Auction Design with Data-Driven Misspecifications*

Philippe Jehiel and Konrad Mierendorff†

November 12, 2021

Abstract

We consider auction environments in which, at the time of the auction, bidders observe noisy signals about their own ex-post value. Novice bidders do not know the signal generating process for others and instead are assumed to build a statistical model relating others' bids to ex-post values based on the data obtained from past similar auctions. Crucially, we assume that only ex post values and bids are accessible while signals observed by bidders in past auctions remain private. We consider steady-states in such environments with a mix of rational and data-driven bidders, and we allow for correlation in the signal distribution. After reviewing the working of the approach in second-price and first-price auctions, we show our main result that inefficiencies must arise when there is a mix of rational and data-driven bidders, signals are correlated, and auction-like mechanisms are considered.

Keywords: Belief Formation, Auctions, Efficiency, Analogy-based Expectations
JEL Classification Numbers: D44, D82, D90

1 Introduction

In the standard model of auctions, bidders hold private information about the value of the object for sale. They commonly know how information is distributed across bidders,

*The first draft of this paper was completed before the tragic passing of the second co-author to whom this paper is dedicated. The content of this paper was presented by the first co-author in a preliminary shape as the Hurwicz lecture at the 2019 Conference on Economic Design in Budapest. Both co-authors thank Dirk Bergemann, Paul Klemperer, Rani Spiegler, Martin Weidner and Bob Wilson for helpful comments as well as the seminar audiences at SAET 2019, Oxford University, TSE, the 14th workshop on Economic Design and Institutions, UCLA, and the Israeli theory seminar. Jehiel thanks the European Research Council for funding (grant no 742816).

†Jehiel: Paris School of Economics and University College London (email: jehiel@enpc.fr); Mierendorff: Department of Economics, University College London (email: k.mierendorff@ucl.ac.uk)

and every bidder i correctly understands how bidder j chooses his bid as a function of his private information. Every bidder best-responds to this correct understanding given the rules of the auction. The resulting strategy profile is a Bayes Nash Equilibrium.

A classic rationale for the correct understanding assumption in Bayes Nash Equilibrium is based on learning (see, for example, Dekel, Fudenberg and Levine, 2004). If similar auctions are played many times (by different subjects in the roles of the various bidders), by looking at previous bids as well as what the corresponding bidders knew (at the time of the auction), one can recover the mapping from information to bids for past observations. If a steady state has been reached, this mapping will correctly describe behavior in the current auction, thereby supporting the rational expectation formulation.

But, access to previous bidders' information is not so natural in practice. If the private information held by previous bidders is not disclosed, it is less clear how a novice bidder (who would only be exposed to data from past similar auctions played by others) would manage to have a correct understanding of other bidders' strategies. A different modeling is required to describe the behavior of novice bidders in such a case.

To make progress in this under-explored direction, we consider one-object auctions in which at the time of the auction, bidder i's private information is a noisy signal about his (own) ex post value for the good assumed to take one of finitely many realizations. We assume that after the auction is completed, what is publicly disclosed is the profile of bids as well as the ex post values of the various bidders, but not the signals observed by the bidders at the time of the auction. In the main part, we restrict attention to two-bidder auctions, but we note that our main insights carry over when there are more bidders.

It should be highlighted that we allow for correlations between signals, which will play a key role in the analysis. Moreover, despite the correlation, the setting is one of private values, since the distribution of bidder i's ex post value is fully determined by bidder i's signal (i.e., it is unaffected by the other bidder's signal, conditional on i's signal). Yet, novice players are assumed to be unaware of the true signal generating process, and thus of the private value character of the auction. Instead, like econometricians would do, they construct a representation of the statistical links between the variables of interest based on the signal they receive as well as the dataset available to them. Specifically, observations from past auctions take the form (b_1, v_1, b_2, v_2) where b_j is the bid previously submitted by a subject in the role of bidder j and v_j is his ex post value.[1] A novice bidder i constructs

[1] We will consider in the discussion how to deal with the cases in which bids are anonymous or cases in which only the winning bid is observed.

from the dataset the empirical distribution describing how b_j is distributed conditional on the various possible ex post values of bidder i. He also uses his own signal θ_i that is informative about the likelihood of his various possible ex post values v_i and combines the two to form a belief about how (v_i, b_j) are jointly distributed given his own signal. He then best-responds to this belief given the rules of the auction.

We will be considering steady state environments in which there is a mixed population of bidders composed of a share of novice bidders (whose expectations are formed as just informally explained) and a complementary share of rational bidders assumed to have a correct understanding of the strategies of the two types of bidders as well as their respective share. Observe that rational bidders can alternatively be viewed as experienced bidders who would have had the opportunity to find out the best strategy in their auction environment (without necessarily an explicit knowledge of how the various types of bidders behave nor of the shares of the various types). We will refer to such steady states as Data-Driven Equilibria (see below for a discussion of how this concept relates to other existing concepts).[2]

We apply this model to understand the efficiency properties of Data-Driven Equilibria, and more particularly, whether by a judicious choice of auction rule, one can implement an efficient allocation.[3] This question resembles classic investigations in mechanism design with the main difference that our solution concept is not the Bayes Nash Equilibrium, but the Data-Driven Equilibrium designed to deal with the presence of novice bidders in our environment. Another departure from classic mechanism design is that for our main result, we will not be considering abstract mechanisms with arbitrary messages to be sent by bidders to the designer before an outcome (allocation and transfer) is decided. Instead, we will focus on what we call auction-like mechanisms defined as mechanisms in which each bidder submits a real-valued bid, and an outcome is chosen as a function of the profile of bids with the restriction that if a bidder submits a higher bid, this bidder has more chance

[2] Our view of novice bidders shares some similarities with the modeling of Doraszelski et al. (2018) with some important differences. The similarity is that in both cases the expectation of novice bidders is assumed to be formed based on the available data. But, while in the long run, bidders bid optimally in Doraszelski et al., this is not so in our approach (which directly analyzes the long run through the focus on steady state). The main reason for this is our premise that the available data in our case do not allow the novice bidders to understand the correct data generating process, thereby leading to mistakes in some cases. Another difference in terms of motivation is that we have in mind novice bidders who would just play once whereas Doraszelski et al. have in mind that the same bidders keep playing repeatedly.

[3] Obviously, given that at the time of the auction, bidders do not know their ex post values, efficiency is defined here at the interim stage, based on the information available to bidders at the time of the auction.

of winning the object.

Our insights are as follows. First, unless the distributions of signals of the two bidders are independent, data-driven bidders rely on a misspecified statistical model, and as a result choose suboptimal bidding strategies. In Section 3, we start illustrating this with Second-Price Auctions (SPA) in the (symmetric) binary case in which there are two possible ex post values. We show that unlike rational bidders, novice bidders do not bid their expected value when there are correlations. As in winner's curse models,[4] novice bidders make inferences about their ex post value from how the other bidder bids. In the case of positive correlation, this leads novice bidders to bid more than their expected value when they receive good signals (because in the neighborhood of large opponent's bids, the own ex post value is more likely to be high) and less than their expected value when they receive bad signals (for a symmetric reason). We provide a numerical characterization of the equilibrium for a parametric class of distributions.

Clearly, the fact that novice and rational bidders do not bid in the same way leads to inefficiencies, unless there is perfect correlation of the signals, or the bidders are all novice or all rational. For our parametric example, we observe that the normalized welfare loss in the data-driven equilibrium of the SPA is U-shaped in the share of novice bidders as well as in the degree of correlation. More generally, we show for the binary mixed population case that as soon as there are correlations, there is some welfare loss in the Second Price Auction. We also consider First-Price Auctions (FPA), for which we also show that there must be inefficiencies whenever there is correlation.[5]

Our main result concerns general auction-like mechanisms when ex post values can take at least three realizations. In Section 4, we provide a general inefficiency result. More precisely, we show in the mixed population case that for generic joint distributions of signals, there is no auction-like mechanism that allows to obtain an efficient outcome with probability one in a Data-Driven Equilibrium. The intuition for this result is as follows. To obtain efficiency among rational bidders, only the Second-Price Auction or a strategically equivalent auction format can be used. This is so because with more than two ex post values there is generically a manifold of signal realizations corresponding to the same expected value for the object, but different beliefs about the signal realization of the other bidder, and if the payment in the auction were to depend on the own bid, then

[4] See Milgrom and Weber (1982) for the classic analysis of such models.

[5] We illustrate through an example that sometimes the welfare loss may be larger in the Second-Price Auction than in the First-Price Auction. The opposite welfare ranking can arise in other examples.

the belief about the opponent would affect the shape of the optimal bid, as in First-Price Auctions. Since in Second-Price auctions, novice bidders do not bid their expected value as also observed in the simplified binary case, we conclude that inefficiencies must occur.

In Section 5, we put our analysis in perspective. First, we discuss alternative specifications of cognitive limitations in auction-like mechanisms either due to different accessibility to data sets from past auctions, or due to more or less sophisticated use of the same limited data sets.[6] We make the simple observation that our main impossibility result would a fortiori hold if we were to consider a mixed population that includes extra cognitive types in addition to those considered in the main part of the paper.

Second, we discuss scenarios in which losing bids would not be accessible from past auctions, thereby considering a natural further restriction on the accessible data sets. We discuss various possible approaches to modeling novice bidders in this context, and suggest that all of them would lead to results similar to those obtained in our main model.

Third, we discuss scenarios in which past bids would be anonymous. After suggesting how the modeling of novice bidders should be modified in this case, we note that similar inefficiency results would be obtained.

Fourth, we briefly discuss more general mechanisms beyond the auction-like mechanisms and note that judicious uses of such mechanisms (direct mechanisms of the scoring rule type) may allow to elicit the beliefs of every bidder i about bidder j's (interim) type. Since in our baseline model, for generic distributions, no two different types have the same belief about their opponent's type, such mechanisms could potentially be used to implement a broad range of allocation rules in the spirit of the work of Crémer and McLean (1988), Johnson, Pratt, and Zeckhauser (1990), MAfee and Reny (1992), and Gizatulina and Hellwig (2017). It should be mentioned though that such mechanisms are not commonly used in market design (perhaps because they require a level of knowledge -how the belief about opponent's type relates to the payoff–relevant type that is rarely available to the designer). This consideration has led us to restrict attention to auction-like mechanisms which seem much more practical from a market design perspective.[7]

[6] We note that the use of data as considered in the main model would not allow to reconstruct the correct signal-generating process due to a fundamental identification constraint. And we suggest our proposed approach can be regarded as corresponding to a sophisticated use of the limited dataset that does not make ad hoc assumptions about bidding behavior in past actions.

[7] It may be mentioned that if we were to consider richer spaces of cognitive types, then one could easily conceive that different types have the same belief, thereby reducing the scope of what can be implemented with such abstract mechanisms (even assuming the designer has the full knowledge required to make the best use of those).

Fifth, we briefly discuss the possibility of more general signal structures, in particular allowing for interdependent value settings rather than private value ones. Section 6 concludes.

Related literature

Our paper relates to different branches of literature. First, the modeling of data-driven bidders is in the spirit of the Analogy-Based Expectation Equilibrium (Jehiel, 2005) to the extent that these bidders aggregate the bid behavior of their opponent according to their own ex post value. Such an aggregation of bidding behavior can be related to the payoff-relevant analogy partition introduced in Jehiel and Koessler (2008).[8] This modeling can also be related to the Bayesian Network Equilibrium (Spiegler, 2016), viewing these agents as believing that their values cause the bid of the opponent, but note that here the reasoning of novice bidders is viewed as a consequence of the nature of the dataset accessible to them, not as a consequence of a subjective wrong causality relation they could have in mind (see Spiegler (2020), and Jehiel (2021), for elaborations of the link between the Analogy-Based Expectation Equilibrium and the Bayesian Network Equilibrium as well as Spiegler (2021) for an extension of the Bayesian Network Equilibrium to settings better suited to deal with the present application).[9]

Our paper is also related to the robust mechanism design literature (Bergemann and Morris, 2005), in the sense that a common motivation in that literature and our approach is that it may be hard to know what the beliefs of agents are. While the robust mechanism design literature uses this observation to motivate the desire to implement outcomes for a large range of (or even all) beliefs, our paper explicitly suggests a method of belief formation for bidders who do not have access to such information from past auctions.[10] Our paper is also mostly concerned with a subclass of mechanisms that we refer to as auction-like mechanisms and how these perform efficiency-wise in the joint presence of data-driven bidders and rational bidders, which has no counterpart in the literature on

[8] See Jehiel (2011) for a different application of ABEE to mechanism design in which unlike in the present setting the designer is assumed to control what is disclosed from past auctions.

[9] At a more general level, one can also relate the Data-Driven Equilibrium to the self-confirming equilibrium (Battigalli, Gilli, and Molinari, 1992; Delel, Fudenberg, and Levine, 2004, as well as recent behavior equilibrium models dealing with misspecified beliefs (see Eyster and Rabin, 2005; Esponda, 2008; Esponda and Pouzo, 2016, among others).

[10] Our paper can thus be viewed as offering a different approach than that of Bergemann and Morris to Wilson's critique calling for a relaxation of the common prior assumption.

robust mechanism design.[11]

Finally, from a technical point of view, our analysis makes use of some results developed in the literature on mechanism design with correlation. In particular, we borrow genericity arguments from Gizatulina and Hellwig (2017).

2 Model

Mechanisms. We consider the allocation of a single object to two bidders $i = 1, 2$ via an auction or more general auction-like mechanism. To simplify notation, when we consider a generic bidder $i \in \{1, 2\}$, we denote the opponent by $j \neq i$. A mechanism $M = [(B_i), q, p]$ consists of three elements: (i) feasible bids B_i for the two bidders. A profile of bids is denoted $b = (b_1, b_2) \in B := B_1 \times B_2$. (ii) an allocation rule $q : B \to [0, 1]^2$, $q(b) = (q_1(b), q_2(b))$, with $q_1(b) + q_2(b) \leq 1$, where $q_i(b)$ is the probability that bidder i gets the object if the bid profile b is submitted. (iii) A payment rule $p : B \to \mathbb{R}^2$, $p(b) = (p_1(b), p_2(b))$, where $p_i(b)$ denotes the payment bidder i has to make if the bid profile b is submitted.

Valuations. Ex-post, the value of the object for bidder i is denoted v_i. It can take values in $V = \left\{v^1, \ldots v^K\right\}$. v_i can also be interpreted as the expected value based on a signal that is known ex-post if the ex-post value is not learned completely. Up to normalization, it is without loss to assume that $0 = v^1 < \ldots < v^K = 1$. When participating in a mechanism, each bidder has an interim type $\theta_i = (\theta_i^1, \ldots, \theta_i^K) \in \Theta := \Delta V$, where θ_i^k denotes the probability that $v_i = v^k$. A profile of types is denoted $\theta = (\theta_1, \theta_2)$. We assume that conditional on θ_i, v_i is independent of θ_j. As a consequence the expected valuation of a bidder only depends on her own interim type: $E[v_i|\theta] = E[v_i|\theta_i]$. In other words, we are considering a setting with *private values*. Interim types are jointly distributed with cumulative distribution function $F(\theta)$ and density $f(\theta)$ defined over Θ^2, and our main interest is in the case where θ_1 and θ_2 are not independent. We assume throughout that the joint distribution is symmetric and has a continuous and positive density. When there

[11] In a related vein, Bergemann et al. (2017) is concerned with providing a lower bound on revenues in first-price auctions when it is not known what bidders' signals are, which can be related to the view that the beliefs held by bidders are not observed by outsiders. However, Brooks et al. have in mind that only the designer or analyst lacks that information, which leads them to consider Bayes Nash Equilibria of the corresponding Bayesian game, in contrast to our modelling of data-driven bidders which has led us to rely on a different notion of equilibrium (steady state).

is no confusion, we slightly abuse notation and denote marginal distributions $F_i(\theta_i)$ and $f_i(\theta_i)$ by $F(\theta_i)$ and $f(\theta_i)$; and conditional distributions $F_i(\theta_i|\theta_j)$ and $f_i(\theta_i|\theta_j)$, by $F(\theta_i|\theta_j)$ and $f(\theta_i|\theta_j)$.

Rational and Misspecified Bidders. We assume that each bidder i is characterized by a *generalized type* $t_i = (\theta_i, s_i)$, where θ_i denotes the *interim type* described before, and $s_i \in \{r, m\}$ specifies the *sophistication* of the bidder. We denote the set of general types by $T = \Theta \times \{r, m\}$. For simplicity we will call θ_i just the *type*. The probability that $s_i = r$ is denoted $\lambda \in (0, 1)$; we assume that it is independent of θ_i and across bidders. $s_i = r$ means that bidder i is *rational*; and $s_i = m$ means that bidder i is *misspecified*. Informally, the rational type correctly understands the environment, whereas the misspecified type holds beliefs that are endogenously determined by past observations of equilibrium outcomes of the mechanism she currently participates in. As we will see, this way of forming beliefs can lead to misspecifications, hence the name of the type.

We now make this precise. Fix a mechanism $M = [(B_i), q, p]$. A strategy of bidder i is a function $b_i : T \to B_i$, where as a shorthand we write $b_i(\theta_i, s_i) = b_i^{s_i}(\theta_i)$—that is, $b_i^r(\cdot)$ is the strategy of the rational type, and $b_i^m(\cdot)$ is the strategy of the misspecified type of bidder i.[12] A strategy profile is denoted by $b = (b_1, b_2) = (b_1^r, b_1^m, b_2^r, b_2^m)$ and we denote the space of all strategy profiles by $\mathcal{B}$.

For a rational type of bidder i, the expected utility of type θ_i when submitting bid $b_i \in B_i$, and assuming that bidder j bids according to $b_j(\cdot)$, is given by

$$U_i^r(b_i, \theta_i | b_j(\cdot)) = \mathbb{E}_f \left[v_i \, q_i(b_i, b_j(\theta_j, s_j)) - p_i(b_i, b_j(\theta_j, s_j)) | \theta_i \right],$$

where $\mathbb{E}_f$ is the expectation with respect to the correct distribution f and the probability λ.

Next consider the misspecified type. We assume that this type forms a belief using past observations from the same mechanism played by similar bidders. Suppose the mechanism is run repeatedly with two (short-lived) bidders whose generalized type profiles are drawn i.i.d., across repetitions. If both bidders play according to a fixed strategy profile, as they would in a steady state, then repeated play generates a data set with observations (b_1, v_1, b_2, v_2). We make the assumption that only bids and ex-post valuations are observable.

[12] We only consider pure strategies in our setting with continuous interim types.

Assumption 1. *For each mechanism we consider, we assume that bidders have access to observations of the form* (b_1, v_1, b_2, v_2) *from the same mechanism. The data about past mechanisms does not include the types* (θ_1, θ_2) *of past bidders.*

The idea behind this assumption is that bids are often disclosed after an auction and as time goes by, the ex-post valuation of the bidders, or an estimate thereof becomes known as well. On the other hand, bidders typically do not have access to the beliefs that past bidders held at the time of bidding.

Past data allow bidders to identify the joint distribution of observable variables. We abstract from issues of estimation, and assume that bidders can recover this distribution without estimation error. The misspecified bidder then forms a simple model that combines relevant information from the empirical distribution of (b_1, v_1, b_2, v_2), and her belief that her own v_i is distributed according to θ_i. To illustrate consider an auction with possible bids $B_1 = B_2 = [0, \infty)$. To assess the payoff from different bids, a bidder needs to know the joint distribution of her own valuation v_i and the opponent's bid b_j, conditional on her own type θ_i. The misspecified bidder combines the distribution of v_i given by her type θ_i with the joint distribution of v_i and the opponent's bid b_j learned from the data in a parsimonious way, taking the joint distribution to be

$$\mathbb{P}_m \left[v_i = v^k, b_j \leq b | \theta_i \right] = \theta_i^k \times H_i(b|v^k) \tag{1}$$

where $H_i(b|v^k)$ is the c.d.f. of b_j conditional on $v_i = v^k$ that is obtained from the data. Throughout, we will use $\mathbb{P}_m$ for probabilities assessed by the misspecified type and $\mathbb{P}_f$ for probabilities computed using the correct probabilistic model (given the density "f"). To see the difference in this particular example, note that

$$\mathbb{P}_f \left[v_i = v^k, b_j \leq b | \theta_i \right] = \theta_i^k \times \mathbb{P}_f \left[b_j \leq b | \theta_i, v_i = v^k \right] = \theta_i^k \times \mathbb{P}_f [b_j \leq b | \theta_i]$$

where the second equality follows from the assumption that θ_j and v_i are independent, conditional on θ_i. Under Assumption 1, $\mathbb{P}_f[b_j \leq b|\theta_i]$ cannot be assessed directly from the data since the types of past bidders are not available. In order to identify $\mathbb{P}_f[b_j \leq b|\theta_i]$ from the data, one would have to make assumptions about the strategies used by past bidders. These assumptions are ad hoc if only data on past bids and ex-post signals are available and a misspecified bidder does not have insight into the type of reasoning used by past bidders. The misspecified type therefore does not attempt to use the data through the lens of such assumptions but just takes the empirical correlation between v_i and b_j as given. One way to interpret the difference between the rational and the misspecified type is that the former is an experienced bidder who understands how bidders reason and is thus able to formulate the correct model to interpret the data where the latter lacks this understanding.[13] In particular, the misspecified type does not know that conditional on her own type θ_i, her own valuation is independent of the opponent's type, and thus that her valuation and the opponents bid are also conditionally independent. The data available from past auctions, however, exhibits a correlation between the valuation v_i and the bid b_j, since conditioning on the unobserved type θ_i is not possible. This is the source of the misspecification of the m-type. As we will see, this gives rise to bidding behavior that is similar to the winner's curse.

To summarize, for a misspecified type of bidder i, the expected utility of type θ_i when submitting bid $b_i \in B_i$, and assuming that bidder $j \neq i$ bids according to $b_j(\cdot)$, is given by

$$\begin{aligned} U_i^m(b_i, \theta_i | b_j(\cdot)) &= \mathbb{E}_m \left[v_i \, q_i(b_i, b_j(\theta_j, s_j)) - p_i(b_i, b_j(\theta_j, s_j)) | \theta_i \right], \\ &= \sum_{k=1}^{K} \theta_i^k \int_{B_j} [v^k \, q_i(b_i, b_j) - p_i(b_i, b_j)] dH_i(b_j | v^k), \end{aligned}$$

where $\mathbb{E}_m$ is the expectation formed according to the model described above. Note that in order to determine $H_i(\cdot|v^k)$, it is enough to specify the strategy $b_j(\cdot)$ since v_i and b_j in the current auction do not depend on the bids placed by the bidder in role i in the past.[14] To understand this better, the example of the second-price auction in the next section will be helpful.

Equilibrium To close the model, we assume that $H_i(\cdot|v^k)$ are equilibrium objects that are generated by the equilibrium strategy profile and the misspecified type best-responds given her beliefs that are captured by $H_i(\cdot|v^k)$.

Definition 1. The strategy profile $b(\cdot)$ is a *"Data-Driven Equilibrium"* of the mechanism $M = [(B_i), q, t]$ if for all $i \neq j$, and for all $\theta_i \in \Theta$,

[13] An alternative equivalent interpretation of a rational bidder is that such a bidder has found out the optimal strategy (possibly through a trial and error process) whereas the misspecified bidder can only rely on the publicly available data from past similar auctions.

[14] Another way to motivate why bidder i does not use the b_i from past auctions is that she is unsure what led a past bidder i to choose his bid, as it could be determined by his information and/or his way of reasoning none of which are accessible to her.

(a) $b_i^r(\theta_i) \in \arg\max_{b_i \in B_i} U_i^r(b_i, \theta_i | b_j(\cdot))$,

(b) $b_i^m(\theta_i) \in \arg\max_{b_i \in B_i} U_i^m(b_i, \theta_i | b_j(\cdot))$, where the distribution $H_i(b_j | v^k)$ used to compute U_i^m is (consistently) derived from $b(\cdot)$, f, and $\text{Prob}[s_i = r] = \lambda$.

3 Standard Auctions

Before considering auction-like mechanisms and presenting the main result of the paper, we apply the model to standard auctions. This illustrates how data-driven beliefs affect bidding behavior. An essential purpose of this Section is to illustrate that data-driven bidders reason as in winner's curse models in the presence of correlation, despite the private value character of the model.

To start with a simple case, we assume here that $|V| = 2$, so that the type of each bidder is one-dimensional. More specifically, we assume that $V = \{0, 1\}$, so that the type can be written as one number $\theta_i \in [0, 1]$, that specifies the probability that bidder i's ex-post valuation is $v_i = 1$. Note that this implies that θ_i is also the interim expected value of bidder i. In the following, we explain the equilibrium logic of our model for two standard auctions formats, the Second-Price Auction and the First-Price Auction. To compute concrete bidding equilibria, we will use a parametric class of joint distributions that allows us to vary the correlation between θ_1 and θ_2.

Example 1. The joint density is given by

$$f(\theta_1, \theta_2) = \frac{2 + \alpha}{2} \left(1 - |\theta_1 - \theta_2|\right)^{\alpha}.$$

The parameter $\alpha \in [0, \infty)$ determines the correlation between the two types where $\alpha = 0$ corresponds to the independent case and $\alpha = \infty$ corresponds to perfect correlation. Figure 1 depicts the joint density for $\alpha \in \{.1, 1, 10\}$

3.1 Second-Price Auction

In a second-price auction, the rational type has a weakly dominant strategy since we are in a private value setting. Hence she bids her interim expected value. We have

$$b^r(\theta_i) = \theta_i,$$

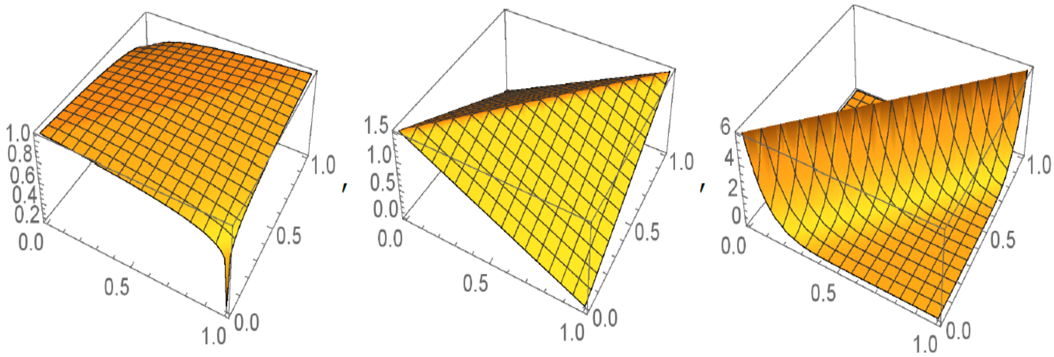

Figure 1: Joint density, $\alpha \in \{.1, 1, 10\}$ (left to right)

where b^r refers to the rational type's strategy. We denote the inverse by $\theta^r(b_i)$, which is of course equal to b_i in this case.

Now consider the misspecified type and consider a symmetric equilibrium, that is $b_i^m(\cdot) = b_j^m(\cdot) = b^m(\cdot)$. Suppose the equilibrium strategy $b^m(\cdot)$ is strictly increasing with inverse $\theta^m(b_i)$. In equilibrium, the cumulative distribution of b_j conditional on $v_i = 1$ is

$$H^{\text{SPA}}(b \mid v_i = 1) = \frac{\mathbb{P}_f\left[b_j \le b, v_i = 1\right]}{\mathbb{P}_f\left[v_i = 1\right]}, \tag{2}$$

where $H^{\text{SPA}}(\cdot)$ refers to this distribution for the SPA. Note that the misspecified type learns the correct joint distribution of v_i and b_j from the data. Hence we have used the correct probabilities $\mathbb{P}_f$ on the right-hand side. In the denominator, we have the unconditional probability of $v_i = 1$ which is given by the (ex-ante) expectation of the random variable $\tilde{\theta}_i$. In the numerator, the probability $\mathbb{P}_f[b_j \le b, v_i = 1]$ is obtained by averaging $\mathbb{P}_f[b_j \le b, v_i = 1|\tilde{\theta}_i]$ over the (ex-ante) random variable $\tilde{\theta}_i$. Since b_j is a function of θ_j and s_j, and the *generalized type* (θ_j, s_j) and v_i are independent *conditional on* $\tilde{\theta}_i$, we have:

$$\begin{aligned} H^{\text{SPA}}(b \mid v_i = 1) &= \frac{\mathbb{E}_{\tilde{\theta}_i}\left[\mathbb{P}_f[b_j \leq b|\tilde{\theta}_i] \times \mathbb{P}_f[v_i = 1|\tilde{\theta}_i]\right]}{\mathbb{E}[\tilde{\theta}_i]} \\ &= \frac{\mathbb{E}_{\tilde{\theta}_i}\left[\left(\lambda\mathbb{P}_f[b^r(\theta_j) \leq b|\tilde{\theta}_i] + (1-\lambda)\mathbb{P}_f[b^m(\theta_j) \leq b|\tilde{\theta}_i]\right) \times \mathbb{P}_f[v_i = 1|\tilde{\theta}_i]\right]}{\mathbb{E}[\tilde{\theta}_i]} \\ &= \frac{1}{\mathbb{E}[\tilde{\theta}_i]}\int_0^1 \left[\lambda F(b \mid \tilde{\theta}_i) + (1-\lambda)F(\theta^m(b) \mid \tilde{\theta}_i)\right] \tilde{\theta}_i\, f(\tilde{\theta}_i)d\tilde{\theta}_i. \end{aligned}$$

In the second line we decomposed the probability $\mathbb{P}_f[b_j \leq b|\tilde{\theta}_i]$ into the probability that a rational and a misspecified type bid below b, conditional on $\tilde{\theta}_i$. If the opponent is rational, the probability of $b_j \leq b$ is given by $\mathbb{P}_f[b^r(\theta_j) \leq b|\tilde{\theta}_i] = F(\theta^r(b) \mid \tilde{\theta}_i) = F(b \mid \tilde{\theta}_i)$, and if the opponent is misspecified it is given by $\mathbb{P}_f[b^m(\theta_j) \leq b|\tilde{\theta}_i] = F(\theta^m(b) \mid \tilde{\theta}_i)$. The term $\tilde{\theta}_i$ in the third line is just $\mathbb{P}_f[v_i = 1|\tilde{\theta}_i]$. We obtain a similar expression for the distribution of b_j conditional on $v_i = 0$:

$$H^{\text{SPA}}(b \mid v_i = 0) = \frac{1}{\mathbb{E}[1-\tilde{\theta}_i]}\int_0^1 \left[\lambda F(b \mid \tilde{\theta}_i) + (1-\lambda)F(\theta^m(b) \mid \tilde{\theta}_i)\right] (1-\tilde{\theta}_i)\ f(\tilde{\theta}_i)d\tilde{\theta}_i$$

where the expectation in the integral differs from that in $H^{\text{SPA}}(b \mid v_i = 1)$ since $P_f[v_i = 0|\tilde{\theta}_i] = (1-\tilde{\theta}_i)$, and outside the integral $\mathbb{E}[1-\tilde{\theta}_i]$ is the unconditional probability $P_f[v_i = 0]$.[15]

In a symmetric equilibrium of the second-price auction, the misspecified type's bid for θ_i solves

$$\max_b \left\{\theta_i H^{\text{SPA}}(b \mid v_i = 1) - \theta_i \int_0^b x dH^{\text{SPA}}(x|v_i = 1) - (1-\theta_i)\int_0^b x dH^{\text{SPA}}(x|v_i = 0)\right\}$$

To obtain an equilibrium we have to determine a bidding strategy b^m and the implied H^{SPA} such that b^m is optimal for the misspecified type give belief H^{SPA}. Taking the first-order condition for b and substituting $H^{\text{SPA}}(b \mid v_i = 1)$ and $H^{\text{SPA}}(b \mid v_i = 0)$, we obtain a differential equation for b^m.

In Example 1, when $\alpha = 0$—the independent case—we have that $H^{\text{SPA}}(b \mid v_i = 1) = H^{\text{SPA}}(b \mid v_i = 0)$ and the first order condition leads to $b^m(\theta_i) = \theta_i$. But, when α differs

[15] Note that $H^{\text{SPA}}(b \mid v_i = 0) = \mathbb{P}_f\left[b_j \leq b, v_i = 0\right] / \mathbb{P}_f\left[v_i = 0\right]$

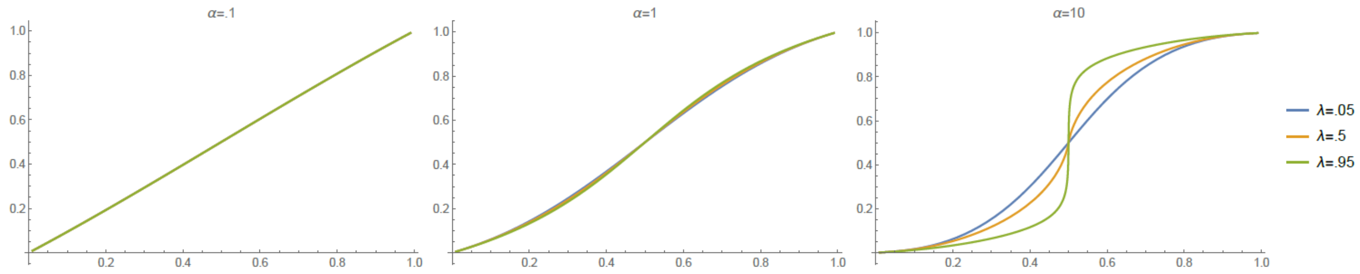

Figure 2: SPA bid-function $b_m(\theta_i)$, $\alpha \in \{.1, 1, 10\}$ (left to right)

from 0, $H^{\text{SPA}}(b \mid v_i = 1)$ differs from $H^{\text{SPA}}(b \mid v_i = 0)$ and $b^m(\theta_i)$ differs from θ_i. Solving the differential equation numerically for the joint distribution from Example 1, we get the bid-functions illustrated in Figure 2 for various values of α.

We see that increasing the correlation leads to stronger deviations from the rational bid. Moreover, the sensitivity of b_m with respect to λ becomes stronger if the correlation is stronger. Generally, for fixed correlation, increasing the share of misspecified types $(1-\lambda)$ leads to smaller deviation from rationality (even if it does not disappear, even as $\lambda = 0$). Bidding against mainly rational types, a misspecified type's behavior exhibits strong deviations from rationality,[16] but in equilibrium, the presence of other misspecified types has a dampening effect.

Intuition. The reasoning leading to the derivation of b^m follows a logic similar to that in classic analysis of winner's curse models (see Milgrom and Weber, 1982). We observe from Figure 2 that the misspecified type overbids for $\theta_i > 1/2$ and underbids for $\theta_i < 1/2$. What explains this behavior? To understand this, it is useful to shut down the (dampening) equilibrium effect of misspecified types and assume that $\lambda \approx 1$. The crucial observation is that the m-type believes that conditional on $v_i = 1$, the opponent's bid distribution is strong. This is because in the data, v_i and b_j are positively correlated: Observations with

[16]Numerical computations indicate that even if $\lambda \to 1$, the slope of b^m remains bounded, where the bound depends on α. In other words, b^m does *not* converge to a step function according to the numerical results.

$v_i = 1$ are more likely generated when $\tilde{\theta}_i$ is high. Due to the positive correlation between θ_i and θ_j, this implies that b_j is also likely to be high. Conversely, the m-type believes that conditional on $v_i = 0$, the opponent's bid distribution is weak.

For an m-type with low θ_i, consider the incentives to decrease the bid below $b = \theta_i$. In this range reducing the bid has a large effect on the winning probability conditional on $v_i = 0$ (the m-type believes that conditional on $v_i = 0$, the opponents bid's are concentrated on a low range) and little effect on the winning probability conditional on $v_i = 1$ (where the m-type believes the opponents bid's are concentrated on a high range). Therefore, the m-type believes that by shading the bid, she can cut the losses from winning with $v_i = 0$, without a strong reduction of the gains from winning when $v_i = 1$.

For a high θ_i, this logic is reversed. Consider the incentives to increase the bid above $b = \theta_i$ when θ_i is high. The bid is now in a range where the m-type believes that increasing the bid mainly effects the winning probability conditional on $v_i = 1$ and has little effect on the winning probability conditional on $v_i = 0$. Hence, she thinks overbidding increases the profits from winning with $v_i = 1$, while only modestly increasing the losses from winning with $v_i = 0$. This leads to bids above θ_i for high types of the misspecified bidder.[17]

Inefficiency of the Second-Price Auction. While the distortions observed in the example are specific to the parametric class of distributions, we can show generally that the SPA is not efficient whenever both rational and misspecified types arise with positive probability, and the types of the two bidders are correlated.[18,19]

Proposition 1. *If $\lambda \in (0, 1)$ and $Corr[\theta_1, \theta_2] \neq 0$, then any data-driven equilibrium of the second-price auction in which the rational types of both bidder play their dominant strategies is inefficient.*

Proof. All omitted proofs can be found in Appendix A. □

[17] The dampening effect of lower values of λ can be understood as follows. Take a value of θ_i larger (resp. smaller) than 0.5. Rational bidders bid less (resp. more) than misspecified bidders. Thus, bidder j ties with the equilibrium bid of a misspecified agent, for a larger (resp smaller) value of θ_j when bidder j is rational than when he is misspecified. Given the correlation between θ_i and θ_j, this in turn gives rise to a bigger winner's curse-like correction when λ is bigger, thereby explaining the dampening effect of decreasing the share of rational types.

[18] Correlation is a sufficient condition for an inefficiency. The careful reader will see from the proof that weaker forms of dependency also lead to inefficiencies. In Section 4 we generalize this proposition to any finite number of valuations (see Lemma 6).

[19] As suggested by Bob Wilson, inefficiencies require the presence of both rational and misspecified types due to our assumed symmetry on the distribution of types. In the absence of symmetry, one would expect inefficiencies to arise in SPA, even if there are no rational types, as long as signals are correlated.

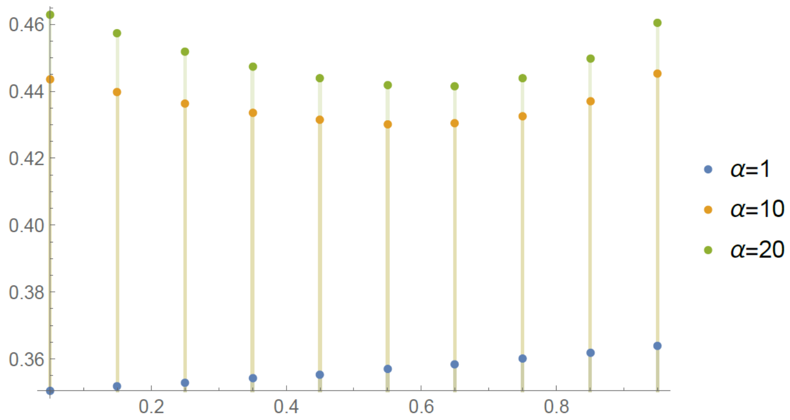

Figure 3: Revenue from the SPA as a function of λ: for $\alpha \in \{1, 10, 20\}$

Revenue and Efficiency. Continuing our illustration for the parametric class in Example 1, we show how revenue and (relative) efficiency of the allocation varies with (a) the share of rational types λ and (b) the correlation between θ_1 and θ_2—that is the parameter α.

Figure 3 plots the revenue as a function of λ for different values of α. Note that the comparison between different values of α with λ held fixed is not very informative since the joint distribution changes in a complicated way as α changes.

We see that for the case of weak correlation ($\alpha = 1$), revenue is increasing in the share of rational type. This suggests that the distortions in the misspecified type's bidding function adversely affect revenue. For highly correlated interim types, the pattern changes and revenue is U-shaped in the share of rational types. The initial decline is intuitive since the distortions in the m-types bid become larger if the share of rational types increases. Profits rise again if the share of rational types becomes so large that presence of m-types becomes unlikely.

Figure 4 shows how efficiency changes depending on λ and α.

To make this comparable across different parameter sets, we normalize efficiency by the expected ex-post value achieved if the object is always allocated to the bidder with the highest interim type. Clearly when $\lambda = 0$ or 1, there is no inefficiency given that bidders of the same sophistication bid in the same way. Moreover, both when $\alpha = 0$ (the independent case) or $\alpha = \infty$ (perfect correlation) there is no inefficiency either. In the parametric example, we observe that the relative efficiency is U-shaped as a function of λ

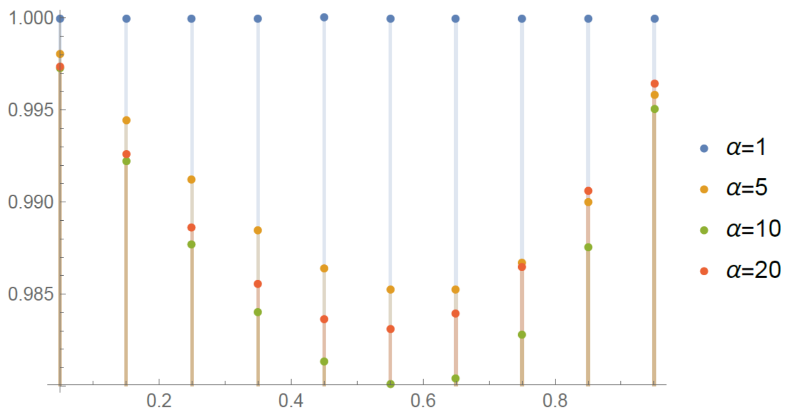

Figure 4: Efficiency of SPA as a function of λ: for $\alpha \in \{1, 5, 10, 20\}$

and α, as shown in Figure 5.

3.2 First-Price auction

In a first price auction, we obtain the misspecified type's belief in a similar way as for the second price auction:

$$H^{\text{FPA}}(b \mid v_i = 1) = \int_0^1 \left[\lambda F(\theta^r(b) \mid \tilde{\theta}_i) + (1-\lambda) F(\theta^m(b) \mid \tilde{\theta}_i)\right] \tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i,$$
$$H^{\text{FPA}}(b \mid v_i = 0) = \int_0^1 \left[\lambda F(\theta^r(b) \mid \tilde{\theta}_i) + (1-\lambda) F(\theta^m(b) \mid \tilde{\theta}_i)\right] (1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\tilde{\theta}_i.$$

$b^r(\cdot)$ and $b^m(\cdot)$ now denote the bidding strategies of the rational and misspecified types in the symmetric equilibrium of the FPA, and their inverses are denoted by $\theta^r(\cdot)$ and $\theta^m(\cdot)$. The misspecified bidder's bid for type θ_i maximizes

$$\max_b \, (1-b)\, \theta_i H^{\text{FPA}}(b \mid v_i = 1) - b(1-\theta_i) H^{\text{FPA}}(b|v_i = 0). \tag{3}$$

Again we obtain a differential equation for $b^m(\theta_i)$. In contrast to the second price auction, however, we cannot assume that rational bidders bid their expected valuations. Instead they maximize

$$\max_b \, (\theta_i - b)\, (\lambda F(\theta^r(b)|\theta_i) + (1-\lambda) F(\theta^m(b)|\theta_i))$$

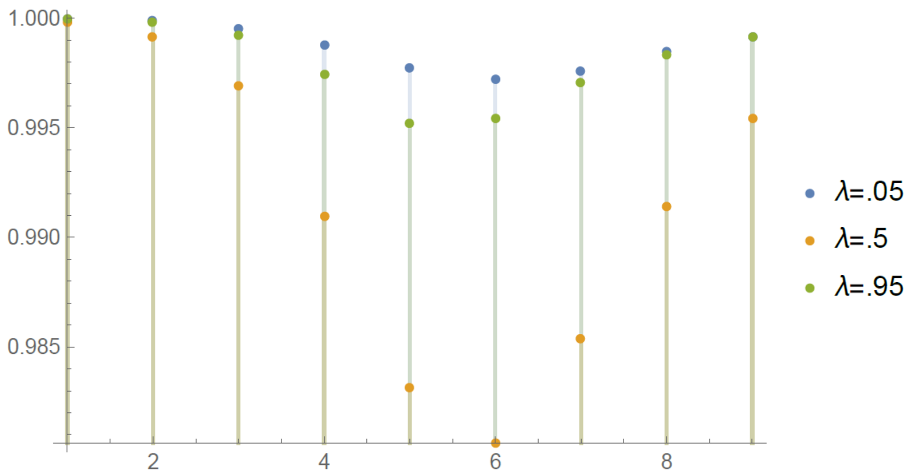

Figure 5: Efficiency of SPA as a function of $\alpha = 1/5 + 5^k$ where $k = 1, \ldots, 9$ is on the horizontal axis. $\lambda \in \{.05, .5, .95\}$.

This optimization problem reflects the complete awareness of the model by rational bidders. They use the correct distribution f, the share of rational types in the population, and the equilibrium bidding strategies of both the rational and the misspecified types when determining their optimal bids. The first-order condition for the rational type's problem yields a second differential equation. To compute an equilibrium, we need to solve the system of two ODEs with the boundary condition $(b^m(0), b^r(0)) = (0, 0)$. This proves challenging even numerically for the distributions in our example, since the system has a singular point at the boundary condition. However, we obtain a similar inefficiency result as we had for the SPA.

Proposition 2. *If $\lambda \in (0,1)$ and $Corr[\theta_1, \theta_2] \neq 0$, then the symmetric equilibrium of the first-price auction is inefficient.*

3.3 Comparison

We can compute the bidding equilibrium for both auction formats for the case of only rational bidders ($\lambda = 1$) and only misspecified bidders $\lambda = 0$. Figure 6 shows the bid functions b_k^s where $k = 1, 2$ denotes first-price or second-price auctions and $s = m, r$ denotes the misspecified or rational type.

To illustrate the role of correlation, the functions are shown for $\alpha \in \{1, 5, 10\}$. Comparing FPA and SPA in the rational case, we see the familiar revenue ranking that the

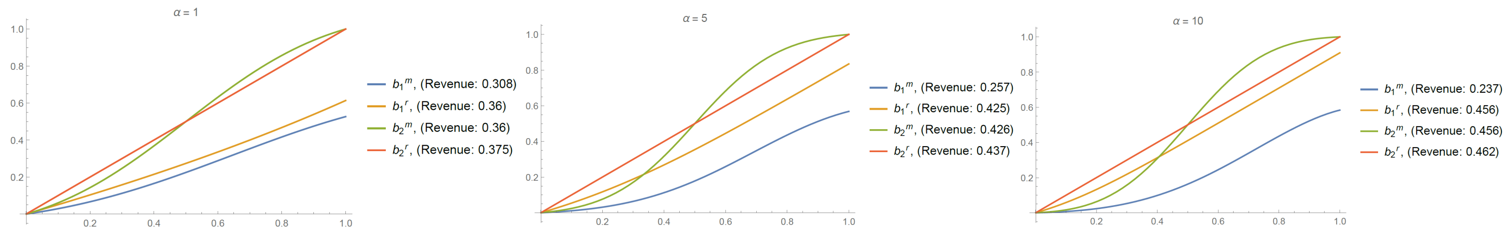

Figure 6: Equilibrium bid functions for $\alpha \in \{1, 5, 10\}$

SPA yields higher revenue than the FPA with correlated types. This revenue ranking is preserved in the case of misspecified bidders. Interestingly, with misspecified bidders, the gap between SPA and FPA becomes more pronounced if values are more correlated. This conforms well with the intuition for the distortions in the bid function: In the SPA low types underbid and high types overbid. In the FPA, the same forces lead the low types to underbid. But this allows the higher types to shade their bids more and the incentive to overbid does not compensate for this force. This leads to much lower bids for misspecified types compared to the rational equilibrium if the correlation is high.

Finally, we want to compare the efficiency of the SPA and FPA. This comparison is not interesting in the purely rational or purely misspecified cases since the symmetric equilibrium implies that both auction formats are fully efficient. A comparison in the mixed case is challenging because we are not able to compute the equilibrium in the FPA. To make progress we consider the best response of a misspecified type to the purely rational equilibrium. This allows us to show how efficiency changes if we inject a small share of misspecified types in a rational population. Figure 7 shows the resulting bidding strategies for $\alpha = 1.5$.

To compare the efficiency we numerically compute how much efficiency is lost in expectation if bidder one uses the purely rational strategy and bidder two uses the misspecified response. This number gives the rate at which efficiency decreases if we decrease λ slightly from $\lambda = 1$. In the example depicted in Figure 7 we have a marginal loss of .0035 for the FPA and .0088 for the SPA. This means that the SPA is less efficient than the FPA.

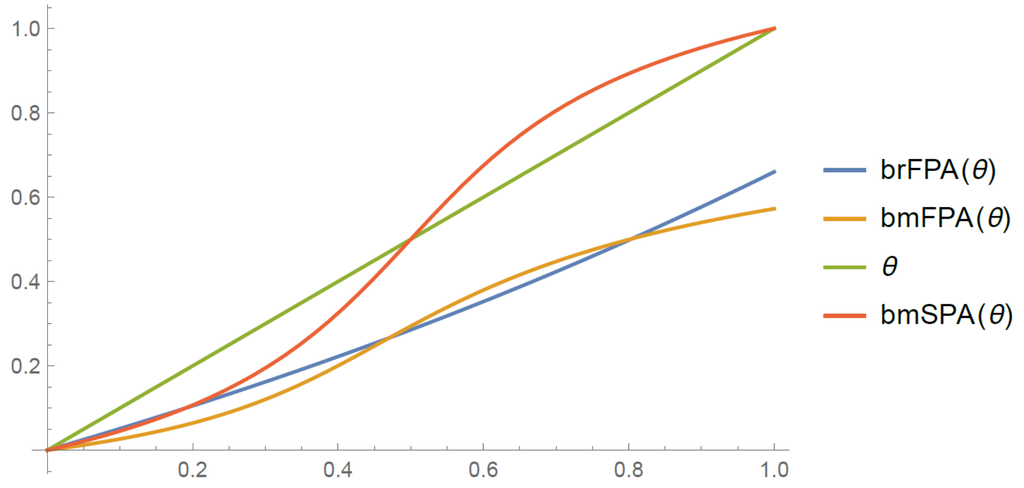

Figure 7: For the FPA ($k = 1$), and the SPA ($k = 2$), $b_k^r(\theta_i)$ is the rational strategy if $\lambda = 1$; $b_k^m(\theta_i)$ is the best response of misspecified type to data generated by the purely rational equilibrium. ($\alpha = 1.5$).

4 Auction-like Mechanisms

We now consider the possibility of implementing an efficient allocation in the presence of both rational and misspecified buyers in the general case. We consider a class of auction-like mechanisms, in which bidders can place a one-dimensional bid $b \in B \subset \mathbb{R}$, and which allocate to the highest bid (possibly adjusted by a bonus or malus). We assume that bidders may choose not to participate in a mechanism in which case their utility is zero.

Definition 2. An *auction-like mechanism* is given by $M = \left[B, (W_i)_{i=1,2}, (L_i)_{i=1,2}, \phi_1\right]$. $B = [\underline{b}, \overline{b}]$ is the set of *feasible bids*. The *allocation rule* $\phi_1 : B \to B$ is a strictly increasing function. The object is allocated to bidder 1 if $b_1 > \phi_1(b_2)$, to bidder two if $b_1 < \phi_1(b_2)$, and with probability 1/2 if $b_1 = \phi_1(b_2)$. We denote the inverse by $\phi_2 = \phi_1^{-1}$. The *payment rules* are $W_i : B \times B \to \mathbb{R}_0^+$, and $L_i : B \times B \to \mathbb{R}_0^+$, which specify the payment bidder i has to make as a function of the bids, if she wins or loses, respectively. We assume that for each $i \in \{1, 2\}$, both functions W_i, L_i are weakly increasing in bidder i's own bid.[20]
An auction-like mechanism is *smooth* if for $i \in \{1, 2\}$, ϕ_i, W_i, and L_i, are continuously differentiable with derivatives that can be continuously extended to the boundary of B.

The smoothness assumption is made for tractability. Many common auction formats

[20] To avoid introducing a participation decision, we assume that when $b_i = 0$, W_i and L_i are 0 so that bidding 0 guarantees a non-negative payoff.

are smooth auction-like mechanisms. Our main result is that if there are at least three possible ex-post valuations, then for generic type distributions, no smooth auction-like mechanism exists that has an efficient data-driven equilibrium.

To make this precise, we reformulate the types of agents. We denote the *interim valuation* of bidder i with type θ_i by

$$w_i(\theta_i) := \mathbb{E}[v_i|\theta_i].$$

Given the normalization $0 = v^1 < \ldots < v^K = 1$, we have $w_i(\theta_i) \in [0,1]$. For each $w_i \in [0,1]$, we denote the set of types θ_i that have interim valuation w_i by

$$\Theta_i(w_i) := \{\theta_i \in \Theta_i | \mathbb{E}[v_i|\theta_i] = w_i\}.$$

For $w_i \in \{0,1\}$ this set is a singleton; and for all $w_i \in (0,1)$, there exists a homeomorphism $x_i(\cdot; w_i) : \Theta_i(w_i) \to [0,1]^{K-2}$, where $K = |V|$ is the number of ex-post valuations. We can therefore write the type of bidder i as $(w_i, x_i) \in [0,1]^{K-1}$. While w_i is the payoff-relevant part of the type, for $w_i \in (0,1)$, x_i can be used to recover the belief $f(\theta_j|x_i^{-1}(x_i; w_i))$ about bidder j's type. Abusing notation we use $f(w_1, x_1, w_2, x_2)$ to denote the joint density of the buyers' types and assume that this density is strictly positive.

Our main result is that for generic distributions, smooth auction-like mechanisms do not have efficient data-driven equilibria. To state this formally, we let $\mathcal{M}_+^d([0,1]^{2K-2})$ be the set of probability measures on $[0,1]^{2K-2}$ that admit continuous and strictly positive densities $f(w_1, x_1, w_2, x_2)$. We endow $\mathcal{M}_+^d([0,1]^{2K-2})$ with the uniform topology for densities. For given V and λ, let $\mathcal{I}(V, \lambda) \subset \mathcal{M}_+^d([0,1]^{2K-2})$ be the set of prior distributions for which all data-driven equilibria of any smooth auction-like mechanism are inefficient.

Theorem 1. *Suppose $K = |V| \geq 3$ and $\lambda \in (0,1)$. Then for generic type distributions, there exists no smooth auction-like mechanism with an efficient data-driven equilibrium. Formally, $\mathcal{I}(V, \lambda)$ is a residual subset of $\mathcal{M}_+^d([0,1]^{2K-2})$, that is, it contains a countable intersection of open and dense subsets of $\mathcal{M}_+^d([0,1]^{2K-2})$.*

The notion of genericity used here is the same as in Gizatulina and Hellwig (2017), who show the genericity of full surplus extraction. The key step in the proof is to show that in the presence of rational bidders, efficiency requires that the mechanism is (strategically equivalent to) a second-price auction. The reason is that to achieve efficiency, the bid in an auction-like mechanism must be a function of w_i only. If there are more than two ex-post valuations, for each $w_i \in (0,1)$, the set $\Theta(w_i)$ is a manifold of dimension $K-2 \geq 1$, and all types in $\Theta(w_i)$ have identical interim expected valuations but different beliefs. We show that for generic distributions, the requirement that the bid is independent of the rational type's belief, implies that the mechanism must be (equivalent to) a second-price auction.[21] We then complete the proof by extending the result of Proposition 1 to more than two ex-post valuations (see Lemma 6 below), showing that in a second-price auction the misspecified type does not bid truthfully. This in turn rules out the possibility of an efficient data-driven equilibrium.

Remark 1 (Two ex-post valuations). With only two ex-post valuations ($K = 2$), our proof does not apply. While the analysis of standard auctions in Section 3 suggests that bid functions of rational and misspecified types in auction-like mechanisms differ, it is an open question whether variations of auction-like mechanisms offer enough flexibility (through the choice of the payment rule) so that types of both sophistication can be incentivized to use an identical bid function when $K = 2$.

Remark 2 (More than two bidders). The restriction to two bidders has been made for simplicity. With more than two bidders, we can consider misspecified types who have access to data from past auctions with observations of the form $(b_1, v_1, \ldots, b_N, v_N)$, where N is the number of bidders. Such bidders will now rely on $h(b_{-i}|v_i)$, the pdf of $b_{-i} = (b_j)_{j\neq i}$ conditional on v_i, to form their beliefs about how variables of interest are distributed. We can define auction-like mechanisms that award the object to the highest bidder and specify payments as a function of all bids. We conjecture that the key argument in our proof—namely that efficiency requires the use of a second-price auction also works with more than two bidders, as long as there are at least three ex-post valuations. Moreover, an analogous result to Proposition 1 and Lemma 6 implies that misspecified types do not use the rational bid function in any equilibrium of the second-price auction.

4.1 Proof of Theorem 1

Regular Equilibria of Simple Mechanisms. First, we show that it suffices to consider *regular equilibria* of *simple mechanisms*. We call a smooth auction-like mechanism *simple* if it is of the form $M = [[0,1], (W_i), (L_i), Id]$, where $\phi = Id$ denotes the identity so that the

[21] The intuition for this is that in any auction in which the payment of the winner would depend non-trivially on the winner's own bid, the optimal equilibrium bid would require some shading that depends non-trivially on the belief, as in the first-price auction. To ensure that the shading is the same for all beliefs as generated by variations of x_i, a second-price auction must be used.

allocation rule is symmetric. We call an equilibrium *regular* if it is symmetric and the bid of each generalized type (w_i, x_i, s_i) is given by a continuous and strictly increasing function $b(w_i)$ with range $b([0,1]) = [0,1]$. In other words, the bid only depends on the interim valuations, but not on the identity, sophistication, or belief x_i, of the bidder. Note that a regular equilibrium of a simple mechanism is efficient. We denote the strictly increasing and continuous inverse of $b(\cdot)$ by $\psi : [0,1] \to [0,1]$.

Lemma 1. *Let $\tilde{M} = [\tilde{B}, (\tilde{W}_i), (\tilde{L}_i), \tilde{\phi}_1]$ be a smooth auction-like mechanism with an efficient data-driven equilibrium $(\tilde{b}_1(w_1, x_1, s_1), \tilde{b}_2(w_2, x_2, s_2))$. Then there exists a simple mechanism $M = [[0,1], (W_i), (L_i), Id]$, with a regular (and hence efficient) data-driven equilibrium.*

Proof. The proofs of all Lemmas can be found in the Appendix. □

In light of Lemma 1, it suffices to consider regular equilibria of simple mechanisms. The intuition behind this result is that in an efficient mechanism with a symmetric allocation rule,[22] all bidders must use the same bids as a function of their interim valuation. The proof shows that mechanisms for which the bidding function has discontinuities, these jumps can be removed in a way that preserves the smoothness of the payment rules. Lemma 1 bears some similarities with but falls short of the revelation principle because the full revelation argument may not preserve the smoothness of the payment rules if the equilibrium of the original mechanism is non-smooth.

Second-Price Auctions. Next we derive a condition on the payment rules and equilibrium bid function that characterizes regular equilibria of the second-price auction. We denote the equilibrium difference in utility between winning and losing of a bidder with bid $b = b(w_i)$, whose bid is tied with the opponent by

$$\delta_i(b) = \psi(b) - (W_i(b,b) - L_i(b,b)) .$$

In a regular equilibrium of the SPA, the rational type bids truthfully ($b(w) = w$), and the payment rules satisfy $W_i(b,b) = b$ and $L_i \equiv 0$, so that $\delta_i(b) = 0$ for all $b \in [0,1]$. The following Lemma shows the converse result.

[22] Clearly, a mechanism with an asymmetric allocation rule can be made symmetric by a simple monotonic transformation.

Lemma 2. *Consider a simple mechanism $M = [[0,1], (W_i), (L_i), Id]$ with a regular data-driven equilibrium. If $\delta_i(b) = 0$, for $i \in \{1,2\}$ and all $b \in [0,1]$, then M is a second-price auction—that is, for all $i \in \{1,2\}$, $L_i(b_i, b_j) = 0$ for all $b_i \leq b_j$ and $W_i(b_i, b_j(w_j)) = w_j$ whenever $b_i \geq b_j(w_j)$.*

Differentiability of the Bidding Strategy. To show that $\delta_i(b) = 0$ for all bids, we derive an implication of $\delta_i(b) > 0$ and show that it is violated generically. In the derivations we will use first-order conditions. The following Lemma shows that the inverse of the bid function, $\psi(b)$ is differentiable if $\delta_i(b) > 0$. The Lemma is based on the proof of Lemma 7 in Persico and Lizzeri (2000).

Lemma 3. *If $\delta_i(b_0) > 0$ for some $b_0 \in [0,1]$, then there exist a non-empty interval $(\alpha, \beta) \subset [0,1]$, with $b_0 \in (\alpha, \beta)$, such that ψ is continuously differentiable on (α, β), and $\psi'(b) > 0$ and $\delta_i(b) > 0$ for all $b \in (\alpha, \beta)$.*

For generic distributions, efficiency requires $M = SPA$. Next, we show that $\delta_i(b) > 0$ implies that a condition similar to the full-surplus extraction condition (McAfee and Reny, 1992) must be violated, and prove results analogous to Gizatulina and Hellwig (2017), to show that for generic prior densities $f(w_1, x_1, w_2, x_2)$, we must have $\delta_i(b) = 0$ for all $b \in [0,1]$, $i \in \{1,2\}$, and any regular data-driven equilibrium of a simple mechanism.

We begin by deriving an implication of $\delta_i(b) > 0$. Fix $b \in (0,1)$ such that $\delta_i(b) > 0$ and consider a rational bidder i with type (w_i, x_i), where $w_i = \psi(b)$ and $x_i \in X$ is arbitrary. In a regular equilibrium, this type maximizes (where we use $j \neq i$ to denote the opponent):

$$\max_{b' \in [0,1]} \int_0^{\psi(b')} \left(\psi(b) - W_i(b', b(w_j))\right) f(w_j|\psi(b), x_i) dw_j - \int_{\psi(b')}^1 L_i(b', b(w_j)) f(w_j|\psi(b), x_i) dw_j$$

Given Lemma 3, we can differentiate the objective function with respect to b', and obtain the first-order condition, which must hold for $b' = b$:

$$f(\tilde{w}_j = \psi(b)|\tilde{w}_i = \psi(b), x_i) = \int_0^1 \frac{\partial P_i(b, b(w_j))/\partial b_i}{\delta_i(b)\psi'(b)} f(w_j|\tilde{w}_i = \psi(b), x_i) dw_j, \tag{4}$$

where we simplify notation by denoting the payment of bidder i as follows

$$P_i(b_i, b_j) := \mathbf{1}_{\{b_i > b_j\}} W_i(b_i, b_j) + \mathbf{1}_{\{b_i < b_j\}} L_i(b_i, b_j).$$

Multiplying (4) by $f(\tilde{w}_i = \psi(b), x_i)/f(\tilde{w}_i = \tilde{w}_j = \psi(b))$, and using

$$f(x_i|\tilde{w}_i = \tilde{w}_j = \psi(b))f(\tilde{w}_i = \tilde{w}_j = \psi(b)) = f(x_i, \tilde{w}_i = \tilde{w}_j = \psi(b)),$$

we obtain for all $x_i \in X_i$:

$$f(x_i|\tilde{w}_i = \tilde{w}_j = \psi(b)) = \int_0^1 m(b, \psi(b), w_j) f(x_i|\tilde{w}_i = \psi(b), w_j) dw_j, \tag{5}$$

where

$$m(b, \psi(b), w_j) = \frac{\partial P_i(b, b(w_j))/\partial b_i f(\tilde{w}_i = \psi(b), w_j)}{\delta_i(b)\psi'(b) f(\tilde{w}_i = \tilde{w}_j = \psi(b))}.$$

Since we consider a simple mechanism and prior densities $f \in \mathcal{M}_+^d([0,1]^{2K-2})$, and $\psi'(b) > 0$, the term $m(b, \psi(b), w_j)$ is finite and non-negative. For fixed b, $m(b, \psi(b), \cdot)$ is in fact a probability density on $[0,1]$.[23]

Condition (5) states that the density $f(\cdot|\tilde{w}_i = \tilde{w}_j = \psi(b))$ can be expressed as a positive linear combination of the densities $f(\cdot|\tilde{w}_i = \psi(b), w_j)$ for $w_j \in [0,1]$, with positive weights on $w_j \neq \psi(b)$. By virtually the same proof as for Theorem 2.4 in GH17, we can show that for generic distributions (5) is violated.

To state the result we need several definitions that mimic GH17. Let $\mathcal{M}_+^d(X)$ be the set of absolutely continuous probabilities measures on X with strictly positive and continuous densities, endowed with the topology induced by the sup-norm for density functions on X; let $\mathcal{C}([0,1], \mathcal{M}_+^d(X))$ be the set of continuous mappings from $[0,1]$ to $\mathcal{M}_+^d(X)$, endowed with the topology of uniform convergence; and let $\mathcal{M}([0,1])$ be the set of probability measures on $[0,1]$, endowed with a topology that is metrizable by a metric that is a convex function on $\mathcal{M}([0,1]) \times \mathcal{M}([0,1])$. Finally let $\mathcal{E}(w_i) \subset \mathcal{C}([0,1], \mathcal{M}_+^d(X))$ be the set of continuous mappings that map $w \in [0,1]$ to densities $g(\cdot|w) \in \mathcal{M}_+^d(X)$ that satisfy the following condition: For all $\mu \in \mathcal{M}([0,1])$:

$$g(x_i|w_i) = \int_0^1 g(x_i|w')\mu(dw'),\ \forall x_i \in X \qquad \Longrightarrow \qquad \mu = \delta_{w_i} \tag{6}$$

where $\delta_{w_i} \in \mathcal{M}([0,1])$ is the Dirac measure with a mass-point on w_i.

Lemma 4 (see Theorem 2.4 in Gizatulina and Hellwig, 2017)**.** *For any $w_i \in (0,1)$, the set $\mathcal{E}(w_i)$ is a residual subset of $\mathcal{C}([0,1], \mathcal{M}_+^d(X))$, that is, it is a countable intersection of*

[23] Integrating both sides of (5) over X we see that $\int_0^1 m(b, \psi(b), w_j) dw_j = 1$.

open and dense subsets of $\mathcal{C}([0,1], \mathcal{M}^d_+(X))$.

The implication of this Lemma is that for fixed $w_i \in (0,1)$, and generic functions $w_j \mapsto f(\cdot|w_i, w_j)$ that map w_j to conditional densities $f(\cdot|w_i, w_j)$, any simple mechanism with a regular data-driven equilibrium must satisfy $\delta_i(b(w_i)) = 0$.

This Lemma is insufficient for our purposes for two reasons. First, we need to show that *for generic priors*, the function that maps w_j to the conditional density $f(x_i|w_i, w_j)$ is an element of $\mathcal{E}(w_i)$, and second we need to show this *for all* w_i. To this end, for $i \in \{1,2\}$ let $\mathcal{W}_i$ be a countable and dense subset of $(0,1)$. We show that for generic prior densities $f(w_1, x_1, w_2, x_2)$, the mapping that maps $w_j \in [0,1]$ to the conditional density $f_i(\cdot|w_i, w_j)$ is an element of $\mathcal{E}_i(w_i)$ for all $w_i \in \mathcal{W}_i$ and all $i \in \{1,2\}$. For the following Lemma, recall that $\mathcal{M}^d_+([0,1]^{2K-2})$ denotes the set of priors with strictly positive and continuous densities.

Lemma 5 (see Theorem 2.7 in Gizatulina and Hellwig, 2017)**.** *For* $i \in \{1,2\}$*, let* $\mathcal{W}_i$ *be a countable and dense subset of* $(0,1)$*. Let* $\mathcal{F}$ *be the set of prior densities in* $\mathcal{M}^d_+([0,1]^{2K-2})$ *such that for all* $i \in \{1,2\}$ *and* $w_i \in \mathcal{W}_i$*, the mapping* $w_j \mapsto f(\cdot|w_i, w_j)$ *is an element of* $\mathcal{E}(w_i)$*. Then* $\mathcal{F}$ *is a residual subset of* $\mathcal{M}^d_+([0,1]^{2K})$*, that is it contains a countable intersection of open and dense subsets of* $\mathcal{M}^d_+([0,1]^{2K})$*.*

This Lemma implies that for generic prior densities $f(w_1, x_1, w_2, x_2)$, any regular equilibrium of a simple mechanism must satisfy $\delta_i(b(w_i)) = 0$ for all $w_i \in \mathcal{W}_i$. Since the functions $b(\cdot)$ and $\delta_i(\cdot)$ are continuous and $\mathcal{W}_i$ is dense, this implies $\delta_i(b) = 0$ for all $b \in [0,1]$. By Lemma 2, this implies that for generic distributions, if a simple mechanism has a regular equilibrium, then it must be the second-price auction.

Bidding Strategy of the Misspecified Type in the Second-Price Auction. So far we have made use of the rational type's first-order condition to show that efficiency cannot be achieved with an auction-like mechanism other than the SPA. To conclude the proof of Theorem 1 we show that for generic distributions, misspecified types do not use $b(w) = w$ in a SPA.

Lemma 6. *Let* $\lambda \in (0,1)$ *and suppose that* $\mathbb{E}_f\left[\theta_i^K | w_j \leq b\right] \neq \frac{\mathbb{E}_f[\theta_i^K]}{\mathbb{E}_f[\theta_i^1]}\mathbb{E}_f\left[\theta_i^1 | w_j \leq b\right]$ *for some* $i \in \{1,2\}$ *and* $b \in [0,1]$*. In any data-driven equilibrium of the second price auction where the rational types bid truthfully, some types* (θ_i, m_i) *place a bid that is different from their interim valuation.*

It is easy to see that the subset of prior densities for which there exists $i \in \{1,2\}$ and $b \in [0,1]$ such that $\mathbb{E}_f\left[\theta_i^K | w_j \leq b\right] \neq \frac{\mathbb{E}_f\left[\theta_i^K\right]}{\mathbb{E}_f\left[\theta_i^1\right]}\mathbb{E}_f\left[\theta_i^1 | w_j \leq b\right]$ is open and dense $\mathcal{M}_+^d([0,1]^{2K})$ so that its intersection with $\mathcal{F}$ is residual by Lemma 5. This concludes the proof of Theorem 1.

5 Discussion and Extensions

In this Section we discuss robustness checks and various extensions. In Subsection 5.1, we discuss alternative models of belief formation from observed data when past bids and ex-post valuations are observable. In Subsection 5.2 we review possible approaches one could take for situations in which data from past auctions only include the winning bid or the payment made in past auctions. In Subsection 5.3, we discuss the case in which past bids would be anonymous. In Subsection 5.4 we discuss general mechanisms that go beyond the auction-like class considered in Section 4. Finally, in Subsection 5.5 we discuss the possible implications of more general signal structures.

5.1 Model of Belief Formation from Observed Data

Our goal in this paper was to model belief formation of novice bidders (the m-types) who lack a complete understanding of the auction environment. Two basic assumptions have guided our modeling choices. First, we have assumed that novice bidders are sophisticated in the sense that they are able to use the empirical joint distribution of observable variables to inform their own bid. Second, we have assumed that novice bidders do not reason about how the bids of past bidders were formed. In particular they do not form a conjecture or model of the information available to past bidders and do not try to analyze how such information drives observed behavior.

Due to missing data about signals (or types) of past bidders, novice bidders are not able to learn the true joint distribution of signals/types, ex-post valuations and opponent's bids. At the same time, a novice bidder knows her own type θ_i, and has access to the empirical distribution of observable variables. She lacks knowledge how these two should be combined, and a priori, many different ways of using the data are conceivable, all of which rely on some implicit or explicit assumptions. Following our second basic assumption, novice bidders do not try to reason about how past bidders have determined their bids. Instead they simply combine the joint distribution of observable variables v_i and b_j with the belief about the distribution of v_i given by their type θ_i to evaluate the expected payoff of different bids. This leads to a misspecified model in which v_i and b_j are correlated even when conditioning on θ_i.

We believe that this simple way of using the data is a plausible model of an inexperienced bidder. But, there may be other ways to think about data-driven belief formation, and other ways of forming beliefs may lead to different misspecifications and deviations from rational behavior. However, it should be clear from the analysis in Section 4, that our main result is robust to the inclusion of various other forms of misspecifications. In fact the presence of various types who differ in their misspecification will make it harder to achieve efficiency since a mechanism would have to account for all types, a task that is already impossible in the presence for rational and a single misspecified type.

One aspect in which bidders could differ is the *precision* of the recorded data. For example, some bidders may only record past bids in broad categories, such as "high" or "low" which would indicate that the bid was above or below some threshold b^*. A bidder with such a coarse record of data would have to make some assumption about the non-identified distribution of the bid on the intervals $[0, b^*]$ and $[b^*, 1]$. A simple starting point could be the uniform distribution but bidders may also make a different assumption. Given such an assumption, our notion of equilibrium can be adapted. While an analysis is beyond the scope of this paper, we conjecture that it would lead to similar conclusions as our present model.

Alternatively, one could consider different degrees of sophistication of novice bidders. We assume that novice bidders do not try to reason about how past bidders determined their bids, nor about the possible joint distribution of v_i, b_j and θ_i. Novice bidders are sophisticated in their ability to analyze data, but boundedly rational in the sense that they do not question their method even though there might be alternative ways of using past data to perform bids.

Thinking about more sophisticated types, we may ask what additional knowledge they would need to have in order to see that their model is misspecified. In the data, one can see that v_i and b_j are independent conditional on b_i since bids are a function of θ_i. However, without further assumptions on how past bids were formed, this does not allow to conclude that v_i and b_j are independent conditional on θ_i. Hence, given the available data, it is not obvious to an observer or the novice bidder, that the m-type in our model uses a misspecified model.[24] A "more sophisticated" type would therefore need to have

[24] The case of two possible ex-post valuations is special. Here, a more sophisticated m-type might make

access to more data that would allow to validate structural assumptions about past bids.

Conversely, we may think of less sophisticated types who have access to the same data as our m-type but do not attempt or are unable to use the statistical link between the bids b_j and the ex post values v_i. For example this could be bidders who are not able to analyze large data sets beyond producing marginal distributions of the opponents bids. Alternatively, the bidder may know her expected valuation $E[v_i|\theta_i]$ but not the full distribution θ_i over ex-post valuations. Such bidders may in some cases actually display less bias in their bidding behavior since they do not use the statistical link between the b_j and v_i that gives rise to a (perceived) conditional correlation. For example, this is the case in second price auctions in which such bidders would bid optimally, in contrast to the m-type bidders we consider.

5.2 Non-observability of losing bids and valuations

In this paper we have assumed that data on past beliefs are not accessible but bidders have full access to past bids and past ex-post valuations. In practice, the ex-post valuations may not be observed precisely, and perhaps only noisy signals of the true valuation are available. To model such a situation, one could formulate the type θ_i as a distribution over such signal realizations and proceed as before.

More importantly, auctioneers may not disclose all bids in an auction. In the following, we discuss possible approaches for the case that only the *winning bid* and the identity of the winner is observable (as well as ex-post valuations of both bidders). Alternatively, the auctioneer may disclose the identity of the winner and the *payment* she has to make. In a first-price auction, this is equivalent to disclosing the winning bid, but in an ascending auction or second-price auction, the payment is equal to the second highest bid and the following discussion has to be modified accordingly.

Consider the case that after an auction with $b_i > b_j$, the data point (i, b_i, v_i, v_j) is observed, so that data about losing bids is not available. In the present paper, we have taken $H(b_j|v_i)$ to be the empirical distribution of the opponent's bid b_j conditional on valuation v_i for bidder i. If losing bids are not observed, this distribution is not directly

the plausible assumptions that (a) past bidders also had a one-dimensional type θ_i and (b) bids are a strictly increasing function of θ_i. Based on these assumptions, the m-type could conclude from the data that v_i and b_j are independent conditional on θ_i, leading her to behave like the rational type. Note however, that with more than two possible ex-post valuations, the bidding strategy cannot be injective, and therefore, without further assumptions about bidding behavior, the m-type cannot conclude from the data that v_i and b_j are independent conditional on θ_i.

accessible. We outline three exemplary models of how a bidder may construct $H(b_j|v_i)$ that reflect different degrees of sophistication. For each approach, the constructed $H(b_j|v_i)$ can be plugged into our equilibrium framework and the analysis would proceed as before.

A *naive bidder* may ignore that the observations about opponent's bids b_j for a given valuation v_i is selected and use the (observable) distribution $H(b_j|v_i, b_j > b_i)$ instead of $H(b_j|v_i)$. This approach will lead bidder i to think that bidder j bids higher than in reality, which creates additional bias.[25]

A *semi-naive bidder* may be aware that for each v_i she only observes a selected sample of opponent's bids b_j which satisfy $b_j > b_i$. For all other observations with a given v_i, b_i is known and she can only infer that $b_j < b_i$. The bidder could then attempt to complete the missing data by assuming some distribution $\tilde{H}(b_j|v_i, b_j < b_i)$. A natural starting point would be the uniform distribution. We call this bidder semi-naive since she makes some ad hoc assumption about $\tilde{H}(b_j|v_i, b_j < b_i)$, but at least she makes an attempt to correct for the selected sample. Given this approach, one could construct a distribution $H(b_j|v_i)$ that combines the empirical distribution $H(b_j|v_i, b_j > b_i)$ and the assumed distributions $\tilde{H}(b_j|v_i, b_j < b_i)$.

Finally, a *sophisticated bidder* may attempt to estimate the distribution of b_j conditional on v_i, using some structural model. Since the correlation between b_1 and b_2 cannot be assessed from the data, a natural starting point is that a bidder takes them to be independent (conditional on v_1 and v_2), and tries to identify the marginal conditional distribution $H(b_j|v_i)$ from the data. An identical identification problem arises in competing risk models. Translated into our context, the results of Tsiatis (1975) show that for any (not necessarily independent) joint distribution of the bids b_1 and b_2, one can construct unique marginal distributions that, under the assumption of independence are consistent with the observed data. The independence assumption is thus not testable and the sophisticated bidder is always able to pursue her approach.

Common to all three approaches is that, the naive, semi-naive, and the sophisticated bidder will deviate from the rational bid in the second-price auction if θ_1 and θ_2 are not independent.[26] This is the case since in all approaches the bidder believes that conditional

[25] Jehiel (2018) uses a similar selection neglect to demonstrate how investor overoptimism can arise if investors only observe realized past projects. Note however, that in the auction context, the direction of the bias crucially depends on what is disclosed. Indeed, if the second highest bid is disclosed, a naive bidder may think that bidder j bids lower than in reality.

[26] Interestingly, only the last approach will have the converse property that when the distributions of types are independent, bidders are behaving optimally. In this sense, it is the closest to the insights developed in

on θ_i, v_i and b_j are correlated. Therefore, the impossibility of an efficient auction-like mechanism continues to hold since the presence of the rational type requires the use of the second-price auction, and as before, the m-type does not bid truthfully in a second-price auction.

5.3 Anonymity of bids

In the above analysis, we have assumed that whether a past bid b came form a bidder in the role of bidder i or j was accessible in the dataset so that bidder i was able to relate the distribution of (past bidder j's bids) b_j to the realizations of (past bidder i's ex post values) v_i. In some cases, past bids would remain anonymous and the datasets would consist of (b, b', v_i, v_j) instead. In such cases, it would not be known whether b or b' was chosen by a bidder in the role of i or j. In the spirit of the analogy-based expectation equilibrium, this would call for considering an analogy partition that is bidder-anonymous in addition to being ex post-payoff relevant (see Jehiel, 2021). That is, for each v_i, bidder i could aggregate the distributions of b and b' (or equivalently of b_i and b_j) conditional on v_i, and best-respond as if bidder j were playing according to such an aggregate distribution when the ex post value of i is v_i.

The analysis would be similar to the one above. In particular, we would still obtain an inefficiency result under the conditions of Theorem 1. But, it should be mentioned that in the anonymous bid case, even when θ_i and θ_j are independently distributed, the resulting steady state would induce some misspecifications on the part of a data-driven bidder i, as it would lead such a bidder i to think that b_j is correlated to v_i when in fact only b_i is.

5.4 General Mechanisms

We have focused on a class of auction-like mechanisms in which bids are one-dimensional and the bidder with highest bid wins. This is a natural class that covers many practically relevant auction formats. Nevertheless, it begs the question whether the designer could achieve efficiency with more elaborate mechanisms. The literature on full surplus extraction with correlated types has demonstrated that under conditions that hold generically in the standard model, (1) the surplus of an incentive compatible allocation rule can be fully extracted (almost fully with continuous types), and (2) with discrete types, any allocation rule can be implemented in a Bayes-Nash equilibrium (Crémer and McLean, 1988; MAfee

this paper.

and Reny, 1992; and Gizatulina and Hellwig, 2017). We are interested in the implementation of an allocation rule that is not known to be incentive compatible given the unusual nature (the cognitive dimension) of a type in our setting. Therefore, a generalization of the second result for continuous types and a mix of rational and misspecified bidders is needed. GH17 show that for generic distributions, the belief $f(\theta_j|\theta_i)$ for given θ_i cannot be expressed as a convex combination of other types' beliefs. Moreover, in an incentive compatible direct mechanism, the belief of an m-type about the type θ_j of the opponent is given by[27]

$$h(\theta_j|\theta_i) = \sum_{k=1}^{K} \theta_i^k f(\theta_j|v_i^k) = \sum_{k=1}^{K} \theta_i^k f(\theta_j|\theta_i = e_k). \tag{7}$$

where $e_k^\ell = \mathbf{1}_{k=\ell}$. Together this implies that no two generalized types (θ_i, s_i) and (θ_i', s_i') have identical beliefs about θ_j. Using a strictly proper scoring rule of the type considered in Johnson, Pratt, and Zeckhauser (1990), we can therefore construct payment rules that (a) require an expected payment equal to zero from an agent who truthfully reports his belief, and (b) lead to a strictly positive expected payment from any misreport. We conjecture that combining such a payment rule with the efficient allocation rule allows to deter any non-local deviations. However, it is an open question if there exists a payment rule that (virtually) implements the efficient allocation rule (or any other allocation rule that is not known to be incentive compatible). As far as we know, an analog result has not been shown with a continuous type space even in the standard model, and exploring this direction is beyond the scope of this paper.

At the same time, even if a positive result can be achieved using scoring rules, the resulting mechanism will be highly unrealistic and schemes like this have been criticized even in the standard model since they rely on detailed knowledge of the environment by the designer (see Wilson's critique). In our context, the designer would not only rely on detailed knowledge about the type distribution, she would also have to know exactly how misspecified types form beliefs based on past data. If, by contrast, there is a rich set

[27] If the designer can use general mechanisms, we could consider direct mechanisms in which the message space coincides with the generalized type-space $M = T$. Any equilibrium of a general mechanism can be replicated by a truthful equilibrium of a direct mechanism. The truthful direct mechanism is a tool in the analysis that should not be taken literally if bidder's are not aware of their sophistication. Therefore if agents are not aware of their sophistication, it is in general not obvious how to replicate the equilibrium of a direct mechanism, using an indirect mechanism in which agents are not asked to report their sophistication directly. However, if we want to replicate a direct mechanism that uses scoring rules as suggested below, an indirect implementation would elicit beliefs from the agents. This does not require agents to report their sophistication.

of possible misspecifications that arise from different ways in which agents may use the available data, it seems unlikely that a designer has the detailed knowledge required to design a suitable incentive scheme. Moreover, for sufficiently rich sets of misspecified types (see discussion in Subsection 5.1), it may not even be possible to separate all types. For example, if there are many different types who observe past data with different coarseness, then the result that no two different types have the same beliefs about the opponent's type is unlikely to hold.

5.5 More general signal structure

In our model, the signal θ_i received by bidder i concerns only i's ex post value v_i. More generally, we could consider that i's signal θ_i defines a belief over the realizations of (v_i, v_j), the ex post values of both bidders. Under general specifications, this would lead us to consider an interdependent value setup in contrast to the private value setup considered above. In this case, the second-price auction may not induce an (interim) efficient allocation, even when only rational bidders are present (see Jehiel and Moldovanu, 2001 for the case in which types are independently drawn). But, under some conditions, the second-price auction would be efficient in such a case (see Dasgupta and Maskin, 2000). However, we conjecture that the additional presence of data-driven bidders (assumed to have access to data of the form (b_i, b_j, v_i, v_j)) would lead to inefficiencies in auction-like mechanisms, as in Theorem 1. A detailed analysis of this is left for future research.

6 Conclusion

Similarly as in the motivation for robust mechanism design, we have assumed that the beliefs of agents were not accessible by others. While the robust mechanism design looks for mechanisms that would work well for a wide range of such unobserved beliefs, our approach suggests a method to model the behaviors of bidders who would miss that information from past auctions. This has led us to observe that, in the presence of correlation, inefficiencies may be inevitable in auction-like mechanisms when there is a mix of novice bidders who reason that way and experienced bidders who behave optimally. Studying the second-best as well offering a welfare and revenue comparison for standard auctions are natural next steps for this research agenda.

A Omitted Proofs

A.1 Proof of Proposition 1

Proof of Proposition 1. If $\lambda \in (0,1)$, efficiency would require that $b^m(\theta_i) = \theta_i$ which implies

$$H^{\text{SPA}}(b \mid v_i = 1) = \int_0^1 F(b \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{\mathbb{E}[\tilde{\theta}_i]} d\tilde{\theta}_i,$$
$$H^{\text{SPA}}(b \mid v_i = 0) = \int_0^1 F(b \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{\mathbb{E}[1-\tilde{\theta}_i]} d\tilde{\theta}_i.$$

Moreover, we must have

$$\theta_i \in \arg\max_b \left\{ \theta_i H^{\text{SPA}}(b \mid v_i = 1) - \theta_i \int_0^b x dH^{\text{SPA}}(x|v_i = 1) - (1-\theta_i) \int_0^b x dH^{\text{SPA}}(x|v_i = 0) \right\}$$

Differentiating the objective function and setting $b = \theta_i$ yields

$$(1-\theta_i)\theta_i \left[H^{\text{SPA}\prime}(\theta_i \mid v_i = 1) - H^{\text{SPA}\prime}(\theta_i \mid v_i = 0) \right]$$

We have

$$\begin{aligned}
& H^{\text{SPA}\prime}(\theta_i \mid v_i = 1) - H^{\text{SPA}\prime}(\theta_i \mid v_i = 0) \\
= & \int_0^1 f(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{\mathbb{E}[\tilde{\theta}_i]} d\tilde{\theta}_i - \int_0^1 f(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{\mathbb{E}[1-\tilde{\theta}_i]} d\tilde{\theta}_i \\
= & \int_0^1 \left[\frac{\tilde{\theta}_i}{\mathbb{E}[\tilde{\theta}_i]} - \frac{1-\tilde{\theta}_i}{1-\mathbb{E}[\tilde{\theta}_i]} \right] f(\theta_i \mid \tilde{\theta}_i) f(\tilde{\theta}_i) d\tilde{\theta}_i \\
= & f(\theta_i) \int_0^1 \left[\frac{\tilde{\theta}_i}{\mathbb{E}[\tilde{\theta}_i]} - \frac{1-\tilde{\theta}_i}{1-\mathbb{E}[\tilde{\theta}_i]} \right] f(\tilde{\theta}_i|\theta_i) d\tilde{\theta}_i \\
= & f(\theta_i) \left[\frac{\mathbb{E}[\tilde{\theta}_i|\theta_j = \theta_i]}{\mathbb{E}[\tilde{\theta}_i]} - \frac{1-\mathbb{E}[\tilde{\theta}_i|\theta_j = \theta_i]}{1-\mathbb{E}[\tilde{\theta}_i]} \right]
\end{aligned}$$

Hence, for bidding θ_i to be optimal for the misspecified type we must have for all θ_i:

$$\frac{\mathbb{E}[\tilde{\theta}_i|\theta_j = \theta_i]}{\mathbb{E}[\tilde{\theta}_i]} - \frac{1-\mathbb{E}[\tilde{\theta}_i|\theta_j = \theta_i]}{1-\mathbb{E}[\tilde{\theta}_i]} = 0$$
$$\iff \mathbb{E}[\tilde{\theta}_i|\theta_j = \theta_i] = \mathbb{E}[\tilde{\theta}_i]$$

If the last line holds for all θ_i we must have

$$\begin{aligned} & \int_0^1 \tilde{\theta}_i f(\tilde{\theta}_i|\theta_j) d\tilde{\theta}_i = \mathbb{E}[\tilde{\theta}_i], && \forall \theta_j, \\ \iff \quad & \int_0^1 \tilde{\theta}_i \theta_j f(\tilde{\theta}_i, \theta_j) d\tilde{\theta}_i = \mathbb{E}[\tilde{\theta}_1]\theta_j f(\theta_j), && \forall \theta_j, \\ \implies \quad & E\left[\tilde{\theta}_i \theta_j\right] = \left(\mathbb{E}[\tilde{\theta}_i]\right)^2. \end{aligned}$$

The last line implies that we must have $\text{Corr}[\theta_1, \theta_2] = 0$ if the misspecified types first-order condition is satisfied for $b = \theta_i$ for all θ_i. Therefore, if $\text{Corr}[\theta_1, \theta_2] \neq 0$, there are types for which a misspecified bidder will not bid θ_i and since $b^r(\theta_j) = \theta_j$ for all types and $\lambda \in (0,1)$, the allocation will be inefficiency for some type profiles. □

A.2 Proof of Proposition 2

Proof of Proposition 2. An efficient allocation requires that $b^r(\theta_i) = b^m(\theta_i) = b(\theta_i)$ for all $\theta_i \in [0,1]$. We denote the inverse of $b(\cdot)$ by θ.

The rational type's bid solves

$$\max_b (\theta_i - b) F(\theta(b)|\theta_i)$$

The FOC yields

$$\begin{aligned} & -F(\theta_i|\theta_i) + (\theta_i - b(\theta_i)) f(\theta_i|\theta_i)\theta_i'(b(\theta_i)) = 0 \\ \iff \quad & b'(\theta_i) = (\theta_i - b(\theta_i)) \frac{f(\theta_i|\theta_i)}{F(\theta_i|\theta_i)}. \end{aligned} \tag{8}$$

The solution with boundary condition $b(0) = 0$ is

$$b(\theta_i) = \int_0^{\theta_i} x e^{-\int_x^{\theta_i} \frac{f(y|y)}{F(y|y)} dy} \frac{f(x|x)}{F(x|x)} dx.$$

The misspecified type maximizes (3)

$$\max_b (1-b)\,\theta_i H^{\text{FPA}}(b \mid v_i = 1) - b(1-\theta_i)H^{\text{FPA}}(b|v_i = 0).$$

with

$$H^{\text{FPA}}(b \mid v_i = 1) = \int_0^1 F(\theta(b) \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i,$$
$$H^{\text{FPA}}(b \mid v_i = 0) = \int_0^1 F(\theta(b) \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\tilde{\theta}_i.$$

This yields

$$\begin{aligned}&\theta_i H^{\text{FPA}}(b(\theta_i) \mid v_i = 1) + (1-\theta_i)H^{\text{FPA}}(b(\theta_i)|v_i = 0)\\ &= (1-b(\theta_i))\,\theta_i H^{\text{FPA}\prime}(b(\theta_i) \mid v_i = 1) - b(\theta_i)(1-\theta_i)H^{\text{FPA}\prime}(b(\theta_i)|v_i = 0)\end{aligned}$$

Using

$$H^{\text{FPA}\prime}(b \mid v_i = 1) = \theta'(b) \int_0^1 f(\theta(b) \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i = \theta'(b) \frac{E[\tilde{\theta}_i|\theta(b)]}{E[\tilde{\theta}_i]} f(\theta(b))$$
$$H^{\text{FPA}\prime}(b \mid v_i = 0) = \theta'(b) \int_0^1 f(\theta(b) \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\tilde{\theta}_i = \theta'(b) \frac{1-E[\tilde{\theta}_i|\theta(b)]}{1-E[\tilde{\theta}_i]} f(\theta(b))$$

we have

$$
\begin{aligned}
&\theta_i \int_0^1 F(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i + (1-\theta_i) \int_0^1 F(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\theta_i \\
&= (1-b(\theta_i))\,\theta_i \theta'(b(\theta_i)) \frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} f(\theta_i) - b(\theta_i)(1-\theta_i)\theta'(b(\theta_i)) \frac{1-E[\tilde{\theta}_i|\theta_i]}{1-E[\tilde{\theta}_i]} f(\theta_i) \\
\iff \quad &\theta_i \int_0^1 F(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i + (1-\theta_i) \int_0^1 F(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\theta_i \\
&= \frac{1-b(\theta_i)}{b'(\theta_i)} \theta_i \frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} f(\theta_i) - \frac{b(\theta_i)}{b'(\theta_i)} (1-\theta_i) \frac{1-E[\tilde{\theta}_i|\theta_i]}{1-E[\tilde{\theta}_i]} f(\theta_i) \\
\iff b'(\theta_i) &= \frac{(1-b(\theta_i))\,\theta_i \frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} f(\theta_i) - b(\theta_i)(1-\theta_i) \frac{1-E[\tilde{\theta}_i|\theta_i]}{1-E[\tilde{\theta}_i]} f(\theta_i)}{\theta_i \int_0^1 F(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i + (1-\theta_i) \int_0^1 F(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\tilde{\theta}_i} \\
&= \theta_i \frac{\frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} f(\theta_i)}{\theta_i \int_0^1 F(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i + (1-\theta_i) \int_0^1 F(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\tilde{\theta}_i} \\
&\quad - b(\theta_i) \frac{\theta_i \frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} f(\theta_i) - (1-\theta_i) \frac{1-E[\tilde{\theta}_i|\theta_i]}{1-E[\tilde{\theta}_i]} f(\theta_i)}{\theta_i \int_0^1 F(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i + (1-\theta_i) \int_0^1 F(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\theta_i} \\
&= (\theta_i - b(\theta_i)) \frac{f(\theta_i|\theta_i)}{F(\theta_i|\theta_i)}
\end{aligned}
$$

Where the last line follows from (8). Matching coefficients, we get

$$
\frac{\frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} f(\theta_i)}{\theta_i \int_0^1 F(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i + (1-\theta_i) \int_0^1 F(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\tilde{\theta}_i} = \frac{f(\theta_i|\theta_i)}{F(\theta_i|\theta_i)}
$$

and

$$
\frac{\theta_i \frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} f(\theta_i) - (1-\theta_i) \frac{1-E[\tilde{\theta}_i|\theta_i]}{1-E[\tilde{\theta}_i]} f(\theta_i)}{\theta_i \int_0^1 F(\theta_i \mid \tilde{\theta}_i)\tilde{\theta}_i \frac{f(\tilde{\theta}_i)}{E[\tilde{\theta}_i]} d\tilde{\theta}_i + (1-\theta_i) \int_0^1 F(\theta_i \mid \tilde{\theta}_i)(1-\tilde{\theta}_i) \frac{f(\tilde{\theta}_i)}{E[1-\tilde{\theta}_i]} d\tilde{\theta}_i} = \frac{f(\theta_i|\theta_i)}{F(\theta_i|\theta_i)}
$$

Combining these we have

$$\frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} = \theta_i \frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} - (1-\theta_i)\frac{1-E[\tilde{\theta}_i|\theta_i]}{1-E[\tilde{\theta}_i]}$$
$$\frac{E[\tilde{\theta}_i|\theta_i]}{E[\tilde{\theta}_i]} = \frac{1-E[\tilde{\theta}_i|\theta_i]}{1-E[\tilde{\theta}_i]}$$

This is the same condition as for the SPA which requires that $\text{Corr}[\theta_1, \theta_2] = 0$. □

A.3 Proof of Lemma 1

Proof of Lemma 1. Consider the equilibrium of the original mechanism $\tilde{M}$. For each bidder i and each $s_i \in \{r, m\}$, we define a (non-empty) correspondence that contains all bids that types with expected valuation w_i use.

$$b_i^{s_i}(w_i) = \tilde{b}_i(w_i, X, s_i)$$

where $X = [0,1]^{K-2}$. We prove the lemma in three steps: (1) we obtain an efficient equilibrium of the original mechanism with single-valued correspondences (or functions) $\hat{b}_i^{s_i}$. (2) We show that these functions satisfy $\hat{b}_i^r(w) = \hat{b}_i^s(w) = \tilde{\phi}_i(\hat{b}_j^r(w)) = \tilde{\phi}_i(\hat{b}_j^r(w))$, and a change of variable allows us to construct a mechanism $\check{M} = \left(B, \left(\check{W}_i\right), \left(\check{L}_i\right), Id\right)$ that has an efficient equilibrium in which $\check{b}_i^r(w) = \check{b}_i(w) = \check{b}_j^r(w) = \check{b}_j^r(w) = \check{b}(w)$. (3) We remove jump continuities in $\check{b}(w)$ and normalize the range of $\check{b}(w)$ to obtain a mechanism $M = ([0,1], (W_i), (L_i), Id)$ so that the (normalized) continuous part of $\check{b}(w)$ is an efficient equilibrium. We show that removing the discontinuities does not destroy the smoothness of the simple mechanism M.

Step 1: First, note that efficiency requires that the correspondences $b_i^{s_i}$ for $i \in \{1,2\}$ must be strictly increasing, meaning any selection must be strictly increasing. We denote the point-wise infimum and supremum of the correspondence by $\underline{b}_i^{s_i}(w) = \inf b_i^{s_i}(w_i)$ and $\overline{b}_i^{s_i}(w) = \sup b_i^{s_i}(w_i)$. Note that the infimum $\underline{b}_i^{s_i}(w)$ is strictly increasing if any selection from $b_i^{s_i}(w)$ is strictly increasing.

Suppose for some w_i, $b_i^r(w_i)$ is not single-valued. Efficiency and the fact that the in requires that for every $b_i \in [\underline{b}_i^{s_i}(w), \overline{b}_i^{s_i}(w)]$, $\left(b_j^{s_j}\right)^{-1}(\phi_2(b_i)) \subset \{w_i\}$, that is, any bid in the closed interval between the between the infimal and supremal bid that bidder i with interim value w_i places in equilibrium is either not placed by bidder j or it is placed by a bidder with the same interim value. We can include the infimum (and supremum) since $w_j \in \left(b_j^{s_j}\right)^{-1}(\phi_2(\underline{b}_i^{s_i}(w)))$ for some $w_j < w_i$ would imply that there exists $w_i' \in (w_j, w_i)$ such that $b_i' < \underline{b}_i^{s_i}(w)$ for some $b_i' \in b_i^{s_i}(w_i')$, which violates efficiency.

Since the probability that $w_j = w_i$ conditional on (w_i, x_i) is zero for all $x_i \in X_i$, the rational type is indifferent between all bids in $[\underline{b}_i^{s_i}(w), \overline{b}_i^{s_i}(w)]$. We set $\hat{b}_i(w_i, x_i, r) := \hat{b}_i^r(w_i) := \underline{b}_i^r(w)$. Similar steps show that we can set $\hat{b}_i(w_i, x_i, m) := \hat{b}_i^m(w_i) := \underline{b}_i^m(w)$.

Since the probability that $E[v_i|\theta_i] = w_i$ is zero, and there are at most countably many discontinuities, this modification of $\tilde{b}_i$ to $\hat{b}_i$ does not change the incentives of bidder j so that we have constructed a new equilibrium in which the correspondences of bidder i are single valued. We can apply the same modification to the strategy of bidder j. Clearly these modification preserve efficiency since $b_j^{s_j}(w_j) < \phi_2(\inf \tilde{b}_i^r(w_i))$ whenever $w_j < w_i$.

Step 2: We have shown in Step 1 that there exists an efficient equilibrium of $\tilde{M}$ that is given by the function $\hat{b}_i^s(w)$, $i \in \{1, 2\}$, $s \in \{r, m\}$. Clearly, efficiency requires that $\hat{b}_i^r(w) = \hat{b}_i^m(w) = \phi_i(\hat{b}_j^r(w)) = \phi_i(\hat{b}_j^m(w)) =: \hat{b}_i(w)$ for almost every w. The only exceptions are a countable set of interim values where all functions have a jump-discontinuity. Here we can redefine $\hat{b}_i^r(w) = \hat{b}_i^m(w) = \hat{b}_i(w) := \lim_{w' \uparrow w} \min\left\{\hat{b}_i^r(w'), \hat{b}_i^m(w'), \phi_i(\hat{b}_j^r(w')), \phi_i(\hat{b}_j^m(w'))\right\}$ for $i \neq j$, so that $\hat{b}_i^r(w) = \hat{b}_i^m(w) = \phi_i(\hat{b}_j^r(w)) = \phi_i(\hat{b}_j^m(w)) = \hat{b}_i(w)$ for every w, and $\hat{b}_i(w)$ is left-continuous.

The bids of bidder i are contained in $\hat{R}_i = [\hat{b}_i(0), \hat{b}_i(1)]$. We now define a new mechanism with $\check{B} = [0, 1]$, $\check{\phi}(w) = w$ and $\check{W}_i, \check{L}_i : [0, 1]^2 \to \mathbb{R}$ given by:

$$\check{W}_i(\check{b}_i, \check{b}_j) = \tilde{W}_i\left(\hat{b}_i(0) + \check{b}_i|\hat{R}_i|,\, \tilde{\phi}_j\left(\hat{b}_i(0) + \check{b}_j|\hat{R}_i|\right)\right),$$
$$\check{L}_i(b_i, b_j) = \tilde{L}_i\left(\hat{b}_i(0) + \check{b}_i|\hat{R}_i|,\, \tilde{\phi}_j\left(\hat{b}_i(0) + \check{b}_j|\hat{R}_i|\right)\right).$$

The new mechanism has an equilibrium given by the functions $\check{b}_i^s(w) = (\hat{b}_i(w) - \hat{b}_i(0))/|\hat{R}_i|$ and $\check{b}_j^s(w) = (\hat{b}_i(w) - b_i(p))/|\hat{R}_i|$. This equilibrium allocates to the bidder with the highest valuation since $\check{b}_i^s(w) > \check{b}_j^s(w)$ if and only if $\hat{b}_i(w) > \phi_i(\hat{b}_i(w))$ and the original mechanism was efficient. This implies that all bidding functions are the same: $\check{b}_i^s(w) = \check{b}_j^s(w) =: \check{b}(w)$ for $s \in \{r, m\}$. Moreover $\check{W}_i$ and $\check{L}_i$ are $\mathcal{C}^1$ since $\tilde{\phi}_j$ is continuously differentiable.

Step 3: The bidding function $\check{b}(w)$ is strictly increasing and can therefore be decomposed as $\check{b}(w) = \check{b}^C(w) + \check{b}^J(w)$, where $\check{b}^C(w)$ is continuous and $\check{b}^J(w)$ is constant except for a countable number of jump-discontinuities. We can modify the definition of $\check{M}$ and obtain a new smooth auction-like mechanism M with a symmetric equilibrium in which

$b(w) = \check{b}^C(w) / \left(\check{b}^C(1) - \check{b}^C(0)\right)$.

The function $b(w_i)$ specifies an equilibrium in the mechanism given by:

$$W_i(b_1, b_2) = \check{W}_i(\check{b}((\check{b}^C)^{-1}(b_1(b^C(1) - b^C(0)))), \check{b}((\check{b}^C)^{-1}(b_2(b^C(1) - b^C(0))))),$$
$$L_i(b_1, b_2) = \check{L}_i(\check{b}((\check{b}^C)^{-1}(b_1(b^C(1) - b^C(0)))), \check{b}((\check{b}^C)^{-1}(b_2(b^C(1) - b^C(0))))).$$

Next, we show that W and L are continuously differentiable. In the mechanism defined in step 2, a rational bidder chooses b_i to maximize

$$\int_0^{\check{b}^{-1}(b_i)} \left(w_i - \check{W}_i(b_i, \check{b}(w_j))\right) dF(w_j|w_i, x_i) - \int_{\check{b}^{-1}(b_i)}^1 \check{L}_i(b_i, \check{b}(w_j)) dF(w_j|w_i, x_i),$$

where $F(w_j'|w, x_i)$ is the probability that $w_j \leq w_j'$, conditional on bidder i's type (w_i, x_i).

Consider a rational bidder with type $w_i = \hat{w} + \varepsilon$, where $\hat{w}$ is a discontinuity in the equilibrium bidding function $\check{b}$ of original mechanism. Placing a bid $b' \in [\check{b}(\hat{w}), \check{b}(\hat{w}_+))$ instead of $\check{b}(w_i)$ must not be profitable:

$$\int_0^{\check{b}^{-1}(\check{b}(w_i))} \left(w_i - \check{W}_i(\check{b}(w_i), \check{b}(w_j))\right) dF(w_j|w_i, x_i) - \int_{\check{b}^{-1}(\check{b}(w_i))}^1 \check{L}_i(\check{b}(w_i), \check{b}(w_j)) dF(w_j|w_i, x_i)$$
$$\geq \int_0^{\check{b}^{-1}(b')} \left(w_i - \check{W}_i(b', \check{b}(w_j))\right) dF(w_j|w_i, x_i) - \int_{\check{b}^{-1}(b')}^1 \check{L}_i(b', \check{b}(w_j)) dF(w_j|w_i, x_i)$$

This can be rewritten as

$$\int_0^{\hat{w}} \left(w_i - \check{W}_i(\check{b}(w_i), \check{b}(w_j))\right) dF(w_j|w_i, x_i) - \int_{\hat{w}}^1 \check{L}_i(\check{b}(w_i), \check{b}(w_j)) dF(w_j|w_i, x_i)$$
$$+ \int_{\hat{w}}^{\hat{w}+\varepsilon} \left(w_i - \check{W}_i(\check{b}(w_i), \check{b}(w_j))\right) dF(w_j|w_i, x_i) + \int_{\hat{w}}^{\hat{w}+\varepsilon} \check{L}_i(\check{b}(w_i), \check{b}(w_j)) dF(w_j|w_i, x_i)$$
$$\geq \int_0^{\hat{w}} \left(w_i - \check{W}_i(b', \check{b}(w_j))\right) dF(w_j|w_i, x_i) - \int_{\hat{w}}^1 \check{L}_i(b', \check{b}(w_j)) dF(w_j|w_i, x_i)$$

The second term in on the left-hand side vanishes as $\varepsilon \to 0$ since $\check{W}_i$ and $\check{L}_i$ are bounded.

Hence we must have

$$\int_0^{\hat{w}} \left(\check{W}_i(b', \check{b}(w_j)) - \check{W}_i(\check{b}(\hat{w}_+), \check{b}(w_j))\right) dF(w_j|\hat{w}, x_i) \\ + \int_{\hat{w}}^1 \left(\check{L}_i(b', \check{b}(w_j)) - \check{L}_i(\check{b}(\hat{w}_+), \check{b}(w_j))\right) dF(w_j|\hat{w}, x_i) \geq 0$$

Since $b' < \check{b}(w_i)$, and $\check{W}_i$ and $\check{L}_i$ are non-decreasing in the first argument, this implies that $\check{W}_i(b', \check{b}(w_j)) = \check{W}_i(b_i, \check{b}(w_j))$ and $\check{L}_i(b', \check{b}(w_j)) = \check{L}_i(b_i, \check{b}(w_j))$ all $b' \in [\check{b}(\hat{w}), \check{b}(\hat{w}_+)]$ and almost every w_j. By continuity of $\check{W}_i$ and $\check{L}_i$ the equalities must hold for all w_j. Hence since $\check{W}_i$ and $\check{L}_i$ are continuously differentiable, $\partial\check{W}_i(b', \check{b}(w_j))/\partial b_i = 0$ and $\partial\check{L}_i(b', \check{b}(w_j))/\partial b_i = 0$ for all w_j and all $b' \in [\check{b}(\hat{w}), \check{b}(\hat{w}_+)]$ and also $\partial\check{W}_i(b', \check{b}(w_j))/\partial b_j = \partial\check{W}_i(\check{b}(\hat{w}), \check{b}(w_j))/\partial b_j = \partial\check{W}_i(\check{b}_+(\hat{w}), \check{b}(w_j))/\partial b_j$ and $\partial\check{L}_i(b', \check{b}(w_j))/\partial b_j = \partial\check{L}_i(\check{b}(\hat{w}), \check{b}(w_j))/\partial b_j = \partial\check{L}_i(\check{b}(\hat{w}_+), \check{b}(w_j))/\partial b_j$ for all $b' \in [\check{b}(\hat{w}), \check{b}(\hat{w}_+)]$ and all w_j. Hence continuous differentiability is preserved by the elimination of the gaps. □

A.4 Proof of Lemma 2

Proof of Lemma 2. We first show that for all i and $b_i, b_j \in [0, 1]$: $\partial W_i(b_i, b_j)/\partial b_i = 0$ if $b_j < b_i$, and $\partial L_i(b_i, b_j)/\partial b_i = 0$ if $b_j > b_i$.

Since $\delta_i(b) = 0$ for all $b \in [0, 1]$ we have that ψ

$$\psi'(b) = \frac{\partial W_i(b, b)}{\partial b_i} + \frac{\partial W_i(b, b)}{\partial b_j} - \frac{\partial L_i(b, b)}{\partial b_i} - \frac{\partial L_i(b, b)}{\partial b_j} < \infty$$

where finiteness follows from the assumption that W_i and L_i are continuously differentiable.

Now suppose that for some $w_i \in (0, 1)$, $\int_0^1 \frac{\partial P_i(b(w_i), b(w_j))}{\partial b_i} f(w_j|w_i, x_i) dw_j > 0$. The same derivation leading to (9) in the proof of Lemma 3, together with $\delta_i(b(w_i)) = 0$ implies that

$$\liminf_{b \nearrow b(w_i)} \frac{\psi(b(w_i)) - \psi(b)}{b(w_i) - b} = \infty.$$

This contradicts $\psi'(b(w_i)) < \infty$. Hence $\int_0^1 \frac{\partial P_i(b(w_i), b(w_j))}{\partial b_i} f(w_j|w_i, x_i) dw_j = 0$ for all $w_i \in [0, 1]$. Since $\partial P(b_i, b_j)/\partial b_i \geq 0$ by assumption, we therefore have $\partial P_i(b_0, b(w_j))/\partial b_i = 0$ for almost every w_j and by continuity of $\partial W_i/\partial b_i$, $\partial L_i/\partial b_i$ and b, this holds for all w_j. Therefore $\partial_{b_i} W_i(b_0, b) = 0$ if $b < b_0$, and $\partial_{b_i} L_i(b_0, b) = 0$ if $b > b_0$.

To conclude the proof, note that individual rationality together with $L_i(b_i, b_j) \geq 0$

requires that $L_i(0, b_j) = 0$ for all b_j.[28] Since $\partial L_i(b_i, b_j)/\partial b_i = 0$ if $b_j > b_i$, this implies that $L_i(b_i, b_j) = 0$ for all $b_i \leq b_j$. Next, $\delta_i(b(w_i)) = 0$ implies $W_i(b_i(w), b_i(w)) = w_i + L_i(b_i(w), b_i(w)) = w_i$, and since $\partial W_i(b_i, b_j)/\partial b_i = 0$, $W_i(b_i, b_j(w_j)) = w_j$ whenever $b_i \geq b_j(w_j)$. □

A.5 Proof of Lemma 3

Proof of Lemma 3. Consider a rational bidder i with types $(w_0, x_i) \in [0,1]^{K-1}$ and any sequence of valuations $w_i^n \nearrow w_0$. w_i^n prefers to bid $b^n = b(w_i^n)$ over bidding $b_0 = b(w_0)$. Therefore

$$\int_0^{\psi(b^n)} \left(w_i^n - W_i(b^n, b(w_j))\right) f(w_j|w_i^n, x_i)dw_j - \int_{\psi(b^n)}^1 L_i(b^n, b(w_j)) f(w_j|w_i^n, x_i)dw_j$$
$$\geq \int_0^{\psi(b_0)} \left(w_i^n - W_i(b_0, b(w_j))\right) f(w_j|w_i^n, x_i)dw_j - \int_{\psi(b_0)}^1 L_i(b_0, b(w_j)) f(w_j|w_i^n, x_i)dw_j$$

$$\iff \frac{1}{b - b^n} \int_0^{\psi(b^n)} \left(W_i(b_0, b(w_j)) - W_i(b^n, b(w_j))\right) f(w_j|w_i^n, x_i)dw_j$$
$$+ \frac{1}{b - b^n} \int_{\psi(b^n)}^1 \left(L_i(b_0, b(w_j)) - L_i(b^n, b(w_j))\right) f(w_j|w_i^n, x_i)dw_j$$
$$\geq \frac{1}{b - b^n} \int_{\psi(b^n)}^{\psi(b_0)} \left(w_i^n - W_i(b_0, b(w_j)) + L_i(b_0, b(w_j))\right) f(w_j|w_i^n, x_i)dw_j$$

Taking the lim sup on both sides we get

$$\int_0^1 \frac{\partial P_i(b_0, b(w_j))}{\partial b_i} f(w_j|w_0, x_i)dw_j \geq \delta_i(b_0) f(w_0|w_0, x_i) \limsup_{n\to\infty} \frac{\psi(b_0) - \psi(b^n)}{b_0 - b^n}$$

[28] Notice that this holds independent of our normalization that $v^1 = 0$, since the lowest type never wins the object in a regular equilibrium.

where $P_i(b_i, b_j) = W_i(b_i, b_j) + L_i(b_i, b_j)$. Similarly, w_0 prefers to bid b_0 over b^n for all $n \in \mathbb{N}$:

$$\int_0^{\psi(b_0)} (w_0 - W_i(b_0, b(w_j))) f(w_j|w_0, x_i) dw_j - \int_{\psi(b_0)}^1 L_i(b_0, b(w_j)) f(w_j|w_0, x_i) dw_j$$
$$\geq \int_0^{\psi(b^n)} (w_0 - W_i(b^n, b(w_j))) f(w_j|w_0, x_i) dw_j - \int_1^{\psi(b^n)} (L_i(b^n, b(w_j))) f(w_j|w_0, x_i) dw_j$$

$$\iff \frac{1}{b_0 - b^n} \int_{\psi(b^n)}^{\psi(b_0)} (w_0 - W_i(b_0, b(w_j)) + L_i(b_0, b(w_j))) f(w_j|w_0, x_i) dw_j$$
$$\geq \frac{1}{b_0 - b^n} \int_0^{\psi(b^n)} (W_i(b_0, b(w_j)) - W_i(b^n, b(w_j))) f(w_j|w_0, x_i) dw_j$$
$$+ \frac{1}{b - b^n} \int_{\psi(b^n)}^1 (L_i(b_0, b(w_j)) - L_i(b^n, b(w_j))) f(w_j|w_i^n, x_i) dw_j$$

Taking the lim inf on both sides we get

$$\delta_i(b_0) f(w_0|w_0, x_i) \liminf_{n\to\infty} \frac{\psi(b_0) - \psi(b^n)}{b_0 - b^n} \geq \int_0^1 \frac{\partial P_i(b_0, b(w_j))}{\partial b_i} f(w_j|w_0, x_i) dw_j.$$

Hence, for $\delta_i(b_0) > 0$ we have

$$\liminf_{n\to\infty} \frac{\psi(b_0) - \psi(b^n)}{b_0 - b^n} \geq \frac{\int_0^1 \frac{\partial P_i(b_0, b(w_j))}{\partial b_i} f(w_j|w_0, x_i) dw_j}{\delta_i(b_0) f(w_0|w_0, x_i)} \geq \limsup_{n\to\infty} \frac{\psi(b_0) - \psi(b^n)}{b_0 - b^n} \tag{9}$$

Notice that so far we have considered the case that $w^n < w_0$. The same steps apply for the case that the sequence satisfies $w^n > w_0$. Hence condition (9) applies for both cases. We have

$$\psi'(b_0) = \psi'_-(b_0) = \psi'_+(b_0) = \frac{\int_0^1 \frac{\partial P_i(b_0, b(w_j))}{\partial b_i} f(w_j|\psi(b_0), x_i) dw_j}{\delta_i(b_0) f(\psi(b_0)|\psi(b_0), x_i)}. \tag{10}$$

Hence $\psi(b_0)$ is differentiable at b_0. Since $\delta_i(b)$ is continuous, there exists ε such that $\delta_i(b) > 0$ for all $b \in B_\varepsilon(b_0)$. Since the right-hand side of (10) is continuous in b_0, ψ is continuously differentiable on $B_\varepsilon(b_0)$. Since ψ is strictly increasing there must be $b' \in B_\varepsilon(b_0)$ such that $\psi'(b') > 0$ and since ψ' is continuous, there exist $\alpha < b' < \beta$ such that $(\alpha, \beta) \subset B_\varepsilon(b_0)$ and ψ is continuously differentiable with $\psi'(b) > 0$ for $b \in (\alpha, \beta)$. □

A.6 Proof of Lemma 4

The proof follows the same steps as the proof of Theorem 2.4 in GH17, except that instead of considering continuous mappings from T_i to the space of all measures on T_{-i}, $\mathcal{M}(T_{-i})$, we consider continuous mappings from $[0,1]$ to the space of all absolutely continuous measures on $X = [0,1]^{K-2}$ with strictly positive and continuous density, which we denoted by $\mathcal{M}_+^d(X)$.

Restricting attention to $\mathcal{M}_+^d(X)$ instead of the space of all measures $\mathcal{M}(X)$, requires a straightforward modification of the constructions of the functions $\boldsymbol{g}$ and the measures $\beta_1, \ldots, \beta_K$ in footnote 20 of GH17. First we take the functions g^k to be functions $g^k : X \to [0,2]$ with $g^k(x^k) = 2$ and $g^k(x) = 0$ for $x \notin B^k$. This allows us to construct perturbations of the measures β_k^0 which need to be elements $\mathcal{M}_+^d(X)$ for our purposes, by setting $\beta_k = (1-\varepsilon)\beta_k^0 + \varepsilon\tilde{\beta}_k$ where the measure $\tilde{\beta}_k$ has a density $\tilde{f}_k$ that satisfies $\tilde{f}_k(x)$ for $x \notin B^k$ and $\int_X g^k(x)\tilde{f}_k(x)dx = 1$. Then, with $\varepsilon \neq -z/(1-z)$ for all negative eigenvalues of the matrix $\left(\int_X g^k(x)\beta_\ell^0(dx)\right)_{k,\ell}$, the vectors $\int_X \boldsymbol{g}(x)\beta_k(dx)$ for $k = 1, \ldots, K$ are linearly independent. The remaining steps in the proof are virtually unchanged.

A.7 Proof of Lemma 5

The proof follows Theorem 2.7 in GH17 and uses results from Section 5.4 in Gizatulina and Hellwig (2014).

First note that for elements of $\mathcal{M}_+^d([0,1]^{2K})$, marginal and conditional densities are defined in the usual way. Moreover, for each w_i, the function that maps w_j to the conditional probability measure on X that is given by the density $f(x_i|w_i, w_j)$, is an element of $\mathcal{C}([0,1], \mathcal{M}_+^d(X))$ (see GH14).

Analog to the proof of Theorem 2.7 in GH17, we let $\mathcal{F}_{w_i}^i \subset \mathcal{M}_+^d([0,1]^{2K})$ be the set of priors such that the function $w_j \mapsto f(\cdot|w_i, w_j)$ is an element of $\mathcal{E}(w_i)$. The key step is to show that the residualness of $\mathcal{E}(w_i)$ in $\mathcal{C}([0,1], \mathcal{M}_+^d(X))$ implies the residualness of $\mathcal{F} = \bigcap_{i\in\{1,2\}, w_i\in\mathcal{W}_i} \mathcal{F}_{w_i}^i$ in $\mathcal{M}_+^d([0,1]^{2K})$. For each $i \in \{1,2\}$ and $w_i \in (0,1)$, let $\psi_{i,w_i} : \mathcal{M}_+^d([0,1]^{2K}) \to \mathcal{M}_+^d([0,1]) \times \mathcal{C}([0,1], \mathcal{M}_+^d(X))$ be the mapping that maps the prior to the conditional distribution $f(w_j|w_i)$ and the function $w_j \mapsto f(x_i|w_i, w_j)$. As shown in the proof of Lemma 5.9 in GH14, the maps ψ_{i,w_i} are continuous and open if $\mathcal{M}_+^d([0,1]^{2K})$ is endowed with the uniform topology for density functions. As in the proof of Theorem 2.7 in GH17, this implies that $\mathcal{F}_{w_i}^i$ is as residual subset of $\mathcal{M}_+^d([0,1]^{2K})$, that is it contains a countable intersection $\bigcap_{n\in\mathbb{N}} H_n(i, w_i)$ of open and dense sets $H_n(i, w_i) \subset \mathcal{M}_+^d([0,1]^{2K})$.

Clearly, $H = \bigcap_{i\in\{1,2\}} \bigcap_{w_i\in\mathcal{W}_i} \bigcap_{n(i,w_i)\in\mathbb{N}} H_{n(i,w_i)}(i,w_i)$ is a subset of $\mathcal{F}$. By a diagonal argument, H is a countable intersection of open and dense subsets of $\mathcal{M}^d_+([0,1]^{2K})$ and hence $\mathcal{F}$ is residual.

A.8 Proof of Lemma 6

Proof of Lemma 6. We have shown this for $|V| = 2$ in Proposition 1. For $|V| \geq 3$, we need to modify the proof. If m-types bid $b(w_i)$, we must have for all θ_i that

$$w_i = \mathbb{E}[v_i|\theta_i] \in \arg\max_b \left\{ \sum_{k=1}^{K} \theta_i^k \left(v_i^k H^{\text{SPA}}(b \mid v_i^k) - \int_0^b z dH^{\text{SPA}}(z|v_i^k) \right) \right\}.$$

The first-order condition is

$$\sum_{k=1}^{|V|} \theta_i^k \left(v_i^k - w_i \right) H^{\text{SPA}\prime}(w_i \mid v_i^k) = 0$$

Considering the type $\theta_i = (1-b, 0, \ldots, 0, b)$ for any $b \in (0,1)$, we have $w_i = b$, and the first-order condition simplifies to

$$H^{\text{SPA}\prime}(b \mid v_i = 1) - H^{\text{SPA}\prime}(b \mid v_i = 0) = 0$$

We have

$$\begin{aligned}
H^{\text{SPA}}(b \mid v_i^k) &= \frac{\mathbb{P}_f\left[b_j \leq b, v_i = v_i^k\right]}{\mathbb{P}_f\left[v_i = v_i^k\right]} = \frac{\int_{\Theta_i} \mathbb{P}_f\left[w_j \leq b|\tilde{\theta}_i\right] \mathbb{P}_f\left[v_i = v_i^k|\tilde{\theta}_i\right] f(\tilde{\theta}_i) d\tilde{\theta}_i}{\mathbb{E}_f\left[\theta_i^k\right]} \\
&= \frac{\int_{\Theta_i} F_{w_j}(b|\tilde{\theta}_i)\tilde{\theta}_i^k f(\tilde{\theta}_i) d\tilde{\theta}_i}{\mathbb{E}_f\left[\theta_i^k\right]} \\
H^{\text{SPA}\prime}(b \mid v_i^k) &= \frac{\int_{\Theta_i} f_{w_j}(b|\tilde{\theta}_i)\tilde{\theta}_i^k f(\tilde{\theta}_i) d\tilde{\theta}_i}{\mathbb{E}_f\left[\theta_i^k\right]}
\end{aligned}$$

Substituting this in the first-order condition, we get for all $b \in B$:

$$\begin{aligned}
&\frac{\int_{\Theta_i} f_{w_j}(b|\tilde{\theta}_i)\tilde{\theta}_i^K f(\tilde{\theta}_i)d\tilde{\theta}_i}{\mathbb{E}_f\left[\theta_i^K\right]} - \frac{\int_{\Theta_i} f_{w_j}(b|\tilde{\theta}_i)\tilde{\theta}_i^1 f(\tilde{\theta}_i)d\tilde{\theta}_i}{\mathbb{E}_f\left[\theta_i^1\right]} = 0 \\
\iff \quad & \int_{\Theta_i}\left[\frac{\tilde{\theta}_i^K}{\mathbb{E}_f\left[\theta_i^K\right]} - \frac{\tilde{\theta}_i^1}{\mathbb{E}_f\left[\theta_i^1\right]}\right] f_{w_i}(\tilde{\theta}_i|w_j = b)f_{w_j}(b)d\tilde{\theta}_i = 0 \\
\iff \quad & \frac{\mathbb{E}_f\left[\theta_i^K|w_j = b\right]}{\mathbb{E}_f\left[\theta_i^K\right]} = \frac{\mathbb{E}_f\left[\theta_i^1|w_j = b\right]}{\mathbb{E}_f\left[\theta_i^1\right]} \\
\iff \quad & \mathbb{E}_f\left[\theta_i^K|w_j \leq b\right] = \frac{\mathbb{E}_f\left[\theta_i^K\right]}{\mathbb{E}_f\left[\theta_i^1\right]}\mathbb{E}_f\left[\theta_i^1|w_j \leq b\right]
\end{aligned}$$

For generic distributions, the last line is violated. □

References

[1] Battigalli, P., M. Gilli, and M. Molinari (1992): "Learning and convergence to equilibrium in repeated strategic interactions: an introductory survey," *Ricerche Econ.*, 46, 335–378.

[2] Bergemann, D., B. Brooks, and S. Morris (2017): "First-Price Auctions With General Information Structures: Implications for Bidding and Revenue," *Econometrica* 85(1), 107-143.

[3] Bergemann, D., and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

[4] Crémer, J., and R. P. McLean (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56(6), 1247–1257.

[5] Dasgupta, P., and E. Maskin (2000): "Efficient Auctions," *Quarterly Journal of Economics* 115, 341-388.

[6] Dekel, E., D. Fudenberg, and D. K. Levine (2004): "Learning to play Bayesian games," *Games and Economic Behavior* 46(2), 282-303.

[7] 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.

[8] Esponda, I. (2008): "Behavioral Equilibrium in Economies with Adverse Selection," *American Economic Review*, 98(4), 126–191.

[9] Esponda, I., and D. Pouzo (2016): "Berk Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," *Econometrica*, 84, 1093–1130.

[10] Eyster, E., and M. Rabin (2005): "Cursed Equilibrium," *Econometrica*, 73(5), 1623–72.

[11] Gizatulina, A., and M. Hellwig (2014): "Beliefs, Payoff, Information: On the Robustness of the BDP Property in Models with Endogenous Beliefs," *Journal of Mathematical Economics*, 51(3), 136–153.

[12] Gizatulina, A., and M. Hellwig (2017): "The generic possibility of full surplus extraction in models with large type spaces," *Journal of Economic Theory*, 170, 385–416

[13] Jehiel, P. (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, 123(2), 81–104.

[14] Jehiel, P. (2011): "Manipulative Auction Design" *Theoretical Economics* 6(2), 185-217.

[15] Jehiel, P. (2018): "Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism," *American Economic Review*, 108(6), 1582–97.

[16] Jehiel, P. (2021): "Analogy-based expectation equilibrium and related concepts: Theory, applications and beyond," mimeo.

[17] Jehiel, P., and F. Koessler (2008): "Revisiting Games of Incomplete Information with Analogy-Based Expectations," *Games and Economic Behavior*, 62, 533–557.

[18] Jehiel, P., and B. Moldovanu (2001): "Efficient Design with Interdependent Valuations," *Econometrica* 69, 1237-1259.

[19] Johnson, S., J. W. Pratt, and R. J. Zeckhauser (1990): "Efficiency Despite Mutually Payoff-Relevant Private Information: The Finite Case," *Econometrica*, 58(4), 873–900

[20] Lizzeri, A., and N. Persico (2000): "Uniqueness and Existence of Equilibrium in Auctions with a Reserve Price," *Games and Economic Behavior*, 30, 83–114.

[21] McAfee, R. P., and P. J. Reny (1992): "Correlated Information and Mechanism Design," *Econometrica*, 60(2), 395–421.

[22] Milgrom, P. and R. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica* 50, 1089-1122.

[23] Spiegler, R. (2016): "Bayesian Networks and Boundedly Rational Expectations," *Quarterly Journal of Economics*, 131, 1243–1290.

[24] Spiegler, R. (2020): "Behavioral Implications of Causal Misperceptions," *Annual Review of Economics* 12, 81-106.

[25] Spiegler, R. (2021): "Modeling Players with Random Data Access," *Journal of Economic Theory* forthcoming.

[26] Tsiatis, A. (1975): "A Nonidentifiability Aspect of the Problem of Competing Risks," *Proceedings of the National Academy of Sciences*, 72(1), 20–22.

[27] Wilson, R. (1987), "Game-Theoretic Approaches to Trading Processes" in Advances in Economic Theory: Fifth World Congress of the Econometric Society, T.Bewley (ed.), Cambridge: Cambridge University Press.