ will be close to the conditional expectation $E[V | z] = \int_0^{\bar{v}} v g(v | z) dv$. This behavior will only hold under fairly restrictive assumptions on the functions $G(v | z)$ and $P'(v)$; however, these assumptions are satisfied by certain classes of conditional distributions commonly assumed in practice. We describe two such sets of conditions below. The first involves conditional invariance of the joint distribution $G(v, z)$ and a linear expected allocation function:

Corollary 6. *Suppose that (i) $G(v, z)$ is translationally invariant, that is, $G(v, z) = \bar{G}(v - z)$; and (ii) $P'(v)$ is constant. Then,*

$$\frac{\check{T}'(z)}{\check{P}'(z)} = E[V | z].$$

While the result in Corollary 6 is stated as an equality, equation (9) is smooth in all of its arguments, so a similar statement about approximate equality holds: if $G(v, z)$ is close to translationally invariant locally at some z , and $P'(\cdot)$ is close to constant in the relevant neighborhood, $\frac{\check{T}'(z)}{\check{P}'(z)}$ will be close to $E[V | z]$.

Our second case where $E[V | z]$ can be approximated well involves the noise in Z being bounded:

Corollary 7. *Suppose $G(v | z)$ has compact support: $G(v | z) \in [\nu(z) - \delta, \nu(z) + \delta]$*

for some $\delta > 0$, where $\nu(z)$ is continuous. Then,

$$\left| \frac{\check{T}'(z)}{\check{P}'(z)} - E(V | z) \right| \leq 2\delta.$$

If the support of $G(v | z)$ is not compact, but the majority of the probability mass is concentrated near $\nu(z)$, Corollary 7 will still hold approximately. Intuitively, this corollary says that if z is a fairly good predictor for v , in the sense that the residual variation in v conditional on observing z is small, the derivatives $\frac{\check{T}'(z)}{\check{P}'(z)}$ will be fairly close to $E(V | z)$.

4.4 Approximating Gains from Trade

Corollary 6 or Corollary 7 can be used to obtain a lower bound on the total expected gains from trade in the game played in the data. We can use the following quantity to obtain an approximate lower bound to the expected gains from trade for all agents. Total expected welfare from trade is:

$$E_V \left[\max_{v'} [V P(v') - T(v') - V\bar{x}] \right],$$

where the maximization is performed pointwise for each realization of V . Using the law of iterated expectations, we can write this as:

$$= E_Z \left[E_V \left[\max_{v'} [V P(v') - T(v') - V\bar{x}] | Z \right] \right].$$

Corollary 3 shows that the convex hull of the $(\check{P}(z), \check{T}(z))$ pairs is an upper bound for the true menu $\{(P(v), T(v))\}$. Hence, for all z ,

$$\begin{aligned} & E_V \left[\max_{v'} [V P(v') - T(v') - V\bar{x}] | Z = z \right] \\ & \geq E_V \left[\max_{z'} [V \check{P}(z') - \check{T}(z') - V\bar{x}] | Z = z \right] \end{aligned} \quad (10)$$

$$\geq \max_{z'} E_V [V | Z = z] \check{P}(z') - \check{T}(z') - E_V [V | Z = z] \bar{x}. \quad (11)$$

This last inequality follows because any value attainable in (11) is also attainable in (10). An intuition for this inequality is that, fixing a feasible expected outcome

$(\check{P}(z), \check{T}(z))$, the gains from trade are linear in values V . Hence, the average maximal utility across agent types is at least the maximal utility of the average agent type. Thus, one can use Corollary 6 or Corollary 7 to identify $E[V | Z = z]$, and then, taking expectations with respect to the empirical distribution of Z , one can obtain (11) as a lower bound on the total gains from trade.

The expression in (11) will correspond exactly to the gains from trade only if the action shifter Z is a one-to-one function of values. If the action shifter Z is not a one-to-one function of V , the inequalities (10) and (11) will hold strictly. Thus, even if we are unsure whether the approximation $E[V | z] \approx \frac{\check{T}'(z)}{\check{P}'(z)}$ is valid, as long as we believe that the assumptions of Corollaries 6 or 7 are even approximately satisfied, this is a generally conservative approach to identifying a bound on the expected gains from trade.

5 Estimation

We now present approaches for estimating valuations, building on the identification arguments presented in Sections 3 and 4. For simplicity of exposition, we describe our estimation results in the context of the fully observed action case, but the same estimation approaches apply to the imperfectly observed actions settings described in Section 4. We discuss estimation with discrete actions in Section 5.1 and with continuous actions in Section 5.2. Throughout this section, we will focus on estimation for a single agent (or a class of agents who are indistinguishable), thus we will omit subscripts i , writing for example $a, v, P(\cdot), T(\cdot)$ instead of $a_i, v_i, P_i(\cdot), T_i(\cdot)$. We will introduce here a subscript j to index an observation in the data, putting this subscript on each object observed in instance j of the game, namely x_j, t_j , and a_j .

5.1 Discrete Actions

Suppose that there are a finite number of possible actions, $\mathcal{A} = \{a^1 \dots a^K\}$, 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

agent's action $a_j^k \in \mathcal{A}$, 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 observations where the action a^k was chosen. We then set, as in Theorem 1,

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

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 $\{(\hat{P}(a^k), \hat{T}(a^k))\}$ 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 $\{(\hat{P}(a^k), \hat{T}(a^k))\}$ graph directly, we take its convex hull, and use the subgradients of the convex hull to estimate values.

For a given collection $\{(\hat{P}(a^k), \hat{T}(a^k))\}$, we define the supporting hyperplane $H(p; \nu)$ of slope ν , as the highest line of slope ν which lies below all points in $\{(\hat{P}(a^k), \hat{T}(a^k))\}$:

$$b(\nu) \equiv \max \left\{ b : \hat{T}(a^k) \geq b + \nu \hat{P}(a^k) \quad \forall a^k \right\}$$

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

We construct the convex hull of $\{(\hat{P}(a^k), \hat{T}(a^k))\}$ at any point p by taking the supremum over all supporting hyperplanes, which we denote $\bar{H}(p) \equiv \sup_{\nu} H(p; \nu)$. We then estimate $\hat{\mathbf{s}}^{-1}(a^k)$ using the set of subgradients of $\bar{H}(p)$ at point $\hat{P}(a^k)$; that is, the set of slopes ν such that $H(p; \nu)$ attains the supremum at point $\hat{P}(a^k)$:

$$\hat{\mathbf{s}}^{-1}(a^k) = \left\{ \nu : H(\hat{P}(a^k); \nu) = \bar{H}(\hat{P}(a^k)) \right\}.$$

$\bar{H}(p)$ 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 menu $\{(P(a^k), T(a^k))\}$ is convex, the inferred $\hat{\mathbf{s}}^{-1}(a^k)$ has the same limit as the first estimator. However, using the convex hull of

$\{(\hat{P}(a^k), \hat{T}(a^k))\}$ ensures that the estimator produces non-overlapping bounds for agents' values in finite samples.

5.2 Continuous Actions

In many cases of interest, the econometrician may be willing to assume that the equilibrium strategies are smooth functions of values, and the equilibrium $P(a)$ and $T(a)$ mappings 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 (or Corollary 5 in the imperfectly observable actions case) 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:

$$\left[\hat{\beta}_0(a), \dots, \hat{\beta}_p(a) \right] = \arg \min_{\beta} \sum_j \left[\left[t_j - \sum_{k=0}^p \beta_k (a_j - a)^k \right]^2 K_h(a_j - a) \right]. \quad (12)$$

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 k th derivative of T . Therefore, an estimate of the first derivative $\hat{T}'(a)$ is given by running 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 (12), $\hat{\beta}_1$. Similarly, an estimate of the first derivative $\hat{P}'(a)$ is given by running 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.⁶

As in Section 5.1, 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.⁷ In addition, it may be desirable in some contexts 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 $\{(P(a), T(a))\}$ menu. In Appendix B, we describe the construction of the quadratic and cubic spline bases in detail. Our estimation approach 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. Constraining the quadratic (cubic) spline coefficients to be nonnegative ensures that the target function is nondecreasing (convex). Note that, by construction, the quadratic splines have two continuous derivatives, and the cubic splines three continuous derivatives. 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.

⁶For the bandwidth, Fan and Gijbels, Chapter 4.2, describes the following rule-of-thumb selection procedure: First, one fits a global quintic polynomial by standard ordinary least squares, $\hat{T}(a) = \alpha_0 + \alpha_1 a + \dots + \alpha_5 a^5$. Let the residual variance from the regression be $\tilde{\sigma}^2$. The rule-of-thumb bandwidth is then equal to the following “variance components”-like formula:

$$\hat{h} = C_{\nu,p}(K) \left[\frac{\tilde{\sigma}^2}{\sum_{i=1}^n (\hat{T}''(a))^2} \right]^{\frac{1}{7}}, \quad (13)$$

where $C_{\nu,p}(K)$ is a kernel-specific constant. Fan and Gijbels, Table 3.2, reports this constant as approximately 0.884 for the Gaussian kernel, 2.275 for the Epanechnikov kernel, and 2.869 for the triweight kernel. This procedure chooses smaller bandwidths for functions that can be fitted better by polynomials.

⁷Menu convexity corresponds exactly to monotonicity of Myerson (1981) marginal revenue, hence “ironing” in this context is exactly equivalent to the Myersonian procedure of enforcing marginal revenue monotonicity for a demand curve.

6 Monte Carlo Simulations

6.1 First-Price Auction Simulations

In this section, we simulate data from a first-price auction and estimate values using the standard Guerre, Perrigne, and Vuong (2000) (GPV) approach as well as our approaches based on local polynomial regressions and convex splines. We use $m = 5$ bidders for each auction instance, using uniform, exponential and lognormal value distributions. We generate datasets of size 4,000, 16,000, and 40,000 for estimation; we solve for the unique equilibrium as in GPV, and then simulate and estimate from each combination of distribution and dataset size 200 times.

The GPV estimator can be written:

$$\hat{v}(a) = a + \frac{\hat{G}(a)}{\hat{g}(a)},$$

where v is the agent's value, a is the agent's bid, and $\hat{G}(\cdot), \hat{g}(\cdot)$ are estimates of the distribution and density functions of the *maximum opposing bid*, rather than the distribution of all bids.⁸ We estimate $\hat{g}(\cdot)$ using a kernel density estimator, using Gaussian kernels with bandwidth chosen using the Silverman (1986) rule of thumb. We estimate $\hat{G}(\cdot)$ using a linearly interpolated empirical CDF. To implement our local polynomial regression method, we run local quadratic regressions to estimate derivatives as described in Section 5.2, using Gaussian kernels with bandwidths equal to half⁹ the Fan-Gijbels rule of thumb in (13). To implement our convex spline method, we first estimate $P_i(\cdot)$ separately for each agent i using a local linear regression, with bandwidth chosen using a rule of thumb that Fan and Gijbels propose for local linear regression. Then, we estimate the function $T_i(P_i)$ for each agent i using convex splines with 10 knots uniformly separated in bid quantile space, minimizing the sum

⁸While we draw i.i.d. bids for simulation, we estimate values allowing for bidders to be potentially drawn from different distributions. The symmetric GPV estimator uses multiple data points from each auction, and could technically be implemented with only observing one auction (with many bidders) whereas our approach requires observing multiple instances of the game. We use the asymmetric version here to hold fixed the effective sample size between our procedure and GPV, comparing only the estimation procedures' performances.

⁹This is an ad-hoc adjustment, because we found that the estimator performed somewhat better with smaller bandwidths, possibly because this reduces bias in the estimator when the target function is relatively nonsmooth.

of squared errors from the (a, x, t) triples observed in data. We did not explicitly apply any procedure for spline knot selection, but we found that results were not very sensitive to using knot counts between 5 and 20.

We show pointwise 95% coverage bands from 200 estimation sample for both estimators, for the case of the exponential distribution with $N = 4,000$ in Figure 3. In Table 1, for all specifications we used, we show average mean squared errors normalized by the true variance in values, for bid values between the 20th and 80th bid quantiles.¹⁰

We find that GPV significantly outperforms both of our estimators: the mean squared error of GPV is less than 1 percent in all of the simulation specifications, and in some cases it is even below 0.1%. In absolute terms, however, both the local quadratic and spline estimators perform fairly well, achieving below 10% MSE in all specifications, and achieving around 1% MSE in specifications with 40,000 observations.

Our proposed spline estimator may have a number of qualitative benefits over the local quadratic regression estimator. First, it is guaranteed to produce convex menus for any given estimation instance, which may be desirable in settings where value and menu estimates are used for economic counterfactuals. Second, as we demonstrate in Section 7.2, the spline estimator allows the researcher to incorporate an unobserved heterogeneity correction procedure. If unobserved heterogeneity is not a serious issue, we suggest using the local quadratic regression estimator, as its finite-sample performance appears to exceed that of the spline estimator. In cases where unobserved heterogeneity or other multi-step procedures that benefit from parametric estimation are needed, or in cases where data is sparse and enforcing menu convexity is important for statistical reasons, our spline procedure may be preferable despite its worse finite-sample performance.

We interpret these first-price auction simulation results as suggesting that if the researcher is well aware that the game generating the data is indeed a first-price auction, the researcher should use the standard GPV approach for identification and

¹⁰While this is a large truncation, as Figure 3 shows, both GPV and our local quadratic regression estimator can function very poorly towards the lower end of the support of the bid distribution. For our estimator, this is essentially because the probability of winning and transfer both become very small; the ratio $\frac{T'(b)}{P'(b)}$ thus becomes numerically unstable. Thus MSE comparisons are more meaningful for bids well inside the interior of the support of the bid distribution.

estimation. If, however, the researcher has data by some unknown bidding procedure or by some unknown equilibrium selection rule, or from a bidding procedure for which no specific, standard tool has yet been derived, and wishes to identify and estimate valuations in a manner that is robust to misspecification of these features of the game, our menu approach can perform well.

6.2 Bargaining Simulations

To further illustrate our method, we choose a setting for which, unlike first-price auctions, there is no existing off-the-shelf method for estimation: a bargaining game with two-sided incomplete information. The theoretical literature on incomplete information bargaining demonstrates that such games in general have multiple equilibria that are difficult to characterize. Here, we generate data from a particular equilibrium from such a game. In order to do so, we adopt the bilateral bargaining game described in Satterthwaite and Williams (1989). Players' roles are denoted $i \in \{B, S\}$. The buyer B and seller S have private valuations and make offers. If the buyer's offer (a_B) exceeds that of the seller (a_S), trade occurs at price $t = ka_S + (1 - k)a_B$, where $k \in [0, 1]$. The parameter k can be considered a bargaining power weight. A game with $k = 1$ corresponds to the seller-optimal mechanism (a take-it-or-leave-it offer by the seller) and a game with $k = 0$ corresponds to the buyer-optimal mechanism (a take-it-or-leave-it offer by the buyer). This game is referred to in the literature as a k double auction (see also Chatterjee and Samuelson 1983). Satterthwaite and Williams (1989) demonstrate that a continuum of equilibria exist in this game, and they provide methods to solve for a class of equilibria. We simulate data following their solution approach, as described in Appendix C.

An observation in the data we generate consists of a final transfer t , an allocation x (i.e. an indicator for whether or not the good sold), and a seller's offer a_S . In applying our estimation approach, we then deliberately tie our hands by treating the buyer's offer and the bargaining power parameter as unknown and treating the seller's action as though arising from an unknown extensive form bargaining game generating the data. Specifically, we treat the seller's offer as a partially observed action vector that is one-to-one with the seller's value as described in Section 4. This simulated data mimicks a type of dataset that a researcher may commonly find:

prices arising from an unknown negotiation process (i.e. the precise extensive form of the game is unknown to the researcher). For example, a researcher might observe housing data with transacted prices, a measure of how many buyers placed offers on the house and failed to win it, and the seller’s initial list price (analogous to the seller offer generated in our data).

We apply our estimators to cases where valuations are distributed according to a uniform or truncated Normal distribution, and for sample sizes of 4,000, 16,000, or 40,000 observations. The distributional parameters for the true distributions are discussed in Appendix C. Figure 4 displays sellers’ offers on the horizontal axis and the estimated and true valuations of sellers on the vertical axis for the uniformly distributed case with 4,000 observations. The estimation procedure performs well in recovering valuations close to the truth. Table 2 displays the mean squared error for the different specifications, normalized by the true variance as in Table 1. The mean squared error of the local quadratic regression estimator is at most 10% in small samples and in larger samples it is below 5%. The convex spline estimator is less accurate in small samples, but also has MSE below 10% in larger samples.

We also remark here that this estimation exercise did not exploit any information about the value of k (the bargaining power), any offer made by the buyer, or the particular equilibrium being played. Recall that a continuum of equilibria exist to this game at any value of k . The generated data in our simulation exercise could have come from any fixed value of k and any fixed equilibrium, or even a sequential bargaining game completely distinct from the static k double auction, and the two-step approach proposed herein would still have returned reasonable estimates of valuations based solely on the observed first seller offer, the allocation, and the transfer.¹¹

7 Extensions: Heterogeneity and Correlated Values

We now provide extensions to our identification arguments provided above to settings of game-level heterogeneity (observed and unobserved heterogeneity) and correlated

¹¹Li and Liu (2015) propose an estimation approach for settings where the researcher instead knows that the data is generated by a k double auction.

values.

7.1 Observable Game-Level Heterogeneity

If realizations of the game differ in a way that is observable to the econometrician, such heterogeneity can easily be controlled for. In particular, the estimation procedures described in Section 5 can be extended to include a vector Y of observable characteristics (that may vary in each realization of the game) as a control in the estimation of the $P_i(\cdot)$ and $T_i(\cdot)$ functions determining the menu. If such heterogeneity enters players' valuations and actions in an additively or multiplicatively separable, single-index form, the econometrician can simply regress observations of actions on observations of the vector Y and treat the residuals from this regression as coming from a "homogenized" game, as in the auction setting of Haile, Hong, and Shum (2003). We discuss this separability condition in more detail below when we discuss unobserved heterogeneity.

Allowing for the equilibrium menu to vary with observable heterogeneity at the game level also allows for the data to represent play from different equilibria of the underlying game: as long as the equilibrium selection can be controlled for using observables, the presence of multiple equilibria poses no threat to identification or estimation. We discuss this further in Section 8.4 below.

7.2 Unobserved Heterogeneity

We focus here on correlation induced in values through game-level heterogeneity that is observed by the players but not by the econometrician, similar to the unobserved auction-level heterogeneity in Krasnokutskaya (2011). We study a class of games that we call *separable heterogeneity games*, which are defined as follows. In the first stage, common component W is drawn from the bounded distribution $F_W(\cdot)$ and commonly observed by all agents $i \in \{1 \dots m\}$ 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 $\tilde{a}_i \in [\underline{a}_i, \bar{a}_i]$, 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 *separable heterogeneity game property* if for all $\tilde{a}_1 \dots \tilde{a}_m, \Delta$, and for all i

$$\mathbf{x}_i(\tilde{a}_1 + \Delta, \tilde{a}_2 + \Delta, \dots, \tilde{a}_m + \Delta) = \mathbf{x}_i(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_m), \quad (14)$$

$$\mathbf{t}_i(\tilde{a}_1 + \Delta, \tilde{a}_2 + \Delta, \dots, \tilde{a}_m + \Delta) = \mathbf{t}_i(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_m) + \Delta(x_i(\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_m) - \bar{x}_i). \quad (15)$$

The separable heterogeneity game property can also be defined multiplicatively. We focus on the additive case here for brevity. This property implies that, if all agents increase bids by a constant amount Δ , the probability of any agent winning is unchanged, and the prices paid by agents change by the same amount by which bids change. Intuitively, this property is appropriate for games which are position-invariant, in the sense that the structure of the game is unchanged when all bids are changed by a given amount. The separable heterogeneity game property may not be appropriate, for example, in bargaining settings with fixed costs, in which case the structure of the game may be qualitatively different depending on the size of the payoff agents are bargaining over.¹²

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 separable heterogeneity 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 a Bayes-Nash equilibrium conditional on any value of w . Fixing a given value of w , the game is identical to that of Section 3. Let A_i^w be a random variable representing i 's equilibrium action when the common component is w . As in Section 3, we define the expected probability and transfer that i achieves when playing a_i in equilibrium as:

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

¹²In such settings, one possible approach is to divide instances of the game into coarse partitions by observable characteristics that may affect payoff size. The separable heterogeneity game property may be a reasonable assumption within each partition, and the approach described in this section could be applied within partitions.

Agent i 's expected utility when her type is v_i and she plays a_i when the common component is w is

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

Analogously to Section 3, equilibrium in a separable heterogeneity game with common component w requires that i 's strategy $\mathbf{s}_i(v_i, w)$ maximize 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

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

In the following proposition, we show that the equilibria of separable heterogeneity games have a “position invariance” property with respect to w —if actions $a_1 \dots a_m$ 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 strategies $\mathbf{s}_1(v_1, w) \dots \mathbf{s}_m(v_m, w)$ constitute an equilibrium of a separable heterogeneity game. Then, for any common component w' , the following strategies constitute an equilibrium:*

$$\begin{aligned} \mathbf{s}_1(v_1, w') &= \mathbf{s}_1(v_1, w) + w' - w \\ &\vdots \\ \mathbf{s}_m(v_m, w') &= \mathbf{s}_m(v_m, w) + w' - w. \end{aligned}$$

Motivated by this proposition, we will define *position-invariant equilibria* by requiring that agents play the same equilibrium for any common component w :

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

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

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

In the equilibrium conditional on $w = 0$, we have, as before in Section 3,

$$\mathbf{v}_i(a_i) \geq \frac{T_i^0(a_i) - T_i^0(a'_i)}{P_i(a_i) - P_i(a'_i)} \quad \forall a'_i : P_i(a_i) > P_i(a'_i),$$

$$\mathbf{v}_i(a_i) \leq \frac{T_i^0(a'_i) - T_i^0(a_i)}{P_i(a'_i) - P_i(a_i)} \forall a'_i : P_i(a_i) < P_i(a'_i).$$

Thus, if we can recover the function $P_i(a_i)$ and $T_i^0(a_i)$, we can bound values or point identify values as in Section 3. Unlike the setting with no unobserved heterogeneity, these objects are not immediately identified from conditional expectations in the data. Rather, the objects that can easily be identified in the data are *noisy* versions of $P_i(a_i)$ and $T_i^0(a_i)$ that include—but cannot condition on—the unobserved heterogeneity term in each realization of the game. We will denote these objects

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

Assuming the econometrician observes multiple independent observations of same underlying separable heterogeneity game in which a position-invariant equilibrium is played, the true menu is identified.

Proposition 4. $P_i(\cdot), T_i^0(\cdot)$ are uniquely identified by $\tilde{P}_i(\tilde{a}_i), \tilde{T}_i(\tilde{a}_i)$ and actions $\{\tilde{a}_i\}$ from a separable heterogeneity game in which the equilibrium is position invariant.

The intuition behind our identification result is as follows. Since the unobserved private value components v_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 that would have been played at a realization of $w = 0$, which we will refer to as F_{A_i} . These distributions are identified from the observed distribution of actions that includes the realized unobserved heterogeneity, $F_{\tilde{A}_i}$, through a deconvolution argument. We then show that the distributions F_W, F_{A_i} allow us to recover $P_i(a_i), T_i^0(a_i)$ from the functions $\tilde{P}_i(\tilde{a}_i), \tilde{T}_i(\tilde{a}_i)$ through a series of deconvolutions and convolutions against $F_W(\cdot)$. In particular, if agent i is the agent whose menu and valuation we wish to identify (for example, suppose i is a seller in a bilateral bargaining game and the other agent is the buyer), then $P_i(a_i)$, the true underlying menu, solves:

$$\tilde{P}_i(\tilde{a}_i) = \frac{\int P_i(a_i) f_{a_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}{\int f_{a_i}(a_i) f_W(\tilde{a}_i - a_i) da_i},$$

and $T_i^0(a_i)$, the homogenized or true underlying transfer function, solves

$$\tilde{T}(\tilde{a}_i) - E(W\Delta P_i | \tilde{a}_i) = \frac{\int T_i^0(a_i) f_{a_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}{\int f_{a_i}(a_i) f_W(\tilde{a}_i - a_i) da_i},$$

where the term $E(W\Delta P_i | \tilde{a}_i)$ is given by

$$E(W\Delta P_i | \tilde{a}_i) = \int (P_i(a_i) - 1) (\tilde{a}_i - a_i) f_{a_i}(a_i) f_W(\tilde{a}_i - a_i) da_i.$$

This identification argument for $P_i(a_i)$ and $T_i^0(a_i)$ immediately leads to an estimation procedure. To describe the estimation, we now re-introduce subscripts j to denote observations in the data. Assume the econometrician observes the allocation x_{ij} , transfer \tilde{t}_{ij} , and action \tilde{a}_{ij} for agent i in $j = 1, \dots, J$ games. The objects \tilde{t}_{ij} and \tilde{a}_{ij} include the unobserved heterogeneity component w_j . In order to perform the initial Krasnokutskaya (2011) deconvolution, the econometrician must also observe the action \tilde{a}_{ij} of at least one other agent. For simplicity here we will assume the econometrician observes all agents' actions, the set $\{\tilde{a}_i\}$, but only one is required. The first step is to estimate f_W and each f_{A_i} (the densities corresponding to F_W and F_{A_i}) through parametric or nonparametric maximum likelihood using observed actions $\{\tilde{a}_i\}$. The log-likelihood function is given by

$$\sum_j \int \left(\prod_i f_{A_i}(\tilde{a}_{ij} - w; \theta_a) \right) f_W(w; \theta_W) dw \quad (16)$$

The objects θ_a and θ_W represent vectors to be estimated, parameterizing the densities f_W and each f_{A_i} to belong to some parametric family or sieve space. The second step is to estimate $\tilde{P}(\cdot)$ and $\tilde{T}(\cdot)$ which can be done through local linear regression just as in the no-unobserved-heterogeneity case in Section 5.2.

Given the estimates of $f_{A_i}(\cdot)$, $f_W(\cdot)$, $\tilde{P}(\cdot)$, $\tilde{T}(\cdot)$, we can solve for $P_i(\cdot)$, $T_i^0(\cdot)$ using minimum weighted distance, using our empirical ironing, shape-constrained spline basis functions suggested in Section 5.2 and described in detail in Appendix B. We propose quadratic spline coefficients θ_P , flexibly parameterizing $P_i(\cdot; \theta_P)$, to

solve the following objective function:

$$\min_{\theta_P} \int \left[\int P_i(a_i; \theta_P) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i - \tilde{P}_i(\tilde{a}_i) \right]^2 \left[\int f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i \right] d\tilde{a}_i. \quad (17)$$

This minimizes the integrated distance between $\int P_i(a_i; \theta_P) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i$ and $\tilde{P}_i(\tilde{a}_i)$, weighting by $\int f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i$, which is the density function of \tilde{a}_i .¹³

We can then estimate $T_i^0(\cdot)$ as a function of $P(\cdot)$, parameterizing it as $T_i^0(P_i(a_i; \theta_P); \theta_T)$, where θ_T are cubic spline coefficients chosen to solve

$$\min_{\theta_T} \int \left[\int T_i^0(P_i(a_i; \theta_P); \theta_T) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i - \left(\tilde{T}_i(\tilde{a}_i) - E(W \Delta P_i | \tilde{a}_i) \right) \right]^2 \left[\int f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i \right] d\tilde{a}_i. \quad (18)$$

The minimum distance problems in (17) and (18) can be computed with standard gradient descent methods. Once $P_i(\cdot), T_i^0(\cdot)$ are estimated, values can then be estimated using the relationship $\mathbf{s}^{-1}(a) = \frac{dT}{dP}$ as in Section 5.2.

7.3 Identification with Non-independent Private Values

In this section, we relax the assumption that values of different agents are independent. Suppose that agents' values $V_1 \dots V_m$ are drawn from some joint distribution $F(v_1 \dots v_m)$, 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 \mathcal{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 Section 3, we can derive bounds on the inverse functions $\mathbf{v}_i(\cdot)$ for each a_i . Our re-

¹³Weighting by the density of \tilde{a}_i is ad-hoc and does not matter asymptotically, but in finite samples it has the effect of lowering the estimation weight on regions where \tilde{a}_i has low density, so $\tilde{P}_i(\tilde{a}_i)$ may not be estimated accurately.

sults in this section are complementary to a contemporaneous paper by Kline (2026), which derives stronger identification results than ours under an additional assumption about equilibrium monotonicity.

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 $G(a_1 \dots a_n)$, derived from $F(v_1 \dots v_n)$ 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 $F(v_{-i} | v_i)$. This conditional distribution of values, combined with the equilibrium strategies of other players \mathbf{s}_{-i} , induces a conditional distribution over opponents' actions $G(a_{-i} | v_i)$. Thus, in equilibrium, if agent i is of type v_i and plays action a'_i , she attains the expected outcome $(P_i^{v_i}(a'_i), T_i^{v_i}(a'_i))$, defined as the expectation over the expected outcomes $\mathbf{x}_i(a'_i, A_{-i}), \mathbf{t}_i(a'_i, A_{-i})$ when $A_{-i} \sim G(a_{-i} | v_i)$. That is,

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

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

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

$$\mathbf{s}_i(v_i) = \arg \max_{a_i} v_i P_i^{v_i}(a_i) - T_i^{v_i}(a_i) - v_i \bar{x}_i. \quad (19)$$

As in Section 3, this allows us to bound $\mathbf{v}_i(a_i)$, the unique type that plays a_i in equilibrium.

Proposition 5. *For each a_i , the unique $v = \mathbf{v}_i(a_i)$ satisfies:*

$$v \geq \frac{T_i^{v_i=\mathbf{v}_i(a_i)}(a_i) - T_i^{v_i=\mathbf{v}_i(a_i)}(a'_i)}{P_i^{v_i=\mathbf{v}_i(a_i)}(a'_i) - P_i^{v_i=\mathbf{v}_i(a_i)}(a_i)} \vee \left\{ a'_i : P_i^{v_i=\mathbf{v}_i(a_i)}(a'_i) < P_i^{v_i=\mathbf{v}_i(a_i)}(a_i) \right\}, \quad (20)$$

$$v \leq \frac{T_i^{v_i=\mathbf{v}_i(a_i)}(a'_i) - T_i^{v_i=\mathbf{v}_i(a_i)}(a_i)}{P_i^{v_i=\mathbf{v}_i(a_i)}(a'_i) - P_i^{v_i=\mathbf{v}_i(a_i)}(a_i)} \vee \left\{ a'_i : P_i^{v_i=\mathbf{v}_i(a_i)}(a'_i) > P_i^{v_i=\mathbf{v}_i(a_i)}(a_i) \right\}. \quad (21)$$

If the distribution $F(v_1 \dots v_n)$ has full support, these bounds collapse to a single point.

Proof. Follows from (19). □

If the distribution $F(v_1 \dots v_n)$ has full support on the rectangle $[\min v_1, \max v_1] \times$

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

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

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

Note that $G(a_{-i} \mid v_i = \mathbf{v}_i(a_i))$, the equilibrium action distribution conditional on $v = \mathbf{v}_i(a_i)$, involves the unknown quantity $\mathbf{v}_i(a_i)$. However, this is equivalent to the conditional distribution $G(a_{-i} \mid a_i)$, which can be derived from $G(a_1 \dots a_n)$. Thus, for any a_i , the econometrician can estimate the functions:

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

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

These functions, plugged into the equations in Proposition 5, allow us to identify the unique $\mathbf{v}_i(a_i)$.

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. 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 Section 3. While our identification argument in this section is quite general, estimation in this general correlated values setting would likely be cumbersome given the number of different conditional objects to be estimated.

Our approach in this section requires that the equilibrium strategy $\mathbf{s}_i(v_i)$ is strictly separating. This assumption is necessary in order to estimate the distribution $G(a_{-i} \mid v_i)$ for the unique $v_i = \mathbf{v}_i(a_i)$ using the observed $G(a_{-i} \mid a_i)$. If \mathbf{s}_i were not invertible, in general $\mathbf{v}_i(a_i)$ would be a set of v_i values; thus, the observed $G(a_{-i} \mid a_i)$ would be a mixture over distributions $G(a_{-i} \mid v_i)$ for different values $v_i \in \mathbf{v}_i(a_i)$

and the researcher would not be able to use $G(a_{-i} | a_i)$ to consistently estimate $P_i^{v_i}(a'_i), T_i^{v_i}(a'_i)$ for any given type v_i .

8 Discussion

8.1 Counterfactuals

As highlighted in the introduction, a number of counterfactual exercises of interest for studying real-world mechanisms involve comparing a mechanism used in practice to some theoretical benchmark that achieves some notion of efficiency, optimality, equity, no-collusion, efficient collusion, etc. Or the counterfactual may involve a change to an alternative mechanism completely, such as a shift from a median price auction to a first-price auction. Counterfactuals such as these can be performed using the objects our approach delivers. Certain types of counterfactuals, however, would not be feasible using our approach. This includes any counterfactual in which the researcher needs to be able to simulate from the game played in the data and solve for an equilibrium. For example, in our application to used-car auctions, it would be infeasible to use our approach to study the question of what would happen if credit constraints were introduced or if players were required to pay a cost for each bargaining offer made. Such counterfactuals would require being able to solve for equilibria of an alternating-offer bargaining game with two-sided incomplete information (which, as explained above, is beyond the state of the current theory literature). Our approach explicitly avoids modeling the game actually played in the data, achieving identification of primitives without solving for an equilibrium. It is an *incomplete model* approach. We see this as a plus for the approach, making it capable of applying to settings where the rules of the game are either unknown or the game's equilibria are complex, but the approach does indeed rule out studying some counterfactuals that the additional structure of a complete model could deliver.

8.2 Beliefs and Rational Expectations

Throughout this paper, we have assumed Bayes-Nash equilibrium, which requires that agents have rational expectations about the distributions of opposing actions,

and thus know the functions $P(a), T(a)$ for any action a that they could play. This assumption is important for our estimation procedure; if agents' beliefs coincide with the true conditional expectations $P(a), T(a)$, we can simply estimate agents' beliefs using expected outcome functions in data. Our approach is in fact robust to some deviations from Bayes-Nash equilibrium: since only agent i 's best response condition is used to estimate her valuation, inference about agent i 's value only requires that agent i has correct beliefs about opponents' actions, not that all agents are rational.¹⁴

Our approach can fail for agents who do not have correct beliefs about the distribution of opponents' actions. These agents can be thought of as optimizing with respect to some $\hat{P}(a), \hat{T}(a)$ functions which may differ from empirical averages, and the econometrician must then take an essentially unverifiable stance on how agents' $\hat{P}(a), \hat{T}(a)$ functions are formed. A number of approaches have been proposed in the literature. Agarwal and Somaini (2018) estimate values assuming rational expectations as well as coarse and lagged expectations. Syrgkanis, Tamer, and Ziani (2017) explore inference in auctions under different information assumptions. In a related literature, Doraszelski, Lewis, and Pakes (2018) show that firms in a new market for frequency response behave in a way that is more consistent with adaptive learning than rational-expectations equilibrium.

Belief formation in multi-agent games is a difficult and important question in the literature; further discussion is beyond the scope of this paper, but we emphasize that the assumption of rational expectations is important for our methods to work. This may be a reasonable assumption for marketplace-like settings, in which traders interact frequently in environments that are relatively stationary over time, but it may be less appropriate for large one-shot mechanisms, or nonstationary settings in which agents are learning and converging towards optimal behavior over time.

¹⁴A similar argument is made in Nekipelov, Syrgkanis, and Tardos (2015), who argue that many learning algorithms should eventually converge to regret-free behavior, which motivates using no-regret conditions—essentially best-response conditions for stationary averages over opponents' actions—in estimation.

8.3 Belief Estimation versus Undominated-Strategy Approaches to Identification

This paper, along with GPV and many of the papers discussed in the subsection above, assumes that agents are optimizing with respect to some beliefs about outcomes, and moreover that the econometrician can estimate these beliefs. If these beliefs can be estimated, revealed preference arguments can be applied to identify agents' values. We view these papers as conceptually distinct from a set of identification and estimation approaches that are based instead on the concept of *rationalizability*, or *undominated strategies*. Examples include Haile and Tamer (2003) (or recent work by Chesher and Rosen (2017)) for estimating bounds on buyer valuations in ascending auctions or Larsen (2021) for estimating bounds on seller valuations in bargaining games. Rather than taking a stance on what agents' beliefs about outcomes are, these approaches rule out only values under which actions observed in the data are dominated—that is, values for which particular actions could not to be rationalizable by *any* belief. In the classic example of Haile and Tamer (2003), an agent bidding some amount y at some point in an ascending auction must have value above y ; if not, there is no possible outcome that the agent could achieve from bidding y that would give the agent positive utility.

The belief estimation approach requires strictly stronger assumptions about agents' beliefs than the rationalizability-based approach; if the econometrician is willing to make these assumptions, the benefit can be powerful identification results with weaker assumptions about game structure. In contrast to rationalizability-based approaches such as Haile and Tamer (2003), the approach we take here allows the econometrician to use essentially no information about the rules of the game; all relevant information about game rules is inferred from observing game outcomes. Moreover, rationalizability-based approaches tend to only allow partial identification of values, whereas our approach can often achieve point identification.

We believe that these two approaches can be used in a complementary manner. In some settings, for example, dominance-based arguments could be used to robustly find bounds on agents' values, and the stronger assumption of Bayes-Nash equilibrium could be used more tentatively to point-identify values.

8.4 Multiple Equilibria

We assume that agents are playing a single equilibrium across all observations. In some marketplace-like settings, agents interact with anonymous trading partners and few observables are available for agents to coordinate on; the assumption that all agents are playing the same equilibrium may be plausible in these settings. In other settings, agents may be able to coordinate to play different equilibria based on various features of the environment that they observe. For example, in timber auctions, agents bidding in auctions for high quality timber could play a certain equilibrium, bidding in a systematically different way than agents bidding in auctions for low quality timber. In such settings, if the econometrician observes all possible characteristics that agents could use for coordination, the problem of multiple equilibria could be addressed by estimating equilibrium menus conditional on observables, rather than pooling all observations into a single menu. This effectively allows a different equilibrium to be played for each possible combination of observables characterizing a trading game. There may be setting in which agents play different equilibria conditional on variables unobserved by the econometrician; this presents an interesting unobservable heterogeneity problem for which we do not yet have a solution.

Examples of features on which equilibrium selection can depend would be a reserve price in an auction or the number of competing bidders in an auction. If reserve prices are present and vary from one instance of the game to the next, and if these reserve prices are observable to agents, the econometrician would also need to observe and condition on reserve prices in order to correctly estimate equilibrium menus for buyers. Similarly, if agents can observe the number of their competitors and this number affects equilibrium behavior (it would do so in a first-price auction, for example, but not in an ascending button auction), then econometrician should observe and condition on this quantity when estimating equilibrium menus. If, however, agents themselves do not observe the number of their competitors, the econometrician need not observe it either in applying the menu approach.

8.5 Affine-Utility Games

We focus in this paper on trading games, in which agents' utility functions depend on an indicator variable x_i for trade and real-valued transfers t_i . However, our arguments

only utilize the fact that agents’ expected utilities are affine functions of expected outcomes, up to some unknown vector of types, and similar arguments apply to other settings in which agents maximize objective functions which are affine in types and observable outcomes.

For example, in an oligopolistic competition setting, our menu estimation approach is related to the Baker and Bresnahan (1988) approach of using firm cost shifters to estimate the *residual demand curve* facing a given firm, and using this to make statements about firms’ price-setting power. Similar arguments apply to other settings in which firms or agents have preferences that are assumed to be linear in some quantity measure $q(p)$, rather than a probability $P(a)$ as in our setting.¹⁵ Applying our approach to such settings could allow for firms to potentially have incomplete information about competitors’ costs without assuming a particular model of conduct by those firms. This could be particularly useful given that Sweeting, Tao, and Yao (2024) point out that small amounts of asymmetric information in oligopoly pricing can have big effects on prices, and price setting in such cases may look similar to collusion if incomplete information is ignored.

Related arguments also apply to settings in which agents maximize affine functions of more than two outcomes. Agarwal and Somaini (2018) analyze school choice problems, in which applicants’ utilities for submitting a given report are affine functions of probabilities of admission to different schools, and uses this to derive identification and estimation arguments. Gentry, Komarova, and Schiraldi (2023) use similar arguments to analyze bidding in combinatorial auctions for highway procurement. While our nonparametric estimation is specialized to the case of two outcomes, menus are also convex in settings with multiple goods, and similar approaches to ours can be applied in such settings. For example, Agarwal and Somaini (2018) show that solving for agents’ values for any preference report requires satisfying a set of revealed-preference inequalities that define the subgradient set of a multidimensional convex function; this is a linear programming problem, so it is fairly straightforward even in high dimensions.

¹⁵For example, consider a monopolist pricing problem, where the monopolist chooses price p to maximize $pq(p) - cq(p)$, where c is the (potentially privately known) marginal cost and $q(p)$ is demand. Defining $T(p) = pq(p)$, we can write monopolist profit as $T(p) - cq(p)$ and obtain c as $\frac{dT}{dq} = c$, which can be re-written to obtain the standard monopolist pricing formula $c = p + qp'(q)$.

8.6 Mechanism Design

The class of quasilinear-utility trading games that we study here is exactly the canonical model of classical mechanism design, as pioneered by Myerson (1981), and many parallels can be drawn between our results and this literature. Similarly to the mechanism design literature, our incentive compatibility constraints in (1) rely only on the quasilinear-utility setup, and are independent of the particular game being played. Indeed, one interpretation of our results from the perspective of mechanism design is that equilibria of arbitrary trading games can alternatively be thought of as revelation mechanisms, in which the abstract actions that agents take are their type reports. The econometrician can estimate the mapping that the revelation mechanism from type reports to outcomes, and use this to estimate the value of an agent making any given type report. The constraint that equilibrium menus are convex is related to the implementability constraint from Myerson, stating that equilibrium allocations must be weakly increasing in types. The “empirical ironing” procedures we propose in Section 5 to enforce menu convexity are thus related to the concept of ironing demand curves to enforce monotonicity of optimal allocation rules in optimal mechanism design.

9 Conclusion

This paper provides a two-step identification argument, based on intuitions from the theoretical mechanism design literature, for general incomplete information trading games, generalizing a number of game-specific approaches in the literature. We propose an method to deal with imperfectly observed actions, derive general estimation procedures, and extend tools for unobserved heterogeneity and correlated private values from auctions setting to general trading games. We believe our results may be useful 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, or where the rules of the game are not necessarily known to the econometrician, such as incomplete-information sequential bargaining games.

References

- AGARWAL, N., P. Z. LI, AND P. J. SOMAINI (2023): “Identification using revealed preferences in linearly separable models,” NBER WP 31868.
- AGARWAL, N., AND P. SOMAINI (2018): “Demand Analysis using Strategic Reports: An application to a school choice mechanism,” *Econometrica*, 86(2), 391–444.
- AMBRUS, A., E. CHANEY, AND I. SALITSKIY (2018): “Pirates of the mediterranean: An empirical investigation of bargaining with asymmetric information,” *Quantitative Economics*, 9(1), 217–246.
- AN, Y., AND X. TANG (2019): “Specification and negotiation in incomplete contracts,” *The RAND Journal of Economics*, 50(1), 57–92.
- ASKER, J. (2010): “A study of the internal organization of a bidding cartel,” *American Economic Review*, 100(3), 724–762.
- ASPELUND, K. M., AND A. RUSSO (2025): “Additionality and asymmetric information in environmental markets: Evidence from conservation auctions,” Working paper, Harvard.
- ATHEY, S., D. COEY, AND J. LEVIN (2013): “Set-asides and subsidies in auctions,” *American Economic Journal: Microeconomics*, 5(1), 1–27.
- ATHEY, S., AND D. NEKIPELOV (2010): “A structural model of sponsored search advertising auctions,” in *Sixth Ad Auctions Workshop*, vol. 15.
- AUSUBEL, L., P. CRAMTON, AND R. DENECKERE (2002): “Bargaining with incomplete information,” *Handbook of Game Theory*, 3, 1897–1945.
- BAGWELL, K., R. W. STAIGER, AND A. YURUKOGLU (2020): “Multilateral trade bargaining: A first look at the GATT bargaining records,” *American Economic Journal: Applied Economics*, 12(3), 72–105.
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): “Estimating dynamic models of imperfect competition,” *Econometrica*, 75(5), 1331–1370.


- BAKER, J. B., AND T. F. BRESNAHAN (1988): “Estimating the residual demand curve facing a single firm,” *International Journal of Industrial Organization*, 6(3), 283 – 300.
- BEST, M. C., J. HJORT, AND D. SZAKONYI (2023): “Individuals and organizations as sources of state effectiveness,” *American Economic Review*, 113(8), 2121–2167.
- CHANG, W.-S., B. CHEN, AND T. C. SALMON (2014): “An investigation of the average bid mechanism for procurement auctions,” *Management Science*, 61(6), 1237–1254.
- CHATTERJEE, K., AND W. SAMUELSON (1983): “Bargaining under incomplete information,” *Operations Research*, 31(5), 835–851.
- CHESHER, A., AND A. ROSEN (2017): “Incomplete English auction models with heterogeneity,” CeMMAP working paper CWP27/17.
- CRAMTON, P., S. ELLERMEYER, AND B. KATZMAN (2015): “Designed to fail: The Medicare auction for durable medical equipment,” *Economic Inquiry*, 53(1), 469–485.
- CRAMTON, P. C., AND J. S. TRACY (1992): “Strikes and holdouts in wage bargaining: Theory and data,” *American Economic Review*, pp. 100–121.
- DECAROLIS, F. (2018): “Comparing public procurement auctions,” *International Economic Review*, 59(2), 391–419.
- DORASZELSKI, U., G. LEWIS, AND A. PAKES (2018): “Just starting out: Learning and equilibrium in a new market,” *American Economic Review*, 108(3), 565–615.
- EINAV, L., T. KUCHLER, J. LEVIN, AND N. SUNDARESAN (2015): “Assessing Sale Strategies in Online Markets Using Matched Listings,” *American Economic Journal: Microeconomics*, 7(2), 215–47.
- 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.
- FUDENBERG, D., D. LEVINE, AND J. TIROLE (1985): “Infinite-horizon models of bargaining with one-sided incomplete information,” in *Game Theoretic Models of Bargaining*, ed. by A. Roth, pp. 73–98. Cambridge: Cambridge University Press.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. Cambridge: MIT Press.
- GENTRY, M., T. KOMAROVA, AND P. SCHIRALDI (2023): “Preferences and performance in simultaneous first-price auctions: A structural analysis,” *The Review of Economic Studies*, 90(2), 852–878.
- GRENNAN, M., AND A. SWANSON (2020): “Transparency and negotiated prices: The value of information in hospital-supplier bargaining,” *Journal of Political Economy*, 128(4), 1234–1268.
- 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.
- HAILE, P., AND E. TAMER (2003): “Inference with an incomplete model of English auctions,” *Journal of Political Economy*, 111(1), 1–51.
- 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.
- KAWAI, K., AND J. NAKABAYASHI (2022): “Detecting large-scale collusion in procurement auctions,” *Journal of Political Economy*, 130(5), 1364–1411.
- KENISTON, D. (2011): “Bargaining and welfare: A dynamic structural analysis,” Working Paper, Yale.
- KLINE, B. (2026): “Identification of incomplete information allocation-transfer games in monotone equilibrium,” *Journal of Econometrics*, 253, 106104.

- KRASNOKUTSKAYA, E. (2011): “Identification and estimation of auction models with unobserved heterogeneity,” *Review of Economic Studies*, 78(1), 293–327.
- LARSEN, B., AND A. L. ZHANG (2021): “Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry,” *Available at SSRN 3990290*.
- LARSEN, B. J. (2021): “The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions,” *The Review of Economic Studies*, 88(2), 851–882.
- LI, H., AND N. LIU (2015): “Nonparametric identification and estimation of double auctions with bargaining,” Working paper, Penn State.
- 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. (1979): “Incentive compatibility and the bargaining problem,” *Econometrica*, pp. 61–73.
- 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.
- PERRY, M. (1986): “An example of price formation in bilateral situations: A bargaining model with incomplete information,” *Econometrica*, pp. 313–321.
- PRESCOTT, J., K. E. SPIER, AND A. YOON (2014): “Trial and settlement: A study of high-low agreements,” *Journal of Law and Economics*, 57(3), 699–746.
- QUINT, D., AND K. HENDRICKS (2018): “A theory of indicative bidding,” *American Economic Journal: Microeconomics*, 10(2), 118–151.
- 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.
- SILVEIRA, B. S. (2017): “Bargaining with asymmetric information: An empirical study of plea negotiations,” *Econometrica*, 85(2), 419–452.
- SILVERMAN, B. W. (1986): *Density Estimation for Statistics and Data Analysis*, vol. 26. CRC press.
- SWEETING, A., X. TAO, AND X. YAO (2024): “Dynamic oligopoly pricing with asymmetric information: Implications for horizontal mergers,” *American Economic Journal: Microeconomics*, 16(3), 345–373.
- SYRGKANIS, V., E. TAMER, AND J. ZIANI (2017): “Inference on auctions with weak assumptions on information,” *arXiv:1710.03830 [cs, econ, math, stat]*, arXiv: 1710.03830.
- VARIAN, H. R. (2009): “Online ad auctions,” *American Economic Review*, 99(2), 430–34.
- WALDINGER, D. (2021): “Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation,” *American Economic Review*, 111(8), 2660–2696.
- YE, L. (2007): “Indicative bidding and a theory of two-stage auctions,” *Games and Economic Behavior*, 58(1), 181–207.

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 .

Figure 2: Example Menu in Imperfectly Observed Action Case




Notes: True menu (P, T) compared to a menu generated in an imperfectly observed action case. The latter menu $(\tilde{P}_i, \tilde{T}_i)$ is based on an action-shifter Z that is a noisy measure of the true underlying valuation.

Figure 3: Estimated Buyer Valuations From First-Price Auction Simulation



Notes: Figure displays the mapping between values (vertical axis) and bids (horizontal axis) from first-price auction simulations. Lines correspond to estimated medians (solid lines) and pointwise 95% quantiles (dotted lines) for the true values (black), GPV (blue), our local quadratic regression method (LQ, red), and our spline method (Spl, green) for exponential value distributions with $N = 4,000$. Means and quantiles are constructed across 200 simulation replications.

Figure 4: Estimated Seller Valuations from Bargaining Simulation



Notes: Figure displays the mapping between values (vertical axis) and bids (horizontal axis)—i.e. seller offers—from bargaining game simulations. Lines correspond to estimated medians (solid lines) and pointwise 95% quantiles (dotted lines) for the true values (black), our local quadratic regression method (LQ, red), and our spline method (Spline, green) for uniform value distributions with $N = 4,000$. Means and quantiles are constructed across 200 simulation replications.

Table 1: Mean Squared Error for First-Price Auction Simulations

Distribution	N	GPV MSE	LQ MSE	Spline MSE
Uniform	4,000	0.316%	2.9%	6.05%
	16,000	0.0959%	1.2%	2.21%
	40,000	0.0446%	0.696%	1.39%
Exponential	4,000	0.148%	3.44%	4.51%
	16,000	0.0485%	1.47%	1.29%
	40,000	0.0211%	0.885%	0.729%
Lognormal	4,000	0.15%	5.37%	5.46%
	16,000	0.0434%	2.45%	1.62%
	40,000	0.022%	1.36%	0.92%

Notes: Mean squared error between bid quantiles 0.2 and 0.8 for GPV, our local quadratic regression method (LQ), and our convexity-constrained spline method (Spline), for uniform, exponential, and lognormal distributed values of various sample sizes. Mean squared errors are in terms of percent of the true variance in bidder values between bid quantiles 0.2 and 0.8 and are based on 200 simulation replications.

Table 2: Mean Squared Error for Bargaining Simulation Results

Distribution	N	LQ MSE	Spline MSE
Uniform	4,000	6.2%	14.2%
	16,000	2.82%	5.45%
	40,000	1.61%	2.89%
Normal	4,000	8.96%	24%
	16,000	4.28%	11.5%
	40,000	2.59%	7.41%

Notes: Mean squared error between true valuations and estimated valuations from bargaining simulations for uniform and Normal distributions of valuations of various sample sizes. LQ refers to the local quadratic estimator and Spline refers to the convexity-constrained spline estimator.

Appendix

A Proofs

A.1 Proof of Proposition 1

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

$$(P_i(a'_i) - P_i(a_i))v_i \leq T_i(a'_i) - T_i(a_i) \quad \forall a'_i. \quad (22)$$

Hence, any type $v_i \in \mathbf{v}_i(a_i)$ is a subgradient of $\{(P_i(a_i), T_i(a_i))\}$ at a_i . Each individual inequality in (22) 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 $\mathbf{s}_i(v_i)$, the union of $\mathbf{v}_i(a_i)$ for all actions a_i is the support of values v_i .

Part 3. Fix some a_i, a'_i and suppose that $P_i(a'_i) > P_i(a_i)$. For any $v_i \in \mathbf{v}_i(a_i)$, by (22), we must have:

$$T_i(a'_i) - T_i(a_i) \geq v_i(P_i(a'_i) - P_i(a_i)).$$

For any $v'_i \in \mathbf{v}_i(a'_i)$, we must have:

$$T_i(a'_i) - T_i(a_i) \leq v'_i(P_i(a'_i) - P_i(a_i)).$$

Since by assumption $P_i(a'_i) > P_i(a_i)$, we have $v'_i \geq v_i$. Moreover, if $v_i = v'_i$, we must have

$$v_i = \frac{T_i(a'_i) - T_i(a_i)}{P_i(a'_i) - P_i(a_i)}.$$

Thus, the intersection of $\mathbf{v}_i(a_i), \mathbf{v}_i(a'_i)$ contains at most a single point.

Part 1. Given a set of points (p_i, t_i) , the graph of (p_i, t_i) is convex if and only if every secant line lies above every point (p_i, t_i) . Formally, for any p_1, p_2, p_3 , if $\alpha p_1 + (1 - \alpha)p_2 = p_3$, and $0 \leq \alpha \leq 1$, then

$$t_3 \leq \alpha t_1 + (1 - \alpha)t_2.$$

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

$$\begin{aligned} T_i(a_i) + (P_i(a'_i) - P_i(a_i))v &\leq T_i(a'_i), \\ T_i(a_i) + (P_i(a''_i) - P_i(a_i))v &\leq T_i(a''_i). \end{aligned}$$

Hence,

$$\begin{aligned} T_i(a_i) + (\alpha P_i(a'_i) + (1 - \alpha) P_i(a''_i) - P_i(a_i))v &\leq \alpha T_i(a'_i) + (1 - \alpha) T_i(a''_i), \\ T_i(a_i) &\leq \alpha T_i(a'_i) + (1 - \alpha) T_i(a''_i), \end{aligned}$$

as desired. □

A.2 Proof of Corollary 2

Proof. Using (6), we have

$$v_i(a_i) = \lim_{\delta \rightarrow 0} \frac{T_i(a_i) - T_i(\mathbf{s}_i(v_i - \delta))}{P_i(a_i) - P_i(\mathbf{s}_i(v_i - \delta))}.$$

Since $\mathbf{s}_i(\cdot)$ is continuous, this is the same as the following, for some $\epsilon > 0$:

$$= \lim_{\epsilon \rightarrow 0} \frac{T_i(a_i) - T_i(a_i - \epsilon)}{P_i(a_i) - P_i(a_i - \epsilon)} = \frac{dT_i}{dP_i} = \frac{T'_i(a_i)}{P'_i(a_i)}.$$

□

A.3 Proof of Corollary 3

Proof. For any z_i , if the true menu is $\{(P_i(a_i), T_i(a_i))\}$, the agent can attain the outcome $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$ by playing a mixed strategy, with distribution over actions equal to the conditional distribution

$$H(a_i \mid Z_i = z_i).$$

This attains the expected outcome:

$$E[(P_i(A_i), T_i(A_i)) \mid Z_i = z_i].$$

which we have shown is equal to $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$. Since any outcome attainable under $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$ is also attainable under $\{(P_i(a_i), T_i(a_i))\}$, maximum utility under $\{(\check{P}_i(z_i), \check{T}_i(z_i))\}$ is a lower bound for maximum utility under the true menu. \square

A.4 Proof of Corollary 5

Proof. Following the proof of Corollary 2, we have:

$$v_i(z_i) = \lim_{\delta \rightarrow 0} \frac{\check{T}_i(z_i) - \check{T}_i(z_i(v_i - \delta))}{\check{P}_i(z_i) - \check{P}_i(z_i(v_i - \delta))}.$$

Since $z_i(\cdot)$ is continuous, this implies, for $\epsilon > 0$,

$$= \lim_{\epsilon \rightarrow 0} \frac{\check{T}_i(z_i) - \check{T}_i(z_i - \epsilon)}{\check{P}_i(z_i) - \check{P}_i(z_i - \epsilon)} = \frac{d\check{T}_i}{dz_i} = \frac{\check{T}'_i(z_i)}{\check{P}'_i(z_i)}.$$

\square

A.5 Proof of Proposition 2

Proof. First, note that we can write

$$\check{P}(z) = E(P(v) \mid z) = \int_0^{\bar{v}} P(v) dG(v \mid z),$$

$$\check{T}(z) = E(T(v) \mid z) = \int_0^{\bar{v}} T(v) dG(v \mid z).$$

We will write $P(v)$, $T(v)$ as:

$$P(v) = \int_0^v P'(\tilde{v}) d\tilde{v},$$

$$T(v) = \int_0^v \frac{dT}{dP} P'(\tilde{v}) d\tilde{v} = \int_0^v \tilde{v} P'(\tilde{v}) d\tilde{v}.$$

Hence, the expectations with respect to z are

$$\begin{aligned} \check{P}(z) &= \int_0^{\bar{v}} P(v) g(v | z) dv = \int_0^{\bar{v}} \int_0^v P'(\tilde{v}) d\tilde{v} g(v | z) dv, \\ \check{T}(z) &= \int_0^{\bar{v}} T(v) g(v | z) dv = \int_0^{\bar{v}} \int_0^v \tilde{v} P'(\tilde{v}) d\tilde{v} g(v | z) dv. \end{aligned}$$

Suppose we have z_2, z_1 , where $z_2 > z_1$, hence $G(v | z_2) >_{FOSD} G(v | z_1)$. The slope of the secant line connecting them is:

$$\frac{\check{T}(z_2) - \check{T}(z_1)}{\check{P}(z_2) - \check{P}(z_1)} = \frac{\int_0^{\bar{v}} \int_0^v \tilde{v} P'(\tilde{v}) d\tilde{v} g(v | z_2) dv - \int_0^{\bar{v}} \int_0^v \tilde{v} P'(\tilde{v}) d\tilde{v} g(v | z_1) dv}{\int_0^{\bar{v}} \int_0^v P'(\tilde{v}) d\tilde{v} g(v | z_2) dv - \int_0^{\bar{v}} \int_0^v P'(\tilde{v}) d\tilde{v} g(v | z_1) dv}. \quad (23)$$

This simplifies to:

$$\frac{\int_0^{\bar{v}} \int_0^v \tilde{v} P'(\tilde{v}) d\tilde{v} [g(v | z_2) - g(v | z_1)] dv}{\int_0^{\bar{v}} \int_0^v P'(\tilde{v}) d\tilde{v} [g(v | z_2) - g(v | z_1)] dv}.$$

Since \tilde{v} is bounded above and below, and both $G(v | z_1)$ and $G(v | z_2)$ are bounded measures, we can change the order of integration for both the numerator and the denominator. For example, for the numerator,

$$\begin{aligned} \int_0^{\bar{v}} \int_0^v \tilde{v} P'(\tilde{v}) d\tilde{v} [g(v | z_2) - g(v | z_1)] dv &= \int_0^{\bar{v}} \int_{\tilde{v}}^{\bar{v}} \tilde{v} P'(\tilde{v}) [g(v | z_2) - g(v | z_1)] dv d\tilde{v} \\ &= \int_0^{\bar{v}} \tilde{v} P'(\tilde{v}) \int_{\tilde{v}}^{\bar{v}} [g(v | z_2) - g(v | z_1)] dv d\tilde{v} \\ &= \int_0^{\bar{v}} \tilde{v} P'(\tilde{v}) [(G(\bar{v} | z_2) - G(\bar{v} | z_1)) - (G(\tilde{v} | z_2) - G(\tilde{v} | z_1))] d\tilde{v}. \end{aligned}$$

Since $G(\bar{v} | z_2) = G(\bar{v} | z_1) = 1$, the term $(G(\bar{v} | z_2) - G(\bar{v} | z_1))$ cancels, and we are left with:

$$= \int_0^{\bar{v}} \tilde{v} P'(\tilde{v}) [G(\tilde{v} | z_1) - G(\tilde{v} | z_2)] d\tilde{v}.$$

For simplicity, we will now write all \tilde{v} 's in the integrand as v 's. Interchanging the

order of integration in both the numerator and denominator in 23, we have:

$$\frac{\check{T}(z_2) - \check{T}(z_1)}{\check{P}(z_2) - \check{P}(z_1)} = \frac{\int_0^{\bar{v}} v P'(v) [G(v | z_1) - G(v | z_2)] dv}{\int_0^{\bar{v}} P'(v) [G(v | z_1) - G(v | z_2)] dv}.$$

Now, suppose that $z_2 = z_1 + \epsilon$ for some $\epsilon > 0$. We then have:

$$\frac{\check{T}(z_1 + \epsilon) - \check{T}(z_1)}{\check{P}(z_1 + \epsilon) - \check{P}(z_1)} = \frac{\int_0^{\bar{v}} v P'(v) [G(v | z_1) - G(v | z_1 + \epsilon)] dv}{\int_0^{\bar{v}} P'(v) [G(v | z_1) - G(v | z_1 + \epsilon)] dv}.$$

Dividing all terms by the difference ϵ , we have:

$$\frac{\frac{\check{T}(z_1 + \epsilon) - \check{T}(z_1)}{\epsilon}}{\frac{\check{P}(z_1 + \epsilon) - \check{P}(z_1)}{\epsilon}} = \frac{\int_0^{\bar{v}} \frac{G(v|z_1) - G(v|z_1 + \epsilon)}{\epsilon} v P'(v) dv}{\int_0^{\bar{v}} \frac{G(v|z_1) - G(v|z_1 + \epsilon)}{\epsilon} P'(v) dv}.$$

Taking the limit as $\epsilon \rightarrow 0$, this becomes:

$$\frac{\check{T}'(z)}{\check{P}'(z)} = \frac{\int_0^{\bar{v}} v \left(-\frac{dG(v|z)}{dz} \right) P'(v) dv}{\int_0^{\bar{v}} \left(-\frac{dG(v|z)}{dz} \right) P'(v) dv},$$

which proves Proposition 2. □

A.6 Proof of Corollary 6

Proof. Let $\check{P}'(v)$ be equal to some constant C for all v . If $G(v, z) = \bar{G}(v - z)$, then $\frac{dG(v|z)}{dz} = g(v | z)$. Hence,

$$\frac{\check{T}'(z)}{\check{P}'(z)} = \frac{\int_0^{\bar{v}} v \left(-\frac{dG(v|z)}{dz} \right) \check{P}'(v) dv}{\int_0^{\bar{v}} \left(-\frac{dG(v|z)}{dz} \right) \check{P}'(v) dv} = \frac{\int_0^{\bar{v}} v g(v | z) C dv}{\int_0^{\bar{v}} g(v | z) C dv} = E[V | z].$$

□

A.7 Proof of Corollary 7

Proof. We have:

$$\frac{\check{T}'(z)}{\check{P}'(z)} = \frac{\int_0^{\bar{v}} v \left(-\frac{dG(v|z)}{dz} \right) \check{P}'(v) dv}{\int_0^{\bar{v}} \left(-\frac{dG(v|z)}{dz} \right) \check{P}'(v) dv} = \frac{\int_{\nu(z)-\delta}^{\nu(z)+\delta} v \left(-\frac{dG(v|z)}{dz} \right) \check{P}'(v) dv}{\int_{\nu(z)-\delta}^{\nu(z)+\delta} \left(-\frac{dG(v|z)}{dz} \right) \check{P}'(v) dv}.$$

We must have $\frac{\check{T}'(z)}{\check{P}'(z)} \in [\nu(z) - \delta, \nu(z) + \delta]$. Likewise, $E(V | z) \in [\nu(z) - \delta, \nu(z) + \delta]$. Hence,

$$\left| \frac{\check{T}'(z)}{\check{P}'(z)} - E(V | z) \right| \leq 2\delta.$$

□

A.8 Proof of Proposition 3

Proof. Fixing common component w , we have

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

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

Suppose strategies $\mathbf{s}_i(v_i, w)$ constitute an equilibrium under w . Then, 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 $P_i^{w'}(a_i), T_i^{w'}(a_i)$ as the expected allocation and transfer i achieves under w' , assuming opponents' actions are distributed as $A_{-i}^{w'}$. That is,

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

$$T_i^{w'}(a_i) = E[t_i(a_i, A_{-i}^{w'})].$$

In order to prove our Proposition, 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 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,

$$(v_i + w') P_i^{w'}(a_i + (w' - w)) - T_i^{w'}(a_i + (w' - w)) - (v_i + w') \bar{x}_i = (v_i + w) P_i^w(a_i) - T_i^w(a_i) - (v_i + w) \bar{x}_i. \quad (24)$$

Thus, if a_i maximizes the right-hand side, $a_i + (w' - w)$ maximizes the left-hand side, and we are done. \square

Proof of (24)

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 separable heterogeneity 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,

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

Thus we have shown that

$$P_i^{w'}(a_i + (w' - w)) = P_i^w(a_i). \quad (25)$$

For transfers, we have

$$\begin{aligned} T_i^{w'}(a_i + (w' - w)) &= E[\mathbf{t}_i(a_i + (w' - w), A_{-i}^w)] \\ &= E[\mathbf{t}_i(a_i + (w' - w), A_{-i}^w + (w' - w))]. \end{aligned}$$

Again using Definition 2, we have

$$\mathbf{t}_i(a_i + (w' - w), A_{-i}^w + (w' - w)) = \mathbf{t}_i(a_i, A_{-i}^w) + (w' - w)(\mathbf{x}_i(a_i, A_{-i}^w) - \bar{x}_i).$$

Hence,

$$\begin{aligned} E[\mathbf{t}_i(a_i + (w' - w), A_{-i}^w + (w' - w))] &= \\ E[\mathbf{t}_i(a_i, A_{-i}^w)] + E[(w' - w)(\mathbf{x}_i(a_i, A_{-i}^w) - \bar{x}_i)]. \end{aligned} \quad (26)$$

The term $E[\mathbf{t}_i(a_i, A_{-i}^w)]$ is equal to $T_i^w(a_i)$ by definition. Using linearity of expectations, the right term simplifies to

$$\begin{aligned} E[(w' - w)(\mathbf{x}_i(a_i, A_{-i}^w) - \bar{x}_i)] &= (w' - w)(E[\mathbf{x}_i(a_i, A_{-i}^w)] - \bar{x}_i) \\ &= (w' - w)(P_i^w(a_i) - \bar{x}_i). \end{aligned}$$

Hence, we have shown that

$$T_i^{w'}(a_i + (w' - w)) = T_i^w(a_i) + (w' - w)(P_i^w(a_i) - \bar{x}_i). \quad (27)$$

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 (25) and (27), this is

$$\begin{aligned} &= [(v_i + w')P_i^w(a_i) - T_i^w(a_i) - (w' - w)(P_i^w(a_i) - \bar{x}_i)] - (v_i + w')\bar{x}_i \\ &= [(v_i + w)P_i^w(a_i) - T_i^w(a_i)] - (v_i + w)\bar{x}_i. \end{aligned}$$

Hence we have proved the equality in (24). □

A.9 Proof of Proposition 4

Proof. We observe actions \tilde{a}_{ij} , allocations x_{ij} , and 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 position-invariant 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 \tilde{a}_{ij} across players. In the remainder of this proof, we omit the subscript j .

We can empirically estimate the functions:

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

The functions involved in the position-invariant equilibrium are the probability of trade and the homogenized $T^0(a_i)$ as a function of the homogenized action a_i :

$$P_i(a_i) = E \left[\mathbf{x}_i(a_i, A_{-i}) \right], \quad T_i^0(a_i) = E \left[\mathbf{t}_i(a_i, A_{-i}) \right].$$

Below, we show that the functions $P_i(a_i), T_i^0(a_i)$ are identified from the functions $\tilde{P}_i(\tilde{a}_i), \tilde{T}_i(\tilde{a}_i), f_W, f_{A_i}$.

Probabilities

First, note that:

$$\begin{aligned} \tilde{P}_i(\tilde{a}_i) &= E \left[\mathbf{x}_i \left(\tilde{a}_i, \tilde{A}_{-i} \right) \right] \\ &= E \left[E \left[\mathbf{x}_i(W + A_i, W + A_{-i}) \mid A_i, W + A_i = \tilde{a}_i \mid W + A_i = \tilde{a}_i \right] \right]. \end{aligned}$$

By Definition 2, we have

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

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

Hence,

$$\begin{aligned} \tilde{P}_i(\tilde{a}_i) &= E[P_i(A_i) \mid W + A_i = \tilde{a}_i] \\ \tilde{P}_i(\tilde{a}_i) &= \frac{\int_{a_i=-\infty}^{\infty} P_i(a_i) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}{\int_{a_i=-\infty}^{\infty} f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}. \end{aligned}$$

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

$$\frac{f_{A_i}(a_i) f_W(\tilde{a}_i - a_i)}{\int_{a_i=-\infty}^{\infty} f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}. \quad (28)$$

Since the distributions of A_i and W both have bounded support, we can set $P_i(a_i)$ to 0 for all a_i outside the support of A_i . Thus, both $P_i(a_i)$ and (28) are in L^1 ; hence the convolution theorem applies, meaning that the convolution of $P_i(a_i)$ and (28) is invertible, hence $P_i(a_i)$ is identified from (28) and $\tilde{P}_i(\tilde{a}_i)$.

Homogenized transfers

First, note that

$$\begin{aligned} \tilde{T}_i(\tilde{a}_i) &= E\left(\mathbf{t}_i\left(\tilde{a}_i, \tilde{A}_{-i}\right)\right) \\ &= E\left[E\left[\mathbf{t}_i(W + A_i, W + A_{-i}) \mid A_i, W + A_i = \tilde{a}_i\right] \mid W + A_i = \tilde{a}_i\right]. \end{aligned}$$

By Definition 2, we have

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

Taking expectations,

$$\begin{aligned} &E[\mathbf{t}_i(W + A_i, W + A_{-i}) \mid W, A_i] \\ &= E[\mathbf{t}_i(A_i, A_{-i}) \mid W, A_i] + E[(P(A_i) - \bar{x}_i) W \mid W, A_i] \\ &= T_i^0(A_i) + (P_i(A_i) - \bar{x}_i) W. \end{aligned}$$

Hence,

$$\tilde{T}_i(\tilde{a}_i) = E \left[(T_i^0(A_i) + (P_i(A_i) - \bar{x}_i)) W \mid W + A_i = \tilde{a}_i \right].$$

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

$$= E \left[(T_i^0(A_i) + (P_i(A_i) - \bar{x}_i)(\tilde{a}_i - A_i)) \mid W + A_i = \tilde{a}_i \right].$$

In integral form, this equation is:

$$\begin{aligned} \tilde{T}_i(\tilde{a}_i) = & \frac{\int_{a_i=-\infty}^{\infty} T_i^0(a_i) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}{\int_{a_i=-\infty}^{\infty} f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i} + \\ & \frac{\int_{a_i=-\infty}^{\infty} (\tilde{a}_i - a_i) (P_i(a_i) - \bar{x}_i) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}{\int_{a_i=-\infty}^{\infty} f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}. \end{aligned} \quad (29)$$

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

$$E(W \Delta P_i \mid \tilde{a}_i) \equiv \frac{\int_{a_i=-\infty}^{\infty} (\tilde{a}_i - a_i) (P_i(a_i) - \bar{x}_i) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}{\int_{a_i=-\infty}^{\infty} f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}. \quad (30)$$

Since we have shown that $P_i(a_i)$ is identified, $E(W \Delta P_i \mid \tilde{a}_i)$ can be calculated for any \tilde{a}_i . We can rearrange (29) to:

$$\tilde{T}_i(\tilde{a}_i) - E(W \Delta P_i \mid \tilde{a}_i) = \frac{\int_{a_i=-\infty}^{\infty} T_i^0(a_i) f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}{\int_{a_i=-\infty}^{\infty} f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}.$$

The term $\tilde{T}_i(\tilde{a}_i)$ can be estimated from the data, so the entire LHS is known. The RHS is a convolution of $T_i^0(a_i)$ against

$$\frac{f_{A_i}(a_i) f_W(\tilde{a}_i - a_i)}{\int_{a_i=-\infty}^{\infty} f_{A_i}(a_i) f_W(\tilde{a}_i - a_i) da_i}.$$

Again, by the convolution theorem, this is invertible, and thus $T^0(a_i)$ is identified. \square

B Spline Construction

Suppose we wish to estimate some function of a variable y using splines. Given a knot sequence $y_1 \dots y_n$, we define the quadratic splines $S_1^q(\cdot) \dots S_{n+1}^q(\cdot)$:

$$\begin{aligned}
 S_1^q(y) &= \frac{2}{y_2 - y_1} \left[|y - y_1|_+ - \frac{|y - y_1|_+^2 - |y - y_2|_+^2}{2(y_2 - y_1)} \right] \\
 &\quad \vdots \\
 S_k^q(y) &= \frac{2}{y_{k+1} - y_{k-1}} \left[\frac{|y - y_{k-1}|_+^2 - |y - y_k|_+^2}{2(y_k - y_{k-1})} - \frac{|y - y_k|_+^2 - |y - y_{k+1}|_+^2}{2(y_{k+1} - y_k)} \right] \\
 &\quad \vdots \\
 S_n^q(y) &= \frac{2}{y_n - y_{n-1}} \left[\frac{|y - y_{n-1}|_+^2 - |y - y_n|_+^2}{2(y_n - y_{n-1})} - |y - y_n|_+ \right] \\
 S_{n+1}^q(y) &= 1.
 \end{aligned}$$

As shown in Figure 5, each $S_k^q(\cdot)$ behaves like a smoothed step function, increasing on the interval $[y_{k-1}, y_{k+1}]$. By constraining the coefficients $\beta_1 \dots \beta_n$ to be nonnegative, we can constrain the target function to be nondecreasing. Moreover, each $S_k^q(\cdot)$ has the property that $\lim_{y \rightarrow \infty} S_k^q(y) = 1$; hence, to constrain the target function to always lie below some bound M , we need only constrain $\sum_{k=1}^{n+1} \beta_k \leq M$. We impose the constraint that $\sum_k \beta_k \leq 1$ in estimating $\hat{P}(\cdot)$.

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

$$\begin{aligned}
 S_1^c(y) &= \int_{-\infty}^y S_1^q(\tilde{y}) d\tilde{y} \\
 &\quad \vdots \\
 S_n^c(y) &= \int_{-\infty}^y S_n^q(\tilde{y}) d\tilde{y} \\
 S_{n+1}^c(y) &= y \\
 S_{n+2}^c(y) &= 1.
 \end{aligned}$$

Constraining the coefficients in this approximation to be nonnegative ensures that the target function is convex.

C Bargaining Simulation Details

For the bargaining simulations, we follow the game studied in Satterthwaite and Williams (1989), referred to as a k double auction. The authors demonstrated that, for any $k \in [0, 1]$, a continuum of strictly increasing, differentiable equilibria exist satisfying the following linked differential equations for the buyer's action a_B and seller's action a_S as a function of their private valuations $V_B \sim F_B$ and $V_S \sim F_S$:


$$a_B^{(-1)}(a_S(v_S)) = a_S(v_S) + ka'_S(v_S) \left(v_S + \frac{F_S(v_S)}{f_S(v_S)} \right), \quad (31)$$

$$a_S^{(-1)}(a_B(v_B)) = a_B(v_B) + (1 - k)a'_B(v_B) \left(v_B - \frac{1 - F_B(v_B)}{f_B(v_B)} \right). \quad (32)$$

Satterthwaite and Williams (1989) provided an approach for solving for equilibria numerically given knowledge of the distributions of player types. A point (v_S, v_B, t) , where t is a final price which v_S and v_B could potentially agree upon, is chosen in the set $\mathcal{P} = \{(v_S, v_B, t) : \underline{v}_S \leq v_S \leq t \leq v_B \leq \bar{v}_B, v_S \leq \bar{v}_S, v_B \geq \underline{v}_B\}$, and then a one-dimensional manifold passing through this point is traced out using the differential equations defined by (31) and (32). This path traces out an equilibrium. 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.


In our bargaining simulations, we focus on two possibilities for valuation distributions: $V_S \sim U[0, 1]$ and $V_B \sim U[0, 1]$, or $V_S \sim N(.2, .5)$ and $V_B \sim N(.6, .3)$. For the Normal case, we truncate the distribution to lie within $[0, 1]$ to simplify computation. We set $k = 1/2$. In each case we choose an equilibrium passing through the point $(v_S, v_B, t) = (0.375, 0.625, 0.5)$. An example of a solution path crossing through this point, for the truncated Normal case, is shown in Figure 6.

Figure 5: Quadratic/Cubic Spline Basis



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

Figure 6: An Equilibrium in the k Double Auction



Notes: A solution to the $k = 1/2$ double auction, lying with the tetrahedron $\mathcal{P} = \{(v_S, v_B, t) : \underline{v}_S \leq v_S \leq t \leq v_B \leq \bar{v}_B, v_S \leq \bar{v}_S, v_B \geq \underline{v}_B\}$. This solution passes through the point $(v_S, v_B, t) = (0.375, 0.625, 0.5)$. 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 $[0, 1]$.