Modeling Bidder Risk Preferences to Optimize Pricing

Robert Zeithammer
University of California, Los Angeles

Lucas Stich
Ludwig-Maximilians-Universität München

October 19, 2019

Abstract: Name-your-own-price selling is a tractable laboratory paradigm for studying bidding behavior because it involves only one bidder per transaction. Using data from an incentive-compatible laboratory experiment that implemented a name-your-own-price seller who charges entry fees, we estimate a flexible model of risk preferences at the individual level, and we find substantial population heterogeneity. The behavior of three quarters of our subjects is more consistent with prospect theory than with expected utility theory in that their estimated utility functions are convex in the loss domain. In several counterfactual simulations, we measure the impact of accounting for the heterogeneity in risk preferences on setting entry and reserve prices in three market institutions that involve bidding, including first-price auctions. We find that when the seller cannot discriminate based on risk preferences, prices set using the simpler homogeneous model (that assumes everyone has the same risk preferences) achieve over 99% of the optimal profit. By contrast, a discriminating seller can benefit from using the heterogeneous model in several situations that we characterize as depending on both the institution and the dispersion in the distribution of valuations.

Keywords: Pricing, Auctions, Risk Preferences, Econometric Modeling

Contact and acknowledgements: Robert Zeithammer, rzeitham@ucla.edu. The authors acknowledge the generous support of the Morrison Family Center for Marketing Studies and Data Analytics

1. Introduction

Bidding behavior has long been recognized as a potentially rich source of information about preferences, because it contains not only ordinal information about utility revealed by the choice to submit the bid, but also cardinal information about the strength of preference revealed by the bid's magnitude. Because most markets that give rise to bidding data are governed by publicly known and strictly enforced rules, empirical researchers and data analysts can infer the bidders' preferences from bidding data by inverting the mapping from preferences to bids implied by the rules (Sutton, 1993; Laffont and Vuong, 1996). A prominent example of this approach is the analysis of first-price auctions with independent private values by Guerre, Perrigne, and Vuong (2000), who invert the optimal bidding strategy to estimate bidders' valuations under the assumption of risk neutrality. Although human bidders are likely not neutral to risk, joint identification of risk preferences and valuations from bids alone is not possible (Guerre, Perrigne, and Vuong, 2009). However, in a laboratory setting with induced values, the analyst knows and controls bidder valuations, so estimation of bidder risk preferences is possible. In this paper, we propose and estimate a flexible individual-level model of risk preferences using bidding data from an incentive-compatible laboratory implementation of a very simple single-bidder auction called "name-your-own-price" (NYOP) selling. Instead of having to compete with other bidders, as would be the case in a standard auction, an NYOP bidder makes a binding price offer to the seller, and the seller's hidden reserve price determines whether the bid is accepted. The econometric inference about bidder preferences is thus particularly tractable because the optimal NYOP bidding strategy is a single-agent decision problem without strategic interactions among different bidders.

Thanks to the tractable laboratory paradigm, this paper is the first to identify individual-level risk preferences from entry and bidding behavior. Our estimates imply a substantial population heterogeneity in risk preferences. Using a set of counterfactual simulations, we explore the impact of this preference heterogeneity on optimal selling strategies in a variety of markets with consumer bidding, including auctions. Compared to accounting for risk preferences only at the population level using a homogeneous model that assumes all customers have the same risk preferences, estimating the model with heterogeneous risk preferences requires considerably more data and effort. We ask the following practical question: When is the additional work involved in empirically identifying the heterogeneous model justified by the additional profit thus attained, and when is the homogeneous model good enough for optimizing seller decisions? We find that sellers who cannot discriminate between buyers gain very little from carefully accounting for heterogeneity in risk preferences, but discriminating sellers can benefit a lot in some market institutions. An important driver of the benefit to discriminating sellers is the variance of the distribution of valuations, which can both enhance the benefit (when the seller uses reserve prices) and reduce it (when the seller uses entry fees).

We now briefly outline our data and modeling assumptions. The data come from an experiment conducted and analyzed by Zeithammer et al. (2019). The experiment varied not only the private valuations of the bidders, but also the entry costs (framed as non-refundable "bidding fees" to bidders). Because our bidders can thus end up with a net loss, we can empirically distinguish between two dominant models of behavior under risk—the expected utility theory (EUT) and prospect theory (PT)—which differ markedly in their model of utility in the loss domain. We measure risk preferences at the individual level using a parametric model with a flexible specification of the utility function that nests both models. Our results suggest substantial heterogeneity in both the best-fitting model as well as in the specific curvatures of the utility function. The behavior of three quarters of our subjects is more consistent with PT than with EUT in that their utility functions seem to be convex in the loss domain. All individual-specific utility functions we estimate are concave in the gain domain, with average relative risk aversion of 0.6 and a wide range from 0 to 0.9.

2. Literature review

Because of the uncertainty inherent in auctions and other bidding markets, individual risk preferences play a key role in modeling bidder behavior, as documented by empirical and experimental studies (e.g., Cox, Smith, and Walker, 1988; Athey and Levin, 2001; Goeree, Holt, and Palfrey, 2002; Abbas and Hann, 2010). However, no general consensus exists concerning which model of risk preferences is best suited to explain commonly observed phenomena, for example, overbidding in auctions (Guerre et al., 2009). We contribute to the literature on modeling bidder behavior by first estimating individual-level risk preferences using a more flexible model specification than prior work on bidding, and then exploring when accounting for the population preference heterogeneity we find is useful for an auctioneer setting reserve prices and entry fees.

The data from the laboratory paradigm we exploit allow us to overcome the lack of nonparametric joint identification of valuations and risk preferences from bidding data alone (Guerre et al., 2009). In a closely related paper that also uses laboratory bidding data, Bajari and Hortaçsu (2005) examined the ability of existing econometric approaches to recover the known valuation distribution from bids, and found the homogeneous risk-averse expected utility model works better than several alternatives. Instead of trying to recover the known valuation distribution using only the observed bids as observations, we consider the known valuations together with the bids in estimation, and infer the shape of bidders' utility functions at the individual level.

Another closely related experimental paper is Chakravarty et al (2011) who also use a bidding task designed to eliminate strategic considerations, and proceed to interpret the observed bids through the lens of a model with constant relative risk aversion. They estimate a population distribution of risk aversion and compare it across two conditions in their study, but they do not conduct counterfactual simulations that use the estimates to optimize selling. In contrast to both Bajari and Hortaçsu (2005) and Chakravarty et al (2011), our model endogenizes entry and is flexible enough to accommodate both expected utility and prospect theories of preferences.

Because of the lack of joint identification of valuations and risk preferences, structural approaches to auctions based on bidding data alone often employ risk neutrality as an identifying assumption (Athey and Haile, 2007). Several papers achieve the joint identification using either parametric assumptions or exogenous variation that gives rise to various exclusion restrictions. Specifically, Guerre et al. (2009) exploit exogenous bidder participation as an exclusion restriction to identify the bidders' utility function and their private value distribution nonparametrically. Campo et al. (2011) exploit heterogeneity across auctioned objects and a parameterization of the bidders' utility function. Lu and Perrigne (2008) use the fact that both ascending and sealed-bid auctions are used in timber sales. Bolotnyy and Vasserman (2019) estimate bidders' risk preferences in infrastructure procurement auctions, identifying a constant absolute risk aversion model of risk preferences using a subset of auctions for which bidding behavior is known a priori to be driven mostly by risk preferences. Campo (2012) exploits a parameterization of the bidders' utility function to show that semi-parametric joint identification of individual-level risk aversion and valuation distribution is possible when the analyst observes a bidder characteristic known to affect the bidder's valuation but not her risk preference.

With the exception of Campo (2012), all of the above existing approaches to identifying risk aversion from bids identify the homogeneous model, in which all bidders have the same risk preferences. We make two contributions to the structural econometric literature on auctions. First, we document when such a homogeneous model is sufficient for the auctioneer's decision-making across several different counterfactual bidding markets. Second, we show the expected-utility model is too restrictive when also modeling entry behavior, as evidenced by the fact that a prospect-theoretic utility model captures the overall entry and behavior of a majority of the subjects in our data better than the expected utility model. Our results reveal that considering the complete entry and bidding behavior has a strong and meaningful impact on inference as compared to considering only observed bids. Palfrey and Pevnitskaya (2008) also study the complete entry and bidding behavior in a laboratory setting, but they do not estimate the underlying preferences, and neither do Zeithammer et al. (2019), whose data we use in this research for our estimations.

The present research also connects to the large literature studying risk preferences from data on individual behavior in non-bidding domains. These papers typically rely on survey and experimental data (e.g., Holt and Laury, 2002; von Gaudecker, van Soest, and Wengström, 2011), but have also exploited economic field data (e.g., Cicchetti and Dubin, 1994; Cohen and Einav, 2007; Sydnor, 2010; Barseghyan, Molinari, O'Donoghue, and Teitelbaum, 2013; Barseghyan, Molinari, and Teitelbaum, 2016; Chiappori, Salanié, Salanié, and Gandhi, 2019). With respect to more methodological contributions in studying risk preferences under PT (e.g., Abdellaoui, 2000; Abdellaoui, Bleichrodt, and Paraschiv, 2007), this paper differs from the extant literature based on lottery choices in that we use data from a setting in which the subjects can decide how much risk they want to face by changing their bids.

Our emphasis on the importance of accounting for heterogeneity further relates this paper to the broader literature on modeling of consumer heterogeneity as summarized, for example, in Allenby and Rossi (1999). That literature focuses on modeling the heterogeneity in preferences for product attributes, and finds both large amounts of heterogeneity and large profit gains for marketers using posted prices who can exploit the heterogeneity through various forms of price discrimination. By contrast, we focus on the modeling of the population heterogeneity in risk preferences, and our counterfactuals study the profit gains of auctioneers facing bidders who are heterogeneous in both risk preferences and valuations. Despite studying different market institutions, our results echo the results from work on the heterogeneity in consumer product preferences in that we also find that a price-discriminating seller can gain substantial profits from using the heterogeneous model instead of relying on the homogeneous one.

3. Data

We analyze data from Experiment 2 in Zeithammer et al. (2019), focusing on the "decision aid" condition. As discussed by Zeithammer et al. (2019), the decision aid makes the task easier and the behavior more stable across rounds of the experiment. 48 subjects participated in this condition, each completing 25 rounds corresponding to a 5×5 within-subject manipulation of valuation drawn from {5, 20, 35, 50, 65} and bidding fee drawn from $\{0,1,6,12,18\}$. Each subject experienced all possible combinations of valuations and fees in random order. The seller was a computer program, which drew its procurement cost c randomly from Uniform[0,70], and accepted bids that exceeded the current cost draw.

In the beginning of each round, the subject learned her valuation and the bidding fee (but not the current cost c), and had to decide whether to participate in the bidding or skip the round. To help the subject with this decision, the bidding interface in the experiment provided a decision aid that calculated and displayed both the probability of acceptance $\frac{b}{70}$ and the monetary payoff $v-b-f$ should the subject's bid be accepted. The subject was required to use the decision aid at least once before submitting her bid, and she could use it as many times as she liked to try out different bid levels. If she chose to bid, she entered the bid amount into a box and pressed "Submit Bid," automatically deducting the bidding fee. Please see Zeithammer et al. (2019) for sample stimuli as well as for a detailed discussion of the effect of the decision aid on bidding and entry behavior.

At the end of each round, payoffs were determined as follows: When the buyer decided to skip a round, she received a payment of 0 points and the round ended. When the buyer submitted a bid and the bid was rejected, she lost the bidding fee f paid. Finally, when the buyer submitted a bid and the bid was accepted, she received a payment equal to her surplus $v-b-f$.

The data were collected at a large European public university. Subjects took a little over a minute per task. They earned, on average, about 16.90 EUR (USD 21.70 at the time of the experiments), which included a show-up fee of 4 EUR (USD 5.10) and another 4 EUR for taking the exit survey.

4. Model

The experiment described above presented the bidder with a costly opportunity to submit a bid in an NYOP setting, where the cost was framed as a non-refundable "bidding fee," deducted upon submission of a bid. From the bidder's perspective, the bidding fee acts exactly like the entry cost in Samuelson (1985). We assume that when considering entry with a bidding fee of f, the buyer with valuation v evaluates the two potential outcomes of offering a bid b as follows. First, if b is accepted, she combines the relevant experimental currency amounts into a net surplus $s=v-b-f$, and experiences her utility of this potential surplus as $u(s)=u(v-b-f)$. Second, if her bid is not accepted, she represents her utility of spending f without getting anything back as $u(-f)$. To decide whether and how much to bid, she solves the following:

$$\max\left\{\underbrace{u(0)}_{V_{not\ enter}},\underbrace{\max_{b\geq 0}\left(\frac{b}{70}\right)u(v-b-f)+\left(1-\frac{b}{70}\right)u(-f)}_{V_{enter}}\right\}, \tag{1}$$

where V represents the buyer's choice-relevant values of entering or not. Without loss of generality, we can set the intercept and scale of the utility function as $u(0)=0$ and $u(-1)=-1$. Given this normalization, the general model in equation 1 nests both classical EUT (with u concave for risk-averse bidders and convex for risk-seeing bidders)[1] and a simple version of PT (with u corresponding to the "value function" concave in gains and convex in losses)[2]. Our goal is to specify a parsimonious yet flexible model of u, and let the data determine which of the above two major theories of risk preferences fits our data better, at the individual level.

To allow for different curvatures in the loss and gain domains, as well as to allow a kink at zero motivated by loss aversion, we adopt the following three-parameter specification:

$$u(s)=\begin{cases} s\geq 0:\gamma s^{r} \\ s<0:-(-s)^{t} \end{cases}, \tag{2}$$

where all three parameters γ, r, and t are positive. When s>0, that is, in the gain domain, equation 2 assumes CRRA with relative risk aversion 1-r. In other words, r<1 implies greater risk aversion

---

[1] Given the $u(0)=0$ intercept, u represents the *change* in the underlying utility from some initial wealth level w_0 to a new wealth level, either $w_0+v-b-f$ in the case of an accepted bid, or w_0-f in the case of a rejected bid. The $u(-1)=-1$ scale normalization simply expresses the changes in the underlying utility as fractions of the change from losing one unit of experimental currency to getting no surplus. Specifically, if the underlying utility over wealth is $\tilde{u}(w)$, we let $u(s)=\frac{\tilde{u}(w_0+s)-\tilde{u}(w_0)}{\tilde{u}(w_0)-\tilde{u}(w_0-1)}$. Because we do not observe the initial wealth level w_0, estimation at the individual level is imperative to preserve the expected-utility interpretation.

[2] PT evaluates uncertain outcomes as gains or losses relative to a reference point. In our NYOP setting, it is natural to let the income state in the beginning of the game be the reference point as implicitly assumed by the $u(0)$=0 normalization: If the buyer does not pay any bidding fees and does not buy any products, she remains at this reference point. When she buys a product, she experiences a net gain of v-b-f. When her bid is rejected, the buyer experiences a loss of the bidding fee f. Also note that compared to the fully general model of Tversky and Kahneman (1992), we assume away the probability weighting function because it is not separately identified in our setting.

in the gains, r=1 implies risk neutrality in the gains, and r>1 implies risk-seeking preferences in the gains. The t parameter has the mirror-image interpretation in the loss domain; for example, t>1 implies risk aversion in the gains. Finally, the γ parameter allows for a kink at zero consistent with loss aversion; that is, a bidder is loss averse when γ<1. Our goal is to estimate three structural parameters $\{r,t,\gamma\}$ for each subject.

To complete the specification of the empirical model, we now describe the econometric error structure that allows observed bids and entry decisions to randomly deviate from theoretical predictions. Let $b(v,f \mid r,t,\gamma)$ be the solution to the inner maximization problem in equation 1, and denote the buyer's n-th observed bid bid_n. We model the observed bid as

$$bid_n = b(v_n, f_n \mid r,t,\gamma) + \varepsilon_n, \tag{3}$$

where ε_n is distributed normally with a mean of 0 and variance σ^2, truncated such that $0 < bid_n < v_n - f_n$. The truncation restricts the observed bid to be low enough that the bidder would receive a net gain should the bid be accepted—a natural rationality restriction.

To model the entry decision, let $V_{enter}(v_n, f_n | r,t,\gamma)$ be the choice-relevant value of entering defined in equation 1. We model the subject's n-th entry choice as a logistic transformation of the choice-relevant value:

$$\Pr_{enter}(v_n, f_n | r,t,\gamma,\tau) = \frac{\exp\left(\tau V_{enter}(v_n, f_n | r,t,\gamma)\right)}{1+\exp\left(\tau V_{enter}(v_n, f_n | r,t,\gamma)\right)}. \tag{4}$$

For tractability, we assume the error term related to entry (with a scale $1/\tau$) is independent of the error term related to bid magnitude (with a scale σ).

5. Parameter Identification and Estimation

We now discuss how the entry and bidding data from the experiment identify the five parameters of the model $\theta \equiv \{r, t, \gamma, \tau, \sigma\}$ at the individual level. The curvature parameter in the gain domain *r* is identified by the magnitude of bids given valuations (larger bids imply smaller *r*) as well as by the entry decisions (more risk-averse buyers enter less). Consider the special case of zero fee. The optimal bidding function in this special case is well known[3] to be $b(v, 0 \mid r, t, \gamma) = \frac{v}{1+r}$, so the ratio of bids to valuations is negatively correlated with the relative risk-aversion parameter. The *t* parameter capturing the shape of *u* in the loss domain is identified both from bids (for a fixed *v* and *f,* a larger *t* implies a larger bid) and from entry decisions (a smaller *t* implies more entry given a fixed *v* and *f*). The loss-aversion parameter is identified by fixing the valuation, and comparing bids with f=0 and f=1. The small increase in fee naturally increases the optimal bid, and it increases it more when γ is larger. Finally, the two standard deviations σ and τ of the errors are identified by departures from predictions.

Note the NYOP framework makes the individual-level identification much more tractable than an auction framework would: When bidders are heterogeneously risk averse, the first-price auction no longer has a symmetric pure-strategy equilibrium. Instead, bidder asymmetries in risk aversion lead to complicated asymmetric equilibria akin to the situation with asymmetric valuation distributions analyzed by Maskin and Riley (2000). Solving for such equilibria in the course of parameter estimation would be much more demanding than solving our equation 1. To appreciate

[3] In standard expected-utility models, risk aversion is known to increase bids in first-price sealed-bid auctions, both in theory (Riley and Samuelson, 1981) and in the laboratory (Cox, Smith, and Walker, 1988). Our model's concavity of the utility function in the gains domain leads to the same result. Bidders bid more because they experience diminishing marginal utility in surplus; therefore (compared to risk-neutral bidders), they prefer increased chances of winning associated with higher bids.

the complexity, please see the counterfactual section for an example of equilibrium bidding when two bidders are heterogeneously risk averse.

To estimate θ at the individual level, we maximize the likelihood of the individual's observations implied by equations 3 and 4. Because bidding more than $v_n - f_n$ is clearly irrational and assumed away in the above model, we use only the 22 observations with $v_n > f_n$ for each bidder. Two of the 48 bidders bid more than $v_n - f_n$ more than once, so we exclude them completely. Six other bidders bid more than or equal to $v_n - f_n$ exactly once, so we exclude these observations and only use the remaining 21 observations in our estimation.

The likelihood maximization is straightforward except for the need to solve the bidding problem in equation 1 at every step, which we accomplish by searching for the optimal bid on a fine grid. To keep the estimates in a realistic range, we conduct a constrained search for the maximum-likelihood estimates, allowing both risk-averse and risk-seeking preferences in gains and losses.[4] To estimate θ at the population level, we pool the observations of all 46 subjects together, and estimate θ using the same procedure under the same constraints.

6. Results: Parameters and Utility Functions

Table 1 shows the maximum likelihood estimates for the homogeneous model as well as several summary statistics about the maximum likelihood estimates for the heterogeneous model. Apparently, the heterogeneous model fits much better: Regarding entry, the average predicted probability of entry when the bidder actually entered is 0.89 in the heterogeneous model versus 0.77 in the homogeneous model. Regarding bidding, the best indication of a better fit is the lower σ estimate shown in Table 1.

---

[4] Specifically, we constrain $r, t \in [0, 2], \sigma \in [0.001, 20], \tau \in [0.1, 10]$, and $\gamma \in [0.5, 2]$.

Table 1: Parameter Estimates	Homogeneous Model Estimates (SE)	Heterogeneous Model Estimates			
		Mean	Median	SD	% > 1
r (curvature in the gain domain)	0.39 (0.01)	0.40	0.35	0.20	0%
t (curvature in the loss domain)	0.51 (0.03)	0.70	0.60	0.46	24%
γ (inverse of loss aversion)	1.65 (0.08)	2.64	1.81	1.77	72%
σ (error in the bids)	6.93 (0.26)	4.42	3.97	2.63	100%
τ (scaling parameter in entry model)	1.47 (0.13)	5.86	5.40	4.07	83%
N (number of observations)	1006	46 subjects, 21 or 22 obs./subject			

To interpret the results shown in Table 1, we begin by discussing the parameters of the homogeneous model. Because its *r* and *t* curvature parameters are both below 1, the homogeneous model involves an *S*-shaped utility function reminiscent of PT. In contrast to most prior calibrations of the prospect-theoretic value function, the homogeneous model does not find loss aversion, because $\gamma > 1$. For a recent survey of other work that also does not find loss aversion, see Walasek, Mullett, and Stewart (2018).

Figure 1: Estimated Utility Function of the Homogeneous Model

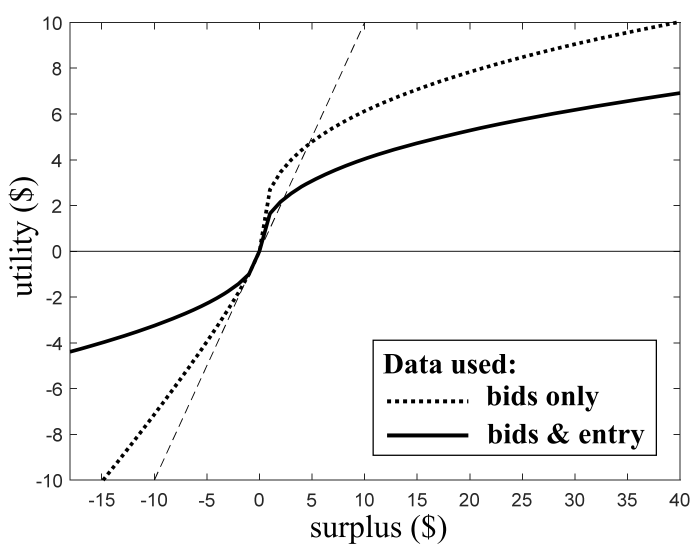

Whereas most prior analyses of bidding behavior consider only the submitted bids, we are able to also add data about the costly entry decisions. We find that considering the complete behavior has a strong and meaningful impact on inference as compared to considering only bids. Figure 1 plots the utility function estimated from the homogeneous model as a solid line, and also plots (as a dotted line) the utility function that using only the bidding data and omitting the observations in which bidders did not enter would suggest. Both functions are S-shaped, but the one based on the complete data is more curved in both losses and gains, and involves lower marginal sensitivity to losses.

We now turn to the parameters of the heterogeneous model. Table 1 does not include asymptotic standard errors for this model, because the small sample size involved in each estimation makes asymptotics invalid. The homogeneous model's estimated coefficients are clearly biased estimates of the respective population averages, but the central tendency of the curvature and gain-liking parameters is qualitatively consistent with the homogeneous model estimates. All coefficients exhibit wide variation in the population. For each parameter, the "% > 1" column shows the proportion of the population with an estimate that exceeds unity. It thus shows that all subjects are risk averse in the gains, but only 24% of them are also risk averse in the losses. Finally, only 28% of the subjects exhibit loss aversion. Figure 2 plots the joint distribution of the two curvatures, with each subject represented by a point, and the vertical line separating subjects whose preferences seem prospect theoretic (left of the line) and subjects whose preferences are consistent with risk-averse EUT (right of the line).

Our findings have three important implications for modeling bidder behavior. First, we clearly reject the risk-neutral model commonly used in the empirical literature on auctions with entry costs. Second, we show the homogeneous risk-averse expected utility model often suggested as a remedy to the known shortcomings of the risk-neutral model (e.g., Bajari and Hortaçsu, 2005; Bolotnyy and Vasserman, 2019) also does not fit our data well, especially when the entry behavior is considered alongside the observed bids. The risk-averse expected utility model does not fit the majority of our subjects even at the individual level. Finally, our results suggest the behavior of about three quarters of our subjects is well captured by a prospect-theoretic model without loss aversion. How important is it for market designers to account for all the heterogeneity we find? The next section outlines our strategy for answering this question.

Figure 2: Joint Distribution of Utility Curvatures *r* and *t* in the Bidder Population

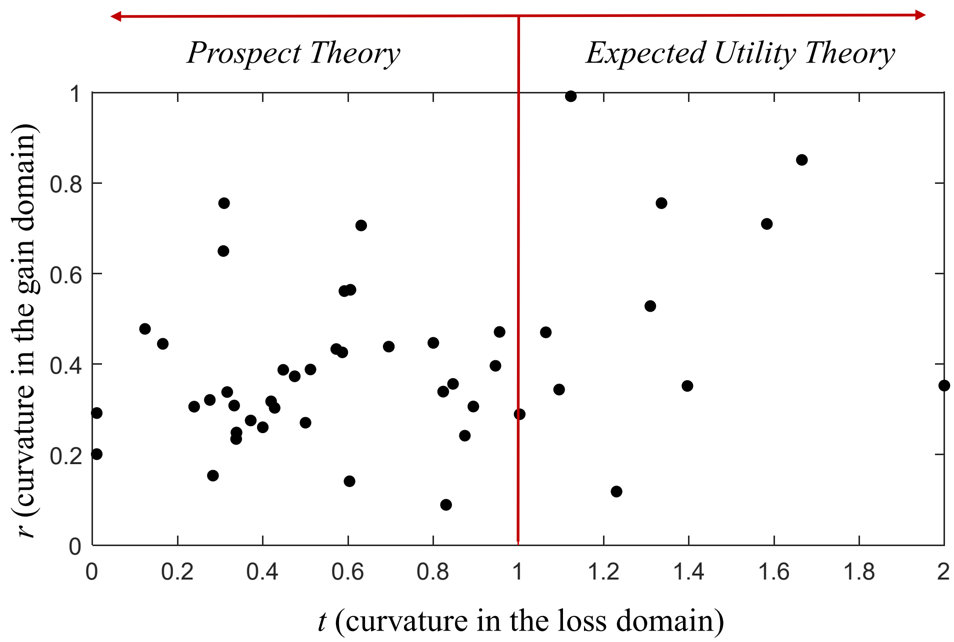

7. How Much Does Accounting for Preference Heterogeneity Increase Profits?

Having summarized the shapes of the utility functions our estimation finds, we now turn to the main question of managerial interest, namely, how important it is for auctioneers and NYOP sellers to account for the heterogeneity in risk preferences we document.

To make the above question concrete, consider an NYOP seller who would like to set the optimal bidding fee, and assume the seller must post this bidding fee publicly and thus cannot price discriminate. Such a seller could obtain a detailed panel dataset such as the one we use, estimate bidder preferences at the individual level, simulate each bidder's response to counterfactual levels of the bidding fee, and then select the most profitable level for the population. Alternatively, the seller could ignore preference heterogeneity and base the counterfactuals on a population-level homogeneous model. How much profit would the seller leave on the table by relying on the homogeneous model to set his fee? To answer this question, we propose the following procedure:

1) Let $\pi_{het}\left(f\right)=\frac{1}{I}\sum_{i=1}^{I}\pi\left(f\middle|\hat{\theta}_i\right)$ be the expected counterfactual profit suggested by the estimated parameters $\left\{\hat{\theta}_i\right\}_{i=1}^{I}$ of the fully heterogeneous model, where *i* indexes bidders and $\pi\left(f\middle|\theta\right)$ is the expected profit from a person with model parameters *θ*. Then, let $f_{het}^{*}=\arg\max_{f\geq 0}\pi_{het}\left(f\right)$.
2) Let $\pi_{hom}\left(f\right)=\pi\left(f\middle|\hat{\theta}_{hom}\right)$ be the expected counterfactual profit suggested by the estimated parameter $\hat{\theta}_{hom}$ of the homogeneous model (estimated on the same data as the heterogeneous model but pooling across bidder identities as in Table 1), and let $f_{hom}^{*}=\arg\max_{f\geq 0}\pi_{hom}\left(f\right)$.
3) Estimate the profit loss from relying on the homogeneous model as $\pi_{het}\left(f_{het}^{*}\right)-\pi_{het}\left(f_{hom}^{*}\right)$.

The key assumption behind step 3) is that we consider the heterogeneous model to be the "truth," so we can use it to estimate the profitability of arbitrary strategies, for example, strategies suggested by the homogeneous model or strategies suggested by the risk-neutral model. Importantly for the above procedure, the data used in the estimation and the specification of the individual-level model are the same under both models; only the model of heterogeneity in risk preferences changes between the objective functions being maximized. In reality, estimation of such a homogeneous model would require less data than the estimation of a heterogeneous model, but we want to isolate the effect of model specification throughout.

When the seller does not have to post the optimal bidding fee publicly, she can engage in first-degree price discrimination and charge higher fees to bidders with higher overall tolerance for risk, effectively setting bidder *i*'s fee to $f_i^{*}=\arg\max_{f\geq 0}\pi\left(f\middle|\hat{\theta}_i\right)$. Obviously, this strategy is only available to the seller who can estimate the heterogeneous model. We can estimate the joint profit improvement from having the heterogeneous model *and* being able to price discriminate in this way as $\sum_{i=1}^{I} \pi\left(f_i^* \middle| \hat{\theta}_i\right) - \pi_{het}\left(f_{hom}^*\right)$. This estimate of the profit gain from first-degree discrimination bounds the profit gain available from any other form of price discrimination, for example, second or third degree.

The key construct of our counterfactual strategy is $\pi\left(z \middle| \theta\right)$—the expected profit from a person with model parameters θ when the seller plays z. This profit depends on the distribution of valuations as follows:

$$\pi\left(f \middle| \theta\right) = E_{v,\varepsilon}\left\{ \Pr_{enter}\left(v, f \middle| r, t, \gamma, \tau\right)\left[f + \frac{\left(b\left(v, f \mid r, t, \gamma\right) + \varepsilon\right)^2}{140} \right]\right\}, \tag{5}$$

where the revenue from received bids is averaged over realizations of the seller's cost (distributed uniformly on [0,70]) that result in the bid being accepted:

$$\Pr\left(b > c\right) E_c\left(b - c \middle| b > c\right) = \int_0^b \left(b - c\right)\left(\frac{1}{70}\right) dc = \frac{b^2}{140}. \tag{6}$$

The experiment we analyze induces a discrete uniform distribution of valuations on {5, 20, 35, 50, 65}, approximating the distribution of seller cost (uniform [0,70]). But in the counterfactuals, we are free to use any distribution of valuations. To explore the impact of population dispersion of valuations on our conclusions, we consider uniform distributions $\left[35 - k, 35 + k\right]$ that all have the same mean but differ in variance of $k^2/3$ (variance of the uniform distribution on [*a,b*] is $\left(b - a\right)^2/12$). In some of our simulations, we also explore the impact of treating the econometric errors as measurement noise, and shutting them down in the profit estimation. Specifically, we sometimes let $\varepsilon = 0$ and $\tau = \infty$, and explore the impact on seller strategy. When the counterfactual setting differs considerably from the experimental setting, for example, in the case of a two-bidder auction with a reserve price, this "zero error" approach is the only one available. When the counterfactual setting is quite similar to the experiment, for example, in the case of an NYOP seller who uses bidding fees, we conduct the analysis both ways (with and without the econometric error terms). Note that we abstract away from our uncertainty about the parameter estimates, and use the plug-in MLE numbers throughout as $\hat{\theta}_i$. This approach is necessary because we do not have good estimates of the MLE estimation errors in our small-sample setting.

Having outlined our strategy for assessing the profit impact of accounting for the heterogeneity in risk preferences on the example of an NYOP seller who sets fees, we note the same strategy can be used to analyze other sellers, for example an NYOP seller who sets a minimum bid level or an auctioneer who sets a public reserve price. Table 2 collects all the simulation results for valuations distributed on [0,70]. We now cover each of these institutions in turn, starting with the bidding fee.

7.1. Results: Bidding Fees in NYOP

The closest counterfactual setting to the experimental situation is an NYOP monopolist who charges bidding fees, faces customers with private valuations distributed uniformly on [0,70], and incurs a production cost also distributed uniformly on [0,70]. The counterfactual simulation needs to interpolate behavior and the implied profit between the experimental levels of the bidding fee.

Figure 3 shows the simulated expected profit of an NYOP retailer for all levels of the fee in increments of 0.1. The solid (red) line shows the expected profit predicted by the heterogeneous model, and finds the fee of about 2.7 to be optimal, as highlighted by the left (red) vertical dashed line. The dashed (green) line shows the expected profit predicted by the homogeneous model, and suggests the fee of about 3.8 highlighted by the right (green) vertical dashed line. To assess the profit lost by a uniform-strategy seller who relies on the homogeneous model, we compare the heterogeneous model's predicted profit at f=3.8 of 6.82 with the maximal profit achieved at f=2.7 of 6.84, and conclude the loss is only 0.4%. To assess the profit lost by a perfectly discriminating seller, we compare the heterogeneous model's predicted profit at f=3.8 suggested by the homogeneous model with the "first degree bound" profit of 7.69 achieved by targeting fees to individual buyers, and find the loss is 11.3%.

Figure 3: Expected Profit of an NYOP Retailer Who Charges Bidding Fees

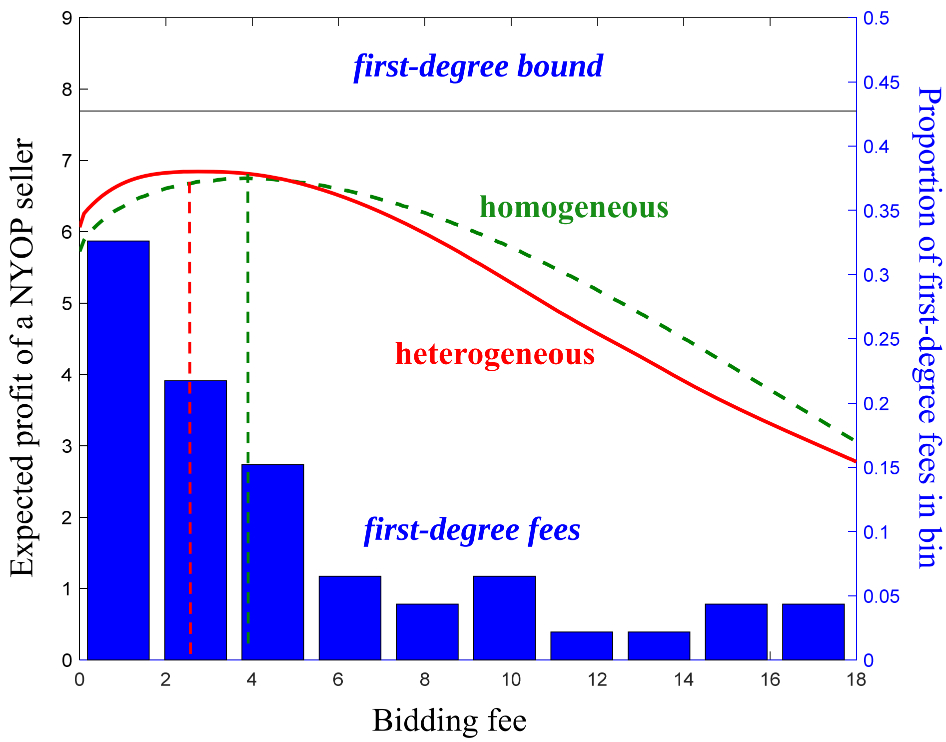

Note to figure: Valuations distributed uniformly on [0,70], econometric error terms included.

Both profit reductions are robust to the exclusion of the econometric error terms from the simulation: When we set the econometric errors to zero, the simulations reveal a 1.2% drop in profit for the uniform-strategy seller and an 11.6% reduction for the discriminating seller. We conclude that when the seller cannot discriminate, the profit loss from using the homogeneous model is minimal. On the other hand, a seller who can discriminate can increase profits by more than 10% by using the heterogeneous model.

Although we find the simplified model with homogeneous risk aversion is sufficient for optimization of uniform selling strategies, further simplification by assuming risk neutrality is not warranted. For example, a seller who relies on the risk-neutral model would select a fee of $40/7 \approx 5.7$, which would result in a 4% profit reduction relative to using the heterogeneous model without the econometric error terms.

Figure 4: Profit Loss from Using the Homogeneous Model to Set Bidding Fees as a Function of Variance of Valuations

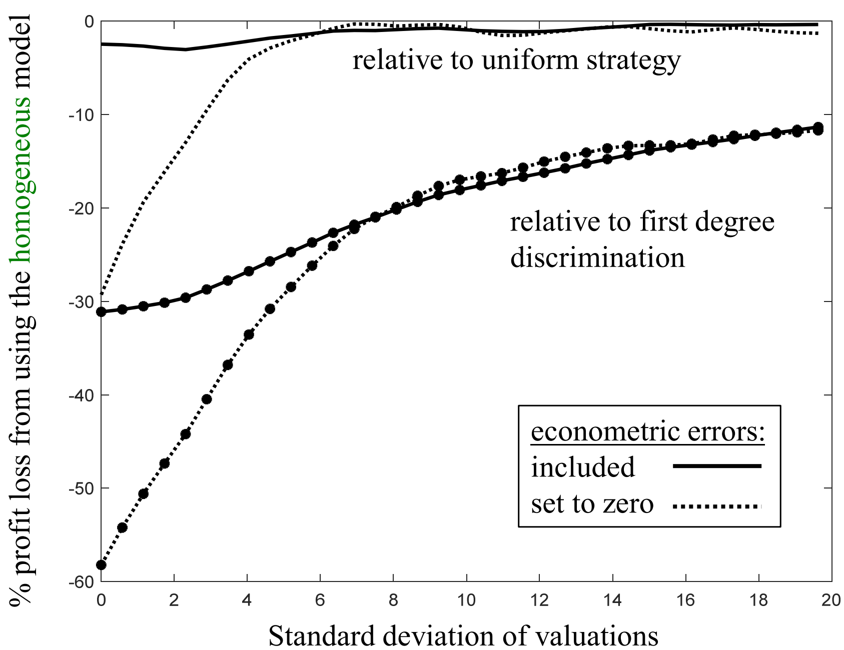

How does the above conclusion depend on the distribution of valuations? Figure 4 shows the percentage profit loss from using the homogeneous model for both types of sellers as a function of the standard deviation of valuations when the distribution of valuations is $[35-k, 35+k]$. In Figure 4, the case shown in Figure 3 is captured by the right endpoints of the solid lines—the simulations with the econometric error terms and largest possible standard deviation. Figure 4 shows the error from using the homogeneous model increases as the distribution of valuations becomes more concentrated around the mean, but remains in low single digits for sellers who use the uniform strategy (i.e., same fee for all buyers). In contrast, sellers who can discriminate leave over 30 percent of profit on the table by using the homogeneous model.

One notable feature of Figure 4 is the divergence between simulations with and without econometric error terms when the standard deviation of valuations is small. The simulations without error terms involve much larger losses because the error-free individual-level (and hence also the homogeneous) model's profit curves look more like Half Dome in Yosemite than the rounded shape in Figure 3 reminiscent of the North Dome just across the famous valley. The half-dome shape with a gentle slope on the left and a sharp drop from the summit as the fee increases beyond the optimal value occurs because each individual enters up to a certain fee level, above which the profit from that individual drops to zero. The summit of the homogeneous model is to the right (higher fee) of the summits of many individual-level models, which means the fee suggested by the homogenous model generates zero profit from many potential customers. Given this topography, a seller with access to the individual-level profit curves tends to select a smaller fee to get some profit from everyone.

7.2. Results: Minimum Bid in NYOP

Published reserve prices, aka "minimum bids," are a prominent feature of the optimal NYOP selling mechanism to risk-neutral bidders (Zeithammer, 2015). We can simulate the profit impact of a minimum bid *m* on the predicted bid by solving $bid\left(v \middle| r,m\right)=\arg\max_{b\geq m} b\left(v-b\right)^{r}$. Because a bidder facing a minimum bid cannot end up with a loss, the prediction does not involve the parameters τ, t, γ of the full model. Optimizing minimum bids is thus simpler than optimizing fees. On the other hand, the minimum bid setting is further from the experimental paradigm, so we need an assumption about the impact of econometric error terms on bidding. To incorporate the estimated error term ε, we truncate its distribution such that $bid\left(v|r,m\right)+\varepsilon \geq m$. We now turn to the results of the simulation.

Figure 5: Expected Profit of an NYOP Retailer Who Charges Bidding Fees

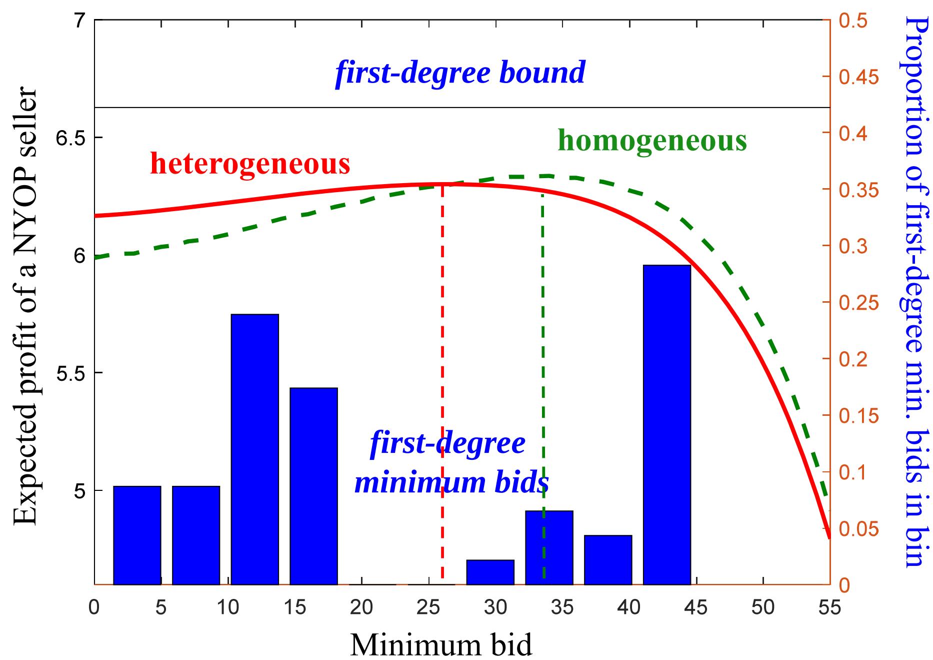

Figure 5 shows the analogue of Figure 3. The homogeneous model suggests a minimum bid of 34—substantially more than the optimal magnitude of 26. However, the heterogeneous-model profit is quite flat around the optimum, so the mistake made by the homogeneous model results in only a 0.5% profit reduction. As in the case of bidding fees, the seller who can tailor minimum bids to individual customers achieves a much higher profit, and charging everyone 34 as suggested by the homogeneous model results in a 5.4% reduction in profit compared to that first-degree bound. Both profit reductions are robust to the exclusion of the econometric error terms from the simulation: When we set ε=0, the simulations reveal no drop in profit for the uniform-strategy seller and a 4.1% reduction for the discriminating seller. The “no drop” result for the uniform-strategy seller arises from the fact that both variants of the model suggest the same zero fee as optimal. Contrary to basic intuition from auction theory, an NYOP seller facing risk-averse bidders can be better off without a bidding fee, because risk aversion increases bids so much that excluding any buyers is not profitable.

As in the case of bidding fees, further simplification by assuming risk neutrality is not warranted. A seller who relies on the risk-neutral model would select a minimum bid of $140/3 \approx 47$, which would result in a 7.7% profit reduction relative to using the heterogeneous model without the econometric error terms.

Figure 6: Profit Loss from Using the Homogeneous Model to Set Minimum Bids as a Function of Variance of Valuations

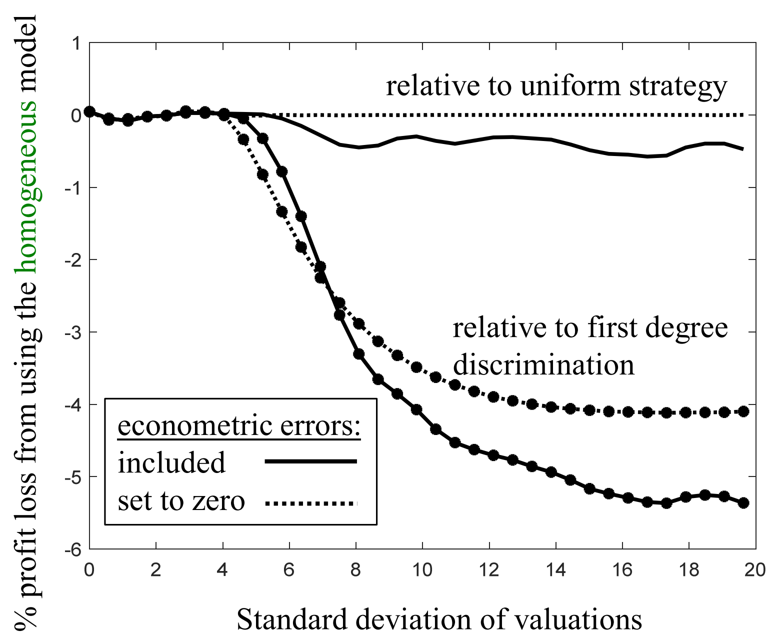

Figure 6 is the analogue of Figure 4. The profit loss is minimal for a uniform-strategy seller, and we find the above 5.4% reduction is the largest across all simulations. The reason the profit loss is zero for small standard deviations is that when valuations are concentrated around the mean, the optimal minimum bid is the bottom of the valuation support, and it is paid by all bidders. In other words, and as also discussed by Zeithammer (2015), a concentrated valuation distribution makes NYOP with a minimum bid equivalent to posted pricing.

7.3. Result: Reserve Price in a First-Price Sealed-Bid Auction

What is the profit foregone by using the simpler homogeneous model instead of the better heterogeneous model to set reserve prices in auctions? The homogeneous model has been analyzed extensively (e.g., Maskin and Riley, 1984; Hu, Matthews, and Zou, 2010), and risk aversion is well known to increase bids and reduce the optimal reserve price. On the other hand, the heterogeneous model has not received much attention to date, perhaps because of the difficulty in characterizing equilibrium bidding. The heterogeneity in risk aversion results in asymmetric equilibria in a fashion analogous to the effect of heterogeneity in valuation distributions analyzed by Maskin and Riley (2000). We numerically solve for equilibria in auctions with two bidders by adapting the approach from Hubbard and Paarsch (2009); please see the Appendix for details of our approach. Solving for bidding equilibrium with three or more bidders, each with a different risk-aversion coefficient, is not tractable to us.

Given the two-bidder limitation described above, we answer our question about the profit impact of correctly accounting for heterogeneity in risk aversion by comparing the profit from the two-bidder homogeneous model with the average profit from all possible pairs of bidders from our population. In terms of the "1-2-3" process introduced in the beginning of this section, i thus indexes pairs of bidders, and π is the auction revenue from a given pair of bidders. In each scenario, we assume the marginal procurement cost is zero (so profit equals revenue), and the valuations of the bidders are distributed uniformly on [0,1]. Because this scenario is the furthest from the experiment setup, more heroic assumptions are needed on the econometric error terms. Because we do not have any data on bidding errors in an auction without two bidders, we consider only the model without econometric error terms. As in the case of minimum bids, only the risk-aversion coefficients in the gain domain play a role here, so each auction setting can be captured by a pair of risk-aversion coefficients. Conceptually, a seller with a uniform strategy can only set the reserve price before knowing the pair of risk aversions participating in the auction. A discriminating seller knows the two risk aversions before setting the reserve. We now turn to the results of the simulation.

Figure 7: Expected Profit of an Auctioneer with Zero Marginal Cost and Two Bidders per Auction

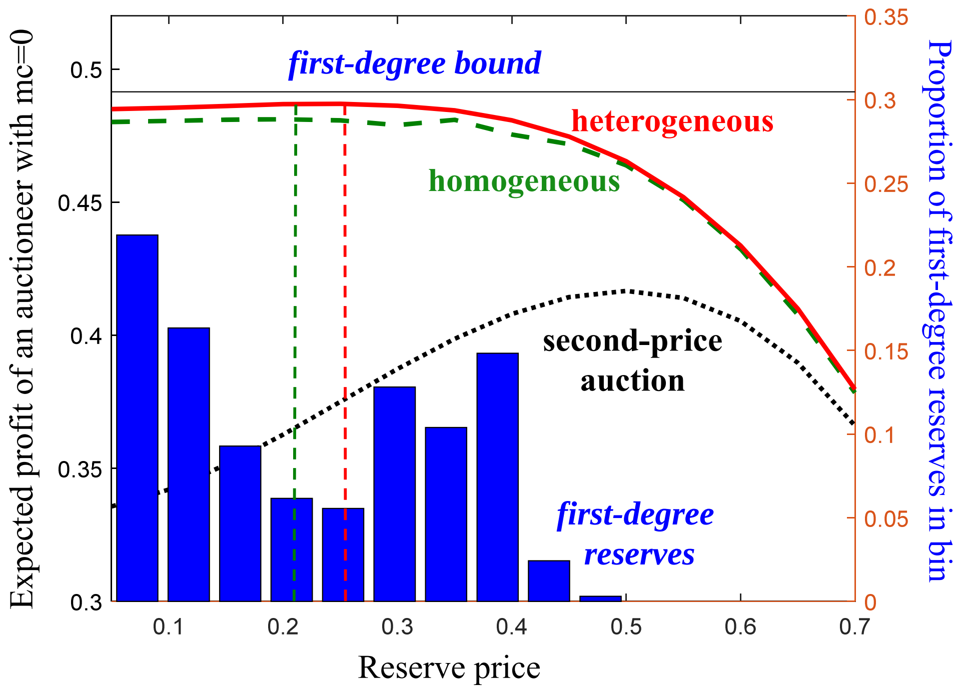

Figure 7 is the analogue of Figures 3 and 5. When we consider the uniform-strategy seller restricted to charging the same reserve price to all bidder pairs, Figure 7 conveys the same qualitative conclusion as Figure 5: The heterogeneous model suggests a reserve of 0.25, which results in an expected profit of 0.487. Relying on the homogeneous model instead (i.e., assuming both bidders have a constant relative risk aversion of 0.38) would make the seller set a smaller reserve price of 0.20 and collect 0.486—less than 0.1% lower than optimal. Again, relying on the risk-neutral model and setting the reserve of ½ would forego a lot more than relying on the homogeneous model, namely, 4.4% of profit.

In contrast to Figure 5, Figure 7 finds that a discriminating seller receives minimal benefit from the heterogeneous model: A seller who can customize the reserve price to the bidder pair would use a range of reserves and collect an average profit of 0.491—only 0.9% more than the seller who relies on the homogeneous model (and is thus restricted to a uniform strategy). One way to explain the difference between Figures 5 and 7 is that compared with NYOP, auctions involve the additional profit-enhancing force of bidder competition, and this additional force is a partial substitute for carefully optimizing the selling strategy (i.e., the reserve price). In fact, one way to interpret Figure 7 is that a seller facing risk-averse bidders can simply stop using the reserve price altogether without facing much downside.

Table 2: Percentage profit loss from using simpler models to optimize strategy

	Simpler model	**homogeneous** estimated		**risk-neutral**
		no errors	estimated errors	no errors
Uniform selling strategies	NYOP with bidding fee	-1.2%	-0.4%	-4.0%
	NYOP with minimum bid	0.0%	-0.5%	-7.7%
	First-price auction with reserve	-0.1%	n/a	-4.4%
Individualized selling strategies	NYOP with bidding fee	-11.6%	-11.3%	
	NYOP with minimum bid	-4.1%	-5.4%	
	First-price auction with reserve	-0.9%	n/a	

Note to table: seller profit attained, relative to strategy suggested by the heterogeneous model, as percent of potential. The results for the homogeneous model are presented both with the estimated econometric errors and without them.

Contrasting the profit curves for the focal first-price auction with those for the second-price (Vickrey) auction, which is not affected by risk aversion, is also useful. As Hu, Matthews, and Zou (2010) show, the optimal reserve in the first-price auction is lower than that in the second-price auction. Figure 7 also extends this ranking to the asymmetric auctions—all the first-degree reserves are below ½—the optimal reserve when valuations are uniform on [0,1]. As predicted by Maskin and Riley (1984), the revenue in the first-price auction is lower than that in the second-price auction. The magnitude of the difference implied by the extent of risk aversion we find is quite large: The optimized first-price auction delivers about 17% higher revenue than the optimized second-price auction.

8. General Discussion

Because bidding markets are typically governed by publicly known and strictly enforced rules, empirical researchers can infer bidders' preferences from bidding data by inverting the mapping from preferences to bids implied by the rules. In this paper, we inform the analysts about what assumptions they should use in formulating the model of individual behavior, and when they should exert the additional effort required to estimate the model at the individual level (or otherwise account for heterogeneity in risk aversion). We propose and estimate a flexible individual-level model of risk preferences using bidding data from an incentive-compatible laboratory implementation of a very simple single-bidder auction called name-your-own-price price (NYOP) selling. Thanks to the tractable laboratory paradigm, this paper is the first to identify individual-level risk preferences from complete bidding behavior (entry decisions *and* bid amounts). Our estimates imply a substantial population heterogeneity in risk preferences. The behavior of three quarters of our subjects is more consistent with prospect theory than with expected utility theory as borne out by their utility functions being convex in the loss domain. All individual-specific utility functions we estimated are concave in the gain domain.

Additionally, we used the estimated individual-level risk preferences for a set of counterfactual simulations in several bidding markets closely related to the paradigm from which the estimates were inferred. Specifically, we considered an NYOP seller who sets bidding fees for the right to submit a bid, an NYOP seller who sets a minimum bid level, and an auctioneer who sets a public reserve price. Another institution we analyzed but not reported in detail is an NYOP seller who offers a buy-it-now price alongside her bidding offer. We have not reported this analysis in detail because the results mirror those for the NYOP seller who uses only bidding fees.

Most empirical models of bidding behavior in auctions and NYOP assume bidder risk neutrality (Athey and Levin, 2007; Guerre, Perigne, and Vuong, 2000). With the exception of Campo (2012), all existing models that acknowledge buyer risk aversion assume a homogeneous model of risk preferences, in which all bidders have the same coefficient of risk aversion (e.g., Bajari and Hortaçsu, 2005; Abbas and Hann, 2010, Campo, Guerre, Perrigne, and Vuong, 2011; Lu and Perrigne, 2008; Bolotnyy and Vasserman, 2019). Our results show future analysts can keep calm and continue to ignore heterogeneity in risk preferences as long as they use their models to optimize non-discriminatory selling strategies, and as long as they allow for prospect-theoretic preferences whenever the bidders face the possibility of a loss. In other words, as long as institutional details prevent price discrimination on the basis of individual risk aversion, the analysts can use a homogeneous model to optimize the seller strategy (e.g., entrance fee or reserve price). The much simpler homogeneous model is sufficient for decision-making in that the strategy it suggests realizes over 99% of potential profit across all three market institutions we studied. See the upper left quadrant of Table 2 for a summary of the relevant profit losses by market institution and assumption about econometric error terms. In the domain of NYOP selling, these profit-loss numbers remain small across all the dispersions of the valuation distribution we studied, maxing out at less than 3 percent for the bidding fee, less than 1 percent for the minimum bid.

Although we find the simplified model with homogeneous risk aversion sufficient for optimization of uniform selling strategies, further simplification by assuming risk neutrality is not warranted. For all selling strategies we considered, the strategy derived from the risk-neutral model foregoes considerably more profit than the strategy derived from the homogeneous risk-averse model: between 4% and 7% depending on the strategy. See the upper right quadrant of Table 2 for a summary of the relevant profit losses by market institution.

In contrast to the homogeneous model's sufficiency for optimizing uniform strategies, we find large profit gains from identifying risk preferences at the individual level when the seller can price discriminate, especially in the market institution of NYOP with bidding fees. The magnitude of gains from price discrimination depends both on the selling strategy and on the dispersion of valuations in the bidder population. For the bidding-fee strategy in the NYOP setting, the profit loss from not discriminating increases to over 30% as the distribution of valuations shrinks to its mean. For the minimum-bid strategy in the NYOP setting, dispersion of valuations acts in the opposite direction, increasing the profit loss from not discriminating to about 5%.

The homogeneous model we find sufficient for decision-making is empirically identified in the auction setting with exogenously varying participation, as shown by Bajari and Hortaçsu (2005). We now briefly outline how their approach can be adapted to the context of NYOP selling without bidding fees by exploiting variance in the probability of bid acceptance: When the probability of an NYOP bid acceptance is $G(b)$ instead of the Uniform[0,70] used in the experiment we analyze, the first-order condition of the bidding problem in equation 1 is $v = b + r\frac{G(b)}{g(b)}$. With two different probabilities of acceptance $G_j(b)$, j=1,2, the α-th percentiles of the two respective bid distributions $b_{j,\alpha}$ can be used as in equation 15 of Bajari and Hortaçsu (2005) to estimate r via the following linear regression:

$$b_{1,\alpha} - b_{2,\alpha} = r\left[\frac{G_2\left(b_{2,\alpha}\right)}{g_2\left(b_{2,\alpha}\right)} - \frac{G_1\left(b_{1,\alpha}\right)}{g_1\left(b_{1,\alpha}\right)}\right] + \varepsilon_\alpha . \tag{7}$$

In settings where the seller has control over $G_j\left(b\right)$, he can select particularly tractable pairs of distributions. For example, letting $G_2\left(b\right)$ be the Uniform[0,*M*] and letting $G_1\left(b\right) = \left(x/M\right)^2$ (the increasing triangle distribution on [0,*M*]) simplifies the expression in the square bracket in equation 7 to $b_{2,\alpha} - \frac{b_{1,\alpha}}{2}$, and suggests a very simple estimator of *r*, namely, $\hat{r} = E_\alpha\left[\frac{2\left(b_{1,\alpha} - b_{2,\alpha}\right)}{2b_{2,\alpha} - b_{1,\alpha}}\right]$, with the expectation taken over different levels of *α*. Of course, the simple linear form of equation 7 is a special feature of the constant relative risk-aversion utility—other functional forms involve a more complicated relationship.

We acknowledge some limitations that provide avenues for future research. In some of our counterfactual simulations, we explored the impact of treating the econometric errors as measurement noise and shutting them down in the profit estimation. However, we could only provide these analyses for counterfactual settings that are more similar to the experiment that generated the data we used, for example, an NYOP seller who uses bidding fees. For counterfactual settings that differed more substantially from the setup of the original experiment (e.g., a two-bidder auction with a reserve price), we are unable to provide the corresponding analyses. This inability is a limitation of the data we use. Examining our counterfactual simulations experimentally would also be very interesting. One could, for example, design an experiment in which, first, sufficient data are generated to estimate the individual risk preferences of the subjects. In a subsequent, second part of the experiment, using the same set of subjects, this information could then be used to test the predictions of our model for the different pricing mechanisms.

References

Abbas, Ali E. and Il-Horn Hann (2010), "Measuring Risk Aversion in a Name-Your-Own-Price Channel," Decision Analysis, 7 (1), 123–136.

Abdellaoui, Mohammed (2000), "Parameter-Free Elicitation of Utility and Probability Weighting Functions," Management Science, 46 (11), 1497–1512.

Abdellaoui, Mohammed, Han Bleichrodt, and Corina Paraschiv (2007), "Loss Aversion Under Prospect Theory: A Parameter-Free Measurement," Management Science, 53 (20), 1659–1674.

Allenby, Greg M. and Peter E. Rossi (1999), "Marketing Models of Consumer Heterogeneity," Journal of Econometrics, 89 (1–2), 57–78.

Athey, Susan and Jonathan Levin (2001), "Information and Competition in U.S. Forest Service Timber Auctions," Journal of Political Economy, 109 (2), 375–417.

Athey, Susan and Philip A. Haile (2007), "Nonparametric Approaches to Auctions," In: James J. Heckman and Edward E. Leamer, eds., Handbook of Econometrics, Vol. 6. Amsterdam: Elsevier.

Bajari, Patrick and Ali Hortaçsu (2005), "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data," Journal of Political Economy, 113 (4), 703–741.

Barseghyan, Levon, Francesca Molinari, and Joshua C. Teitelbaum (2016), "Inference Under Stability of Risk Preferences," Quantitative Economics, 7, 367–409.

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum (2013), "The Nature of Risk Preferences: Evidence from Insurance Choices," American Economic Review, 103 (6), 2499–2529.

Bolotnyy, Valentin and Shoshana Vasserman (2019), "Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement," Working Paper.

Campo, Sandra (2012), "Risk Aversion and Asymmetry in Procurement Auctions: Identification, Estimation and Application to Construction Procurements," Journal of Econometrics, 168 (1), 96–107.

Campo, Sandra, Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong (2011), "Semiparametric Estimation of First-Price Auctions with Risk-Averse Bidders," Review of Economic Studies, 78 (1), 112–147.

Chakravarty, Sujoy, Harrison, Glenn W., Haruvy, Ernan E., and Rutström, E. Elisabet (2011), "Are You Risk Averse over Other People's Money?" Southern Economic Journal, 77 (4), 901– 913.

Chiappori, Pierre-André, Bernard Salanié, François Salanié, Amit Gandhi (2019), "From Aggregate Betting Data to Individual Risk Preferences," Econometrica, 87 (1), 1–36.

Cicchetti, Charles J. and Jeffrey A. Dubin (1994), "A Microeconometric Analysis of Risk Aversion and the Decision to Self-Insure," Journal of Political Economy, 102 (1), 169–186.

Cohen, Alma and Liran Einav (2007), "Estimating Risk Preferences from Deductible Choice," American Economic Review, 97 (3), 745–788.

Cox, James C., Vernon L. Smith, and James M. Walker (1988), "Theory and Individual Behavior of First-Price Auctions," Journal of Risk and Uncertainty, 1 (1), 61–99.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2002), "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," Journal of Economic Theory, 104 (1), 247–272.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong (2000), "Optimal Nonparametric Estimation of First-Price Auctions," Econometrica, 68 (3), 525–574.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong (2009), "Nonparametric Identification of Risk Aversion in First-Price Auctions Under Exclusion Restrictions," Econometrica, 68 (3), 525–574.

Holt, Charles A. and Susan K. Laury (2002), "Risk Aversion and Incentive Effects," American Economic Review, 92 (5), 1644–55.

Hu, Audrey, Steven A. Matthews, and Liang Zou (2010), "Risk Aversion and Optimal Reserve Prices in First- and Second-Price Auctions," Journal of Economic Theory, 145 (3), 1188–1202.

Hubbard, Timothy P. and Harry J. Paarsch (2009), "Investigating Bid Preferences at Low-Price, Sealed-Bid Auctions with Endogenous Participation," International Journal of Industrial Organization, 27 (1), 1–14.

Kaplan, Todd R. and Shmuel Zamir (2012), "Asymmetric First-Price Auctions with Uniform Distributions: Analytic Solutions to the General Case," Economic Theory, 50 (2), 269–302.

Laffont, Jean-Jacques and Quang Vuong (1996), "Structural Analysis of Auction Data," American Economic Review, 86 (2), 414–420.

Lu, Jingfeng and Isabelle Perrigne (2008), "Estimating Risk Aversion from Ascending and Sealed-Bid Auctions: The Case of Timber Auction Data," Journal of Applied Econometrics, 23 (7), 871–896

Maskin, Eric and John Riley (1984), "Optimal Auctions with Risk-Averse Buyers," Econometrica, 52 (6), 1473–1518.

Maskin, Eric and John Riley (2000), "Asymmetric Auctions," Review of Economic Studies, 67 (3), 413–438.

Palfrey, Thomas R. and Svetlana Pevnitskaya (2008), "Endogenous Entry and Self-Selection in Private Value Auctions: An Experimental Study," Journal of Economic Behavior & Organization, 66 (3–4), 731–747.

Riley, John G. and William F. Samuelson (1981), "Optimal Auctions," American Economic Review, 71 (3), 381–392.

Samuelson, William F. (1985), "Competitive Bidding with Entry Costs," Economics Letters, 17 (1–2), 53–57.

Sutton, John (1993), "Much Ado About Auctions: Some Introductory Remarks," European Economic Review, 37 (2–3), 317–319.

Sydnor, Justin (2010), "(Over)insuring Modest Risks," American Economic Journal: Applied Economics, 2 (4), 177–199.

Todd R. Kaplan and Shmuel Zamir (2012), "Asymmetric First-Price Auctions with Uniform Distributions: Analytic Solutions to the General Case," Economic Theory, 50 (2), 269–302.

Tversky, Amos and Daniel Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty," Journal of Risk and Uncertainty, 5 (4), 297–323.

von Gaudecker, Hans-Martin, Arthur van Soest, and Erik Wengström (2011), "Heterogeneity in Risky Choice Behavior in a Broad Population," American Economic Review, 101 (2), 664–694.

Walasek, Lukasz, Timothy L. Mullett, and Neil Stewart (2018), "A Meta-Analysis of Loss Aversion in Risky Contexts," Working Paper.

Zeithammer, Robert (2015), "Optimal Selling Strategies When Buyers Name Their Own Prices," Quantitative Marketing and Economics, 13 (2), 135–171.

Zeithammer, Robert, Lucas Stich, Martin Spann, and Gerald Häubl (2019), "Consumer Behavior in Bidding Markets with Participation Costs," Working Paper.

Appendix: Bidding Equilibrium in a First-Price Sealed-Bid Auction with a Reserve and Two Bidders with Different Risk Aversions

Two bidders exist with valuations distributed *iid* uniformly on the [0,1] interval. The two bidders differ in their relative risk aversion r_i. Similar to the approach in Kaplan and Zamir (2012), let $v_i(b)$ be the inverse bidding function of bidder *i*, and look for a pure-strategy equilibrium in which both bidders bid inside an interval $\left[R,\bar{b}\right]$, with initial conditions $v_i(R)=R$. Bidder *i* solves

$$\max_{b\geq R} v_{-i}(b)(v_i-b)^{r_i}. \tag{8}$$

When the $b \geq R$ constraints are not binding, first-order conditions are necessary for optimality, and they imply the two inverse bidding functions must satisfy the following system of differential equations:

$$\begin{aligned} v_1'(b)\left(v_2(b)-b\right)-r_2 v_1(b)=0 \\ v_2'(b)\left(v_1(b)-b\right)-r_1 v_2(b)=0 \end{aligned}. \tag{9}$$

To compute expected revenue, we only need the product of the two inverse bidding functions, because $G(x) = \Pr(\beta_1(v_1) < x \,\&\, \beta_2(v_2) < x) = v_1(x)v_2(x)$ is the cumulative distribution function (cdf) of the revenue. Given *G*, the expected revenue with reserve price *R* is

$$\begin{aligned} \int_R^{\bar{b}} x dG(x) &= -\int_R^{\bar{b}} x(1-G(x))' \, dx = -\bar{b}\underbrace{\left(1-G(\bar{b})\right)}_{0} + R\underbrace{\left(1-G(R)\right)}_{\Pr(v_i<h)} + \int_R^{\bar{b}} 1-G(x)dx \\ &= R\left(1-R^2\right) + \int_R^{\bar{b}} 1 - v_1(x)v_2(x)dx \end{aligned} , \qquad (10)$$

where the inverse bidding functions implicitly also depend on *R*.

For each $\{r_1, r_2, R\}$, we use Chebyshev polynomials of degree 5 to approximate the pair of inverse bidding functions as follows: Given a guess of the upper bound of bid support $\bar{b}$, we represent each inverse bidding function on a fine grid of bids $\{b_1 = R, \ldots, b_k, \ldots \bar{b}\}$, and we minimize the sum of squares implied by the above first-order condition:

$$SSQ\left(r_1, r_2 \middle| \bar{b}\right) = \sum_i \sum_k \left[v_i'(b_k)\left(v_{-i}(b_k) - b_k\right) - r_{-i} v_i(b) \right]^2 . \qquad (11)$$

To find the upper bound $\bar{b}$, we try every possible value on a fine grid, and select the one that minimizes the $SSQ\left(r_1, r_2 \middle| \bar{b}\right)$ over all possible $\bar{b}$. With estimates of equilibrium inverse bidding functions in hand, we then estimate the key integral $\int_R^{\bar{b}} 1 - v_1(x)v_2(x)dx$ in equation 10 by the appropriate sum.