

Do Auctioneers Pick Optimal Reserve Prices?

Andrew M. Davis, Elena Katok, Anthony M. Kwasnica

Smeal College of Business, Pennsylvania State University, University Park, Pennsylvania 16802
{adavis@psu.edu, ekatok@psu.edu, kwasnica@psu.edu}

We investigate how auctioneers set reserve prices in auctions. A well-established theoretical result, assuming risk neutrality of the seller, is that the optimal reserve price should not depend on the number of participating bidders. In a set of controlled laboratory experiments, we find that seller behavior often deviates from the theoretical benchmarks. We extend the existing theory to explore three alternative explanations for our results: risk aversion, anticipated regret, and probability weighting. After fitting our data to each of these models through parameter estimation techniques on both an aggregate and individual level, we find that all three models are consistent with some of the characteristics of our data, but that the regret model provides a slightly more favorable fit overall.

Key words: reserve prices; procurement auctions; behavioral operations

History: Received August 10, 2009; accepted May 23, 2010, by Teck-Hua Ho, decision analysis. Published online in *Articles in Advance*.

1. Introduction

What about the sellers? Although there has been extensive effort devoted toward understanding whether buyers (bidders) in auctions approximate the predictions of the positive theory of auction bidding, little effort has been devoted toward understanding whether sellers (bid takers) follow the normative prescriptions for good auction design. In almost all instances of experimental observation of auction behavior, the role of the seller is played by the experimenter and the rules of the auction are predetermined. This is despite the fact that there are many well-known instances in which observed seller behavior appears to run contrary to normative theoretical prescriptions. For example, consider the case of secret reservation prices, the use of which cannot be justified by standard theoretical results. In spite of mounting evidence that secret reserve prices are detrimental to sellers (Katkar and Reiley 2005), Anwar et al. (2006) report that over 6% of eBay auctions in their large-scale study used secret reserves.

Mechanism design, specifically optimal auction design, provides strong recommendations about how a seller, who is interested in maximizing her expected profits, should conduct an auction. In seminal papers, Myerson (1981) and Riley and Samuelson (1981) characterize the optimal design of auctions. In addition to such important principles as revenue equivalence, both papers observe that, in general, an expected revenue maximizing seller should set a *nonnegligible* reserve price, one that exceeds her own value for the object, whenever the bidders have independent private values (IPV). An important implication of this

result is that the auction that maximizes the expected revenue for the seller will sometimes be inefficient, meaning that the object will occasionally not be sold despite the potential for mutually beneficial trade. Remarkably, whenever there are IPVs, this optimal reserve price should not depend upon the number of bidders at the auction.

Whereas the optimal reserve does not depend upon the number of bidders, the number of bidders does impact the importance of the reserve price decision. When there is a large number of bidders, the reserve price is not likely to matter; therefore, the seller's choice of a reserve price has a vanishingly small impact on his or her actual expected revenue from the auction. On the other hand, when the number of bidders is small, the selection of an optimal reserve price can be critical to obtaining substantial revenue. For example, if there is only one bidder, it is the reserve price that determines the sale price at the auction. Given that many auctions are conducted with very few real bidders (Anwar et al. 2006 report that about 17% of eBay auctions in their sample did not attract any bids), the selection of a suboptimal reserve price can mean substantial foregone revenue.

The objective of this study is to better understand how human sellers set reserve prices in auctions. Specifically, we wish to compare the behavior of sellers in our laboratory experiment to normative theoretical benchmarks. In practice, there are many reasons why a seller may select seemingly suboptimal reserve prices. For example, if, as has been thoroughly documented in the literature, bidders are not following the predictions of game theory, then the seller's best

reply may be to respond with a reserve price that is different from what the standard theory suggests (see Crawford et al. 2009). In addition, the real-world auction environment may involve a number of complications that are not a part of the standard model. For example, competition between sellers can result in efficient (lowered) reserve prices (McAfee 1993), and nonindependent (affiliated) values can make it optimal for the seller's reserve price to converge to their personal valuation as more bidders arrive (Levin and Smith 1996). Because it is nearly impossible to disentangle all of these competing reasons for reserve prices that do not match predictions of the standard theory in practice, we turn to the laboratory for a more controlled examination. By conducting auctions in the laboratory with well-defined value distributions, payoffs, and rules, we create an environment that mimics the model presented in the standard theory as closely as possible. Furthermore, by using simulated bidders who are programmed to follow the game theoretic predictions of risk-neutral bidders, we can eliminate any strategic complications caused by the interaction between a human seller and human bidders. Specifically, we conducted second-price sealed bid auctions (described to the seller as English auctions) with simulated bidders, and systematically varied the number of bidders, the identical value distributions of the bidders, and the information provided to the sellers.

Whereas a majority of the experimental auction literature focuses on buyers (see Kagel 1995 for a comprehensive survey of laboratory auction research), our research centers on the seller's perspective. There have only been a select few works that investigate research similar to this. Greenleaf (2004) considers regret and rejoicing in English auctions with secret reserve prices and finds evidence that this behavioral model influences how reserve prices are set. Chen et al. (2005) consider the effect of ambiguity on reserve prices with a fixed number of bidders and varying auction formats. Our work extends and contributes to the literature in two ways. First, we systematically examine how sellers set reservation prices in the laboratory when faced with a different number of bidders and different distributions of bidder values. Second, we develop new theory that organizes our data better than the standard theory does.

Our laboratory results are twofold: We find evidence that, qualitatively, the data matches certain aspects of the standard theory well. Importantly, when a value distribution dictates lower optimal reservation prices, we observe significantly lower average reservation prices in our experiment. Other features of our data, however, are inconsistent with standard theoretical benchmarks. First, overall reservation prices tend to be lower than risk-neutral optimal. Although this feature can be potentially explained by risk aversion, we

also find that reservation prices increase with the number of bidders, to the extent that for a large number of bidders, average reservation prices are (contrary to what should be observed with risk-averse sellers) above risk-neutral optimal levels. Moreover, although we find a substantial amount of individual heterogeneity in behavior, we observe the above regularities on both an aggregate and individual level.

We proceed to explore several models of seller behavior that might rationalize the experimental data we observe. We consider the impact of risk aversion, anticipated regret, and probability weighting, on how sellers set reservation prices, and find that on an aggregate and individual level, all three models explain some characteristics of our data, but that the anticipated regret model appears to be favored overall.

In the following section we provide an overview of our experimental design and theoretical benchmarks. In §3, we provide a summary of the results of our experiment. In §4, we investigate two alternative models and, in §5, we conclude with a final summary and future research.

2. Experimental Design

We begin by providing a brief background of the reserve price benchmarks for risk-neutral and risk-averse sellers. We then follow this with a summary of our experimental procedure and implementation.

2.1. Risk-Neutral Benchmarks

We investigate the following auction setting. There are n bidders with values $v_i \in [\underline{v}, \bar{v}]$ drawn independently from the identical distribution F with density given by $F' = f$. We assume that $\phi(v) = (1 - F(v))/f(v)$ is decreasing for all v or $\phi(v)$ is concave.¹ Bidders place bids (drop-out prices) to purchase a single indivisible object from the seller in a second-price sealed bid auction (English). Let $b = (b_1, \dots, b_n)$ be the vector of bids. It is a weakly dominant strategy for all bidders to bid $b(v_i) = v_i$ if $v_i \geq r$ and 0 otherwise. The seller must determine the reserve price r in the auction given that bidders are assumed to follow this strategy.

The optimal reserve price under risk neutrality (which we will refer to as the standard theory hereafter) is given by

$$r^* = v_0 + \frac{(1 - F(r^*))}{f(r^*)}, \quad (1)$$

where v_0 is the seller's personal valuation for the object (Myerson 1981, Riley and Samuelson 1981).

¹ These assumptions are sufficient to ensure that the second-order conditions are satisfied for an interior maximum in the standard theory.

In our experiment some sellers always faced bidders whose values were drawn from the cumulative distribution given by $F(v) = (v/100)^{1/3}$, which we denote as the *Cuberoot* treatment; the remaining sellers faced bidders whose values were drawn from the cumulative distribution given by $F(v) = (v/100)^3$, which we denote as the *Cube* treatment. We assume that the personal value of the object to the seller is equal to zero. Therefore, the optimal reserve prices under our two distributions are given by 42 (Cuberoot) and 63 (Cube).

We chose these distributions because they provided valuations in which the majority of values were low (Cuberoot, mean of 25, standard deviation of 28.4) and in which the majority of values were high (Cube, mean of 75, standard deviation of 19.4). Furthermore, these distributions were simple for subjects to comprehend, and, importantly, the expected values of each distribution did not match the optimal reserve price (as is the case with the uniform distribution). To ensure that subjects understood the value distribution, subjects were shown a table of 100 values from the distribution along with a histogram showing the frequencies based on 10,000 randomly generated values. See the online appendix (provided in the e-companion) for sample instructions.²

2.2. Risk-Aversion Benchmarks

Let $u(x)$ be the Bernoulli utility function for the seller. Assume without loss of generality that $u(0) = 0$. The seller's ex post utility from the auction is given by

$$v(b; r) = \begin{cases} u(b_{(2)}) & \text{if } b_{(2)} > r, \\ u(r) & \text{if } b_{(1)} > r > b_{(2)}, \\ 0 & \text{if } r > b_{(1)}, \end{cases} \quad (2)$$

where the notation $x_{(k)}$ is taken to indicate the k th highest element (bid or value) of an n -dimensional vector. Given (2) and the bidders' strategy, the seller's expected utility for an auction with reserve price r is given by

$$Ev(r) = u(r)G(r) + \int_r^{\bar{v}} u(v)h(v)dv, \quad (3)$$

where $G(v) = n(1 - F(v))F^{n-1}(v)$ is the probability that one value exceeds v and all others fall below v , $g(v) = G'(v)$ is the associated density, and $H(v) = F^n(v) + G(v)$ is the distribution of the second highest value with $h(v) = H'(v)$ the density. Then, we have the following necessary first-order condition:

$$\frac{dEv(r^*)}{dr} = u'(r^*)G(r^*) + u(r^*)g(r^*) - u(r^*)h(r^*) = 0. \quad (4)$$

Simplifying we obtain the following:

$$\frac{u(r^*)}{u'(r^*)} = \frac{(1 - F(r^*))}{f(r^*)}. \quad (5)$$

Note this shows that the same basic relationship holds for risk-averse preferences (concave $u(x)$) as it does for risk neutrality. In particular, the optimal reserve price for a risk-averse seller is also independent of the number of bidders n . When $u(x)$ is linear, this condition simplifies to the well-known standard optimal reserve price condition given in (1). Let r_u^* be the optimal reserve price associated with a seller with Bernoulli utility function u , where r^* is assumed to refer to the risk-neutral seller's optimal reserve price.

PROPOSITION 1. For all (strictly) risk-averse u , $r_u^* < r^*$.

PROOF. Let $m(r) = u(r)/u'(r)$ and note that $m(0) = 0$ by assumption. Also note $m'(r) = 1 - (u(r)u''(r)/u'(r)^2)$. Because we assume u to be strictly risk averse, we know that u is strictly concave so $u''(r) < 0$, and it must be that $m'(r) > 1$. But this obviously implies that $m(r)$ intersects $(1 - F(r))/f(r)$ (and thus satisfies the necessary condition given by (5) at a lower r than the intersection implied by (1) or $r_u^* < r^*$). This proposition follows from Waehrer et al. (1998). \square

Constant relative risk-aversion (CRRA) preferences have been used extensively to study risk aversion in both laboratory settings and empirical studies (Binswanger 1980, Chen and Plott 1998, Campo et al. 2011). CRRA utility functions are given by $u(x) = x^{1-\alpha}$ where $0 \leq \alpha < 1$ is the measure of relative risk aversion. If values are drawn from the Cuberoot distribution, then the risk-averse optimal reserve price is given by

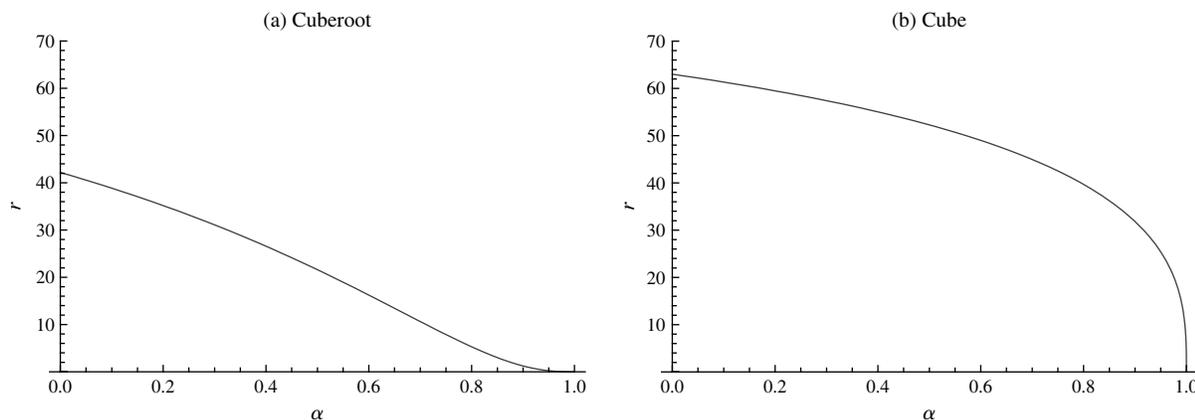
$$r^* = \left[\frac{27(1 - \alpha)^3}{(4 - 3\alpha)^3} \right] \times 100, \quad (6)$$

and when values are drawn from the Cube distribution, then the risk-averse optimal reserve is given by

$$r^* = \left[\frac{(1 - \alpha)^{1/3}}{(4 - \alpha)^{1/3}} \right] \times 100. \quad (7)$$

In Figure 1 we plot the optimal reserve price under the assumption of CRRA preferences for different values of α . Note that the risk-aversion parameter impacts the choice of an optimal reserve price in a nonlinear fashion. Although previous studies have shown estimates for the level of risk aversion (α) to be between 0.45 and 0.67 and to vary depending on the setting (Cox and Oaxaca 1996, Goeree et al. 2002, Goeree and Holt 2004), all levels of risk aversion for which $0 \leq \alpha < 1$ result in optimal reserves being below risk-neutral theoretical optimums. Additionally, in a range of risk aversion typically observed in the laboratory, a risk-averse seller's reserve price will vary less from the optimal value under the Cube distribution than under the Cuberoot distribution.

² An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Figure 1 Optimal Reserves with CRRA Utility Differing Levels of Relative Risk Aversion α 

2.3. Experimental Implementation

Recall that in our experiments each subject acted as a seller of a single, indivisible object with a value to the seller of zero, and the object was offered for sale via a second-price (or English) auction with n potential buyers. The seller's task was to determine a reserve price $r \in [0, 100]$ below which the object would not be sold.³

The n buyers were computerized and programmed to follow the weakly dominant strategy of setting their bid equal to their valuation contingent upon their value being equal to or above the chosen reserve price. They had a private value $v \in [0, 100]$ drawn from an identical distribution (Cube or Cuberoot).

Each seller participated in 60 separate auction periods. In each auction period, the seller faced a new set of bidders with valuations drawn independently from the same distribution. The seller was also informed of the number of computerized bidders he would face in advance of the auction period. We manipulated the number of bidders that a seller faced so that in one treatment $n \in \{1, 2, 3, 4\}$ and in the other treatment $n \in \{1, 4, 7, 10\}$. For every period, we selected the number of bidders n uniformly (with replacement) from one of these sets. We chose the first treatment because the reserve price decision is most salient when the seller is facing only a small number of bidders. We conducted the second treatment to determine if the effects we observed for small numbers of bidders persisted in auctions with more bidders. Under both treatments, the standard theory for either risk-neutral or risk-averse sellers predicts that the reserve price is invariant in n .

We described the second-price auction to the human sellers as an English auction. Because subjects, in order to maximize their profit, had to perform a number of complex calculations on their own regarding order statistics and probabilities for this auction,

³ Reserve prices were restricted to be integer valued.

we also conducted two additional treatments where we provided the subjects with a testing field. In this field, they could enter a reserve price and observe the following information for that specific reserve price: the probability of not selling the object, the probability of selling at the reserve price, the probability of selling above the reserve price, the average selling price if sold, and the average revenue from the stated reserve price. For these two treatments, we investigated both the Cuberoot and Cube distributions with $n \in \{1, 4, 7, 10\}$. We will refer to these treatments as *FullInfo* and the former as *NoInfo*.

For all treatments, following the determination of a reserve price in each round, we provided the following information to the sellers: the drop-out prices for the losing bidders (if the drop-out price was equal to or above r), whether some buyers were forced to drop out because their values were below r , the winning bid (r or the highest drop-out price above r), and the seller's resulting profit.⁴

After participants completed 60 periods of the auction, but before subjects received their cash earnings, we used a modified version of the instrument introduced by Holt and Laury (2002) to elicit a measure of each subject's level of risk aversion. The instrument involved 10 lottery pairs. Option A in each pair was a 50/50 chance of winning \$4.50 or \$5.50. Option B in each pair involved a 50/50 chance of winning \$1 or \$X, where X varied from \$9 in pair 1 to \$20 in pair 10 (see the online appendix for instructions from this phase). Subjects were asked to select one option for each pair. Immediately following this choice, a single chosen lottery was selected at random and played by the computer. The earnings from the lottery outcome were added to the subject's earnings from the auctions. A more risk-averse person should select option

⁴ Note that the highest drop-out price is equal to the second highest valuation above r so the auctioneer did not learn the value of the winning bidder.

Table 1 Experimental Design and Sample Sizes

	Cuberoot	Cube	Total
$n \in \{1, 2, 3, 4\}$	20	20	40
$n \in \{1, 4, 7, 10\}$	20	20	40
$n \in \{1, 4, 7, 10\}$ (FullInfo)	20	20	40
Total	60	60	120

Note. The number of seller subjects in each particular treatment combination.

A in more of the pairs only switching to option B as the expected value of the “risky” lottery exceeded some threshold.

Table 1 summarizes the design of the experiment. Twenty subjects participated in each treatment, and, because participants did not interact with each other, each subject constituted a single independent observation, which we use as the main unit of statistical analysis in this study.

We conducted all sessions at the Laboratory for Economic Management and Auctions at the Pennsylvania State University, Smeal College of Business, during the Spring and Summer of 2008. Participants in all six treatments were students, mostly undergraduates, from a variety of majors. Before each session, subjects were allowed a few minutes to read over the instructions themselves. Following this, we read the instructions aloud and answered any questions (we did this separately for both phases of the experiment so that subjects would not “read ahead”). Each individual participated in a single session only and was recruited through an online recruitment system. Cash was the only incentive offered, where subjects were paid a \$5 show-up fee plus an additional amount that was based on their personal performance for all 60 decisions in phase 1 and their decisions in phase 2. Average compensation for the participants, including the show-up fee, was \$20. Each session lasted approximately 45 minutes, and the software was programmed using the zTree system (Fischbacher 2007).

3. Results

We begin by providing summary statistics for all treatments. In Figure 2, we provide the mean reserve price for each bidding distribution evaluated for each n . We also display how these average reserve prices compare to risk-neutral theoretical optimal reserve levels (42 for Cuberoot; 63 for Cube).

We can make a number of observations from Figure 2: (1) reservation prices are clearly not independent of the number of bidders, they increase in n , contrary to standard theory (assuming either risk neutrality or risk aversion); (2) reservation prices are higher in the Cube condition than in the Cuberoot condition, consistent with standard theory; (3) whereas reserve prices are generally below the

optimal reserve price for low n , reserve prices are often above the optimal value when n is large (contrary to the assumption of risk aversion for any value of α); and (4) all three of these results continue to hold when sellers are provided with detailed information about reserve price outcomes (FullInfo). To examine these and other regularities more rigorously, we conducted a regression analysis. We ran all regressions with random effects for individual subjects and used the observed reserve price in a period as the dependent variable.⁵ The explanatory variables, their descriptions, and model estimates for the NoInfo treatments are listed in Table 2.

We centered the n variable by subtracting the average number of bidders faced by all sellers in order to be able to interpret the *Constant* as the average reservation price in the Cuberoot condition with the average number of bidders (we did the same with the decision period). Below we summarize our major results based on the regression models. First, consider Model (1)—the baseline model that does not consider the independent risk-aversion measure from the second phase of the experiment or variables that relate to more complex distributional effects.

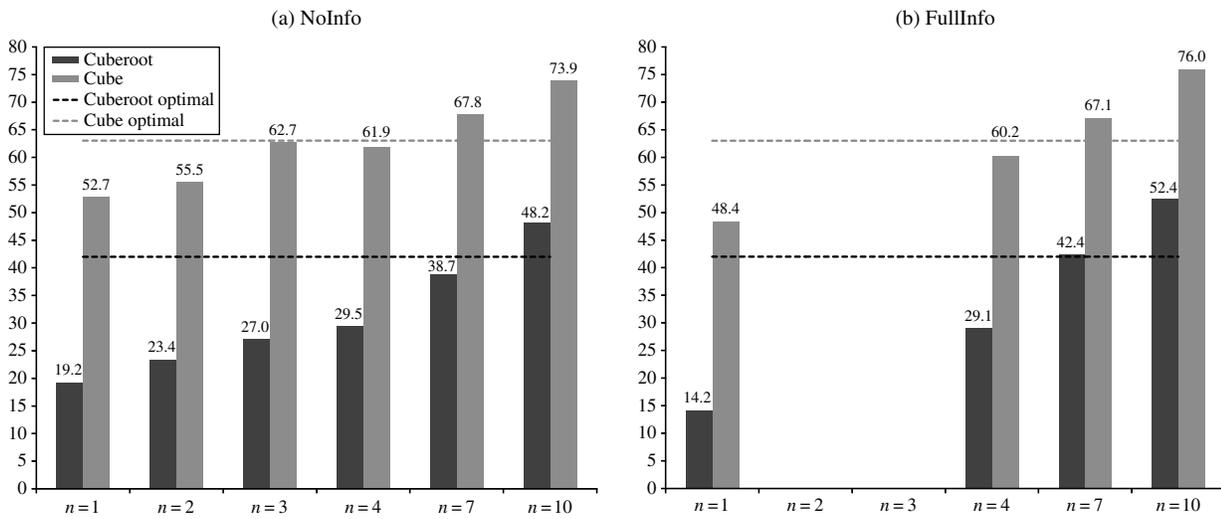
As is evidenced by the value of the *Constant* coefficient, subjects select reserve prices in the Cuberoot treatment that are significantly below the optimal level ($p < 0.01$). In the Cube treatment, subjects set reserve prices that are significantly greater than in the Cuberoot treatment as is evidenced by the positive and significant coefficient on the *Cube* indicator variable. In fact, subjects set reserve prices that, on average, are not significantly different from the optimal reserve price of the standard theory; 63, the optimal reserve price under the standard theory, is inside the joint 95% confidence interval for the sum of the *Constant* and *Cube* coefficients. Although the absolute value of the reserve price in the Cuberoot treatment is not consistent with the standard theory, these results indicate that subject behavior with respect to the distribution manipulation is consistent with the standard theory.

The other independent variables, however, indicate how seller behavior deviates from theoretical predictions. Sellers increase reservation prices with the number of bidders; the coefficient on $(n - \bar{n})$ is positive and significant, indicating that, on average, sellers in the Cuberoot treatment increase the reservation price by 2.98 for each additional bidder.

To account for the possibility of learning, we included the centered period number as an independent variable. Other factors constant, sellers increase

⁵ We conducted a Hausman (1978) test to ensure that random effects provided consistent estimates for each model; the interested reader is referred to the online appendix, where we provide all results with fixed effects.

Figure 2 Average Reserve Prices for Each n



reserve prices slightly over time, as evidenced by the positive and significant coefficient on the $(Per - \bar{Per})$ variable. The time trend, although significant, is slight, and is insufficient to bring average reserve prices in the Cuberoot condition to standard optimal levels, even at the end of the session.

To provide a direct test of the impact of risk aversion on reserve price choice, we conducted a risk-aversion elicitation exercise that was a modified version of the instrument introduced by Holt and Laury (2002). A more risk-averse seller should select fewer of the risky (option B) lotteries and also charge a lower reserve price. Model (2) from Table 2 adds the number of risky options ($\#B$) chosen as an independent variable. We center the $\#B$ variable by looking at the difference between the observed number of option Bs chosen by a subject and the average number

of option Bs chosen ($\overline{\#B}$) so that the other regression coefficients can have the same original interpretation as describing the reserve price setting of the “average” risk-averse seller versus the average number of bidders. The coefficient for Model (2) is significant and positive, indicating that a less risk-averse seller charges higher reserve prices, thus agreeing with the theoretical predictions. Although risk aversion does provide some explanation of the observed data, it does not account for the pronounced increases in the reserve price as the number of bidders increases or the fact that reserve prices above the risk-neutral optimal are frequently observed for large n . We will discuss Model (3) in later sections.

We estimated the same model for the FullInfo treatments, and present these results in Table 3. In focusing on Model (4), one can see the value of the *Constant*

Table 2 Regression Results for the NoInfo Treatments

Variable	Description	Model (1)	Model (2)	Model (3)
<i>Constant</i>	Intercept	29.45* [1.83]	29.52* [1.77]	29.53* [1.76]
$n - \bar{n}$	No. of bidders minus the average number of bidders	2.98* [0.08]	2.98* [0.08]	3.24* [0.11]
$Per - \bar{Per}$	No. of decision period minus the average number of periods	0.04* [0.011]	0.04* [0.011]	0.03* [0.011]
<i>Cube</i>	1 if bids from Cube distribution	31.78* [2.59]	31.64* [2.51]	31.64* [2.49]
$\#B - \overline{\#B}$	No. of B choices selected minus the average number of B choices		1.40* [0.49]	1.40* [0.49]
$(n - \bar{n}) \times \text{Cube}$	Interaction for $(n - \bar{n})$ and <i>Cube</i>			-0.55* [0.15]
R^2		0.499	0.519	0.522

* $p < 0.01$. Standard errors are in brackets.

Table 3 Regression Results for the FullInfo Treatments

Variable	Description	Model (4)	Model (5)	Model (6)
<i>Constant</i>	Intercept	34.86* [3.64]	34.12* [3.63]	34.12* [3.65]
$n - \bar{n}$	No. of bidders minus the average number of bidders	3.57* [0.08]	3.57* [0.08]	4.26* [0.12]
$Per - \bar{Per}$	No. of decision period minus the average number of periods	-0.10* [0.016]	-0.10* [0.016]	-0.10* [0.016]
<i>Cube</i>	1 if bids from Cube distribution	28.29* [5.15]	29.78* [5.19]	29.76* [5.22]
$\#B - \overline{\#B}$	No. of B choices selected minus the average number of B choices		-1.30 [0.97]	-1.38 [0.97]
$(n - \bar{n}) \times \text{Cube}$	Interaction for $(n - \bar{n})$ and <i>Cube</i>			-1.38* [0.16]
R^2		0.449	0.464	0.469

* $p < 0.01$. Standard errors are in brackets.

term is slightly lower than the Cuberoot optimal value of 42 (42 is just outside the confidence interval). Combining the coefficient on *Cube* with the *Constant* results in the optimal reserve price of 63 being within the joint confidence interval. Similar to the NoInfo treatments, subject behavior with respect to the bidding distributions is consistent with the standard theoretical predictions.

As with the NoInfo treatments, the other independent variables in Model (4) indicate how our results differ with the standard theory. Once again, the coefficient on $(n - \bar{n})$ is positive and significant, suggesting that subjects increase reserve prices by 3.57 for each additional bidder, even when presented with additional information about their reserve price choice. Another way that the FullInfo treatments relate to the NoInfo treatments pertains to learning effects. Whereas both coefficients are significant, in the FullInfo treatments the coefficient is small and negative (compared to small and positive for NoInfo). In Model (5), the coefficient on this risk-aversion parameter is not significant at any level, implying that risk aversion does not play a significant role when sellers are provided with information about reserve price outcomes. We will discuss Model (6) in later sections. The previous analysis shows how aggregate behavior varies with the treatment conditions. We next turn to examine whether these aggregate relationships were consistent with the behavior of the individuals.

We completed regressions on reserve prices for each of the 120 participants for all 60 periods of each treatment, with $(n - \bar{n})$ as the only independent variable. We conducted this analysis to determine what portion of the subjects systematically increased, decreased, or did not change their reserve prices as n changed. If the coefficient was statistically significant at the 95% level and positive, we classified the n effect as “positive” on reserve prices. Similarly, if the coefficient was statistically significant at the 95% level and negative, we classified it as “negative,” and if the coefficient was not significant at the 95% level, we categorized it as “neutral.” By reviewing the constant term and its corresponding 95% confidence interval, we classified subjects as on average setting reserve prices above, below, or at theoretical optimum levels. Table 4 displays the proportion of individuals in each of the nine categories based on the 120 individual regression models, where all results are pooled together.⁶

In Table 4, one can see that a large portion of subjects increased reserve prices as n went up (73.3%). Additionally, average reserves prices were generally below risk-neutral theoretical predictions (56.7%), but also a large portion set reserves above risk-neutral

Table 4 Classification of How Subjects Set Reserve Prices on Average, and How They Set Reserve Prices with Respect to Different Numbers of Bidders

n effect	Average r (%)			Total (%)
	Below	Neutral	Above	
Negative	9.2	0.8	0.0	10.0
Neutral	13.3	2.5	0.8	16.7
Positive	34.2	8.3	30.8	73.3
Total	56.7	11.7	31.7	100.0

theoretical predictions (31.7%). One last observation from this table is that the center cell, neutral/neutral, is very rarely seen in our data, suggesting serious deviations from the standard theory. We will investigate these deviations further in §4.

Because a seller’s expected profits are affected most by reserve prices when facing a small number of bidders, we examine more closely the data when subjects faced $n = 1$ bidder. The plot of the cumulative distributions for average reserve prices when facing one bidder is shown in Figure 3 for both distributions and information conditions. The figure clearly shows that in the Cuberoot treatment, for both conditions, subjects set reserve prices far below optimal; approximately 25% of the time they set reserves between 0 and 10 in the NoInfo treatment, and approximately 50% of the time they set reserves between 0 and 10 in the FullInfo condition. In contrast, we find that in the Cube treatment, a large percentage of reserve prices are in the neighborhood of the optimal value of 63.

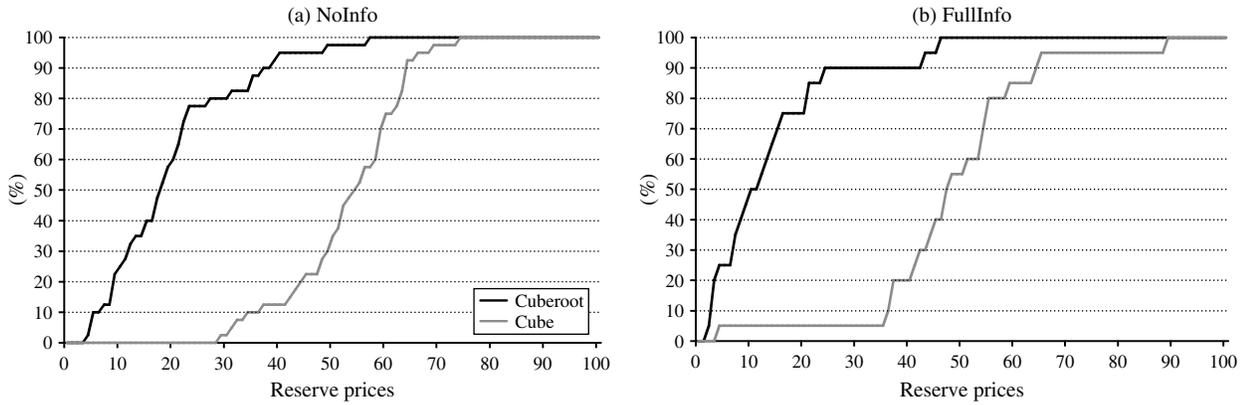
4. Alternative Models

As we have seen, the experimental data has a number of features that are inconsistent with the standard theory and risk aversion. Participants routinely set lower than optimal reserve prices when the number of bidders in the auction is low, but as the number of bidders increases, sellers increase the reserve price until, with high numbers of bidders ($n = 10$), the reserve price is often above the risk-neutral optimum. In an attempt to explain some of this behavior, we consider two separate models that have been previously advanced to explain why auction bidders deviate from the risk-neutral Nash equilibrium:

1. Winner and loser anticipated regret.
2. Probability weighting.

We consider each model independently in order to understand the potential effect of incorporating each feature into a model of seller behavior. Although actual behavior may be more fully described by a combination of some of these (and other) models, the approach we take here is to attempt to understand what features are necessary and sufficient to produce the qualitative results observed in the laboratory experiments.

⁶ The interested reader is referred to the online appendix, where Cube and Cuberoot are separated between NoInfo and FullInfo.

Figure 3 Cumulative Distribution of Reserve Prices for $n = 1$ 

4.1. Anticipated Regret

A model that has shown to consistently describe the behavior of bidders in auctions is that of anticipated regret (Filiz-Ozbay and Ozbay 2007, Engelbrecht-Wiggans and Katok 2008). In these types of models, the agent feels bad (experiences some disutility) when his or her ex ante action fails to capture all of the potential earnings that are available (and known) ex post. For the seller in an auction, this is likely to happen in two ways: (1) the seller fails to sell the object even though there were bidders willing to pay for the object, and thus she regrets those foregone revenue; (2) the price exceeds the reserve price and the seller realizes that a higher reserve price might have extracted more revenues from the winning bidder. Following Bell (1982) and Engelbrecht-Wiggans (1989), we term the first type of regret *loser's regret* and the second type *winner's regret*. A model of anticipated regret assumes that the agent correctly anticipates these feelings of regret and modifies his or her ex ante actions in order to maximize his or her expected payoffs modified by this regret.

Let $w(x)$ be the amount of winner's regret. We assume winner's regret to be a function of only the observed sales price or revenue x . Let $l(r)$ be the amount of loser's regret given the chosen reserve price r . We assume that these terms enter into the seller's payoff function in a linearly separable fashion; the seller's ex post utility function is then given by

$$v(b; r) = \begin{cases} u(b_{(2)}) - w(b_{(2)}) & \text{if } b_{(2)} > r, \\ u(r) - w(r) & \text{if } b_{(1)} > r > b_{(2)}, \\ 0 - l(r) & \text{if } r > b_{(1)}. \end{cases} \quad (8)$$

Given these payoffs and bidders' equilibrium strategies, the seller's expected utility for a second-price auction with reserve price r is given by

$$Ev(r) = [-l(r)]F^n(r) + [u(r) - w(r)]G(r) + \int_r^{\bar{v}} [u(v) - w(v)]h(v) dv. \quad (9)$$

Then, we have the following necessary first-order condition:

$$\begin{aligned} \frac{dEv(r^*)}{dr} &= -l'(r^*)F^n(r^*) + [-l(r^*)]nF^{n-1}(r^*)f(r^*) \\ &\quad + [u'(r^*) - w'(r^*)]G(r^*) + [u(r^*) - w(r^*)]g(r^*) \\ &\quad - [u(r^*) - w(r^*)]h(r^*) = 0. \end{aligned} \quad (10)$$

Divide by $nF^{n-1}(r^*)$ and rearrange terms to obtain

$$\begin{aligned} -l'(r^*)\frac{1}{n}F(r^*) + [-l(r^*)]f(r^*) + [u'(r^*) - w'(r^*)] \\ \cdot (1 - F(r^*)) - [u(r^*) - w(r^*)]f(r^*) = 0, \end{aligned} \quad (11)$$

which can be rearranged to obtain the following condition:

$$\frac{u(r^*) - w(r^*) + l(r^*)}{u'(r^*) - w'(r^*) - l'(r^*)(1/n)(K(r^*))} = \frac{(1 - F(r^*))}{f(r^*)}, \quad (12)$$

where $K(r^*) = F(r^*)/(1 - F(r^*))$. To consider the potential effect of anticipated regret on optimal reserve prices, we assume that the seller is risk neutral⁷ and posit regret functions that are consistent with the information available to the seller after the auction is completed. Let $W(r) = w(r) - \phi(r)w'(r)$ be the term's associated winner's regret from the first-order conditions and let $L(r) = l'(r)(1/n)\psi(r) + l(r)$ be the term's associated loser's regret from the first-order conditions where $\psi(v) = F(v)/f(v)$. Note that, despite not appearing as an argument of the function, $w(v)$ and $l(r)$ may be functions of the number of bidders n . Thus, particular formulations of these functions may result in the optimal reserve price increasing with n as was observed for many experimental subjects. The following are sufficient conditions for particular directional changes in r^* with respect to n .

⁷Qualitatively similar results can be derived in the event that the seller is risk averse, yet winner and loser regret enter into the seller's utility function in a similar fashion to that modeled here.

PROPOSITION 2. If $\phi'(r^*) + W'(r^*) - L'(r^*) < 1$, then

1. the optimal reserve price r^* is increasing in n whenever $\partial L(r^*)/\partial n - \partial W(r^*)/\partial n < 0$;
2. the optimal reserve price r^* is decreasing in n whenever $\partial L(r^*)/\partial n - \partial W(r^*)/\partial n > 0$.

PROOF. Given the notation presented above, the necessary first-order conditions can be expressed as follows:

$$\phi(r^*) + W(r^*) - L(r^*) - r^* = 0. \quad (13)$$

We differentiate with respect to n at the optimal solution to obtain the following:

$$[\phi'(r^*) + W'(r^*) - L'(r^*) - 1] \frac{\partial r^*}{\partial n} + \frac{\partial W(r^*)}{\partial n} - \frac{\partial L(r^*)}{\partial n} = 0. \quad (14)$$

Rearranging terms we obtain

$$\frac{\partial r^*}{\partial n} = \frac{\partial L(r^*)/\partial n - \partial W(r^*)/\partial n}{[\phi'(r^*) + W'(r^*) - L'(r^*) - 1]}. \quad (15)$$

Note that the denominator is the second-order condition that must be less than or equal to zero if r^* is a local maximum. Obviously, $\phi'(r^*) + W'(r^*) - L'(r^*) < 1$ is sufficient to guarantee that the denominator is strictly negative.⁸ Then it follows that if $\partial L(r^*)/\partial n - \partial W(r^*)/\partial n < 0$ we have that $\partial r^*/\partial n > 0$ and if $\partial L(r^*)/\partial n - \partial W(r^*)/\partial n > 0$ we have that $\partial r^*/\partial n < 0$. \square

To understand the conditions of this proposition, note that, in general, loser's regret will encourage lower reserve prices whereas winner's regret will lead to higher reserve prices. Thus, whenever the change in winner's regret due to an increase in n is bigger than the change in loser's regret, reserve prices should increase, and whenever the change in loser's regret due to an increase in n is bigger than the change in winner's regret, reserve prices should decrease.

The appropriate choice of winner's and loser's regret functions likely depends upon the frame of the auction institution as well as the information available to the seller ex post. The following is an example of winner's and loser's regret that matches the qualitative features of our data and is consistent with the information conditions of the auctions we conducted. When the auction has ended with a sale, the seller's information may be of two different forms:

1. If the item sold at the reserve price r , then the seller knows that the second highest value $v_{(2)}$ (and all lower values) was below r and the highest value $v_{(1)}$ was above r .

2. If the item sold above the reserve price, then the seller knows that the price was equal to the second highest value $v_{(2)}$ and the highest value $v_{(1)}$ was (trivially) above $v_{(2)}$.

We assume that winner's regret is proportional to the difference between the obtained price (r or higher) and expected high value given the seller's information. In the first case above, this implies that

$$w(r) = \alpha_w [E(v_{(1)} | v_{(1)} > r, v_{(2)} < r) - r], \quad (16)$$

and in the second case when the sale price is $v > r$, we have that

$$w(v) = \alpha_w [E(v_{(1)} | v_{(2)} = v) - v], \quad (17)$$

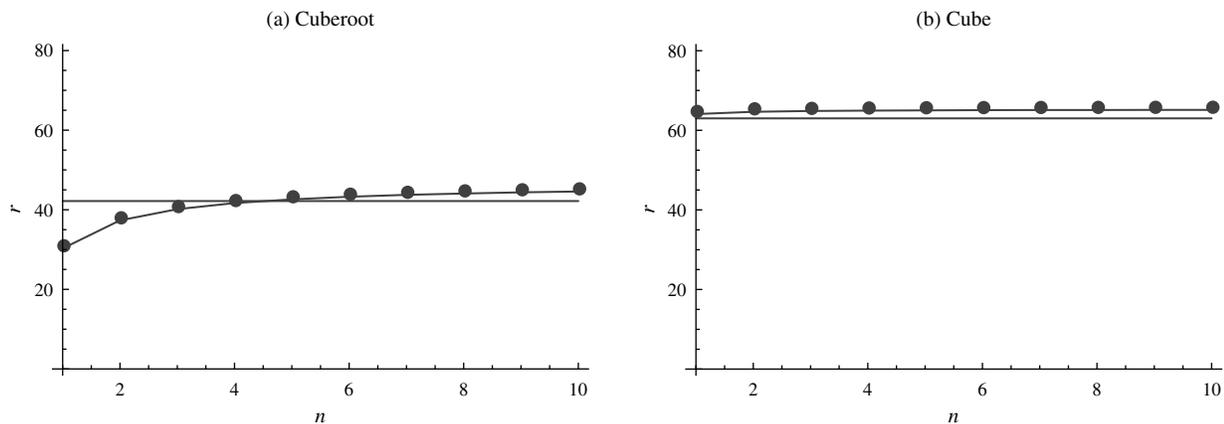
where $\alpha_w > 0$. However, because $E(v_{(1)} | v_{(1)} > r, v_{(2)} < r) = E(v_{(1)} | v_{(2)} = r)$ we can merge these two information conditions into a common winner's regret function given by (17) for all $v \geq r$. See the technical appendix for the derivation of these two expectations and their equivalence.

Next, let loser's regret be proportional to the reserve price or $l(r) = \alpha_l r$ or that the seller generally feels worse for setting a higher reserve price that did not result in a sale. In our experiments, a loser's regret functional form of this type seems particularly reasonable given the general lack of information provided when there is no sale; when all values are below the reserve price the seller simply observes that fact and does not observe the realizations of values. Therefore, she cannot calculate exactly how much money was left on the table by not selling in that instance.⁹

Then, substituting these functional forms for $w(x)$ and $l(r)$ into $W(r)$ and $L(r)$, we have that $W(r) = \alpha_w \phi(r)$ so $\partial W(r^*)/\partial n = 0$. We also have that $L(r) = \alpha_l(1/n)\psi(r) + \alpha_l r$ so $\partial L(r^*)/\partial n = -\alpha_l(1/n^2)\psi(r^*)$, which is clearly negative. Therefore, we have that $\partial L(r^*)/\partial n - \partial W(r^*)/\partial n < 0$ and the optimal reserve price should increase with n according to Proposition 2. Figure 4 depicts the optimal reserve prices for sellers who anticipate winner's and loser's regret as specified above for $\alpha_l = 0.3$, $\alpha_w = 0.5$, and number of bidders n from 1 to 10. The effect of α_l is to change the importance of loser's regret relative to the winner's regret. When loser's regret is relatively low (α_l small), the seller is induced to charge higher than (standard theory) optimal reserve prices in order to avoid selling at too low a reserve price where the winning bidder (if there is one) is likely to have substantial profits. This family of intuitive and simple regret functions approximate our data reasonably well. For example, in both cases, the reserve price is increasing, and for particular

⁸ A more obvious but slightly less tight sufficient condition is that $\phi'(r^*) \leq 0$, $L'(r^*) \geq 0$, and $W'(r^*) \leq 0$. These conditions are almost always satisfied for most models of regret and distributions we have examined.

⁹ Greenleaf (2004) examines models where this information is assumed to be known and arrives at results regarding the optimal reserve price.

Figure 4 Optimal Reserve Prices in the Regret Model with $\alpha_l = 0.3$ and $\alpha_w = 0.5$ as a Function of n 

Note. The horizontal line represents the standard theory optimal r^* .

choices of α_l , the reserve price goes from being below the risk-neutral optimal reserve price to being above it.

One can immediately recognize from these figures that the effect of anticipated regret is more dramatic with the Cuberoot distribution than the Cube distribution. In Model (3) in Table 2, and Model (6) in Table 3, we explore whether the effect of the number of bidders is different across the two value treatments by adding a variable that interacts the distribution treatment Cube and the number of bidders ($n - \bar{n}$). We find that the coefficient on the interaction term between the number of bidders and the indicator variable for the Cube treatment is negative and significant at the 99% level (coefficient of -0.55) for the NoInfo condition, and negative and significant at the 99% level for the FullInfo condition (coefficient of -1.38). This illustrates that our data agrees with the qualitative direction of the regret theory; the impact of regret has a more pronounced effect in the Cuberoot treatment compared to the Cube distribution treatment. Additionally, it appears as though the positive relationship between reserve prices and n that we observed in our experiment agrees with the theoretical predictions from the regret model.

4.2. Probability Weighting

Although anticipated regret models change the payoffs associated with particular outcomes in order to depart from standard models of expected utility, it still maintains that a human seller calculates correctly and maximizes expected utilities given those payoffs. An alternative model is that the human seller might fail to calculate expected utilities correctly. In these models, the agent might behave as if some events are more likely and others are less likely. It has been well documented that human subjects often behave as if small probability events have a greater probability of occurring than their given probabilities and high probability events are less likely to occur than their

given probabilities (Kahneman and Tversky 1979). To incorporate this potential behavioral model into the reserve price setting decision of sellers, we assume that the cumulative probability of certain events are distorted by a *probability weighting* function $w(p)$ that takes the actual probability of occurrence and maps it to some other probabilistic value (value between zero and one). In particular, we assume the probability weighting function has the following features: (1) $w(0) = 0$ and $w(1) = 1$, (2) $w'(p) > 0$ for all p , and (3) there exists a p^* such that for all $0 < p < p^*$ $w(p) > p$ and $w''(p) < 0$ and for all $1 > p > p^*$ $w(p) < p$ and $w''(p) > 0$.

These are standard assumptions made in the literature on probability weighting (Wu and Gonzalez 1996). The first assumption means that events of probability zero and one are easily discernible and are therefore not subject to distortions. The second assumption means that greater probabilities still imply greater weights. The final assumption indicates the empirical regularity that low probability events are typically overweighted and high probability events typically underweighted. In addition, these features allow the probability weighting model to be incorporated into the calculus of expected utility calculations with a minimal amount of complications.¹⁰

A closely related explanation for probability weighting is that the sellers, particular in the NoInfo treatments, had to calculate these probabilities themselves.

¹⁰ However, it should be noted that these probability weighting functions essentially imply a nonadditive probability measure, so this is a special case of the broad class of nonexpected utility models such as those where subjects are ambiguity averse. Chen et al. (2005) examine bidder and auctioneer behavior in settings where ambiguity aversion is likely to be present. It is possible that some of the observed behavior in the NoInfo treatment may also be due to ambiguity aversion. In the FullInfo treatment ambiguity should not be an issue, but we observe substantially similar behavior.

Although all the necessary information was provided, these are nontrivial calculations that may have involved substantial errors. By using a probability weighting model we are assuming the errors in calculation take a particular form. The seller's expected utility calculation is given by

$$Ev(r) = u(r)w(G(r)) + \int_r^{\bar{v}} u(v)dw(H(v)). \quad (18)$$

The following necessary first-order condition results:

$$\frac{\partial Ev(r^*)}{\partial r} = u'(r^*)w(G(r^*)) + u(r^*)w'(G(r^*))g(r^*) - u(r^*)w'(H(r^*))h(r^*) = 0, \quad (19)$$

which can be simplified to yield

$$\frac{u(r^*)}{u'(r^*)} = \frac{w(G(r^*))}{w'(H(r^*))h(r^*) - w'(G(r^*))g(r^*)}. \quad (20)$$

This condition is equivalent to the first-order conditions from the standard theory if $w(p) = p$ or there is no distortion due to probability weighting. Also note that this first-order condition is only defined for $n \geq 2$ because $g(v)$ is only appropriately defined when there are at least two bidders. When $n = 1$ the first-order condition is easily found to be

$$\frac{u(r^*)}{u'(r^*)} = \frac{w(1 - F(r^*))}{w'(1 - F(r^*))f(r^*)}. \quad (21)$$

To provide a more concrete example, we assume sellers are risk neutral and consider a probability weighting function extensively utilized in the literature. The following is a one-parameter weighting function used by Tversky and Kahneman (1992), Camerer and Ho (1994), and Wu and Gonzalez (1996):¹¹

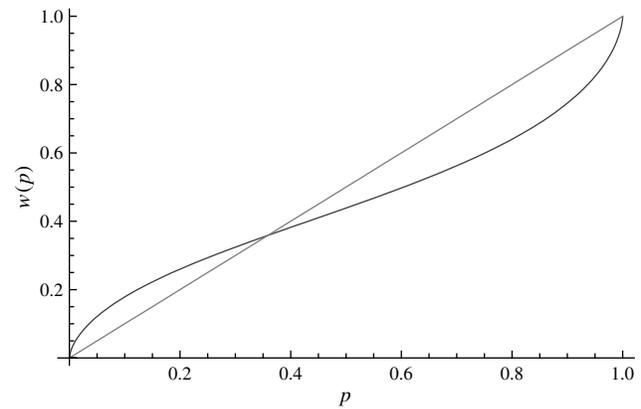
$$w(p) = \frac{p^\beta}{(p^\beta + (1 - p)^\beta)^{1/\beta}}. \quad (22)$$

Figure 5 depicts the probability weighting function for the parameter $\beta = 0.65$. Previous studies have found a wide range in the values of β for different scenarios, 0.57 to 0.94 (Tversky and Kahneman 1992), 0.28 to 1.87 (Camerer and Ho 1994), and a slightly different probability weighting function utilized by Prelec showed the parameter to vary from 0.03 to 0.95 (Prelec 1998).

Although an exact analytical solution given the distributions of values used and the weighting function is not readily available, the results show that for $n \geq 2$ under both distributions the optimal reserve price is varying with the number of bidders. Figure 6 depicts the optimal reserve price for both value distributions as a function of the number of bidders.

¹¹ It should be noted that for this probability weighting function, (21) is only satisfied uniquely when β is sufficiently large.

Figure 5 One Parameter Probability Weighting Function $\beta = 0.65$



Although this figure demonstrates that we might expect similar behavior under the two distribution treatments, substantially different behavior is expected when $n = 1$. The optimal reserve price is approximately 64 and 61 under the Cube and Cuberoot distributions, respectively. Importantly, under the Cube distribution the optimal reserve price for $n = 1$ is lower than that for higher numbers of bidders ($n = 2$ optimal reserve is 69). Under the Cuberoot distribution treatment the optimal reserve price for $n = 1$ is substantially higher than it is for higher numbers of bidders; the optimal reserve falls to 43 for $n = 2$. This behavior holds for many parameters and alternative probability weighting functions.¹² Additionally, Figure 6(b) shows that for the Cube treatment when n exceeds 6, sellers should set slightly lower reserve prices as the number of participating buyers increases. Comparing these figures to our experimental data results in some key differences. Specifically, for the Cuberoot condition, when n is less than or equal to two bidders, we should observe a negative relationship between n and reserve prices. In reality, what we observe is the opposite of this. Likewise, there is no evidence that the effect of n decreases for larger n under the Cube condition.

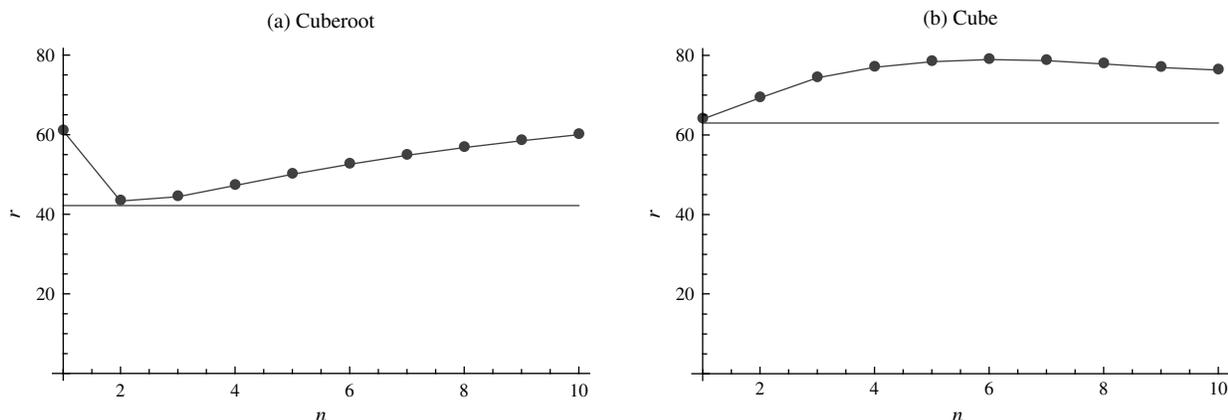
In many cases the predictions of probability weighting models are virtually indistinguishable from those of models that change the payoffs of associated outcomes (such as models of regret) (Goeree and Holt 2004). Interestingly, the predictions of (at least these particular implementations) of probability weighting and anticipated regret are substantially different for the treatments we have considered. At least qualitatively, the data appears to favor a model of anticipated regret over probability weighting.

4.3. Model Estimation

So far we have seen that some features of these models are consistent with the qualitative behavior of

¹² This feature is also true for the uniform distribution.

Figure 6 Optimal Reserve Prices Given One Parameter Probability Weighting Function $\beta = 0.65$



Note. The horizontal line represents the standard theory optimal r^* .

the sellers in our experimental sessions. However, the models also differ from the empirical regularities of the data. In this section, we estimate parameters for the different alternative models (including CRRA) and use maximum likelihood to assess the quality of their fit to the observed data. We then allow for bidder heterogeneity in both choice of parameters and behavioral models.

Each behavioral theory outlined above predicts some optimal reserve price as a function of parameters of the model (which we generically label θ) and (potentially) the number of bidders n , and the value distribution F . Let $r^*(\theta, n, F)$ be the optimal reserve price under a given model for a particular selection of θ , n , and F . We begin by positing a simple normal likelihood function for the choice of an observed reserve price r_i or $f(r_i | \theta, n, F) = N(r^*(\theta, n, F), \sigma^2)$ or that each subject selects the observed reserve price r_i with some mean zero, normal error around the optimal reserve price predicted by theory. Because n and F are controlled and observed in the experiment, for each given model, the objective is to select a θ that maximizes the likelihood function:

$$\max_{\theta \in \Theta} L(\theta, \sigma^2 | r, n, F) = \prod_{i=1}^m f(r_i, \sigma^2 | \theta, n, F). \quad (23)$$

In addition to considering the risk-neutral model, using standard numerical maximum likelihood techniques, we estimate a single parameter class of each of the three alternative models. Specifically, we examine the following:

- For the risk-aversion model, we estimate $\alpha \in [0, 1]$ the coefficient of relative risk aversion in the CRRA model.
- For the anticipated regret model, we estimate the model where it is assumed that $\alpha_l = 1 - \alpha_w$ and $\alpha_l \in [0, 1]$.
- For the probability weighting model, we estimate the weighting function given in (22), where $\beta \in (0, 1)$.

In Table 5 we provide a summary of the maximum likelihood estimates. Focusing on the NoInfo treatments first and considering the pooled data where all individuals are assumed to have the same parameter value for the particular model being evaluated, we see that the anticipated regret model has the highest log likelihood of the three models. Whereas the differences in the log likelihoods do not appear very substantial, a Vuong test (Vuong 1989) for comparing nonnested models reveals that the models are statistically different from one another, favoring the model of anticipated regret. We also applied a likelihood-ratio test to compare the standard theory as a nested version of the other models. Results from those tests agree with the Vuong test at even higher levels of significance.

The estimated parameters provide insights into the ability of the model to fit the data for reasonable values of the parameters. For the NoInfo condition, the CRRA parameter estimate of $\alpha = 0.29$ is somewhat lower than observed in previous studies suggesting that sellers in our experiment appear to act less risk averse. The anticipated regret parameter suggests that, on average, subjects seem to put slightly higher emphasis on loser regret but that the effect of the two are roughly similar. This weighting rationalizes the increase in reserve price with respect to n because

Table 5 Maximum Likelihood Estimates for Standard Theory, Risk Aversion, Anticipated Regret, and Probability Weighting Aggregated Between NoInfo and FullInfo

Behavioral model	NoInfo treatments		FullInfo treatments	
	Log(L) for pool	Parameter for pool	Log(L) for pool	Parameter for pool
Standard theory	-21,556	—	-11,092	—
Risk aversion	-21,141	$\alpha = 0.29$	-11,050	$\alpha = 0.17$
Anticipated regret	-20,851	$\alpha_l = 0.55$	-10,880	$\alpha_l = 0.48$
Probability weighting	-21,571	$\beta = 0.99$	-11,093	$\beta = 0.99$

Copyright: INFORMS holds copyright to this *Articles in Advance* version, which is made available to institutional subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

Table 6 Parameter Estimates for Risk Aversion, Anticipated Regret, and Probability Weighting When Allowing for Subject Heterogeneity

Behavioral model	NoInfo treatments		FullInfo treatments	
	Log(L) for individuals	Mean parameter value	Log(L) for individuals	Mean parameter value
Standard theory	-261	—	-264	—
Risk aversion	-242	$\alpha = 0.28$	-253	$\alpha = 0.20$
Anticipated regret	-237	$\alpha_j = 0.54$	-251	$\alpha_j = 0.42$
Probability weighting	-253	$\beta = 0.89$	-254	$\beta = 0.82$

the loser regret term matters less as the number of bidders increases. Finally, the probability weighting parameter (0.99) is far from that observed in other experiments and suggests that a risk-neutral nonprobability weighting model (e.g., $\beta = 1$) is a better fit to the data. In fact, if we removed the restriction that $\beta < 1$, we would have found that the best fit would have been obtained via a model in which subjects appear to act as if $\beta > 1$ or they place *lower* weight on low probability events and *higher* weight on high probability events; this is counter to most other experimental observations of probability weighting.

For the FullInfo treatments, the results are much the same. The regret model generates the largest log likelihood, with the Vuong test showing it to be significantly higher than the other models and thus favoring it overall (the only models not statistically different from each other were standard theory and probability weighting).

As discussed earlier, the high degree of subject heterogeneity obviously affects the quality of the fit of the pooled model and substantially lowers the log likelihood. Therefore, we next allow for different parameter values for every subject and examine the behavioral models in turn. As one can see from Table 6, in the NoInfo and FullInfo treatments, the regret model is once again slightly favored. A description of the estimated parameter values when within model subject heterogeneity is allowed for is provided in Table 6 as well. Whereas the mean parameter values for the risk aversion and anticipated regret models do not change significantly, the mean probability weighting estimate is now slightly more consistent with previous studies.

Because of the difficulty in drawing conclusions about individual behavior using average log likelihoods, we now take a different approach. Specifically, it may be that individual heterogeneity is not only in terms of parameters within a model but also is in terms of different subjects using different models; whereas some may be risk averse, others may be driven primarily by regret or probability weighting. This provides a more accurate depiction of which

Table 7 Proportion of Largest Log-Likelihood Function for Each Subject by Behavioral Model

Behavioral model	NoInfo treatments		FullInfo treatments	
	% of subjects best fit	Mean parameter estimate	% of subjects best fit	Mean parameter estimate
Standard theory	0.00	—	2.50	—
Risk aversion	33.75	$\alpha = 0.52$	12.50	$\alpha = 0.77$
Anticipated regret	40.00	$\alpha_j = 0.54$	47.50	$\alpha_j = 0.51$
Probability weighting	26.25	$\beta = 0.67$	37.50	$\beta = 0.58$

model is favored by subjects, compared to the results in Table 6. We summarize the proportion of subjects for whom each of the three models provides the best fit in Table 7. Interestingly, each model, except for the standard theory, fits a considerable percentage of the subjects for the NoInfo treatments with anticipated regret being the best fit with 40%. A similar situation exists in the FullInfo treatments, where 47.5% were best fit by anticipated regret. Furthermore, when considering the estimated parameter values for each behavioral model for *only* those subjects best fit by that model, the mean parameter estimates for the risk aversion and probability weighting models are more consistent with the values observed in previous experiments.

4.4. Model Predictive Power

We have demonstrated that the regret model tends to fit our data slightly better than the risk aversion and probability weighting models. However, this result does not prove that the regret model, or any of the alternative models, can fully account for the variation of reserve prices we observe with respect to n . One possible way to address this issue is to run a regression of reserve prices on the number of bidders and an independent variable that represents the predicted reserve price, \hat{r}_i , for each subject's best model and maximum-likelihood estimate (for example, if a subject's maximum log likelihood of all four models was the probability weighting model with $\beta = 0.5$, then \hat{r}_i is the predicted reserve price for probability weighting with that parameter estimate). The regression model is

$$r_i = \gamma_0 + \gamma_1(n - \bar{n}) + \gamma_2\hat{r}_i + \epsilon. \quad (24)$$

If the coefficient on $(n - \bar{n})$ becomes insignificant, and the coefficient on \hat{r}_i is 1, then the behavioral models fully explain the n effect. However, introducing a regressor that is highly dependent on other independent variables, such as \hat{r}_i with $(n - \bar{n})$, presents a number of econometric issues. Therefore, we can impose a restriction on the coefficient of \hat{r}_i so that $\gamma_2 = 1$ and estimate

$$(r_i - \hat{r}_i) = \gamma_0 + \gamma_1(n - \bar{n}) + \epsilon. \quad (25)$$

Table 8 Regression Results Comparing the Predictive Power of the Alternative Models for the n Effect in the NoInfo and FullInfo Treatments

Dependent variable	Independent variable	Description	NoInfo		FullInfo	
			Cuberoot	Cube	Cuberoot	Cube
r_i	Constant	Intercept	29.46* [0.28]	61.67* [0.26]	34.85* [0.34]	63.14* [0.44]
	$n - \bar{n}$	No. of bidders minus the average number of bidders	3.24* [0.11]	2.70* [0.11]	4.28* [0.10]	2.90* [0.13]
$(r_i - \hat{r}_i)$	Constant	Intercept	1.93* [0.27]	-3.18* [0.25]	2.42* [0.31]	-7.69* [0.39]
	$n - \bar{n}$	No. of bidders minus the average number of bidders	1.88* [0.11]	1.73* [0.10]	2.80* [0.09]	1.54* [0.12]

Note. Note how the coefficient on $(n - \bar{n})$ is reduced, but not entirely eliminated, when the alternative models' predicted reserve prices are incorporated.

* $p < 0.01$. Standard errors are in brackets.

If the coefficient on $(n - \bar{n})$ is significant, then the behavioral models do not fully account for the n effect that we see in our data. However, if we see a reduction in the coefficient on $(n - \bar{n})$ compared to a regression that omits \hat{r}_i , then the models partially explain the n effect.

Table 8 compares regressions with dependant variable r_i and the dependant variable $r_i - \hat{r}_i$ from (25). We estimated these models using fixed effects, where Cuberoot and Cube are separated to better interpret the constant.¹³

As we reported in §3, when we do not incorporate any alternative models, there is a largely positive and significant coefficient on $(n - \bar{n})$. When we incorporate each subject's best prediction, based on their most favorable model, we find that the coefficient on $(n - \bar{n})$ is reduced, and that for all four regressions, the differences in the coefficient on $(n - \bar{n})$ are all statistically significant at 99% based on a Wald test. Nevertheless, the coefficient on $(n - \bar{n})$ remains positive and significant. This result indicates that the alternative models we consider explain part but not all of the n effect that we see in our data.

5. Conclusion and Discussion

We examine, in the laboratory, the standard normative model of setting the optimal reservation price in an auction with bidders who have independent private values. We find that the data is consistent with certain aspects of the theory but not with others. When we vary the distribution of values of the bidders, we find that the actual reservation prices our sellers set shift in the direction predicted by the model. However, we also find that reservation prices increase with the number of bidders, even in situations where sellers are provided with detailed information about reserve price outcomes. This is contrary to theoretical

predictions that state the optimal reservation price is independent of the number of bidders.

To examine whether an alternative model organizes the data better, we proceed to extend the standard theory along three separate dimensions: risk aversion, anticipated regret, and probability weighting. We find that whereas risk aversion is consistent with average reservation prices our sellers set for a small number of bidders being below theoretical optimums, it does not predict a dependence between the reservation prices and the number of bidders. The model in which sellers anticipate regret fits most of the qualitative aspects of our data by predicting that the optimal reserve price increases with the number of bidders, and may even exceed the risk-neutral optimum level when the number of bidders is large enough. The model in which sellers distort probabilities does not fit particularly well on aggregate, but it does seem to fit a reasonable proportion of our subjects when we fit the model on an individual level. Overall, we find that all three alternative models fit some portion of our subjects, whereas only a single individual is best described by the standard theory. Many experiments tend to focus on a single behavioral explanation, but we feel it is important to recognize that human actors may exhibit highly heterogeneous behavior and that such heterogeneity may persist in the population, so attempting to fit one model to such individual data is dangerous.

One practical implication of our work is that sellers are subject to behavioral biases at least as much, if not more, than are the bidders. Even in our simple laboratory setting, sellers are not able to set the reservation prices optimally, which leads us to conjecture that in more complex auctions in the field, sellers may well be foregoing a substantial amount of potential profit by not setting the reservation prices in a profit-maximizing way. Further research into sellers' behavior in other auctions such as first-price auctions or auctions with affiliated values is needed in order to provide better insights into how our findings might translate to other auction environments where

¹³ The Hausman (1978) test for consistency of random effects models was not always satisfied, hence fixed effects is appropriate.

changes in the information structure may vary the predictions of the models. Finally, the introduction of human bidders might pull auctioneer behavior further from standard theoretical predictions, because issues such as equity and fairness may now become salient. All these factors need to be considered when providing analysis or decision-support tools associated with real-world auctions.

6. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Acknowledgments

Financial support was provided by the Smeal College of Business Small Grants Program. The authors thank seminar participants at the 2008 International Economic Science Association meetings, the 2008 INFORMS conference, the Fall 2009 Midwest International Economics Group meeting, and the Caltech Social and Information Sciences Laboratory Seminar Series. The authors also thank an anonymous referee and associate editor, who helped improve this paper.

Technical Appendix

To consider the expected values needed for the computation of winner's regret first consider the expectation given in (16). The joint density of the highest and second highest values is given as follows (see Krishna 2002):

$$f_{(1),(2)}(y_1, y_2) = n(n-1)f(y_1)f(y_2)F(y_2)^{n-2}, \quad (26)$$

so the conditional joint density is given by

$$f_{(1),(2)}(y_1, y_2 | v_{(1)} > r, v_{(2)} < r) = \frac{1}{nF(r)^{n-1}(1-F(r))} f_{(1),(2)}(y_1, y_2) \quad (27)$$

$$= \frac{1}{nF(r)^{n-1}(1-F(r))} (n(n-1)f(y_1)f(y_2)F(y_2)^{n-2}) \quad (28)$$

$$= \frac{n-1}{F(r)^{n-1}(1-F(r))} f(y_1)f(y_2)F(y_2)^{n-2} \quad (29)$$

for $y_1 > r$ and $y_2 < r$ and 0 otherwise. The expected value is then given by

$$E(v_{(1)} | v_{(1)} > r, v_{(2)} < r) = \int_r^{\bar{v}} \int_v^r y_1 f_{(1),(2)}(y_1, y_2 | v_{(1)} > r, v_{(2)} < r) dy_2 dy_1 \quad (30)$$

$$= \frac{n-1}{F(r)^{n-1}(1-F(r))} \int_r^{\bar{v}} \int_v^r y_1 f(y_1)f(y_2)F(y_2)^{n-2} dy_2 dy_1 \quad (31)$$

$$= \frac{n-1}{F(r)^{n-1}(1-F(r))} \int_r^{\bar{v}} y_1 f(y_1) \left(\int_v^r f(y_2)F(y_2)^{n-2} dy_2 \right) dy_1. \quad (32)$$

Consider the y_2 integral (inside parentheses). Via integration by parts, we know that

$$\int_v^r f(y_2)F(y_2)^{n-2} dy_2 = F(y_2)^{n-2}F(y_2)|_v^r - (n-2) \int_v^r F(y_2)^{n-3} f(y_2)F(y_2) dy_2 \quad (33)$$

$$= F(r)^{n-1} - (n-2) \int_v^r f(y_2)F(y_2)^{n-2} dy_2 \quad (34)$$

$$(n-1) \int_v^r f(y_2)F(y_2)^{n-2} dy_2 = F(r)^{n-1} \quad (35)$$

$$\int_v^r f(y_2)F(y_2)^{n-2} dy_2 = \frac{F(r)^{n-1}}{n-1}. \quad (36)$$

Substitute into (32) to obtain the following:

$$E(v_{(1)} | v_{(1)} > r, v_{(2)} < r) = \frac{n-1}{F(r)^{n-1}(1-F(r))} \int_r^{\bar{v}} y_1 f(y_1) \frac{F(r)^{n-1}}{n-1} dy_1 \quad (37)$$

$$= \frac{1}{1-F(r)} \int_r^{\bar{v}} y_1 f(y_1) dy_1. \quad (38)$$

Now consider the expectation given in (17). The density of the highest value given that the second highest value is known to be v is given by

$$f_{(1)}(y_1 | y_2 = v) = \frac{f_{(1),(2)}(y_1, v)}{f_{(2)}(v)}, \quad (39)$$

where the joint density of the highest and second highest values is given by (26) and the denominator is the marginal density of the second highest value evaluated at v and is given by (see Krishna 2002) $f_{(2)}(v) = n(n-1)(1-F(v))F(v)^{n-2}f(v)$, so we obtain the following:

$$f_{(1)}(y_1 | y_2 = v) = \frac{n(n-1)f(y_1)f(v)F(v)^{n-2}}{n(n-1)(1-F(v))F(v)^{n-2}f(v)} \quad (40)$$

$$= \frac{f(y_1)}{1-F(v)} \quad (41)$$

for $y_1 > v$ and 0 otherwise. Then calculating the expected value of the highest value we have that

$$E(v_{(1)} | v_{(2)} = v) = \int_v^{\bar{v}} y_1 f_{(1)}(y_1 | y_2 = v) dy_1 \quad (42)$$

$$= \frac{1}{1-F(v)} \int_v^{\bar{v}} y_1 f(y_1) dy_1. \quad (43)$$

Notice that for $r = v$ we have that (38) and (43) are identical.

The derivative of winner's regret with respect to the observed price is given by

$$w'(v) = \alpha_w \left[\frac{f(v)}{(1-F(v))^2} \int_v^{\bar{v}} y_1 f(y_1) dy_1 - \frac{1}{1-F(v)} v f(v) - 1 \right] \quad (44)$$

$$= \alpha_w \frac{f(v)}{1-F(v)} \left[\frac{1}{1-F(v)} \int_v^{\bar{v}} y_1 f(y_1) dy_1 - v - \frac{1-F(v)}{f(v)} \right] \quad (45)$$

$$= \alpha_w \frac{f(v)}{1-F(v)} [E(v_{(1)} | v_{(2)} = v) - v - h(v)]. \quad (46)$$

References

- Anwar, S., R. McMillan, M. Zheng. 2006. Bidding behavior in competing auctions: Evidence from eBay. *Eur. Econom. Rev.* 50(2) 307–322.
 Bell, D. 1982. Regret in decision making under uncertainty. *Oper. Res.* 30(5) 961–981.

- Binswanger, H. P. 1980. Attitude toward risk: Experimental measurement in rural India. *Amer. J. Agricultural Econom.* **62**(3) 395–407.
- Camerer, C., T.-H. Ho. 1994. Violations of the betweenness axiom and nonlinearity in probability. *J. Risk Uncertainty* **8**(2) 167–196.
- Campo, S., E. Guerre, I. Perrigne, Q. Vuong. 2011. Semiparametric estimation of first-price auctions with risk averse bidders. *Rev. Econom. Stud.* Forthcoming.
- Chen, K.-Y., C. R. Plott. 1998. Nonlinear behavior in sealed bid first-price auctions. *Games Econom. Behav.* **25**(1) 34–78.
- Chen, Y., P. Katuscak, E. Ozdenoren. 2005. Sealed bid auctions with ambiguity: An experimental study. Working paper, Center for Economic Research and Graduate Education–Economics Institute, Prague, Czech Republic.
- Cox, J. C., R. L. Oaxaca. 1996. Is bidding behavior consistent with bidding theory for private value auctions? *Res. Experiment. Econom.* **6** 131–148.
- Crawford, V., T. Kugler, Z. Neeman, A. Pauzner. 2009. Behaviorally optimal auction design: Examples and observations. *J. Eur. Econom. Assoc.* **7**(2–3) 377–387.
- Engelbrecht-Wiggans, R. 1989. The effect of regret on optimal bidding in auctions. *Management Sci.* **35**(6) 685–692.
- Engelbrecht-Wiggans, R., E. Katok. 2008. Regret and feedback information in first-price sealed-bid auctions. *Management Sci.* **54**(4) 808–819.
- Filiz-Ozbay, E., E. Y. Ozbay. 2007. Auctions with anticipated regret: Theory and experiment. *Amer. Econom. Rev.* **97**(4) 1407–1418.
- Fischbacher, U. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experiment. Econom.* **10**(2) 171–178.
- Goeree, J. K., C. A. Holt. 2004. A model of noisy introspection. *Games Econom. Behav.* **46**(2) 365–382.
- Goeree, J. K., C. A. Holt, T. R. Palfrey. 2002. Quantal response equilibrium and overbidding in private-value auctions. *J. Econom. Theory* **104**(1) 247–272.
- Greenleaf, E. 2004. Reserves, regret, and rejoicing in open English auctions. *J. Consumer Res.* **31**(2) 264–273.
- Hausman, J. A. 1978. Specification tests in econometrics. *Econometrica* **46**(6) 1251–1271.
- Holt, C. A., S. K. Laury. 2002. Risk aversion and incentive effects. *Amer. Econom. Rev.* **92**(5) 1644–1655.
- Kagel, J. H. 1995. A survey of experimental research. J. H. Kagel, A. E. Roth, eds. *Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 501–585.
- Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47**(2) 263–291.
- Katkar, R., D. H. Reiley. 2005. Public versus secret reserve prices in eBay auctions: Results from a Pokemon field experiment. Working paper, Stanford University, Stanford, CA.
- Krishna, V. 2002. *Auction Theory*, 1st ed. Academic Press, San Diego.
- Levin, D., J. L. Smith. 1996. Optimal reservation prices in auctions. *Econom. J.* **106**(438) 1271–1283.
- McAfee, R. P. 1993. Mechanism design by competing sellers. *Econometrica* **61**(6) 1281–1312.
- Myerson, R. 1981. Optimal auction design. *Math. Oper. Res.* **6**(1) 58–73.
- Prelec, D. 1998. The probability weighting function. *Econometrica* **66**(3) 497–527.
- Riley, J. G., W. F. Samuelson. 1981. Optimal auctions. *Amer. Econom. Rev.* **71**(3) 381–392.
- Tversky, A., D. Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertainty* **5** 297–323.
- Vuong, Q. H. 1989. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* **57**(2) 307–333.
- Wahrer, K., R. M. Harstad, M. H. Rothkopf. 1998. Auction form preferences of risk-averse bid takers. *RAND J. Econom.* **29**(1) 179–192.
- Wu, G., R. Gonzalez. 1996. Curvature of the probability weighting function. *Management Sci.* **42**(12) 1676–1690.