Identifying Unobservables With Experiments:

Risk Attitudes and Bidding Behavior

by

Glenn W. Harrison †

May 2005

ABSTRACT

One of the perennial challenges of testing economic theory is that predictions from theory often depend on unobservables. Experimental methods offer a significant methodological advance in such settings, since one can often design an experiment to identify the previously unobservable variable. But one confounding unobservable has remained latent in a wide range of experiments: risk attitudes. Recent advances in the direct elicitation of risk attitudes in experiments offer the prospect that one can finally design experiments that build in controls for this unobservable. An excellent testing ground for this strategy is provided by bidding in a first-price sealed-bid auction characterized by private and independent values. In a series of laboratory experiments data is collected on observed valuations and bids, using standard procedures. However, information is also elicited that identifies the risk attitudes of the same subject, since that is a critical characteristic of the predicted bid under the standard model. It is then straightforward to specify a joint likelihood function for the observed risk aversion responses and bids, estimate the risk aversion characteristic, and test if the implied Nash Equilibrium bid differs from the observed bid systematically. The results are striking. In the simplest possible case, when there are only two bidders, received theory does a wonderful job of characterizing behavior when one controls for the risk attitudes of the individual bidder.

† Department of Economics, College of Business Administration, University of Central Florida. E-mail contact: GHARRISON@BUS.UCF.EDU. All data and statistical software are available from the ExLab Digital Archive located at http://exlab.bus.ucf.edu.

One of the perennial challenges of testing economic theory is that predictions from theory often depend on unobservables. Experimental methods offer a significant methodological advance in such settings, since one can often design an experiment to identify the previously unobservable variable (Manski [2002]). But one confounding unobservable has remained latent in a wide range of experiments: risk attitudes. Recent advances in the direct elicitation of risk attitudes in experiments, due primarily to Holt and Laury [2002], offer the prospect that one can finally design experiments that build in controls for this unobservable.[1] To understand the important of doing so, consider the vast array of canonical experiments that force subjects to face risk in some form when one allows minimal off-equilibrium behavior: public good contribution games, common pool resource extraction games, ultimatum bargaining games, trust or investment games, labor gift-exchange games, battle of sexes games, coordination games, centipede games, and so on.

To illustrate the potential importance of controlling for the risk attitude confound, consider an important case in which there has been considerable debate over the ability of received theory to account for behavior: bidding in a first-price sealed-bid auction characterized by private and independent values.[2] Auction theory is very rich, and has been developed specifically for the parametric cases considered in experiments (e.g., Cox, Roberson and Smith [1982] and Cox, Smith and Walker [1988]). In a new series of laboratory experiments data are collected on observed valuations and bids, using standard procedures. However, information is also elicited that identifies the risk attitudes of the same subject, since that is a critical characteristic of the predicted bid under

[1] As noted crisply by Camerer [2003; p.154], there "... are three general strategies for controlling preferences in experiments: assume, measure, or induce." The first approach is to assume the problem away by assuming risk neutrality, perhaps on the grounds that the stakes are "small." This argument will simply not survive the weight of data showing moderate levels of risk aversion by typical subjects. The third approach is to use experimental designs that induce specific risk attitudes, such as the "lottery procedures" of Roth and Malouf [1979] and Berg, Daley, Dickhaut and O'Brien [1986]. There have been debates over the validity of these procedures, although much of that debate has been at cross purposes. The second approach is to directly measure risk attitudes and control for them, and is the approach employed here.

[2] See Kagel [1995] and Harrison [1989][1990] for a flavor of the debates.

the standard model (e.g., Harrison [1990]). It is then straightforward to specify a joint likelihood function for the observed risk aversion responses and bids, estimate the risk aversion characteristic, and test if the implied Nash Equilibrium bid systematically differs from the observed bid.

The results are striking. In the simplest possible case, when there are only two bidders (N=2), received theory does a wonderful job of characterizing behavior when one controls for the risk attitudes of the individual bidder.[3]

In section 1 we review theoretical predictions for this class of auction. In section 2 we describe our procedures. In section 3 we examine the results, and present the statistical model of behavior. Section 4 draws some conclusions.

1. Theoretical Predictions

Cox, Roberson and Smith [1982] develop a model of bidding behavior in first-price sealed bid auctions that assumes that each agent has a CRRA utility function $U(y) = y^r$, where U is the utility of experimental income y and $(1-r_i)$ is the Arrow-Pratt measure of risk aversion (RA). Each agent has their own r_i, so each agent is allowed to have distinct risk attitudes. However, r_i is restricted to lie on the closed interval [0,1], where $r_i = 1$ corresponds to risk neutrality (RN). Hence this model allows (weak) risk aversion, but does not admit risk-loving behavior.[4] Each agent in the

---

[3] Harrison, List and Tra [2005] show, however, that when auctions consist of more and more bidders, received theory does increasingly poorly in terms of characterizing "one shot" behavior. Their evidence suggests that received theory is relevant for "small auctions" but not for "large auctions." Thus if one were testing received theory it would matter on what domain the data were generated. Cox, Roberson and Smith [1982] reported different results, with the smallest of their auctions (N=3) generating the data that seemed to most obviously contradict the risk-averse Nash Equilibrium bidding model. However, this could have been due to collusion. In *all* of their experiments the same N bidders participated in multiple rounds, facilitating coordination of collusive under-bidding strategies that wreak havoc with the one-shot predictions of the theory.

[4] Cox, Smith and Walker [1988] offer a generalization that admits of some degrees of risk-loving behavior. Since we do not observe much risk-loving in the population used in these experiments, college students in the United States, this extension is not needed for present purposes.

model knows their own risk attitude, their own valuation v_i, that everyone's risk attitudes are drawn from the closed interval [0,1], and that everyone's valuation is drawn from a uniform distribution over the interval $[v_0, v^1]$. It can then be shown that the symmetric Bayesian Nash Equilibrium (NE) implies the following bid function:

$$b_i = v_0 + [(N-1)/(N-1+r_i)]\ (v_i - v_0)$$

where there are N active bidders. In the RN case in which $v_0=0$, $v^1=1$ and $r_i=1$, this model is the one derived by Vickrey [1961], and calls for bidders to choose their optimal bid using a simple rule: take the valuation received and shade it down by (N-1)/N. When N=2, the RN NE bidding rule is therefore particularly simple: bid one-half of the valuation. Thus one might expect the N=2 case to provide a particularly compelling test of the general RA NE bidding rule, since the optimal RN NE bid is also an arithmetically simple heuristic.[5]

2. Experimental Design and Procedures

Each subject in our experiment participated in a single session consisting of two tasks. The first task involved a sequence of choices designed to reveal each subject's risk preferences. In the second task, subjects participated in a series of ten first-price auctions against random opponents, followed by a small survey designed to collect individual characteristics.[6]

Subjects were recruited to the Behavioral Research Lab at the College of Business

[5] That is, ½ is arguably more focal than ⅔ or ¾, and so on for N>2. It is certainly easier to implement arithmetically, absent calculating aids.

[6] All tasks were undertaken with the instructions and procedures provided on the V*e*con Lab software developed by Charles Holt. See http://veconlab.econ.virginia.edu/guide.htm for an introduction. An obvious extension of these experiments would be to reverse the order of the risk attitude and bidding tasks, since there is some evidence of order effects in the elicitation of risk attitudes (Harrison, Johnson, McInnes and Rutström [2005]).

Administration of the University of Central Florida.[7] A total of 58 subjects participated over three sessions. The smallest number of subjects in one session was 16, so there was little chance that the subjects would (rationally) believe that they could establish reputations over the 10 rounds of bidding against a random opponent.

Each subject was told that they would be privately assigned induced values between $0 and $8, using a uniform distribution. Cox, Roberson and Smith [1982] show that for RN subjects the expected earning of each subject in a first price auction is $(v^1-v_0)/N(N+1)$, where v^1 and v_0 are the upper and lower bound for the support of the induced values. Thus expected RN earnings were $1.33 per subject in each period.

Subjects in each session were informed (a) of the number of other bidders in the auction; (b) that the other bidders' induced values were, like their own, drawn from a uniform support with bounds given above; and (c) that their earnings in the auction would equal their induced value minus their bid if they have the highest bid, or zero otherwise.

Holt and Laury [2002] (HL) devise a simple experimental measure for risk aversion using a multiple price list design, and we employed their procedure. Each subject is presented with a choice between two lotteries, which we can call A or B. Table 1 illustrates the basic payoff matrix presented to subjects in their experiments. The first row shows that lottery A offered a 10% chance of receiving $2 and a 90% chance of receiving $1.60. The expected value of this lottery, EV^A, is shown in the third panel as $1.64.[8] Similarly, lottery B in the first row has chances of payoffs of $3.85 and $0.10, for an expected value of $0.48. Thus the two lotteries have a relatively large difference in

[7] Recruitment was undertaken with the *ExLab* software available at http://exlab.bus.ucf.edu. The University of Central Florida is located in Orlando, and has over 44,000 undergraduates eager to earn cash in experiments. All subjects were recruited with the lure of a $5 participation fee and the expectation of additional earnings, although no information about the task or expected earnings was provided.

[8] The EV columns of Table 1 were not presented to subjects.

expected values, in this case \$1.17. As one proceeds down the matrix, the expected value of both lotteries increases, but the expected value of lottery B becomes greater than the expected value of lottery A. In our experiments we scaled these prizes up by a factor of 2, so that the largest prize was \$7.70 and the smallest prize was \$0.20.[9] The prizes in these lotteries effectively span the range of possible incomes in the auction, which range from \$8.00 to zero.

The subject is asked to choose A or B in each row of the counterpart of Table 1 presented to them. One row is later selected at random for payout for that subject. The logic behind this test for risk aversion is that only risk-loving subjects would take lottery B in the first row, and only risk-averse subjects would take lottery A in the second-to-last row. Arguably, the last row is simply a test that the subject understood the instructions, and has no relevance for risk aversion at all. A RN subject should switch from choosing A to B when the EV of each is about the same, so a RN subject would choose A for the first four rows and B thereafter.

3. Results

Figure 1 displays observed bidding behavior. The induced value is displayed on the bottom axis, a 45° line is shown and corresponds to the subject just bidding their value, and then the RN bid prediction is shown under that 45° line. The standard behavior from a long series of such experiments is observed: subjects tend to bid higher than the RN prediction, to varying degrees.

The statistical model consists of a likelihood of observing the risk aversion responses *and* the observed bidding responses.

The likelihood of the risk aversion responses is modeled with a probit choice rule defined over the 10 binary choices that each subject made, conditional on a value for the r coefficient in the

[9] Such scalar increases in all prizes do not change the CRRA values implied by choices in this task.

power utility function. Thus the likelihood function evaluates the expected utility of each lottery for a given candidate value of r, uses that to predict which lottery will be chosen, and infers the likelihood of the observed choices in the risk aversion task.

To allow for subject heterogeneity with respect to risk attitudes, the parameter r is modeled as a linear function of observed individual characteristics of the subject. For example, assume that we only had information on the age and sex of the subject, denoted Age (in years) and Female (0 for males, and 1 for females). Then we would estimate the coefficients α, β and γ in $r = \alpha + \beta \times \text{Age} + \gamma \times \text{Female}$. Therefore, each subject would have a different estimated r, $\hat{r}$, for a given set of estimates of α, β and γ to the extent that the subject had distinct individual characteristics. So if there were two subjects with the same sex and age, to use the above example, they would literally have the same $\hat{r}$, but if they differed in sex and/or age they would generally have distinct $\hat{r}$. In fact, we use 12 individual characteristics in our model. Apart from age and sex, these include binary indicators for race (NonWhite), a Business major, rich (parental or own income over $80,000 in 2003), high GPA (above 3.75), low GPA (below 3.25), college education for the father of the subject, college education for the mother of the subject, whether the subject works, whether the subject is a Catholic, and whether the subject is some other Christian denomination.

The likelihood of the bidding responses is then modeled as a multiplicative function of the predicted bid conditional on the estimated risk attitude for the subject. Thus we estimate a coefficient b which scales up or down the predicted NE bid: if b = 1 then the observed bid exactly tracks the predicted bid *for that subject*. The predicted NE bid for each subject i depends, of course, on the $\hat{r}_i$ for that subject, as well as the parameters N, v_0, v^1 and v_i. Thus if we observe two subjects with the same v_i but different bids, it is perfectly possible for this to be consistent with the predicted NE bid if they have distinct individual characteristics and hence distinct $\hat{r}_i$. The coefficient b is also modeled as a linear function of the same set of individual characteristics as the coefficient r.

The full specification of the likelihood function for bidding allows for heteroskedasticity with respect to individual characteristics. Thus the specification is $(b \times b^{NE}) + e$, where e is again a linear function of the individual characteristics. Thus we obtain information from the coefficients of b on which types of subjects deviate systematically from the NE prediction, and we obtain information from the coefficients on e on which types of subjects exhibit more noise in their bidding.

The overall likelihood consists of the likelihood of the risk aversion responses plus the likelihood of the bidding responses, conditional on estimates of r, b and e. In turn, these three parameters are linear functions of a constant and the individual characteristics of the subject. Since each subject provides 10 binary choices in the risk aversion task, and 10 bids in the auction task, we allow for the responses of the same subject to be correlated due to unobserved individual effects.

Table 2 displays the maximum likelihood estimates.[10] The intercept for r is estimated to be 0.612, consistent with evidence from comparable experiments of risk aversion.[11] The intercept for b is 1.02, consistent with bids being centered on the RA NE bid conditional on the estimated risk aversion for each subject. The top panel of Figure 2 shows the distribution of predicted values of b for each of the 58 subjects. Some subjects have estimates of b as low as 0.8, or as high as 1.35, but the clear majority seem to be tracked well by the RA NE bidding prediction. The bottom panel of Figure 2 displays a distribution of comparable estimates when we use the RN NE bidding prediction instead of the RA NE bidding prediction, and re-estimate the model. Observed bids are abut 25%

---

[10] Estimation was undertaken using the flexible "ml" procedure in *Stata*. All code is available at the *ExLab* archive for this project.

[11] For example, Holt and Laury [2002][2005] and Harrison, Johnson, McInnes and Rutström [2005a][2005b] for comparable subject pools.

higher than predicted if one assumes, counter-factually, that subjects are all RN.[12]

Figure 3 displays the predicted risk attitudes from this estimation exercise. One might be concerned that the full model fits the RA NE bidding model simply because it has a "free parameter r_i" to fit the bidding data to. In some sense this is true, since the joint likelihood of the data includes the effect of different $\hat{r}_i$'s on bids, and the estimates seek $\hat{r}_i$ values that explain the bidding data best. But it is not true entirely, since the joint likelihood must also explain the risk attitude choice data as well. So Figure 3 shows the difference between the distribution of predicted risk attitudes if one only uses the risk aversion tasks (top panel), and the distribution that is generated if one uses all data simultaneously (bottom panel). The two distributions are virtually identical. Kendall's τ statistic can be used to test for rank correlation; it has a value of 0.82, and leads one to reject the null hypothesis that the two sets of estimates of risk attitudes are independent at *p*-values below 0.0001.

4. Conclusion

Economists are increasingly recognizing that the power of the experimental method can be extended to generate observations on previously unobserved confounds of tests of theory.[13] Apart from risk attitudes, subjective beliefs are an obvious candidate in any strategic setting, and one can jointly estimate models of observed choices and elicited beliefs (e.g., Nyarko and Schotter [2002]). Similarly, preferences over the distribution of payoffs might also be a factor influencing behavior in strategic games, so one might elicit those preferences in non-strategic games as a control (e.g., Cox [2004]).

[12] In this case the joint likelihood of the risk aversion and bidding data collapses into two completely independent likelihoods, since the estimates from the former do not "feed into" the likelihood of the latter.

[13] This general point is, in fact, the major methodological innovation of experimental economics, the concept of "induced values." Binswanger [1982; p. 393] was among the first to see the broader implications for latent variables such as risk attitudes and subjective beliefs.

There are several new methodological issues involved in such designs, such as the effect that the additional elicitation has on observed behavior in the "target" task for which is a latent variable. For example, eliciting subjective beliefs about the likely behavior of opponents in a game could well influence behavior in that game, by encouraging subjects to think more like a game theorist.[14] However, these potential issues themselves should be amenable to appropriate experimental and statistical design to control for such effects.[15]

[14] For example, see Erev, Bornstein and Wallsten [1993], Croson [2000] and Rutström and Wilcox [2005].

[15] One could draw two samples at random from the same population, and have one sample undertake the belief elicitation task and the other sample undertake the game theory task. Manski [2002; §4] makes a similar suggestion.

Table 1: Design of the Holt and Laury Risk Aversion Experiments

Lottery A				Lottery B						
p($2)		p($1.60)		p($3.85)		p($0.10)		EV^A	EV^B	Difference
0.1	$2	0.9	$1.60	0.1	$3.85	0.9	$0.10	$1.64	$0.48	$1.17
0.2	$2	0.8	$1.60	0.2	$3.85	0.8	$0.10	$1.68	$0.85	$0.83
0.3	$2	0.7	$1.60	0.3	$3.85	0.7	$0.10	$1.72	$1.23	$0.49
0.4	$2	0.6	$1.60	0.4	$3.85	0.6	$0.10	$1.76	$1.60	$0.16
0.5	$2	0.5	$1.60	0.5	$3.85	0.5	$0.10	$1.80	$1.98	-$0.17
0.6	$2	0.4	$1.60	0.6	$3.85	0.4	$0.10	$1.84	$2.35	-$0.51
0.7	$2	0.3	$1.60	0.7	$3.85	0.3	$0.10	$1.88	$2.73	-$0.84
0.8	$2	0.2	$1.60	0.8	$3.85	0.2	$0.10	$1.92	$3.10	-$1.18
0.9	$2	0.1	$1.60	0.9	$3.85	0.1	$0.10	$1.96	$3.48	-$1.52
1	$2	0	$1.60	1	$3.85	0	$0.10	$2.00	$3.85	-$1.85

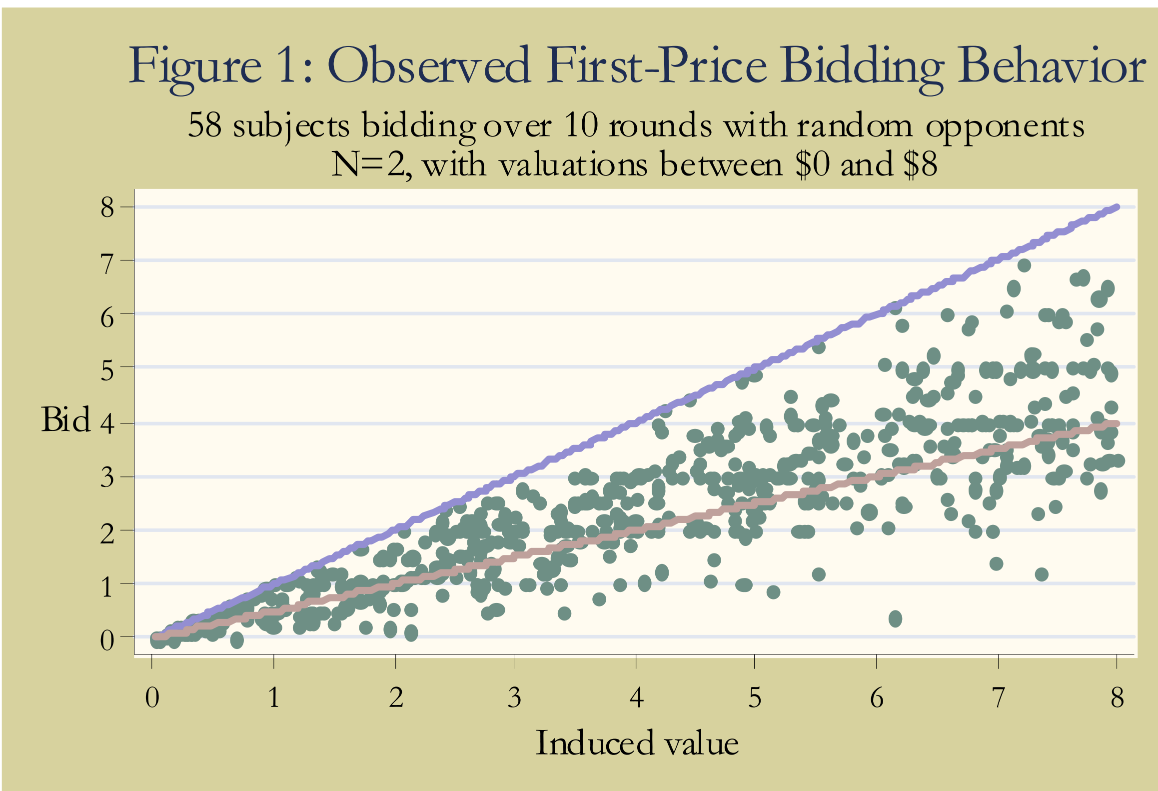

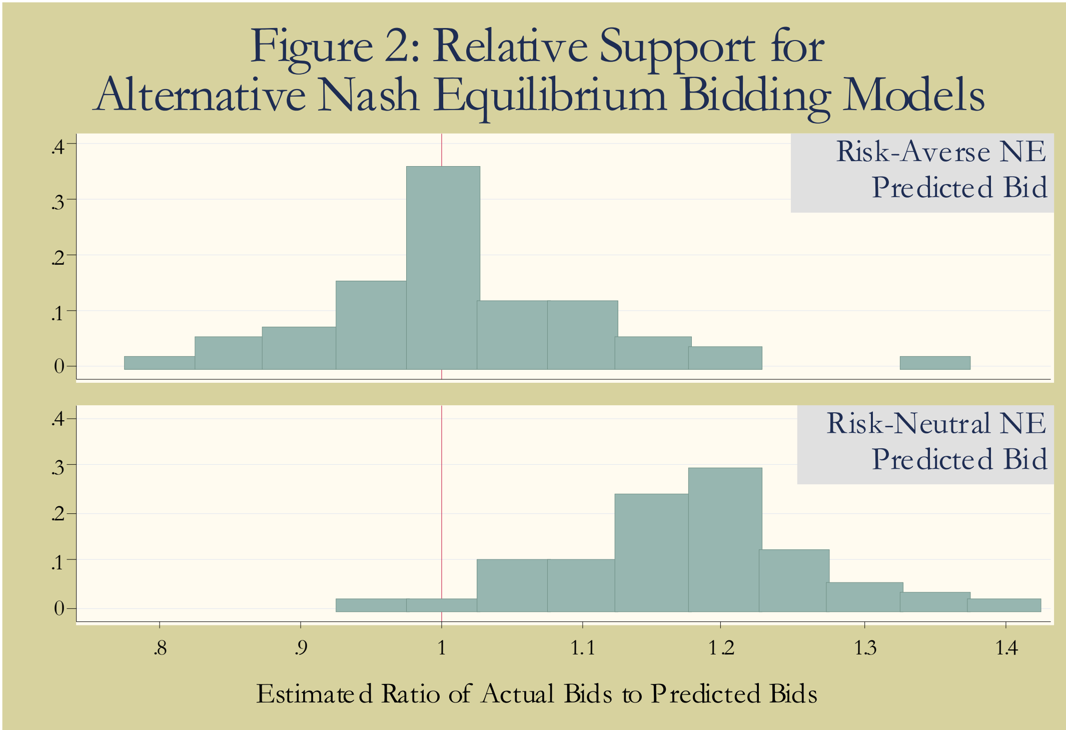

Figure 2: Relative Support for
Alternative Nash Equilibrium Bidding Models
Risk-Averse NE
Predicted Bid
Risk-Neutral NE
Predicted Bid
0
.1
.2
.3
.4
.8
.9
1
1.1
1.2
1.3
1.4
Estimated Ratio of Actual Bids to Predicted Bids

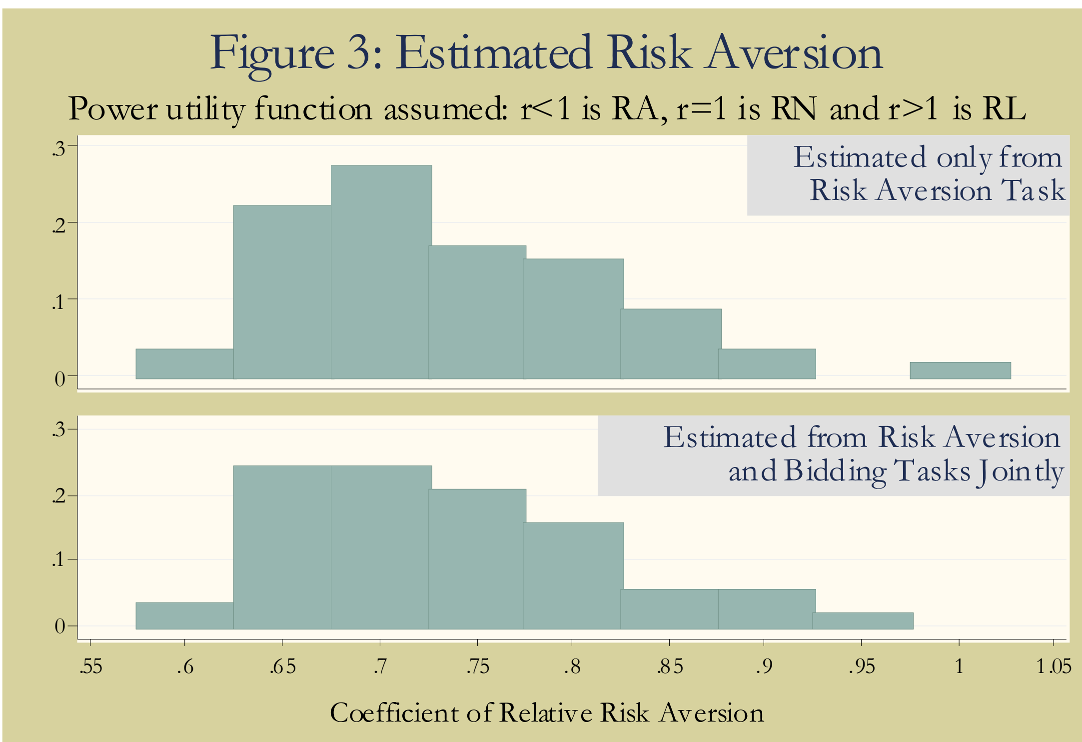

Figure 3: Estimated Risk Aversion
Power utility function assumed: r<1 is RA, r=1 is RN and r>1 is RL
Estimated only from
Risk Aversion Task
Estimated from Risk Aversion
and Bidding Tasks Jointly
0
.1
.2
.3
.55
.6
.65
.7
.75
.8
.85
.9
.95
1
1.05
Coefficient of Relative Risk Aversion

Table 2: Maximum Likelihood Estimates for Model of Bidding Behavior

Parameter	Variable	Point Estimate	Standard Error	*p*-value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
r	Constant	0.612	0.320	0.06	-0.02	1.24
	Age	0.003	0.015	0.82	-0.03	0.03
	Female	-0.052	0.079	0.51	-0.21	0.10
	Non-white	0.011	0.081	0.89	-0.15	0.17
	Major is in business	0.058	0.086	0.50	-0.11	0.23
	Father completed college	-0.004	0.083	0.96	-0.17	0.16
	Mother completed college	0.003	0.093	0.97	-0.18	0.19
	Income over $80k in 2003	0.036	0.073	0.62	-0.11	0.18
	Low GPA (below 3.24)	-0.024	0.098	0.81	-0.22	0.17
	High GPA (greater than 3.75)	0.190	0.113	0.09	-0.03	0.41
	Work full-time or part-time	-0.022	0.074	0.77	-0.17	0.12
	Catholic religious beliefs	-0.046	0.130	0.72	-0.30	0.21
	Other Christian religious beliefs	0.040	0.080	0.62	-0.12	0.20
b	Constant	1.021	0.721	0.16	-0.39	2.43
	Age	-0.007	0.030	0.81	-0.07	0.05
	Female	0.019	0.084	0.82	-0.14	0.18
	Non-white	-0.059	0.079	0.45	-0.21	0.09
	Major is in business	0.023	0.083	0.78	-0.14	0.19
	Father completed college	0.054	0.068	0.43	-0.08	0.19
	Mother completed college	-0.023	0.085	0.79	-0.19	0.14
	Income over $80k in 2003	0.078	0.074	0.29	-0.07	0.22
	Low GPA (below 3.24)	0.001	0.079	0.99	-0.15	0.16
	High GPA (greater than 3.75)	0.210	0.124	0.09	-0.03	0.45
	Work full-time or part-time	-0.019	0.068	0.78	-0.15	0.11
	Catholic religious beliefs	0.035	0.095	0.72	-0.15	0.22
	Other Christian religious beliefs	0.157	0.080	0.05	0.00	0.31
e	Constant	0.096	1.093	0.93	-2.05	2.24
	Age	-0.011	0.046	0.81	-0.10	0.08
	Female	0.008	0.156	0.96	-0.30	0.32
	Non-white	-0.077	0.162	0.63	-0.39	0.24
	Major is in business	0.123	0.133	0.36	-0.14	0.38
	Father completed college	-0.330	0.139	0.02	-0.60	-0.06
	Mother completed college	0.078	0.199	0.69	-0.31	0.47
	Income over $80k in 2003	0.044	0.119	0.71	-0.19	0.28
	Low GPA (below 3.24)	0.116	0.111	0.29	-0.10	0.33
	High GPA (greater than 3.75)	0.044	0.191	0.82	-0.33	0.42
	Work full-time or part-time	-0.144	0.148	0.33	-0.43	0.15
	Catholic religious beliefs	-0.341	0.157	0.03	-0.65	-0.03
	Other Christian religious beliefs	-0.077	0.149	0.61	-0.37	0.21

References

Berg, Joyce E.; Daley, Lane A.; Dickhaut, John W.; and O'Brien, John R., "Controlling Preferences for Lotteries on Units of Experimental Exchange," *Quarterly Journal of Economics*, 101, May 1986, 281-306.

Binswanger, Hans, "Empirical Estimation and Use of Risk Preferences: Discussion," *American Journal of Agricultural Economics*, 64, May 1982, 391-393.

Camerer, Colin F., *Behavioral Game Theory* (Princeton: Princeton University Press, 2003).

Cox, James C., "How To Identify Trust and Reciprocity," *Games and Economic Behavior*, 46(2), February 2004, 260-281.

Cox, James C.; Roberson, Bruce, and Smith, Vernon L., "Theory and Behavior of Single Object Auctions," in V. L. Smith (ed.), *Research in Experimental Economics* (Greenwich, CT: JAI Press, Volume 2, 1982).

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

Croson, Rachel, "Thinking Like a Game Theorist: Factors Affecting the Frequency of Equilibrium Play," *Journal of Economic Behavior & Organization*, 41(3), 2000, 299-314.

Erev, Ido; Bornstein, Gary, and Wallsten, Thomas S., "The Negative Effect of Probability Assessments on Decision Quality," *Organizational Behavior & Human Decision Processes*, 55, 1993,78-94.

Harrison, Glenn W., "Theory and Misbehavior of First-Price Auctions," *American Economic Review*, 79, September 1989, 749-762

Harrison, Glenn W., "Risk Attitudes in First-Price Auction Experiments: A Bayesian Analysis," *Review of Economics & Statistics*, 72, August 1990, 541-546.

Harrison, Glenn W.; Johnson, Eric; McInnes, Melayne M., and Rutström, E. Elisabet, "Risk Aversion and Incentive Effects: Comment," *American Economic Review*, 95, 2005 forthcoming.

Harrison, Glenn W.; List, John A., and Tra, Constant, "Statistical Characterization of Heterogeneity in Experiments," *Working Paper*, Department of Economics, College of Business Administration, University of Central Florida, 2005.

Holt, Charles A., and Laury, Susan K., "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), December 2002, 1644-1655.

Kagel, John H., "Auctions: A Survey of Experimental Research," in J. H. Kagel and A.E. Roth (eds.), *The Handbook of Experimental Economics* (Princeton: Princeton University Press, 1995).

Manski, Charles F., "Identification of Decision Rules in Experiments on Simple Games of Proposal and Response," *European Economic Review*, 46(4-5), 2002, 880-891.

Nyarko, Yaw, and Schotter, Andrew, "An Experimental Study of Belief Learning using Elicited Beliefs," *Econometrica*, 70, 2002, 971-1005.

Roth, Alvin E., and Malouf, Michael W. K., "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, 86, 1979, 574-594.

Rutström, E. Elisabet, and Wilcox, Nathaniel T., "Learning and Belief Elicitation: Observer Effects," *Working Paper*, Department of Economics, College of Business Administration, University of Central Florida, 2005.

Vickrey, William S., "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16, 1961, 8-37.