2006s-15

Overbidding in Independant Private-Values Auctions and Misperception of Probabilities

Olivier Armantier, Nicolas Treich

Série Scientifique
Scientific Series

Montréal
Août 2006

CIRANO
Centre interuniversitaire de recherche
en analyse des organisations

Overbidding in Independant Private-Values Auctions and Misperception of Probabilities [*]

Olivier Armantier[†], Nicolas Treich[‡]

Résumé / *Abstract*

Nous menons une expérience pour évaluer si le fait qu'un sujet évalue mal ses probabilités de gagner peut être une hypothèse alternative à l'aversion au risque pour expliquer les surenchères lors d'enchères indépendantes privées au premier prix. Les résultats expérimentaux montrent, en effet, que les sujets sous-estiment leurs probabilités de gagner l'enchère et ont tendance à surenchérir. Cependant, lorsqu'on leur présente plus de précisions sur leurs prédictions, les sujets apprennent d'abord à prédire correctement leurs probabilités de gagner, puis à limiter considérablement la surenchère. L'estimation de différents modèles du comportement suggère que i) les sujets sont hétérogènes par rapport à leurs préférences du risque et leurs perceptions des probabilités, ii) les sujets choisissent leur meilleure réponse conditionnellement aux croyances qu'ils révèlent, et iii) bien que nécessaire pour expliquer pleinement le comportement des sujets, l'aversion au risque semble jouer un rôle moins important que prévu. Finalement, nos résultats expérimentaux sont consistants avec un modèle théorique d'enchères standard qui combine l'aversion au risque et la mauvaise perception des probabilités.

Mots clés : aversion au risque, enchères, mauvaise perception des probabilités, surenchère

We conduct an experiment to test whether probability misperception may be a possible alternative to risk aversion to explain overbidding in independent first-price private-values auctions. The experimental outcomes indicate that subjects underestimate their probability of winning the auction, and indeed overbid. Yet, when provided with feed-back on the precision of their predictions, subjects learn first to predict their probability of winning correctly, and second to curb-down significantly overbidding. The structural estimation of different behavioral models suggests that i) subjects are heterogenous with respect to risk preferences and probability perceptions, ii) subjects tend to best-respond to their stated beliefs, and iii) although necessary to explain fully behavior, risk aversion appears to play a lesser role than previously believed. Finally, our experimental findings are shown to be consistent with a standard theoretical auction model combining risk aversion and misperception of probabilities.

Keywords: auctions, misperception of probabilities, overbidding, risk-aversion

Codes JEL : C70, C92, D44, D81

1. Introduction

There is a wealth of evidence indicating that subjects in independent first-price private-values auctions tend to bid above the risk neutral bayesian Nash equilibrium (hereafter RNBNE). Although often rationalized by risk aversion, there does not seem to be a consensus in the literature around the cause(s) of overbidding.[1] In particular, Goeree, Holt and Palfrey (2002) (hereafter GHP) suggest probability misperception as a possible alternative to risk aversion.[2] The object of the present paper is twofold: first, we conduct an experiment to establish whether or not subjects perceive correctly their probability of winning the auction; second, we verify whether misperception of probabilities may be considered a driving force behind overbidding.

The probability misperception hypothesis may be considered appealing to explain overbidding. Indeed, it has been repeatedly shown in the psychology literature on judgment that individuals have biased perceptions of probabilities. Typically, psychologists have observed a specific misperception pattern: individuals overestimate low probabilities, while they underestimate high probabilities.[3] Likewise, several individual decision experiments suggest that agents exhibit comparable probability distortion patterns when making risky choices.[4] It appears therefore natural to expect that agents may suffer from similar biases in games where probabilities are involved.[5]

---

[1]Other potential explanations include a "Joy of winning" (see e.g. Cox, Smith and Walker 1983, 1988), a lack of monetary incentives (the "flat maximum critique" of Harrison 1989), bidding errors (Kagel and Roth 1992), and asymmetric costs of deviating (Friedman 1992). See also the December 1992 issue of the American Economic Review for a flavor of the debate pertaining to the causes of overbidding.

[2]Cox, Smith and Walker (1985) were in fact the first to propose to study first-price auctions with a utility function exhibiting non-linearity in the probabilities. The authors however, only consider a power probability weighting function, and they conclude without further analysis that it is observationally equivalent to a model with risk aversion.

[3]See e.g. Sanders (1973), and Murphy and Winkler (1984) for probability forecasting in meteorology; Lichtenstein, Slovic, Fischhoff, Layman and Combs (1978), Viscusi, Hakes and Carlin (1997), and Benjamin, Dougan and Buschena (2001) for predictions of lethal risks; Viscusi and O'Connor (1984), and Gerkins, DeHaan and Schulze (1988) for perception of job related hazards; as well as Hurley and Shogren (2004) for evidence in sterile laboratory experiments.

[4]See the numerous references in Camerer (1995) for probability distortion in lottery choices; Ali (1977), and Golec and Tamarkin (1998) for horse track betting evidence; as well as Wakker, Thaler, and Tversky (1997) for probability distortion in insurance decisions.

[5]It has been suggested that probability distortion may reflect both the misperceptions that agents may have on the probabilities they face, as well as intrinsic preferences over probabilities when making a risky choice (see e.g. Kahneman and Tversky 1979). Following GHP's suggestion, we concentrate exclusively throughout the paper on probability misperception.

It is well known however, that the identification of probability perceptions separately from preferences is not trivial. Indeed, different combinations of beliefs and preferences may generate the same observed behavior. For instance, GHP's econometric estimation suggests that a model with a sensible constant relative risk aversion parameter, and a model in which risk neutral agents misperceive their probability of winning the auction, both fit bidding choices equally well. GHP therefore conclude that subjects behave "*as if*" risk averse, but their auction data do not enable them to establish unambiguously whether this should be attributed to probability misperception or risk aversion.[6] Following Manski (2002, 2004), we circumvent this identification problem by simultaneously eliciting choices and subjective probabilities. As a result, we can test directly which of the probability misperception or risk aversion hypothesis, is the most relevant to explain overbidding.

The experiment we propose therefore consists in a first-price independent private-values auction similar to the one in GHP. In addition, we also ask subjects to predict their probability of winning the auction in order to evaluate the extent of probability misperception. To promote truthful revelation, subjects are rewarded according to the accuracy of their predictions in addition to their auction profits. We conduct two different treatments differentiated by the feed-back provided to subjects at the end of each round on the precision of their predictions. No information is revealed in treatment 1, while in treatment 2 subjects are informed of the quality of their predictions.

The experimental outcomes in treatment 1 indicate that subjects overbid, and underestimate their probability of winning the auction. However, after observing their objective probability of winning, subjects in treatment 2 not only learn to make more accurate predictions, but they also drastically curb-down their tendency to overbid. In fact, bidding above the RNBNE virtually disappears once subjects have learned to predict correctly their probability of winning.

To explain the experimental outcomes, we estimate several structural models of noisy behavior. The estimation results suggest that subjects are heterogenous with respect to risk preference and probability perception. In addition, and according with Nyarko and Schotter (2002), we find that actions seem to be consistent with beliefs, as subjects appear to "best-respond" to their stated beliefs on their probability of winning. The structural models also suggest that probability

---

[6]Dorsey and Razzolini (2003) face a similar identification problem. Indeed, their experiment suggests the presence of probability misperception in private-values auctions, but they cannot formally confirm this hypothesis as they do not elicit beliefs.

misperception is a main source of overbidding, and that risk aversion may play a lesser role than previously believed. In fact, the estimated constant relative risk aversion parameter drops from 0.6 to 0.2 when one accounts for heterogeneity and for the probability misperceptions revealed by subjects.

Finally, we confirm theoretically the link observed in the experiment between probability misperception and overbidding. To do so, we consider a standard independent private-values model combining risk aversion and probability misperception. We show that a class of probability misperception functions, encompassing the distortion identified in our experiment, induces overbidding compared to the RNBNE under perfect perception.

The paper is structured as follows: the experimental design is presented in section 2 and discussed in section 3; the experimental outcomes are commented in section 4; different noisy models of behavior are estimated and compared in section 5; we present in section 6 our theoretical results derived from a standard private-values auction model with probability misperception and risk aversion; finally section 7 concludes.

2. The Experimental Design

We present in this section the different experimental treatments. The choices made when designing the experiment are then discussed in a subsequent section. The experiment was conducted with volunteers at the State University of New York at Stony Brook. There were eight experimental sessions, four for each treatment, and each session included 10 subjects and 15 rounds. No subject participated in more than one session. At the beginning of a session, players were assigned to an isolated computer. Subjects were told in advance how many rounds would be played, and they knew that the experiment would not exceed one hour and thirty minutes. Instructions were then read aloud, followed by participants' questions, and a brief training with the computer software.[7]

As further discussed in section 3, the experimental design is essentially motivated by the following three objectives: first, the auction has to be similar to the one in GHP; second, the subjects' behavior in the auction has to be independent from the elicitation of their beliefs about their probability of winning; third, the remuneration scheme must lead subjects to reveal their beliefs as precisely as possible.

---

[7]The complete list of instructions is available in Appendix A.

Before describing how the experiment unfolds, we summarize the discrete auction model in GHP. Two players participate in a sealed bid independent private-values auction. Each player receives a randomly determined prize-value, which is equally likely to be $0, $2, $4, $6, $8, or $11. The players must simultaneously make a sealed bid, which is constrained to be an integer dollar amount. The prize is awarded to the highest bidder (with ties decided by the flip of a coin), for a price equal to his bid. GHP prove that the unique RNBNE is to bid $0, $1, $2, $3, $4, and $5, for values of respectively $0, $2, $4, $6, $8, and $11. The equilibrium strategy yields subjects an average profit of $1.9 per round.

In each round, subjects are randomly matched in pairs. To avoid reputation building, the subjects are informed that the assignment is such that it is not possible to identify the other member of the pair. The problem in each round may be decomposed in three phases. In phase 1, the GHP's auction is implemented in strategy form. In other words, subjects are asked to identify the bid they would make for each possible prize value. They are told that a prize value will be assigned to them at the end of the round, and that their conditional choice in phase 1 corresponding to that value would determine what we called their "effective bid". Subjects are also informed that the outcome of the auction is decided by comparing the effective bids of the two members of the pair. We emphasized in the instructions that subjects should make careful decisions in phase 1, as their effective bid directly influences their auction payoffs.

In phase 2, subjects are asked to predict their own probability of winning the auction for a given list of bids. A coin flip decides whether all subjects must make their predictions for a list of even bids (i.e. $0, $2, $4, $6, $8 and $10), or uneven bids (i.e. $1, $3, $5, $7, $9 and $11).[8] Subjects are asked to express as a frequency (i.e. an integer between 0 and 100) their chances of winning the auction for each of the six bids in the list.[9] We carefully explained to participants that their probability of winning with a given bid is equal to the probability that this bid is higher than the effective bid of an average person in the room. We also made them aware that the experiment was designed in such a way that their choices and payoffs in phase 1 are completely independent of their choices and

---

[8] GHP also conduct a "low cost" treatment with prize values of $0, $2, $4, $6, $8, and $12. Although this "low cost" treatment was likely to generate slightly more overbidding, it would have required to ask predictions for either six uneven bids, or seven even bids. To avoid such asymmetry, we preferred to adopt GHP's "high cost" treatment.

[9] In a series of pilot experiments, we experienced different mechanisms for subjects to express their beliefs (e.g. a scale, a pie, a comparison with other random events). No substantial difference was identified in the way subjects reported their beliefs (i.e. the distributions of the predictions in the various treatments were not statistically different).

payoffs in phase 2 (and vice-versa).

Once every prediction has been submitted, we determine in phase 3 the auction and prediction payoffs. To do so, we randomly match subjects in pairs. Then, we successively stop by each subject computer station, and roll a 6-sided die to assign a prize value to each participant. As previously mentioned, we determine a subject's effective bid by matching his prize value with his corresponding conditional choice in phase 1. The effective bids of the member of each pair are then compared, and the auction payoffs are calculated. To determine the prediction payoffs, a 6-sided die is first rolled once in public. The outcome of the draw determines which of the six integer bids in the list given in phase 2 will be used to measure the accuracy of the subjects' predictions. The predictions of both members of a pair are compared with an "objective probability" of winning the auction for the bid randomly selected. We carefully explained that this objective probability is calculated precisely by looking at the bid decisions made in phase 1 of that round by all the participants in the room, other than the members of the pair to which the subject belongs. As further discussed in section 3, the objective probability differs across pairs, but it is the same for both members of a pair. The prediction payoff is $4 to the member of the pair with the closest prediction to the objective probability for the bid randomly selected. We told subjects that in the event of a tie, the allocation of the $4 would be decided by the flip of a coin. Note that, to promote a similar level of introspection when selecting bids and making predictions, we have arranged so that the number of choices and the expected payoffs are roughly the same in phases 1 and 2.

At the end of the round, the auction outcomes (i.e. the bids of both members and their own payoffs) are revealed to each member of the pair. This is the only information revealed in treatment 1. In treatment 2, participants are also informed of the quality of their absolute predictions. More precisely, we presented both on a graph and in a table the subject's predictions along with its objective probabilities of winning for each of the six integer bids in the list given in phase 2. To promote absolute precision, rather than relative precision, we did not reveal to the subjects in treatments 1 and 2 the predictions payoffs, nor the accuracy of any other participant. In other words, subjects do not know until the end of the session whether their prediction in a given round was more accurate.

Following GHP, subjects were paid in cash at the end of the session half of their accumulated earnings. The auction earnings in treatment 1 are comparable to those in GHP (respectively $10.58 versus $10.7 in GHP), but auction payoffs in treatment 2 are slightly higher ($11.49). The auction profits increase over time in

both treatments (from $9.38 in the first three periods to $10.25 in the last three periods of treatment 1, and from $9.00 to $14.06 in treatment 2). Except for the last three periods of treatment 2, auction profits remain significantly smaller than the RNBNE expected profit of $14.25.

3. Comments on the Experimental Design

We briefly justify in this section some of the choices concerning the experimental design. The first objective of the paper is to evaluate whether subjects misperceive their probability of winning the auction. To do so, we compare a player's predictions with its objective probability, which can only be derived if we observe his (potential) opponents full bidding strategy. Two features of the design allows us to infer the subjects strategies: first, the possible private-values are discrete and finite; second the auction is implemented in its strategic form. As a result, the design allows us to observe the full bidding strategy for any bidder, and therefore we can calculate his opponents' objective probability of winning the auction. In addition, this approach allows us to collect a much larger sample of data in each round than GHP. This additional information will prove helpful in section 5 when we estimate different behavioral models. Note that it has been argued, that the strategy method may generate different behaviors in certain experimental games (see e.g. Roth 1995). For instance, it is conceivable that a subject may regret his bid selection in phase 1 when he observes the actual prize value assigned to him. To prevent such a problem, we emphasized to subjects that they had to select their bids carefully in phase 1, as one of them would directly influence their auction payoffs. Moreover, we shall see in section 4 that a comparison with the experimental outcomes obtained by GHP under the extensive implementation of the game does not indicate any significant treatment effect.[10]

How to elicit subjective probabilities is a question that has been often debated among psychologists, statisticians and economists. Two desired properties are difficult to reconcile when devising a mechanism: the ease of implementation and the theoretical properties. Most psychologists (and some economists) believe that, if asked, subjects will reveal their best estimates.[11] Economists are in general concerned that such a simple approach may not incite subjects to report their

---

[10]Selten and Buchta (1999), as well as Pezanis-Christou and Sadrieh (2003) adopted an essentially comparable strategy method in an auction experiment. According with our results, they did not identify the presence of a significant treatment effect.

[11]See e.g. Lichtenstein et al. (1978), Viscusi et al. (1997), as well as Manski (2004) and the references therein.

true probabilities. Instead, economists often prefer to use "proper" scoring rules for which payoffs are maximized by truthful revelation (see e.g. De Finetti 1965, Murphy and Winkler 1970, Savage 1971). It is well known however, that this procedure is not incentive compatible when subjects' expected utilities are non-linear in payoffs and/or in probabilities, as it will turn out to be the case in our experiment. In addition, scoring rules do not necessarily encourage a subject to assess as precisely as possible his subjective probability since i) most scoring rules are essentially flat around the optimum, and ii) untruthful assessments can secure higher minimum payoffs. Scoring rules however, have been mostly criticized for their complexity to be fully comprehended by non-professional forecasters (Hogarth 1987, Wilcox and Feltovich 2000, or Read 2003).[12] More importantly, recent evidence suggest that eliciting beliefs during the course of a repeated game experiment with intrusive procedures such as scoring rules, may alter the way subjects play the game (Croson 2000, Camerer, Ho and Chong 2001, Rutström and Wilcox 2003). Since no conclusive evidence indicates that scoring rules generate significantly more accurate predictions (Camerer and Hogarth 1999, or Sonnemans and Offerman 2001), some economists have preferred to abandon this approach in favor of simpler methods more transparent to subjects (e.g. Dufwenberg and Gneezy 2000, Croson 2000, Wilcox and Feltovich 2000, Charness and Dufwenberg 2003). Likewise, we adopt a simple method consisting in comparing subjects predictions, and rewarding only the closest to the objective probability. This technique is commonly used in practice to obtain accurate estimates, and it has been shown to be incentive compatible for subjects with uniformed priors (Ottaviani and Sørensen 2003).[13] In fact, only 3 out of the 80 participants declared in a post-experiment survey that they did not systematically try to report their best estimates.[14] In addition, we will see that our elicitation procedure is further validated by the fact that the probabilities reported by subjects are consistent with their actions.[15]

---

[12] In fact, several experimental studies based on scoring rules merely ask subjects to trust the experimenter that truthful revelation is their best strategy (see e.g. Offerman, Sonnemans and Schram 1996, Offerman 2002, Costa-Gomez and Weizsäcker 2004).

[13] Examples of practical implementations of this method include the National Collegiate Weather Forecasting Contest, the Earthquake Prediction Contest, the Federal Forecasters Forecasting Contest, or the Wall Street Journal semi-annual forecasting survey.

[14] In addition, 79 out of 80 subjects declared that they perfectly understood the prediction payment mechanism. In contrast, only 28 out of 40 subjects declared that they systematically provided their best estimates, and 18 out of 40 subjects declared that they perfectly understood the prediction payment mechanism, in a comparable experiment conducted with a quadratic scoring rule.

[15] For the sake of completeness, we conducted a series of additional sessions with a quadratic scoring rule. The results obtained, although slightly noisier, are not substantially different from the one presented here. In other words, our conclusions appear to

Let us now discuss how subjects' predictions are induced independently from their bidding behavior. A possible strategy may consist in comparing every subject's prediction with a single objective probability of winning, which in this case would be the average probability of winning calculated with the bids of all participants. However, the decisions and payoffs of a player in phases 1 and 2 are not independent in this situation, since his bidding decisions in phase 1 influence the objective probability of winning. In this context, one could conceivably imagine that a player may submit absurd bids in phase 1, to skew knowingly the probability of winning and improve his chances of winning the prediction reward in phase 2. The payment scheme adopted in our experiment do not suffer from this drawback. Indeed, a subject's prediction is compared to an objective probability based on the actions of all participants except the members of the pair to which he belongs. Therefore, the members of a given pair cannot influence with their bidding decision in phase 1 the objective probability to which they will be compared. In other words, an individual's payoffs and actions in phase 1 cannot affect his payoffs and actions in phase 2 (and vice versa).

We also strived to find an appropriate balance between collecting as much information as possible, and requiring a similar amount of introspection in each of the bidding and prediction phase. As a result, we do not ask subjects to make a prediction for each of the twelve bids between $0 and $11. Instead, we opted to ask subjects to submit six bids in phase 1, and to make a prediction for a list of six even or uneven bids in phase 2. A drawback of this approach is that we do not observe the full probability weighting function (hereafter PWF) for each subject. However, note that in each period, the unobserved predictions (e.g. the list of uneven bids) are equally spaced between the answers provided by subjects (e.g. the list of even bids). Therefore, we can easily infer unobserved predictions in our subsequent analysis by applying simple interpolation techniques.

We did not want subjects to infer any information from the questions asked during a round. For instance, we could have asked subjects to predict their probability of winning for a single bid. This approach would have yielded much less information, but more importantly, we were concerned that subjects may interpret this as a signal from the experimenter that they should have selected this bid. Instead, we preferred to ask predictions for either even or uneven bids, which we believed would be less likely to send any signal.

---

be robust to the elicitation method employed. A summary of the experimental outcomes under the quadratic scoring rule may be found on one of the authors' website at http://www.sceco.umontreal.ca/liste_personnel/armantier/index.htm.

Finally, the objective of the paper is not only to test for the presence of misperception, but also to compare the explanatory power of the probability misperception and risk aversion hypotheses. When designing treatment 2, our contention was that, given the financial incentives, providing feed-back on the accuracy of their predictions should help subjects correct their perception bias (if any). This contention was motivated by experiments in psychology showing that training may significantly reduce judgement errors (see e.g. Fischhoff 1982). If predictions biases could be totally eradicated, then we could perfectly distinguish between probability misperception and risk aversion. Indeed, the probability misperception hypothesis would have little credit, if overbidding remains prevalent even though subjects have correct estimates of their probability of winning. However, if overbidding is eliminated when perceptions are accurate, then this would indicate that risk aversion may not be the primary source of overbidding as previously believed. Although we did not expect to eradicate totally perception biases, we believed that a partial correction would help us better disentangle risk aversion from probability misperception.

4. Experimental Outcomes

4.1. Treatment 1

The experimental outcomes observed in treatment 1 are summarized in Tables 1 to 3, as well as Figures 1, 3 and 5. Figure 1 displays on the X-axis a subjects' objective probabilities of winning the auction, and on the Y-axis the subjects' predictions in treatment 1 (averaged across periods and subjects). The figure actually consists of twelve consecutive dots, representing the predictions for integer bids ranging from \$0 to \$11. For instance, the first dot on the left corresponds to the predicted probability of winning with a bid $b = 0$. The diagonal has also been plotted to guide the eye. If subjects make unbiased predictions, then their stated probabilities should fall around the diagonal. Figure 1 indicates that subjects systematically underestimate their probability of winning for any bid between \$0 and \$11. The underestimation is in fact quite significant for a large number of bids. For instance subjects believe that their probability of winning with a bid $b = 3$ and $b = 5$ (the RNBNE bids for a private-value of $v = 6$ and $v = 11$) are respectively 19.66% and 44.86%, while the actual probabilities are 47.38% and 74.08%.

It is interesting to note that, although we find a slight evidence of overesti-

mation for the lowest bid (i.e. $b = 0$) during the first three periods, subjects' predictions did not produce the S-shape pattern of misperception commonly observed in the psychology literature (see e.g., Camerer, 1995). This observation is consistent with GHP whose results also suggest that the PWF may be convex everywhere. In other words, it appears that subjects in our experiment are pessimists, and do not exhibit the traditional bias (i.e. overestimate small probabilities and underestimate large probabilities). Similar evidence has been detected in other experimental studies (e.g. Schotter and Sopher 2001 identify pessimistic beliefs in an ultimatum game experiment), as well as field studies (Giordani and Soderlin 2003 show that professional forecasters have been historically pessimistic in their predictions of GDP). Our experimental results may therefore raise the question of the relevance of S-shaped PWF to apply to private-values auctions, and may be more generally to game situations.

No evidence of learning may be detected in Figure 1. Indeed, the predictions in the first and last three periods appear indistinguishable. This observation is only partially confirmed by the estimation of an econometric model in which the predictions stated by subjects are assumed to be drawn from a normal distribution truncated on $[0, 100]$.[16] The expected predictions are modeled as

$$E\left[\widehat{P}_{i,b,t}\right] = \lambda_b \left(1 + \delta_b^5 P_{>5} + \delta_b^{10} P_{>10} + \eta_i\right) \quad, \tag{4.1}$$

where $\widehat{P}_{i,b,t}$ is the prediction of player $i = 1, ..., 10$ in period $t = 1, ..., 15$ regarding his probability of winning the auction with a bid equal to $b = \$0, ..., \$11$; $(P_{>5}, P_{>10})$ are two dummy variables defined such that $P_{>5}$ (respectively $P_{>10}$) equals one when $t > 5$ (respectively $t > 10$), and equals zero otherwise; $\left(\lambda_b, \delta_b^5, \delta_b^{10}\right)$ are parameters to be estimated;[17] and η_i is a normally distributed individual random effect with mean zero and variance σ_η^2.[18] Finally, to account for possible heteroskedasticity across predictions for different bid values and/or different time periods, we model the standard deviation of the predictions as

$$Std\left(\widehat{P}_{i,b,t}\right) = \sigma \left|\gamma_1 - b\right|^{\gamma_2} t^{\gamma_3} \quad. \tag{4.2}$$

---

[16] The truncation reflects the fact that subjects' predictions must necessarily lay between 0 and 100.

[17] The parameter λ_b may be interpreted as the initial average prediction for a bid b, while δ_b^5 and δ_b^{10} represent the average percentage deviation in the predictions stated for a bid b after respectively 5 and 10 periods.

[18] To identify the presence of learning, a non-linear specification of the form (4.1) was preferred to a more traditional linear regression with time trends, because it enables i) to compare the speed of learning between early and late periods, and ii) to compare the speed with which subjects adjust their bids, and the speed with which subjects adjust their predictions.

11

This specification of the variance is quite flexible and it allows in particular i) for subjects to learn to make more homogenous predictions with time (i.e. $\gamma_3 < 0$); and ii) for predictions to be more homogenous for low and high bids (for which most predictions are likely to be close to 0 and 100) than for intermediate bids (i.e. $\gamma_1 \in [0, 11]$ and $\gamma_2 < 0$ in which case the variance has an inverse U shape).

The results of the maximum likelihood estimation presented in Table 1 indicate no systematic adjustment, as most estimated parameters $\left(\widehat{\delta}_b^5, \widehat{\delta}_b^{10}\right)$ are not significantly different from zero.[19] Subjects however appear to learn to lower their predictions for small bids after the first five periods, since $\widehat{\delta}_b^5$ is significantly smaller than zero for $b = 0, ..., 4$. In fact, the reduction in the stated probabilities may be quite consequent. For instance, when asked to predict their probability of winning for a bid $b = 0$ (respectively $b = 1$), subjects lower their predictions by 55.1% (respectively 43.9%) after the first five periods, and an additional 24.7% (respectively 16.5%) after period 10. This decrease is consistent with reinforcement learning, as most subjects did not witness during the experiment any auction won by a bidder submitting a bid of 0 or 1.[20] In other words, the only evidence of learning we find is that subjects quickly reduce their predictions for small bids.

The estimates of the parameters controlling the variance of the predictions are presented in Table 2. We find that $\widehat{\gamma}_1$ is almost centrally located in $[0, 11]$, while $\widehat{\gamma}_2$ is significantly lower than zero. As expected, this implies that the variance of the predictions has a nearly symmetrical inverse U-shape. In other words, predictions are more homogenous for bids with which bidders are almost certain to loose or to win the auction (i.e. low and high bids). The variance of the predictions, however, does not seem to contract with time, as $\widehat{\gamma}_3$ is found to be insignificant. Finally, we identify substantial heterogeneity within the subjects' pool as $\widehat{\sigma}_\eta^2$, the variance of the individual random effect, is found to be significant and relatively large. For instance, a subject with an individual effect equal to one standard deviation $(\eta_i = \widehat{\sigma}_\eta)$ will submit predictions 12.2% larger than the average bidder (for whom $\eta_i = 0$).

Let us now turn to the bids submitted in treatment 1. The average bids actually submitted for each possible private-value (red line) may be compared in

---

[19]Throughout the paper the estimates' standard deviations, and the distributions of the test statistics have been evaluated by bootstrap, in order to control for the finiteness of the sample (see Shao and Tu 1995). Note also that the regressions in this section have also been conducted with dummy variables identifying the session in which the subjects participated. No significant session effect has been detected.

[20]Out of the 300 auctions conducted during the four sessions of treatment 1, 2 were won by a bidder submitting a bid of 0, and 9 by a bidder submitting a bid of 1.

Figure 3 with the RNBNE (black line). The figure confirms that except for low private-values (i.e. $v = 0$ and $v = 2$) subjects tend to bid significantly higher than the RNBNE. This result does not necessarily imply that subjects are either risk averse, or not perfectly rational. Indeed, even though subjects do not submit the RNBNE, it is entirely possible that they best-respond to their opponents actions, or that their choices are consistent with their beliefs. To explore these hypotheses, two additional bid functions are also plotted in Figure 3. The blue line represents a subject's risk neutral objective best-response. In other words, it represents what a risk neutral agent should bid if he knew or could infer correctly the other participants actions. The orange line represents a subject's risk neutral subjective best-response. In other words, it represents what a risk neutral agent should bid conditional on her stated beliefs regarding her probability of winning the auction.[21] Figure 3 indicates that subjects overbid compared to their objective best-responses. Their actions however, appear to be quite consistent with their beliefs. Indeed, the average bids submitted for the different private-values are only slightly higher than the subjective best-responses.[22] The previous results confirm that subjects cannot be assimilated to perfectly rational risk neutral agents since they do not comply with either the RNBNE, or their objective risk neutral best-responses. Subjects however, may not be far from risk neutral utility maximizers, since their actions are close to their subjective risk neutral best-responses. Nevertheless, the slight overbidding remaining may still be explained by risk aversion, although it would suggest that risk aversion may play a lesser role than previously believed. This conjecture will need to be statistically confirmed in section 5 when we estimate noisy behavioral models.

Figure 5 indicates that, if anything, subjects learn to increase slightly their bid over time. This observation is confirmed by the estimation of an econometric model in which the bids are assumed to follow a normal distribution truncated on $[0, 11]$, and the expected bids are modelled as

$$E\left[B_{i,v,t}\right] = \lambda_v \left(1 + \delta_v^5 P_{>5} + \delta_v^{10} P_{>10} + \eta_i\right) \quad , \tag{4.3}$$

where $B_{i,v,t}$ is the bid submitted by player $i = 1, ..., 10$ in period $t = 1, ..., 15$ for a

---

[21]Unlike previous private-values experiments we can compare subjects' choices with their objective and subjective best-responses. Indeed, two features specific to our design enable the calculation of these best-responses: first, the auction is implemented in strategic form; second, we elicit subjects' beliefs about their probability of winning.

[22]This observation is consistent with Nyarko and Schotter (2002), as well as Bellemare, Kröger and van Soest (2005) who find that subjects' actions appear to me more compatible with their subjective rather than their objective best-responses.

private-value $v \in \{0, 2, 4, 6, 8, 11\}$.[23] The remaining parameters and variables are defined as in equation (4.1). Finally, the standard deviation of the bids is also modeled as

$$Std\,(B_{i,v,t}) = \sigma\,|\gamma_1 - v|^{\gamma_2}\,t^{\gamma_3} \quad . \tag{4.4}$$

The maximum likelihood outcomes in Table 1, indicate that subjects only learn to increase their bids significantly for private-values equal to $v = 4$ and $v = 6$. In other words, we cannot find conclusive evidence of systematic strategy adjustment. The estimates of the standard deviation parameters in Table 2 indicate that i) the variance of the bids increases with the private-value v (i.e. $\widehat{\gamma}_1$ is insignificant and $\widehat{\gamma}_2$ is positive), ii) subjects' behavior does not become more homogenous with time (i.e. $\widehat{\gamma}_3$ is insignificant), and iii) there is significant heterogeneity between subjects in their bid selection (i.e. $\widehat{\sigma}_\eta$ is significantly larger than zero). In particular, a change of one standard deviation in the random effect results in a 14.8% variation in the bids of an individual across all possible values.

The frequency with which subjects submitted the correct RNBNE bid for a given private-value is reported in Table 3. This table shows that although subjects select the RNBNE bid 98.67% and 94.67% of the time for the lowest private-values of $v = 0$ and $v = 2$, they very rarely submit the RNBNE bid for the highest private-values of $v = 8$ and $v = 11$ (respectively 12.17% and 5.67%). These percentages vary only slightly between the first and last three periods, which further confirms that subjects do not learn to play the RNBNE bids. The last column of Table 3 represents the frequency with which subjects submit the entire RNBNE bid function (i.e. $B(0) = 0$, $B(2) = 1$, $B(4) = 2$, $B(6) = 3$, $B(8) = 4$ and $B(11) = 5$). The correct RNBNE strategy has been chosen only 1.00% of the time overall, but there seems to be a slight increase between the first and last 3 periods from 0.83% to 4.17%.

Finally, note that the auction outcomes in treatment 1 appear to be both qualitatively and quantitatively consistent with those in GHP. In other words, implementing the auction in strategy form, and asking subjects to predict their probability of winning did not appear to introduce any significant treatment effect. This result is consistent with Rutström and Wilcox (2003) who suggest that eliciting beliefs with a non-intrusive procedure does not affect significantly the way subjects play a game (see also Nyarko and Schotter 2002).

---

[23]Note that although bidding is not constrained to $[0, 11]$, we did not observe any bid outside this interval during the experiment.

4.2. Treatment 2

The experimental outcomes observed in treatment 2 are summarized in Tables 1 to 3, as well as Figures 2, 4 and 6. Figure 2 shows that providing feed-back has a dramatic effect on the accuracy of subjects predictions. Indeed, although predictions in the first three periods are roughly similar to treatment 1, we see in Figure 2 that subjects make nearly unbiased estimates of their probability of winning during the last three periods of treatment 2. This observation is confirmed by the estimation of an econometric model specified as in (4.1) and (4.2). Indeed, Table 1 indicates that subjects' predictions for all possible bids increase rapidly after the first five periods, and keep increasing, although at a slower pace, after period 10. For instance, when asked to predict their probability of winning for a bid $b = 2$, subjects increase their predictions by 41.2% after the first five periods, and an additional 12.2% after period 10. In other words, from the feed-back provided, subjects in treatment 2 appear to learn to correct their probability misperception over time.[24] This result seems to contrast sharply with the outcomes in treatment 1.

To confirm statistically the presence of a treatment effect, we estimate a model in which the predictions stated by subjects in treatments 1 and 2 are assumed to be drawn from a normal distribution truncated on $[0, 100]$. The standard deviations of the predictions are specified as in (4.2), while the means are now defined as

$$
E\left[\widehat{P}_{i,b,t}\right] = \left(\lambda_b + \widetilde{\lambda}_b T_2\right)\left[1 + \left(\delta_b^5 + \widetilde{\delta}_b^5 T_2\right) P_{>5} + \left(\delta_b^{10} + \widetilde{\delta}_b^{10} T_2\right) P_{>10} + \eta_i\right] \quad,
$$
(4.5)

where T_2 is a dummy variable equal to 1 when the observation was collected in treatment 2. The model is estimated by maximum likelihood with the joint sample consisting of the data collected in treatments 1 and 2. The estimation results in Table 4 indicate that subjects' predictions in the early periods of treatments 1 and 2 are nearly indistinguishable. Indeed, except for $b = 1$ and $b = 2$, the

---

[24]Although subjects were explicitly warned that the objective probabilities were likely to change in successive rounds as other players modify their bidding behavior, a possible strategy for a subject in treatment 2 could be to report the probabilities revealed to him by the experimenter in the previous round. The plausibility of this hypothesis is in part affected by the fact that subjects have a 50% to make predictions for a different list of bids (either even or uneven bids) in two consecutive rounds. Inspection of the data also reveals that when a subject has to make predictions for the same list of bids in rounds t and $t-1$ (respectively, $t-2$), his prediction for a given bid in round t equals the probability revealed to him by the experimenter in round $t-1$ (respectively, $t-2$) only 7.2% (respectively, 3.8%) of the time. In other words, it appears that subjects do not simply copy the feed-back provided by the experimenter, but instead use this information to learn to correct their misperceptions.

parameters $\widetilde{\lambda}_b$ are not significantly different from zero. With time however, subjects in treatment 2 learn to make substantially higher predictions than subjects in treatment 1, since most of the parameters $\widetilde{\delta}_b^5$ and $\widetilde{\delta}_b^{10}$ are significantly greater than zero.

Let us now turn to the bidding behavior of subjects in treatment 2. Figure 4 indicates that on average subjects still bid above the RNBNE, although overbidding is slightly less prominent than in treatment 1. Figure 4 however, only tells one part of the story. Indeed, Figure 6 shows a dramatic reduction of overbidding over time. Although the bids submitted during the first three periods are comparable in treatments 1 and 2, subjects in treatment 2 learned to reduce drastically their bids. In fact, during the last three periods of treatment 2 only the bids submitted for high private-values exceed slightly (on average) the RNBNE. For instance, for the highest private-value $v = 11$, subjects submitted instead of the RNBNE $b = 5$, an average bid of 5.29 during the last three periods of treatment 2, compared to 6.57 in treatment 1. This treatment effect is confirmed by a maximum likelihood estimation similar to the one conducted in Section 4.1 where the expected bids are now defined as

$$E\left[B_{i,v,t}\right] = \left(\lambda_v + \widetilde{\lambda}_v T_2\right)\left[1 + \left(\delta_v^5 + \widetilde{\delta}_v^5 T_2\right)P_{>5} + \left(\delta_v^{10} + \widetilde{\delta}_v^{10} T_2\right)P_{>10} + \eta_i\right] \quad, \tag{4.6}$$

the standard deviations are specified as in (4.4), and the exogenous variables are defined as in (4.5). Table 4 corroborates the fact that subjects in treatment 2 learned to make lower bids than subjects in treatment 1, since most of the parameters $\widetilde{\delta}_v^5$ and $\widetilde{\delta}_v^{10}$ are significantly smaller than zero.

The last two columns of Table 1 also indicate that subjects in treatment 2 adjust their strategy at a slower, but more constant pace than their predictions. Indeed, we have just seen that subjects in treatment 2 essentially adjust their predictions within the first five periods (the average percentage increase in the predictions is 20.8% within the first five periods and an additional 6.2% during the last 5 periods). In contrast, the magnitude of the strategy adjustment is basically comparable after 5 and 10 periods (the average percentage decrease in the bids submitted is 8.6% within the first five periods and an additional 6.8% during the last 5 periods). In other words, subjects tend to first correct their misperceptions, and then adjust their bidding behavior accordingly. The asymmetry between the prediction and strategy learning speeds may explain why, in contrast with treatment 1, subjects do not appear to best-respond to their beliefs in treatment 2. Indeed, Figure 4 indicates that on average subjects in treatment

2 bid above their risk neutral subjective best-response. Note however, that after the adjustment period (i.e. by the end of the session) subjects actions, subjective and objective risk neutral best-responses become consistent, since perceptions are almost unbiased, and behavior nearly conforms with the RNBNE.

Finally, one of the most striking difference between treatments 1 and 2 may be found in Table 3. Indeed, we can see that during the last three periods of treatment 2, subjects submitted much more frequently the RNBNE bids for each possible private-value. For instance, subjects submitted the RNBNE bid $b = 3$ for a private-value $v = 6$, 88.33% of the time in treatment 2, versus only 13.33% of the time in treatment 1. Even more remarkable, subjects precisely complied with the RNBNE bid function 41.67% of the time in the last three periods of treatment 2, versus 4.17% of the time in treatment 1.

To summarize, subjects overbid and underestimate their probability of winning in treatment 1. From the feed-back provided in treatment 2, subjects first learn to correct their misperceptions, and then nearly eliminate their tendency to overbid. This treatment effect is not as surprising as it may first appear. Indeed, it seems natural that the feed-back and the financial stimulus lead subjects to correct their misperceptions. Then, we can reasonably expect a subject to lower his bids, once he realizes that his probability of winning with any given bid is higher than he previously believed. For instance, when a subject realizes that he is virtually guaranteed to win the auction with a bid of $b = 5$ or $b = 6$, he should be very unlikely to submit a higher bid.

5. Noisy Behavioral Models

Although the experimental outcomes just presented suggest that risk aversion may play a lesser role than previously believed, we can only confirm this conjecture statistically by estimating a behavioral model. We adopt the Quantal Response Equilibrium concept (QRE hereafter) developed by McKelvey and Palfrey (1995) to model "noisy" decision making. This equilibrium concept has been shown to be powerful to organize behavior in numerous experimental settings.[25] The QRE approach is based on two key principles: behavior is random, and the probability to choose an action increases with the expected utility this action may yield. In other words, agents are not expected to select systematically their best-responses, but they play these strategies with a higher frequency. A Nash-like condition on

---

[25]See e.g. McKelvey and Palfrey (1995, 1998), Yi (2001), Capra, Goeree, Gomez, and Holt (2002), Anderson, Goeree and Holt (2002), as well as GHP.

the consistency of actions and beliefs is then imposed to determine the QRE choice probabilities. We present in the following two versions of the model adapted to our auction experiment. The first version assumes that agents are homogenous, while the second takes into consideration the possibility that agents may be heterogenous with respect to risk preference and probability perception.

5.1. Quantal Response Models with Homogenous Agents

Consider the power function probabilistic choice rule adopted by GHP. Under the assumption that weakly dominated strategies (e.g. bidding above one's own private-value) are excluded, the probability that an agent i with a private-value v selects a bid $b \in \{0, 1, ..., v\}$ when facing an opponent j may be written

$$P_i\left(b \mid v\right) = \frac{\left\{E\left[U\left(b, v\right) \mid P_j\right]\right\}^{1/\mu}}{\sum_{\widetilde{b}=0}^{v-1}\left\{E\left[U\left(\widetilde{b}, v\right) \mid P_j\right]\right\}^{1/\mu}} \quad , \tag{5.1}$$

where $U\left(.\right)$ denotes the individual indirect utility function, P_j is the vector of choice probabilities selected by agent i's opponent, and $\mu > 0$ is a "noise parameter" reflecting the sensitivity of the choice probabilities to expected utilities. A large μ yields essentially random behavior, while μ close to zero implies Nash-like behavior since best-response strategies are chosen with a probability close to one.

Following GHP, we assume that subjects are homogenous and exhibit constant relative risk aversion:
$$U\left(b, v\right) = \frac{\left(v - b\right)^{1-r}}{1 - r} \quad ,$$

where $r \in [0, 1[$ is the Arrow-Pratt coefficient of relative risk aversion common to all subjects.[26] In this model the actual probability that bidder i wins the auction with a bid b against an opponent j may be written

$$P_i^W\left(b\right) = \frac{1}{6}\sum_{v\in\{0,2,4,6,8,11\}}\sum_{\widetilde{b}<b}P_j\left(\widetilde{b} \mid v\right) + \frac{1}{12}\sum_{v\in\{0,2,4,6,8,11\}}P_j\left(b \mid v\right) \quad ,$$

---

[26]The GHP discrete auction model does not have a bayesian Nash equilibrium in pure strategy for every possible value of the risk aversion parameter r. Therefore, it will not be possible to estimate a bayesian Nash equilibrium model under risk aversion, as a possible alternative to the QRE.

where the first term represents the probability that agent j bids below b when receiving one of the six possible private-values, and the second term represents a favorable coin flip in the event of a tie.

Following GHP, we also assume that subjects misperceive their probability of winning the auction homogeneously according to the PWF proposed by Prelec (1998):

$$\Phi\left(P_i^W\left(b\right)\right) = \exp\left(-\beta\left(-\ln\left(P_i^W\left(b\right)\right)\right)^\alpha\right) \quad , \tag{5.2}$$

where $\alpha > 0$ and $\beta > 0$ are parameters to be estimated. Note that this PWF can display the typical S-shape pattern for specific values of (α, β), and agents have perfect perceptions when $(\alpha, \beta) = (1, 1)$.

Agent i expected utility when bidding b for a private-value v is then defined as

$$E\left[U\left(b, v\right) \mid P_j\right] = U\left(b, v\right)\Phi\left(P_i^W\left(b\right)\right) \quad .$$

The symmetric QRE choice probabilities are then such that

$$P_i\left(b \mid v\right) = P_j\left(b \mid v\right) = P^*\left(b \mid v\right) \quad \forall v \in \{0, 2, 4, 6, 8, 11\}, \ \forall b \leq v.$$

Since the private-values and bids are discrete and finite, we can find the QRE choice probabilities for any combination of (μ, r, α, β) by replacing P_i and P_j by P^* in (5.1), and solving numerically the resulting fixed-point problem. The different structural parameters may then be estimated by standard maximum likelihood techniques, in which a subject's actions are compared with their corresponding QRE probabilities of choice:

$$L\left(\mu, r, \alpha, \beta\right) = \prod_{i,v,t,b} \left[P^*\left(b \mid v\right)\right]^{\mathbb{I}\left[B_{i,v,t}=b\right]} \quad ,$$

where $B_{i,v,t}$ is the bid actually submitted by bidder i for the private-value v at period t, and $\mathbb{I}_{[B_{i,v,t}=b]}$ is the indicator function verifying $\mathbb{I}_{[B_{i,v,t}=b]} = 1$ when $B_{i,v,t} = b$, and $\mathbb{I}_{[B_{i,v,t}=b]} = 0$ otherwise.

To compare our results with those of GHP, we start by estimating the QRE model using only the data collected in the auction phase. In other words, we ignore for the moment the predictions made by subjects about their probability of winning. Table 6 reports the estimates of the three models considered by GHP. Namely, we estimate a QRE model with risk aversion and perfect perception, a QRE model with risk neutrality and perfect perception, and a QRE model with probability misperception and risk neutrality. As noted in the introduction, the

risk aversion parameter r cannot be identified separately from the parameters of the PWF (α, β) when one relies only on the auction data.[27] At this point, this prevents us from estimating a more general QRE model including simultaneously probability misperception and risk aversion. Such a model will be estimated next when we utilize the subjects' predictions in addition to the auction data.

Let us first concentrate on the estimation results for treatment 1. Table 6 indicates that the estimation of the QRE model with risk aversion is consistent with GHP's results. Indeed, we estimate the noise and risk aversion parameters to be $(\widehat{\mu}, \widehat{r}) = (0.067, 0.611)$, versus $(0.08, 0.55)$ in GHP. Our results are also compatible with recent estimates of the relative risk aversion parameter in private-value auctions.[28] According with GHP, we find that a QRE model under risk aversion generates a considerably higher log-likelihood (-2,856 versus -5,017 under risk neutrality), thereby clearly rejecting the risk neutrality hypothesis.[29] Finally, the estimation of a risk neutral QRE model with probability misperception yields a PWF of the form $\Phi(P) = P^{2.45}$, closely matching the quadratic PWF obtained by GHP. The similitude between the estimates in Table 6 and in GHP reinforces the conjecture that implementing the auction in strategy form, and asking subjects to make predictions did not introduce any significant treatment effect. It also confirms the ability of the QRE to explain behavior with a parsimonious model in two different experimental analyses.

The estimated results in treatment 2 vary slightly compared to treatment 1, but remain within the same order of magnitude. For instance, the noise parameter in the risk averse QRE model increases slightly in treatment 2 (0.117), while the risk aversion parameter becomes lower (0.497). These differences may be partially explained by the dual adjustment process taking place in treatment 2, as subjects first learned to correct their predictions, and then to adjust their strategies.

We now turn to the estimation of a QRE model combining probability misperception and risk aversion. The econometric analysis proceeds in two steps. In step 1, we use the probabilities of winning predicted by the subjects during the experiment to estimate by maximum likelihood the PWF parameters (α, β) in the

---

[27]Indeed, observe that the two sets of parameters (μ, r, α, β) and $(k\mu, k(r-1)+1, \alpha, k\beta)$, where $k > 0$, yield the same choice probability in (5.1).

[28]See e.g. Cox and Oaxaca (1996), Chen and Plott (1998), Campo, Perrigne and Vuong (2000), as well as Pezanis-Christou and Romeu (2003) who structurally re-estimate the risk aversion parameters from several published experimental studies.

[29]A formal likelihood ratio test for nested hypotheses yields a test statistics of 4,321.653, much larger than 3.841, the asymptotic critical value corresponding to the significance level 0.05. In fact, any test of the risk neutrality hypothesis in Tables 6 to 10 yields P-values inferior to 10^{-50}. Therefore, these P-values and the corresponding test statistics will not be reported in the text.

model

$$\widehat{P}_{i,b,t} = \exp\left(-\beta\left(-\ln\left(P_{i,t}^{W}(b)\right)\right)^{\alpha}\right) + \epsilon_{i,b,t} \quad , \tag{5.3}$$

where $\widehat{P}_{i,b,t}$ is the prediction of player i in period t regarding his probability of winning the auction with a bid equal to b, $P_{i,t}^{W}(b)$ is the actual probability that bidder i wins the auction given the bids submitted by the other subjects in round t, and $\epsilon_{i,b,t}$ is a normal error term with variance σ^2. In step 2, using the PWF estimated in step 1, we estimate the structural parameters (μ, r) in a QRE model with probability distortion. As previously shown however, subjects in treatment 2 learned to correct their misperceptions over time.[30] Therefore, the shape of the PWF is likely to differ significantly between the first and last periods. To account for this possibility, the PWF has been successively estimated in step 1 with data collected during the entire experiment (i.e. 15 periods), and with the data collected during the last five periods only. The estimation results in step 1 are presented in Table 7. The parameters (α, β) are significantly lower than $(1, 2)$, which implies that, although of similar shape, the PWF is different from the one suggested by GHP. As expected, the PWF estimated with the last five periods of treatment 2 is flatter, reflecting the fact that subjects make nearly unbiased estimates. The estimated parameters however, are significantly different from $(\alpha, \beta) = (1, 1)$ which corresponds to the absence of probability distortion. In other words, although nearly unbiased, subjects' predictions in the last five periods of treatment 2 remain imperfect.

The parameters estimated in step 2 are summarized in Table 8.[31] The first element to note is that the log-likelihoods of the QRE models with risk aversion and risk neutrality are smaller in Table 8 than in Table 6. To test formally which of these models prevail we adopt the specification test proposed by Vuong (1989) for non-nested hypotheses. A test therefore consists in comparing the specification of one of the models estimated with the data collected during all 15 periods in Table 8, against the specification of its analog in Table 6. For the risk aversion models in treatments 1 and 2, the adjusted test statistics are respectively 3.929

---

[30]Although the estimated values of (α, β) may change significantly between the early and late periods, the distribution of $\epsilon_{i,b,t}$ appears to remain roughly constant over time. In particular, a model in which the variance is specified as in (4.2) is rejected in favor of the homoscedastic specification proposed in (5.3) on the basis of a likelihood ratio test (the P-value is 0.256 in treatment 1 and 0.146 in treatment 2). The data, however, provide evidence of heterogeneity across subjects which we will account for in the next sub-section.

[31]Note that the standard deviations in table 8, as well as the test statistics, are robust to the errors generated by the estimation of the PWF in step 1.

and 4.260, which correspond to P-values of $1.36E-5$ and $6.83E-6$.[32] For the risk neutral models in treatments 1 and 2, the adjusted test statistics are respectively 18.322 and 6.057, which correspond to P-values of $2.39E-32$ and $6.71E-10$. Therefore, one can clearly reject at the usual significance level the pure QRE specifications in Table 6, in favor of the corresponding QRE models with probability misperception in Table 8. In other words, taking into consideration the subjects' predictions improves significantly the fit of the QRE model. Next, observe that the estimations of the risk aversion parameter decrease sharply compared to Table 6. For instance, the risk aversion parameter estimated with the entire sample collected in treatment 1 drops from 0.611 to 0.375 when accounting for the stated predictions. Even more remarkably, this parameter is estimated at 0.234 during the last five periods of treatment 2. This does not imply, however, that agents behaved as if risk neutral. Indeed, a comparison of the log-likelihoods in Table 8 indicates that the data still strongly reject risk neutrality. In contrast with the risk aversion coefficients, the estimated "noise parameter" μ is comparable in Tables 6 and 8, except for the last five periods of treatment 2 where it becomes closer to zero. This result reflects the fact that, as previously mentioned, subjects behavior nearly conforms with the RNBNE during the last periods of treatment 2.

Finally, we relax the assumption of equilibrium behavior to estimate the risk aversion parameter. Instead, inspired by the results in section 4, we assume that subjects best-respond to their stated beliefs about their probability of winning. In other words, we compare the subjects' actions in the experiment, with their subjective best-responses. To account for noisy behavior, we consider a Quantal Best-Response approach in which subjects do not necessarily select their best-responses with probability one. Instead, according with the basic principles underlying the QRE approach, we assume that the probability that subject i selects a bid $b = 0, ..., v$ when receiving a private-value v at period t is

$$P_{i,t}\left(b \mid v\right) = \frac{\left\{\widehat{E}_{i,t}\left[U\left(b,v\right)\right]\right\}^{1/\mu}}{\sum\limits_{\widetilde{b}=0}^{v-1}\left\{\widehat{E}_{i,t}\left[U\left(\widetilde{b},v\right)\right]\right\}^{1/\mu}} \quad \text{where} \quad \widehat{E}_{i,t}\left[U\left(b,v\right)\right] = \frac{\left(v-b\right)^{1-r}}{1-r}\widehat{P}_{i,b,t} \ ,$$

where $\widehat{P}_{i,b,t}$ is subject i's stated probability of winning with a bid b in period

---

[32]Recall that the P-values are estimated by Bootstrap, and therefore they differ slightly from the asymptotic P-values.

22

t. Recall that we observe $\widehat{P}_{i,b,t}$ for each agent and each period. Therefore, the only parameters to estimate in this model are (μ, r). Table 9 indicates that the Quantal Best-Response model generates significantly lower log-likelihoods than the QRE models in Tables 6 and 8. In fact, pairwise comparisons on the basis of the Vuong specification test systematically lead to reject the models in Tables 6 and 8 in favor of the corresponding models in Table 9 (the P-values range from $1.06E - 3$ to $8.39E - 128$). These tests therefore confirm that subjects' behavior may be best explained by the fact that they tend to best-respond to their beliefs. The likelihoods reported in Table 9 also shows that the risk neutrality hypothesis is still strongly rejected in favor of risk aversion. Moreover, note that, unlike the models estimated in Tables 6 and 8, the parameters $(\widehat{\mu}, \widehat{r})$ estimated in treatment 1, and in the last five periods of treatment 2, are now quite similar. This result further confirms the ability of the Quantal Best-Response approach to capture behavior equally well in the two different treatments. The difference observed when estimating the risk aversion parameter with the entire sample of data collected in treatment 2 may again be attributed to the fact that subjects in treatment 2 learned to adjust their strategies and predictions at a slightly different rate. Finally, observe that the risk aversion parameter is significantly lower than previously estimated. Indeed, we now find that r lays between 0.24 and 0.27, instead of the 0.5 to 0.6 range suggested by Table 6, as well as GHP and other studies.

5.2. Quantal Response Models with Heterogenous Agents

To explain overbidding, the leading risk aversion model, the so-called CRRAM (see e.g. Cox et al. 1988), assumes that agents may have heterogenous preferences. To give the risk aversion hypothesis its best chance to organize the data collected in the experiment, we explore in this section the possibility that agents may be heterogenous with respect to risk preference and/or probability misperception. This approach is also consistent with Section 4 where we detected significant heterogeneity across subjects in the way they bid and predict probabilities.

Following Cox et al. (1988), we assume that heterogenous risk preferences may be represented by a utility function of the form:

$$U_i\left(b, v\right) = \frac{\left(v - b\right)^{1-r_i}}{1 - r_i} \quad .$$

In addition, to model heterogenous probability perception, we generalize the PWF

proposed by Prelec (1998):

$$\Phi_i\left(P\right) = \exp\left(-\beta_i\left(-\ln\left(P\right)\right)^{\alpha_i}\right) \quad . \tag{5.4}$$

Following Cox et al. (1988), we assume that an agent faces uncertainty about the risk preferences and probability perceptions of others. In other words, agent i observes $\theta_i = (r_i, \alpha_i, \beta_i)$, but he only knows the distribution from which the specific risk aversion and the probability perception parameters of his opponents are generated. Since (α_i, β_i) must be strictly positive, we assume that the vectors $(r_i, \ln(\alpha_i), \ln(\beta_i))$ are identically and independently distributed across agents from a multi-normal distribution with mean $\left(r, \widetilde{\alpha}, \widetilde{\beta}\right)$, and variance-covariance matrix Σ.[33]

The analysis of the model is somewhat similar to the homogenous case, except that we must now distinguish the choice probability of agent i conditional on θ_i,

$$P_i\left(b \mid v, \theta_i\right) = \frac{\left\{E\left[U_i\left(b, v\right) \mid P_j, \theta_i\right]\right\}^{1/\mu}}{\sum\limits_{\widetilde{b}}^{v-1}\left\{E\left[U_i\left(\widetilde{b}, v\right) \mid P_j, \theta_i\right]\right\}^{1/\mu}} \quad , \tag{5.5}$$

from the unconditional (or average) choice probability of agent i, $P_{\theta,i}\left(b \mid v\right) = E_\theta\left[P_i\left(b \mid v, \theta_i\right)\right]$.

Agent i's probability of winning and expected utility may now be written

$$P_i^W\left(b\right) = \frac{1}{6}\sum_{v \in \{0,2,4,6,8,11\}}\sum_{\widetilde{b} < b}P_{\theta,j}\left(\widetilde{b} \mid v\right) + \frac{1}{12}\sum_{v \in \{0,2,4,6,8,11\}}P_{\theta,j}\left(b \mid v\right) \quad ,$$

$$\text{and} \quad E\left[U_i\left(b, v\right) \mid P_j, \theta_i\right] = U_i\left(b, v\right)\Phi_i\left(P_i^W\left(b\right)\right) \quad .$$

Although agents are endowed with different θ_i, the model is ex-ante symmetric, and therefore, the QRE choice probabilities $P^*\left(b \mid v, \theta\right)$ are also symmetric across agents. Note however, that unlike the homogenous case, the QRE choice probabilities are now a function of a multi-dimensional continuous variable θ. Consequently, unless one imposes arbitrary parametric restrictions on the shape

---

[33]Although the distribution of the parameters has been specified in terms of $(\ln(\alpha_i), \ln(\beta_i))$, we will report in the remainder the descriptive statistics for (α_i, β_i). In particular, (α, β) will denote the expectation of (α_i, β_i), σ_α^2 and σ_β^2 will denote the variance of α_i and β_i, and $\rho_{\alpha,\beta}$ will denote the coefficient of correlation between α_i and β_i.

of the choice probabilities as a function of θ, the heterogenous QRE model cannot be solved even numerically. In contrast, the set of unconditional QRE choice probabilities $P_\theta^* (b \mid v)$ is discrete and finite, and it may be determined with the same fixed-point approach as in the homogenous case. Indeed, observe that

$$P_{\theta,i} (b \mid v) = E_\theta \left[\frac{\{E\left[U_i\left(b,v\right) \mid P_j, \theta_i\right]\}^{1/\mu}}{\sum_{\widetilde{b}}^{v-1} \left\{ E\left[U_i\left(\widetilde{b},v\right) \mid P_j, \theta_i\right]\right\}^{1/\mu}} \right] , \qquad (5.6)$$

only involves the unconditional choice probabilities $P_{\theta,i}$ and $P_{\theta,j}$, which in equilibrium are both equal to $P_\theta^* (b \mid v)$.[34] The Likelihood associated with the heterogenous model may then be written:

$$L\left(\mu, r, \alpha, \beta, \Sigma\right) = \prod_{i,v,t,b} \left[P_\theta^* (b \mid v)\right]^{\mathbb{I}\left[B_{i,v,t}=b\right]} .$$

The heterogenous analogs to the various homogenous models estimated in Tables 6 to 9, are reported in table 10. Four points are particularly worth noting.

First, the average values of the noise, risk aversion and probability misperception parameters (μ, r, α, β), change slightly compared to the homogenous case (Tables 6 to 9), but they remain of comparable magnitude. Observe however, that the risk aversion parameter is in general lower in Table 10. In particular, the coefficient of relative risk aversion is found to be slightly below 0.2 when the Quantal Best-Response model accounting for the subjects' stated beliefs is estimated with the data collected during the last five periods of treatments 1 and 2.

Second, most of the standard deviation parameters $(\sigma_r, \sigma_\alpha, \sigma_\beta)$ are significantly larger than zero, and relatively substantial compared to their respective means. In other words, our estimates suggest significant heterogeneity across subjects both in terms of risk aversion and probability perception. The first part of this result is consistent with a number of experimental analyses of first-price auctions, in which heterogeneity in risk preference is often identified (see e.g. Cox et al. 1988, Cox and Oaxaca 1996, or Chen and Plott 1998). To the best of our knowledge,

---

[34]The expectation in (5.6) does not have a closed form, and it must be evaluated numerically. To speed-up the numerical integration we rely on a quasi Monte Carlo sampling method consisting in generating extensible lattice points modified by the "baker's transformation" (see Hikernell et al. 2000).

however, the presence of heterogeneity in probability perception has never been previously detected in experimental economics.

Third, there appears to be a strong correlation between risk preference and probability misperception. Indeed, $\rho_{r,\beta}$, the coefficient of correlation between an individual risk aversion parameter r_i and his probability perception parameter β_i, is significantly greater than zero. Therefore, although observationally equivalent, probability distortion and risk aversion do not appear to act as substitutes for an individual. Instead, our estimation results suggest that they complement each other, as highly risk averse subjects also appear to be more pessimistic and/or less accurate when predicting probabilities.[35] This result is somewhat consistent with Bellemare et al. (2005) who also identify a correlation between preferences and beliefs.[36]

Four, a series of statistical tests indicate that i) models accounting for heterogeneity systematically dominate their homogenous analogs in Tables 6 to 9; ii) within the class of heterogenous models, the risk neutrality hypothesis is rejected in favor of risk aversion; iii) the experimental data are best described by the heterogenous Quantal Best-Response model in which agents play with a probability close to one their best-response conditional on their stated beliefs.

To summarize, the structural estimations suggest that subjects in our experiment may be best characterized as utility maximizers with heterogenous levels of risk aversion and probability misperception. Moreover, the risk aversion parameters estimated, although non-negligible, are far too low to explain alone the subjects' tendency to bid above the RNBNE. Instead, the structural estimations confirm that probability misperception is one of the main determinant of overbidding in our experiment.

6. Probability Misperceptions in a Private Values Auction Model

Our experimental results suggest that probability misperception may complement risk aversion to explain overbidding in an independent private values auction. The theoretical effect of risk-aversion on the bayesian Nash equilibrium bid function

---

[35]Given the estimated values of (α, β), β_i essentially controls the degree of convexity of the PWF. Therefore, an individual with a high level of risk aversion r_i is more likely to have a large β_i which corresponds to a more severely convex PWF.

[36]In an analysis of the ultimatum game, Bellemare et al. (2005) find that proposers who are optimistic about the acceptance rates of responders also tend to have significantly higher levels of inequity aversion.

is well-known. Milgrom and Weber (1982) showed that the bid function is higher when bidders are risk-averse compared to risk-neutral. To the best of our knowledge, no equivalent result exists on the effect of probability misperception. The objective of this section is to examine whether the class of PWF identified in our experiment may lead to overbidding in a standard private values auction. To simplify, we only consider a continuous auction model with homogenous agents.

Let us consider the standard first-price independent private-values auction model. There are N agents with identical Von Neuman-Morgenstern utility function $u(.)$. They participate in an auction where they each submit a sealed bid for an indivisible object. Agent $i = 1, ..., N$ has a private-value v_i for the object. This private-value is drawn independently from a distribution with cumulative $F(.)$, density $f(.)$ and support $[\underline{v}, \overline{v}]$. The highest bidder gets the object. His payoffs is equal to his own valuation of the object minus his bid, $v_i - b_i$. The other bidders receive no payoffs.

Let us denote $B(.)$ the symmetric bayesian Nash equilibrium of this game. This strategy is a solution to the following optimization and fixed point problems,

$$B(v_i) = \underset{b_i}{ArgMax} \left\{ \begin{array}{c} G[b_i > B(v_j), \forall j \neq i]u(w + v_i - b_i) + \\ (1 - G[b_i > B(v_j), \forall j \neq i])u(w) \end{array} \right\} \quad \begin{array}{l} \forall v_i \in [\underline{v}, \overline{v}] \text{ and} \\ \forall i = 1, ..., N \end{array},$$

(6.1)

where $G[.]$ is the subjective probability of winning the auction and w is initial wealth. In other words, $B(v_i)$ is bidder i's best-reply when the other bidders' select the equilibrium strategy $B(.)$, given that all bidders have the same preferences and beliefs.

Although $G[.]$ is assumed to be common knowledge, we do not require bidders to perceive accurately winning probabilities. Instead, we assume that they all distort their probability of winning according to a continuous and increasing PWF $\Phi(.)$ verifying $\Phi(0) = 0$ and $\Phi(1) = 1$ such that

$$G[.] = \Phi(P[.]),$$

and where $P[.]$ is the objective probability of winning.

Let us introduce the indirect utility function U as

$$U(x) = u(w + x) - u(w). \tag{6.2}$$

Suppose then that the equilibrium strategy $B(.)$ is a monotonically increasing function (see Maskin and Riley, 2000). Therefore, if $B^{-1}(.)$ stands for the inverse

of $B(.)$, and maximizing (6.1) is equivalent to maximizing over b_i

$$\Phi(F(B^{-1}(b_i))^{N-1})U(v_i - b_i).$$

Differentiating with respect to b_i yields

$$\Phi'(F(B^{-1}(b_i))^{N-1})(N-1)F(B^{-1}(b_i))^{N-2}\frac{f(B^{-1}(b_i))}{B'(B^{-1}(b_i))}U(v_i - b_i)$$

$$-\Phi(F(B^{-1}(b_i))^{N-1})U'(v_i - b_i).$$

Setting this expression equal to zero gives the first order condition

$$B'(v_i) = (N-1)f(v_i)\frac{\Phi'(F(v_i)^{N-1})F(v_i)^{N-2}}{\Phi(F(v_i)^{N-1})}\frac{U(v_i - B(v_i))}{U'(v_i - B(v_i))} \ \forall v_i \in [\underline{v}, \overline{v}] \ . \quad (6.3)$$

Together with the boundary condition $B(\underline{v}) = \underline{v}$, this differential equation characterizes the bayesian symmetric Nash equilibrium bidding behavior in an auction where bidders have a PWF $\Phi(.)$.

Let us now introduce a class of PWF that includes as a special case the function identified in GHP's and in our experiment. This is the class of "star-shaped" PWF.

Definition 6.1. *Let a PWF $\Phi(p)$ with $\Phi(0) = 0$ and $\Phi(1) = 1$; then $\Phi(p)$ is said to be star-shaped if $\Phi(p)/p$ is increasing in p.*

The term star-shaped is taken from Landsberger and Meilijson (1990), and relates to their definition of a star-shaped *utility* function.[37] We have plotted on Figure 7 a typical star-shaped PWF. Observe that the class of star-shaped PWF provides a formal generalization of the PWF identified in our experiment (see Figure 1 and 2). A star-shaped PWF means that the chord to the PWF drawn from 0 to p must lay above $\Phi(p)$ for any p. This is equivalent to assuming that the slope of the chord is lower than the slope of the tangent to $\Phi(p)$ at p, namely $\Phi(p)/p \leq \Phi'(p)$. It is also immediate that any convex PWF is star-shaped. This is of interest since the convexity of Φ has been shown to be equivalent to "risk-aversion" in Rank Dependent models in the sense that a risk-averse agent must decline any mean-preserving spread (Yaari 1986, Chew, Karni and Safra 1987).

---

[37]The exact term, as it is defined in Landsberger and Meilijson (1990), and more generally in Mathematics, would be that Φ is star-shaped at zero.

Furthermore, observe that Φ star-shaped implies underestimation everywhere, i.e. $\Phi(p) \leq p$.[38]

We now show that a star-shaped PWF leads to an unambiguous comparative statics analysis. Let us compare $B(x)$, the symmetric equilibrium bidding strategy under probability distortion as given by the general condition (6.3), to the corresponding equilibrium condition under perfect perception denoted $B_0(x)$. Assume that Φ is star-shaped, or equivalently that $\Phi'(p) \geq \Phi(p)/p$ for any p. Then, we get

$$
\begin{aligned}
& B'(v) - B_0'(v) \\
=\ & (N-1)f(v)F(v)^{N-2}\Big[\frac{\Phi'(F(v)^{N-1})}{\Phi(F(v)^{N-1})}\frac{U(v-B(v))}{U'(v-B(v))} - \frac{1}{F(v)^{N-1}}\frac{U(v-B_0(v))}{U'(v-B_0(v))}\Big] \\
\geq\ & (N-1)f(v)F(v)^{N-2}\Big[\frac{U(v-B(v))}{U'(v-B(v))} - \frac{U(v-B_0(v))}{U'(v-B_0(v))}\Big],
\end{aligned}
$$

by assumption and since U, U', Φ and Φ' are positive. From the last inequality, we have that, for any v,

$$B(v) = B_0(v) \text{ implies } B'(v) \geq B_0'(v).$$

We thus have a single crossing property. This property means that B can only cross B_0 from below. Since $B(\underline{v}) = B_0(\underline{v}) = \underline{v}$, the function $B(v)$ will always be larger than $B_0(v)$ for any v such that $v \geq \underline{v}$. Therefore, agents increase their bids when the PWF Φ is star-shaped.

Proposition 6.2. *Suppose that Φ is star-shaped. Then the PWF Φ leads agents to overbid compared to the case of perfect perception.*

Proposition (6.2) implies that a star-shaped PWF could explain by itself the overbidding commonly observed in private-values auctions. An alternative interpretation of Proposition (6.2), is that explaining overbidding in the presence of a star-shaped PWF may required a lower level of risk aversion than previously believed. The latter interpretation is relevant for two reasons: first, it is consistent with the conclusions of our experimental analysis; second, it has been argued that the levels of risk aversion estimated to explain overbidding in experimental

---

[38]Suppose it is not the case. Then there exists p_0 such that $\Phi(p_0) > p_0$. As a result, $\Phi(p_0)/p_0 > 1 = \Phi(1)/1$, which contradicts Φ star-shaped.

auctions may be considered unreasonably high (see e.g. Harrison 1990, or Rabin 2000).[39]

To illustrate Proposition (6.2), consider the special case of power functions both for the utility and PWF, namely $u(w) = w^{1-r}/(1-r)$ with a positive constant relative risk aversion parameter r, and $\Phi(p) = p^\beta$ with $\beta \geq 1$. Observe that Φ is star-shaped since it is convex, and that it basically corresponds to the type of PWF identified both in GHP's and our experiments. As it is usual in auction models, assume a zero initial wealth so that $U(x) = x^{1-r}/(1-r)$ together with $r \in [0,1[$. Assume also a uniform distribution $F(v) = v$ over the support $[0,1]$. With these functional forms, we can get a closed-form solution to the differential equation (6.3), which simply reduces to the linear strategy

$$B(v) = \frac{\beta v}{1 + \beta - r},$$

for the two-bidder case. Consistent with Proposition (6.2), the equilibrium bid can increase indifferently with either β the curvature of the PWF, or r the risk-aversion coefficient.[40] Even this simple auction model makes it clear that there is an identification problem and that the source of overbidding may be either risk aversion or probability distortion.

Finally, we conclude this section by providing some intuition as to why a star-shaped PWF leads to overbidding. In fact, two underlying effects are at play. To see that, consider a simpler individual decision-making problem that parallels the one faced by participants in an auction.[41] Consider an agent who maximizes over b an objective function

$$\Phi(p(b))U(v - b), \tag{6.4}$$

where $p(b)$ is the probability of getting a prize of value v, and $\Phi(.)$ is a PWF. To simplify, we will assume that this program is concave in b.

---

[39]This point actually needs to be explained in the context of first-price auctions. Indeed subjects are not found to be extremely risk-averse in standard experiments on first-price auctions (usually r si estimated around 0.6). But recall that participants' wealth is usually assumed to be zero. Risk aversion estimates may arguably increase rapidly with wealth, up to reach unreasonable values. Risk-aversion estimates then have to be very close to zero (as we found) in order to be consistent with the introduction of sensible levels of wealth.

[40]It is easy to generalize Milgrom and Weber (1982)'s result to show that more risk-aversion (in the sense of Arrow-Pratt) always leads to increase the equilibrium bid function, even in the presence of probability distortion.

[41]Further intuitions on why the comparative statics analysis is often similar in individual decision-making problems and in corresponding strategic contexts are given in Gradstein, Nitzan and Slutsky (1992). See also Milgrom and Roberts (1994).

First, if $\Phi(p) = p$ the solution is simply given by b^* solving

$$p'(b^*)U(v - b^*) - p(b^*)U'(v - b^*) = 0. \tag{6.5}$$

It is optimal to equate the marginal benefit of increasing the probability of winning $p'(b^*)U(v-b^*)$ to the marginal cost of reducing the value of the winning $p(b^*)U'(v-b^*)$. Our objective is to compare the maximizer of (6.4) to b^*. It is easy to see that introducing a PWF leads to overbidding compared to the case of no misperceptions if and only if the slope of the tangent of $\Phi(p(b))U(v - b)$ at b^* is positive. This is equivalent to

$$\Phi'(p(b^*))p'(b^*)U(v - b^*) - \Phi(p(b^*))U'(v - b^*) \geq 0. \tag{6.6}$$

Using the first-order condition (6.5) we have $p'(b^*) = \frac{p(b^*)U'(v-b^*)}{U(v-b^*)}$ so that this latter inequality is also equivalent to $\Phi'(p(b^*))p(b^*) \geq \Phi(p(b^*))$, that is for any $p \in [0, 1]$

$$\Phi'(p) \geq \Phi(p)/p. \tag{6.7}$$

Hence, a star-shaped PWF leads to increase the value of b.

Let us now interpret this result. Observe that the result derives from the comparison of (6.5) to (6.6). These conditions represent the trade-off mentioned above between an increase in the probability of winning versus a decrease of the payoff contingent on winning. How does the probability distortion affect each of these two effects? First, there is the effect of the change in the probability of winning. Directly comparing the first terms of (6.5) and (6.6) shows that this effect is controlled by Φ' compared to 1. This means that when $\Phi'(p)$ is larger than 1, increasing b by one unit is *perceived* as relatively more profitable at the margin, so that the misperception in the probability leads to increase b. Second, there is the effect related to the decrease of the payoff. Comparing the second terms of (6.5) and (6.6) shows that this second effect is controlled by $\Phi(p)$ compared to p. When the probability is under-estimated $\Phi(p) < p$, this effect leads to reduce the *perceived* cost associated with a reduction of the payoff, so that this effect leads to increase b as well. Obviously, there does not exist any continuous probability distortion function with $\Phi(0) = 0$ and $\Phi(1) = 1$ such that $\Phi'(p) < 1$ together with $\Phi(p) < p$ for any $p \in [0, 1]$. Hence, there is no hope of finding a probability distortion for which both effects go in the same direction for any p. Yet, our results show that the aggregate effect, that simply condenses the two effects mentioned above, critically depends on condition (6.7), namely on whether the PWF Φ is star-shaped.

7. Conclusion

The objectives of this paper were i) to determine whether bidders in a private-values auction misperceive their probability of winning, and ii) to verify whether probability misperception can play a significant role in explaining overbidding. To identify beliefs separately from preferences, we followed Manski's (2004) suggestion, and we conducted a two treatments auction experiment in which subjects are asked, in addition to bidding, to state their beliefs about their probability of winning the auction. The experimental outcomes in treatment 1 indicate that subjects overbid, and underestimate their probability of winning. In treatment 2, we find that providing feed-back on the accuracy of their predictions, leads subjects to correct their predictions, and then to learn to curb-down drastically their tendency to overbid. The estimation of noisy behavioral models suggests that bidders best-respond to their stated beliefs about their probability of winning. In addition, the econometric analysis indicates that the risk aversion parameter drops significantly (from 0.6 to 0.2) when one accounts for probability distortion and heterogeneity across agents. We also identify an interesting correlation between beliefs and preferences, as we find that highly risk averse bidders are more likely to underestimate their probability of winning the auction. Finally, we show theoretically that our experimental findings are consistent with the predictions of a standard independent private-values auction model combining risk aversion and probability misperception. In summary, we find that bidders underestimate their probability of winning, which appears to be a main source of overbidding. In contrast, although still necessary to explain fully behavior, risk aversion is found to play a lesser role than previously believed to explain overbidding.

These results tend to support the views of Kagel and Roth (1992) who argue that risk aversion may be one of the determinant of overbidding in private-values auctions, but not necessarily the most important one. The present paper suggests that probability misperception may complement risk aversion to explain overbidding. Our analysis also appears to be consistent with Isaac and Walker (1985), Ockenfels and Selten (2004), and Neugebauer and Selten (2003) who find that overbidding is less prominent when subjects observe their opponents' highest bid at the end of each round. Indeed, subjects may have used this information to predict more precisely their probability of winning the auction. Finally, our estimations, although not totally immune from Rabin (2000)'s critique, appears more compatible with the low financial incentives provided in the experiment. Indeed, the econometric analysis yields significantly lower, and arguably more reasonable

levels of risk aversion than previous studies.

It has to be noted that the conclusions of this paper should be interpreted with caution, as we need to acknowledge the limitations of our experiment. Indeed, to infer the subjects' probability distortion, we had to conduct a discrete auction which does not necessarily possess a bayesian Nash equilibrium in pure strategy for any risk aversion coefficient. This prevented us from directly comparing the probability misperception hypothesis with the leading risk aversion model, the bayesian Nash equilibrium model with constant relative risk aversion (see e.g. Cox, Smith, and Walker 1988). Moreover, we have assumed throughout the paper that probability misperception was the only possible source of probability distortion. We have ignored in particular the probability weights which may reflect agents' preferences over probabilities when facing uncertain decisions (see e.g. Kahneman and Tversky 1979). It is unclear at this point whether such probability weights would complement or partially offset the effect of risk aversion in explaining overbidding. We found it non-trivial to modify our experiment to account for probability weights. The design of such an experiment remains an open question left for future research.

We also realize that the relevance of the probability misperception hypothesis will only be fully assessed by its ability to organize behavior beyond the private-values auction model. To do so, additional experiments with different auction formats will need to be conducted. However, we can already conjecture that, unlike risk aversion, probability misperception may be a valid candidate to explain the persistence of bidding above the dominant strategy in second-price auction, or the regular failures of the binary lottery procedure to induce risk neutral bidding. Note also that, unlike the "joy of winning" hypothesis, probability misperception is consistent with the absence of overbidding in English auctions. Again, further experimental analyses are needed to confirm these conjectures.

One may also question the usefulness of establishing whether overbidding is generated by probability misperception or risk aversion. Indeed, these two factors may be considered behaviorally equivalent, since they lead bidders to behave "*as if*" risk averse. Our experiment illustrates why it is of importance to distinguish between the two hypotheses. Indeed, risk aversion is usually assumed to be an intrinsic individual characteristic that is time and context independent. In contrast, our experiment indicates that probability misperception may be corrected, in which case individual behavior may be significantly affected. Therefore, as argued by Camerer (1995), the identification of systematic judgment errors may have important policy implications. For instance, a government seeking to limit

overbidding in its public auctions could provide feed-back to reduce perception biases. Alternatively, a government could select a specific auction format, such as the English auction, for which behavior appears to be less dependent on probability perceptions.

To conclude, one may comment on the general identification problem between beliefs and preferences. In practice, the usual way to address this problem is to impose some beliefs (e.g. rational or naive) on all agents, and then focus on the estimation of preferences. This procedure may be flawed, as an incorrect specification of the beliefs may lead to a misrepresentation of preferences. As discussed in Manski (2004), the combination of observed choices with data on beliefs should mitigate this identification problem and improve our ability to predict behavior. Our experiment illustrates how this approach may provide new insights into preferences. Indeed, our estimates of risk aversion based on elicited beliefs differ notably from previous studies relying solely on choice data. In addition, the elicitation of probability perception in our experiment allowed us to address additional questions such as the evolution of beliefs over time, the consequences of heterogeneity in beliefs, and the correlation that may exist between beliefs and preferences.

References

Ali M., 1977, "Probability and Utility Estimates for Racetrack Bettors," Journal of Political Economy, vol. 85(4), 803-15

Anderson S., Goeree J. and C. Holt, 2002, "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies," Southern Economic Journal, 69(1), 21-47.

Bellemare C., Kröger S., and A. van Soest, 2005, "Actions and Beliefs: Estimating Distribution-Based Preferences Using a Large Scale Experiment with Probability Questions on Expectations," mimeo, Tilburg University.

Benjamin D., Dougan W. and D. Buschena, 2001, "Individuals' Estimates of the Risks of Death: Part II–New Evidence," Journal of Risk and Uncertainty, 22, 35-57.

Camerer C., 1995, "Individual Decision Making," In: Kagel, J.H. & Roth, A.E. (eds.): Handbook of Experimental Economics, Princeton, NJ: Princeton University Press, 587-703.

Camerer C., Ho T. and J. Chong, 2001, "Behavioral Game Theory: Thinking, Learning, and Teaching," Paper presented at the Nobel Symposium.

Camerer C. and R. Hogarth, 1999, "The Effects of Financial Incentives in Economics Experiments: A Review and Capital-Labor-Production Framework," Journal of Risk

and Uncertainty, 7-42.

Campo S., Perrigne I. and Q. Vuong, 2001, "Asymmetry in First-Price Auctions with Affiliated Private Values," Journal of Applied Econometrics, Vol., 2, 179-207.

Capra M., Goeree J., Gomez R. and C. Holt, 2002, "Learning and Noisy Behavior in an Experimental Study of Imperfect Price Competition," International Economic Review, 43, no 3.

Charness G. and M. Dufwenberg, 2003, "Promises & Partnership," mimeo, University of California at Santa Barbara.

Chen KY. and C. Plott, 1998, "Nonlinear Behavior in Sealed Bid First-Price Auctions," Games and Economic Behavior, 25, 34-78.

Chew C.S., Karni E. and Z. Safra, 1987, "Risk-Aversion in the Theory of Expected Utility with Rank Dependent Probabilities," Journal of Economic Theory, 42, 370-381.

Costa-Gomes M. and G. Weizsäcker, 2004, "Stated Beliefs and Play in Normal-Form Games," mimeo, Harvard University.

Cox J. and R. Oaxaca, 1996, "Is Bidding Behavior Consistent with Bidding Theory for Private Value Auctions," in R. M. Isaac, ed., Research in Experimental Economics, Vol. 6, Greenwich, Conn.: JAI Press, 131-148.

Cox J., Smith V. and J. Walker, 1983, "Tests of a Heterogeneous Bidder's Theory of First-Price Auctions," Economics Letters, 12:3-4 207-212.

Cox J., Smith V. and J. Walker, 1985, "Experimental Development of Sealed-Bid Auction Theory; Calibrating Controls for Risk Aversion," American Economic Review, 75, 160-165.

Cox J., Smith V. and J. Walker, 1988, "Theory and Individual Behavior of First-Price Auctions," Journal of Risk and Uncertainty, 1:1, 61-99.

Croson R., 2000, "Thinking Like a Game Theorist: Factors Affecting the Frequency of Equilibrium Play," Journal of Economic Behavior and Organization, 41, 299-314.

De Finetti B., 1965, "Methods of Discriminating Levels of Partial Knowledge Concerning a Test Item," British Journal of Mathematical and Statistical Psychology, 18, 87-123.

Dorsey R. and L. Razzolini, 2003, "Explaining Overbidding in First Price Auctions using Controlled Lotteries," Experimental Economics, 6(2), 123-140.

Dufwenberg M. and U. Gneezy, 2000, "Measuring Beliefs in an Experimental Lost Wallet Game," Games and Economic Behavior 30, 163-182.

Fischhoff B., 1982, "Debiasing," In Judgment under Uncertainty: Heuristics and Biases, Kahneman, D., Slovic, P. & Tversky, A. (Eds.) London: Cambridge University Press.

Friedman D., 1992, "Theory and Misbehavior of First-Price Auctions: Comment,"

American Economic Review, 82:5, 1374-1378.

Gerking S., DeHaan M. and W. Schulze, 1988, "The Marginal Value of Job Safety: A Contingent Valuation Study," Journal of Risk and Uncertainty, 1, 18-99.

Giordani P. and P. Söderlind, 2003, "Is there Evidence of Pessimism and Doubt in Subjective Distributions? A Comment on Abel," mimeo, Stockholm School of Economics.

Goeree J., Holt C., and T. Palfrey, 2002, "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," Journal of Economic Theory, 104(1), 247-272 .

Golec J. and M. Tamarkin, 1998, "Bettors Love Skewness, Not Risk, at the Horse Track," Journal of Political Economy, 106 (1), 205-225.

Gradstein M., Nitzan S. and S. Slutsky, 1992, "The Effect of Uncertainty on Interactive Behavior," Economic Journal, 102, 554-61.

Harrison G., 1989, "Theory and Misbehavior of First-Price Auctions," American Economic Review, 79:4, 749-762.

Harrison G., 1990, "Risk Attitudes in First-Price Auction Experiments: A Bayesian Analysis," Review of Economics and Statistics, 72:3, 541-546.

Hickernell F.J., Hong H.S., L'Ecuyer P. and C. Lemieux, 2000, "Extensible Lattice Sequences for Quasi-Monte Carlo Quadrature", SIAM Journal on Scientific Computing, 22, 1117–1138.

Hogarth R., 1987, "Judgment and Choice: The Psychology of Decision," 2nd ed. New York:Wiley.

Hurley T. and J. Shogren, 2004, "A Note on Capturing Beliefs with Predictions," mimeo, University of Wyoming.

Isaac M. and J. Walker, 1985, "Information and Conspiracy in Sealed Bid Auctions," Journal of Economic Behavior and Organization, 6, 139-159.

Kagel J. and A. Roth, 1992, "Theory and Misbehavior in First-Price Auctions: Comment," American Economic Review, 82:5, 1379-1391.

Kahneman D. and A. Tversky, 1979, "Prospect Theory: an Analysis of Decision under Risk," Econometrica, 47, 263-291.

Landsberger M. and I. Meilijson, 1990, "Lotteries, Insurance and Star-shaped Utility Functions," Journal of Economic Theory, 52, 1-17.

Lichtenstein S., Slovic P., Fischhoff B., Layman B. and B. Combs, 1978, "Judged Frequency of Lethal Events," Journal of Experimental Psychology: Human Learning and Memory, 4, no. 6, 551-78.

Manski C. F., 2002, "Identification of Decision Rules in Experiments on Simple Games of Proposal and Response," European Economic Review, 46, 880–891.

Manski C. F., 2004, "Measuring Expectations," Econometrica, 72(5), 1329–1376.

Maskin E. and J. Riley, 2000, "Equilibrium in Sealed High Bid Auctions," Review of Economic Studies, 67, 439-54.

McKelvey R. and T. Palfrey, 1995, "Quantal Response Equilibria for Normal Form Games," Games and Economic Behavior, 10, 6-38.

McKelvey R. and T. Palfrey, 1998, "Quantal Response Equilibria in Extensive Form Games," Experimental Economics, 1(1), 9-41.

Milgrom P. and J. Roberts, 1994, "Comparing Equilibria," American Economic Review, 84, 441-459.

Milgrom P. and R. Weber, 1982, "A Theory of Auctions and Competitive Biddings," Econometrica, 50, 1089-1122.

Murphy A. and R. Winkler, 1970, "Scoring Rules in Probability Assessment and Evaluation," Acta Psychologica, 34, 273-386.

Murphy A. and R. Winkler, 1984, "Probability Forecasting in Meteorology," Journal of the American Statistical Association, 79, 489-500.

Neugebauer T. and R. Selten, 2003, "Individual Behavior of First-Price Sealed-Bid Auctions: The Importance of Information Feedback in Experimental Markets," mimeo University of Bonn.

Nyarko Y. and A. Schotter, 2002, "An Experimental Study of Belief Learning Using Elicited Beliefs," Econometrica, 70, 971-1005.

Ockenfels A. and R. Selten, 2004, "Impulse Balance Equilibrium and Feedback in first-price Auctions," Games and Economic Behavior, forthcoming.

Offerman T., 2002, "Hurting Hurts More Than Helping Helps," European Economic Review, 46, 1423-1437.

Offerman T., Sonnemans J. and A. Schram, 1996, "Value Orientation, Expectations and Voluntary Contributions in Public Goods", Economic Journal, 106, 817-45.

Ottaviani M. and P. Sørensen, 2003, "Forecasting and Rank-Order Contests," mimeo London School of Economics.

Pezanis-Christou P. and A. Romeu, 2003, "Structural Inferences from First-Price Auction Experiments," mimeo, Institut d'Anàlisi Econòmica.

Pezanis-Christou P. and A. Sadrieh, 2003, "Elicited Bid Functions in (A)Symmetric First-Price Auctions," mimeo, Institut d'Anàlisi Econòmica.

Prelec D., 1998, "The Probability Weighting Function," Econometrica, 66:3, 497-527.

Rabin M., 2000, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," Econometrica, 68(5), 1281-1292.

Read D., 2003, "Monetary Incentives, What Are they Good For?", Mimeo London School of Economics.

Roth A., 1995, "Bargaining Experiments," In: Kagel, J.H. & Roth, A.E. (eds.): Handbook of Experimental Economics, Princeton, NJ: Princeton University Press, 253-348.

Rutström E. and N. Wilcox, 2003, "Learning, Belief Elicitation and Heterogeneity," Mimeo, University of Houston.

Sanders F., 1973, "Skill in Forecasting Daily Temperature and Precipitation: Some Experimental Results," Bulletin of the American Meteorological Society, 54, 1771-9.

Savage L., 1971, "Elicitation of Personal Probabilities and Expectations," Journal of the American Statistical Association, 66, 783-801.

Schotter A. and B. Sopher, 2001, "Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach," mimeo New-York University.

Selten R. and J. Buchta, 1999, "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Function," in: Budescu, D., I. Erev, and R. Zwick (eds.), Games and Human Behavior - Essays in Honor of A. Rapoport, Lawrence Erlbaum Associates, Inc., Mahwa, NJ, USA.

Shao J. and D. Tu, 1995, "The Jackknife and Bootstrap," New York: Springer-Verlag.

Sonnemans J. and T. Offerman, 2001, "Is the Quadratic Scoring Rule Behaviorally Incentive Compatible?", mimeo, CREED.

Viscusi K., Hakes J. and A. Carlin, 1997, "Measures of Mortality Risks," Journal of Risk and Uncertainty, 14, 213-33.

Viscusi K. and C. O'Connor, 1984, "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?", American Economic Review, 74, 942-56.

Vuong Q., 1989, "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," Econometrica, 57, 307-333.

Wakker, P., Thaler R., and A. Tversky, 1997, "Probabilistic Insurance," Journal of Risk and Uncertainty, 15, 7-28.

Wilcox N. and N. Feltovich, 2000, "Thinking Like a Game Theorist: Comment," mimeo University of Houston.

Yaari M.E., 1987, "The Dual Theory of Choice Under Risk," Econometrica, 55, 95-115.

Yi K., 2001, "Quantal-Response Equilibrium Models of the Ultimatum Bargaining Game," mimeo, Hong Kong University of Science and Technology.

Table 1
Evolution of Strategies and Predictions
Estimates of the Parameters Affecting the Mean

	Treatment 1			Treatment 2		
	λ_b	δ_b^5	δ_b^{10}	λ_b	δ_b^5	δ_b^{10}
b=0	4.421* (0.961)	-0.551* (0.047)	-0.247* (0.067)	8.389* (0.712)	0.073* (0.036)	-0.005 (0.028)
b=1	8.758* (1.516)	-0.439* (0.032)	-0.165* (0.058)	16.281* (2.670)	0.270* (0.024)	0.107* (0.019)
b=2	13.806* (2.186)	-0.224* (0.019)	-0.071* (0.030)	24.838* (2.854)	0.412* (0.015)	0.122* (0.019)
b=3	24.802* (2.860)	-0.136* (0.024)	0.002 (0.027)	32.973* (3.008)	0.398* (0.017)	0.131* (0.017)
b=4	35.121* (3.965)	-0.076* (0.019)	-0.002 (0.021)	42.883* (4.553)	0.353* (0.013)	0.146* (0.016)
b=5	47.431* (4.934)	-0.048 (0.026)	-0.005 (0.015)	51.396* (5.867)	0.320* (0.016)	0.127* (0.014)
b=6	58.412* (5.920)	-0.028 (0.017)	0.002 (0.017)	61.340* (7.701)	0.287* (0.013)	0.073* (0.013)
b=7	71.667* (6.806)	-0.011 (0.012)	0.001 (0.012)	75.044* (6.693)	0.195* (0.010)	0.025* (0.011)
b=8	80.128* (7.327)	-0.003 (0.009)	0.001 (0.010)	88.605* (7.859)	0.102* (0.008)	0.013 (0.008)
b=9	91.424* (8.045)	-0.002 (0.004)	-0.003 (0.007)	95.585* (8.459)	0.054* (0.006)	0.001 (0.005)
b=10	96.700* (5.600)	-0.002 (0.003)	-0.002 (0.005)	98.314* (6.263)	0.018* (0.003)	0.001 (0.004)
b=11	98.672* (3.318)	-0.002 (0.002)	-0.001 (0.002)	99.132* (2.836)	0.008* (0.002)	0.001 (0.001)
	λ_v	δ_v^5	δ_v^{10}	λ_v	δ_v^5	δ_v^{10}
v=0	0.018 (0.015)	-0.001 (0.005)	0.000 (0.005)	0.021 (0.013)	-0.001 (0.004)	0.000 (0.004)
v=2	1.016* (0.123)	-0.018 (0.011)	0.001 (0.008)	1.034* (0.081)	-0.025* (0.006)	0.001 (0.006)
v=4	2.214* (0.326)	0.169* (0.026)	0.050 (0.028)	2.406* (0.342)	-0.112* (0.020)	-0.098* (0.021)
v=6	3.634* (0.584)	0.111* (0.016)	0.013 (0.015)	3.740* (0.329)	-0.131* (0.020)	-0.095* (0.017)
v=8	4.912* (0.800)	0.033 (0.021)	-0.005 (0.014)	5.178* (0.664)	-0.146* (0.018)	-0.080* (0.019)
v=11	6.542* (0.734)	0.024 (0.017)	-0.010 (0.012)	6.765* (0.692)	-0.098* (0.014)	-0.137* (0.017)

* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 2
Evolution of Strategies and Predictions
Estimates of the Parameters Affecting the Variance

	Treatment 1					Treatment 2				
Model	σ	γ_1	γ_2	γ_3	σ_η	σ	γ_1	γ_2	γ_3	σ_η
Predictions	7.845* (0.539)	6.386* (0.749)	-0.673* (0.156)	0.024 (0.038)	0.122* (0.045)	8.247* (0.781)	5.435* (0.928)	-0.507* (0.180)	-0.098* (0.033)	0.084* (0.036)
Bids	0.232* (0.042)	0.786 (1.752)	0.895* (0.205)	-0.008 (0.024)	0.148* (0.053)	0.192* (0.036)	0.420 (0.977)	0.977* (0.248)	-0.054* (0.019)	0.111* (0.043)

* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 3
Frequency with which the Correct RNBNE Bid or Strategy is Played

Treatment 1

| | Frequency of RNBNE Bid played for Value Equals to | | | | | | Frequency of RNBNE Strategy |
	v=0	v=2	V=4	v=6	v=8	v=11	
All Periods	98.67%	94.67%	38.33%	15.33%	12.17%	5.67%	1.00%
First Three Periods	93.33%	87.50%	46.67%	26.67%	18.33%	6.67%	0.83%
Last Three Periods	100.00%	98.33%	24.17%	13.33%	10.00%	5.83%	4.17%

Treatment 2

| | Frequency of RNBNE Bid played for Value Equals to | | | | | | Frequency of RNBNE Strategy |
	v=0	v=2	v=4	v=6	v=8	v=11	
All Periods	99.17%	95.67%	67.83%	51.33%	39.17%	24.17%	12.50%
First Three Periods	95.83%	87.50%	34.17%	14.17%	17.50%	6.67%	0.00%
Last Three Periods	100.00%	99.17%	95.83%	88.33%	70.00%	50.00%	41.67%

Table 4
Estimates of the Parameters Affecting the Mean

	λ_b	$\tilde{\lambda}_b$	δ_b^5	$\tilde{\delta}_b^5$	δ_b^{10}	$\tilde{\delta}_b^{10}$
b=0	5.212* (0.792)	2.871 (1.838)	-0.488* (0.051)	0.513* (0.023)	-0.277* (0.062)	0.286* (0.039)
b=1	10.849* (1.707)	4.304* (2.037)	-0.377* (0.037)	0.647* (0.017)	-0.154* (0.054)	0.269* (0.029)
b=2	17.431* (2.391)	6.622* (2.157)	-0.161* (0.021)	0.564* (0.012)	-0.083* (0.031)	0.203* (0.017)
b=3	28.922* (2.878)	2.089 (2.103)	-0.096* (0.027)	0.366* (0.017)	0.002 (0.021)	0.146* (0.015)
b=4	37.505* (4.382)	3.733 (1.926)	-0.056* (0.016)	0.373* (0.013)	0.008 (0.016)	0.133* (0.019)
b=5	48.136* (5.064)	1.059 (1.619)	-0.032* (0.013)	0.295* (0.015)	-0.021 (0.023)	0.148* (0.012)
b=6	61.861* (5.935)	-0.645 (1.904)	-0.042* (0.014)	0.263* (0.021)	-0.005 (0.014)	0.083* (0.010)
b=7	72.225* (7.369)	2.041 (1.546)	-0.026* (0.012)	0.241* (0.023)	0.007 (0.018)	0.016* (0.007)
b=8	83.986* (7.483)	3.623 (1.916)	-0.011 (0.010)	0.086* (0.018)	-0.003 (0.008)	0.013 (0.009)
b=9	92.056* (8.108)	2.212 (2.313)	0.004 (0.007)	0.056* (0.017)	-0.008 (0.010)	-0.017 (0.010)
b=10	97.758* (5.835)	-0.385 (1.431)	-0.003 (0.004)	0.024* (0.0011)	0.000 (0.004)	0.001 (0.005)
b=11	98.811* (3.535)	-0.446 (1.308)	-0.002 (0.003)	0.014 (0.012)	-0.001 (0.003)	0.002 (0.001)

	λ_v	$\tilde{\lambda}_v$	δ_v^5	$\tilde{\delta}_v^5$	δ_v^{10}	$\tilde{\delta}_v^{10}$
v=0	0.023 (0.022)	-0.005 (0.054)	-0.001 (0.005)	0.000 (0.003)	0.000 (0.004)	0.000 (0.002)
v=2	1.024* (0.166)	-0.009 (0.061)	-0.020* (0.004)	-0.007 (0.005)	0.001 (0.007)	0.000 (0.003)
v=4	2.186* (0.344)	0.210 (0.113)	0.155* (0.026)	-0.282* (0.013)	0.046* (0.020)	-0.120* (0.016)
v=6	3.588* (0.626)	0.087 (0.094)	0.070* (0.022)	-0.204* (0.024)	0.007 (0.014)	-0.099* (0.015)
v=8	5.117* (0.852)	-0.020 (0.124)	0.035* (0.013)	-0.169* (0.018)	-0.015 (0.014)	-0.076* (0.018)
v=11	6.508* (0.741)	0.229* (0.105)	0.025 (0.014)	-0.108* (0.017)	-0.024 (0.016)	-0.100* (0.014)

* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 5
Estimates of the Parameters Affecting the Variance

Prediction Model					Bid Model				
σ	γ_1	γ_2	γ_3	σ_η	σ	γ_1	γ_2	γ_3	σ_η
7.974*	5.744*	-0.569*	-0.037	0.109*	0.206*	0.548	0.928*	-0.029	0.130*
(0.586)	(0.815)	(0.179)	(0.042)	(0.042)	(0.047)	(1.221)	(0.279)	(0.021)	(0.049)

* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 6
Maximum Likelihood Estimates of the QRE Models in GHP

Model	Treatment 1					Treatment 2				
	μ	r	α	β	Ln(L)	μ	r	α	β	Ln(L)
QRE RA	0.067* (0.002)	0.611* (0.004)	───	───	-2,856	0.117* (0.003)	0.497* (0.004)	───	───	-3,304
QRE RN	0.554* (0.002)	───	───	───	-5,017	0.298* (0.007)	───	───	───	-4,435
QRE PWF	0.148* (0.002)	───	1.018* (0.009)	2.455* (0.031)	-2,860	0.197* (0.004)	───	0.778* (0.015)	2.111* (0.026)	-3,261

QRE RA stands for Quantal Response Equilibrium with Risk Aversion.
QRE RN stands for Quantal Response Equilibrium with Risk Neutrality.
QRE PWF stands for Quantal Response Equilibrium with Probability Weighting Function.
* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 7
Probability Weighting Function Estimation from Subjects' Predictions

	# of Periods	α	β	σ	Ln(L)
Treatment 1	All	0.635* (0.004)	1.849* (0.009)	0.082 (0.004)	-7,796
	Last 5	0.639* (0.005)	1.891* (0.013)	0.080 (0.003)	-2,645
Treatment 2	All	0.717* (0.006)	1.221* (0.006)	0.078 (0.005)	-8,117
	Last 5	0.730* (0.005)	1.103* (0.006)	0.060 (0.005)	-3,328

* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 8
Maximum Likelihood Estimates of the QRE Model with Misperceptions

Model	# of Periods	Treatment 1			Treatment 2		
		μ	r	Ln(L)	μ	r	Ln(L)
QRE RA	All	0.112* (0.003)	0.375* (0.005)	-2,732	0.107* (0.002)	0.447* (0.005)	-3,145
	Last 5	0.105* (0.003)	0.387* (0.005)	-856	0.047* (0.002)	0.234* (0.004)	-662
QRE RN	All	0.181* (0.003)	—	-3,666	0.276* (0.003)	—	-4,100
	Last 5	0.179* (0.004)	—	-1,135	0.043* (0.001)	—	-798

QRE RA stands for Quantal Response Equilibrium with Risk Aversion.
QRE RN stands for Quantal Response Equilibrium with Risk Neutrality.
* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 9
Maximum Likelihood Estimates of the Quantal Best-Response Model

Model	# of Periods	Treatment 1			Treatment 2		
		μ	r	Ln(L)	μ	r	Ln(L)
QBR RA	All	0.064* (0.003)	0.270* (0.004)	-2,229	0.104* (0.003)	0.438* (0.005)	-2,914
	Last 5	0.055* (0.004)	0.244* (0.007)	-634	0.053* (0.003)	0.254* (0.009)	-628
QBR RN	All	0.121* (0.007)	——	-3,021	0.278* (0.009)	——	-3,918
	Last 5	0.100* (0.011)	——	-859	0.089* (0.005)	——	-755

QBR RA stands for Quantal Best-Response under Risk Aversion.
QBR RN stands for Quantal Best-Response under Risk Neutrality.
* denotes a parameter larger than zero at a 5% significance level.
Numbers in parenthesis refer to the standard deviations of the estimates

Table 10
Maximum Likelihood Estimates of the QRE Models with Heterogeneity

Model	Periods	μ	r	α	β	σ_r	σ_α	σ_β	$\rho_{r,\alpha}$	$\rho_{r,\beta}$	$\rho_{\alpha,\beta}$	Ln(L)
Treatment 1												
QRE RA	All	0.059* (0.004)	0.691* (0.004)	—	—	0.103* (0.013)	—	—	—	—	—	-2,616
QRE PWF	All	0.145* (0.005)	—	0.993* (0.012)	2.224* (0.025)	—	0.054* (0.023)	0.181* (0.038)			-0.127 (0.106)	-2,597
QRE RA With PWF	All	0.102* (0.005)	0.338* (0.007)	0.708* (0.010)	1.789* (0.029)	0.063* (0.010)	0.041* (0.018)	0.160* (0.029)	0.201 (0.256)	0.735* (0.176)	0.054 (0.139)	-2,582
	Last 5	0.101* (0.004)	0.351* (0.009)	0.721* (0.009)	1.801* (0.022)	0.056* (0.012)	0.044 (0.026)	0.172* (0.025)	0.321 (0.198)	0.525* (0.206)	-0.112 (0.242)	-692
QRE RN With PWF	All	0.173* (0.005)	—	0.746* (0.016)	1.866* (0.031)	—	0.050* (0.022)	0.179* (0.043)			-0.192 (0.132)	-3,192
	Last 5	0.169* (0.008)	—	0.761* (0.014)	1.834* (0.026)	—	0.059* (0.019)	0.196* (0.058)			-0.025 (0.183)	-902
QBR RA With PWF	All	0.060* (0.004)	0.237* (0.007)	0.785* (0.018)	1.838* (0.019)	0.065* (0.009)	0.038 (0.016)	0.192* (0.035)	0.098 (0.145)	0.498* (0.152)	-0.128 (0.176)	-2,076
	Last 5	0.053* (0.004)	0.211* (0.010)	0.782* (0.015)	1.843* (0.024)	0.071* (0.013)	0.042 (0.023)	0.207* (0.044)	0.276 (0.210)	0.637* (0.183)	0.022 (0.149)	-542
QBR RN With PWF	All	0.124* (0.010)	—	0.748* (0.020)	1.876* (0.030)	—	0.035* (0.014)	0.197* (0.029)			-0.084 (0.166)	-2,816
	Last 5	0.096* (0.014)	—	0.762* (0.017)	1.790* (0.031)	—	0.044* (0.020)	0.186* (0.040)			-0.050 (0.175)	-796
Treatment 2												
QRE RA	All	0.098* (0.006)	0.527* (0.007)	—	—	0.092* (0.019)	—	—	—	—	—	-3,135
QRE PWF	All	0.178* (0.012)	—	0.804* (0.016)	2.217* (0.021)	—	0.093* (0.038)	0.227* (0.043)			0.077 (0.139)	-3,088
QRE RA With PWF	All	0.090* (0.006)	0.392* (0.008)	0.764* (0.017)	1.156* (0.024)	0.064* (0.014)	0.068 (0.038)	0.172* (0.033)	0.165 (0.200)	0.602* (0.166)	-0.017 (0.154)	-2,945
	Last 5	0.041* (0.004)	0.196* (0.007)	0.780* (0.012)	1.069* (0.020)	0.054* (0.014)	0.040 (0.030)	0.138* (0.024)	0.230 (0.184)	0.563* (0.195)	-0.086 (0.180)	-588
QRE RN With PWF	All	0.237* (0.008)	—	0.778* (0.020)	1.269* (0.023)	—	0.071* (0.028)	0.163* (0.044)			-0.184 (0.177)	-3,812
	Last 5	0.050* (0.004)	—	0.815* (0.022)	1.191* (0.026)	—	0.053 (0.038)	0.145* (0.035)			-0.108 (0.190)	-701
QBR RA With PWF	All	0.084* (0.008)	0.322* (0.010)	0.817* (0.016)	1.148* (0.025)	0.058* (0.012)	0.047 (0.026)	0.166* (0.040)	0.127 (0.225)	0.688* (0.142)	-0.057 (0.223)	-2,777
	Last 5	0.049* (0.006)	0.204* (0.008)	0.807* (0.016)	1.109* (0.028)	0.049* (0.013)	0.045 (0.024)	0.129* (0.035)	0.258 (0.193)	0.716* (0.168)	-0.008 (0.173)	-543
QBR RN With PWF	All	0.235* (0.014)	—	0.770* (0.022)	1.204* (0.025)	—	0.050 (0.032)	0.194* (0.037)			-0.055 (0.148)	-3,593
	Last 5	0.077* (0.013)	—	0.791* (0.018)	1.172* (0.023)	—	0.052* (0.025)	0.162* (0.046)			-0.086 (0.181)	-656

QRE RA stands for Quantal Response Equilibrium with Risk Aversion.
QRE PWF stands for Quantal Response Equilibrium with Probability Weighting Function.
QRE RA with PWF stands for Quantal Response Equilibrium with Risk Aversion and Probability Weighting Function.
QRE RN with PWF stands for Quantal Response Equilibrium with Risk Neutrality and Probability Weighting Function.
QBR RA with PWF stands for Quantal Best-Response with Risk Aversion and Probability Weighting Function.
QBR RN with PWF stands for Quantal Best-Response with Risk Neutrality and Probability Weighting Function.
* denotes a parameter larger than zero at a 5% significance level. Numbers in parenthesis refer to the standard deviations of the estimates
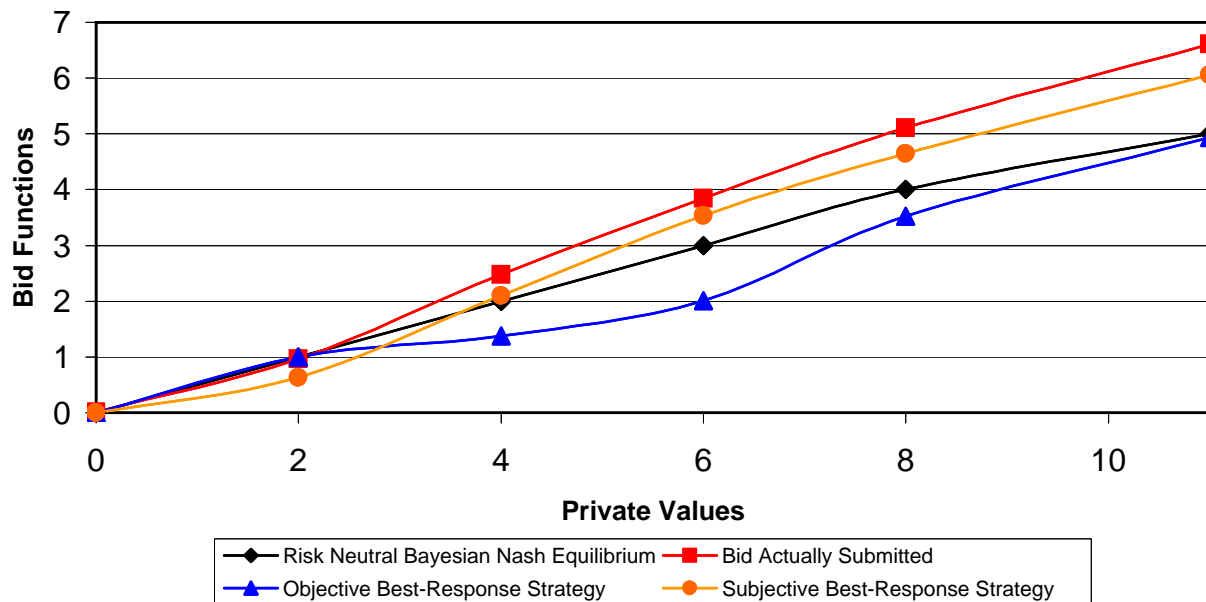
Figure 1
Subjects' Predictions of their
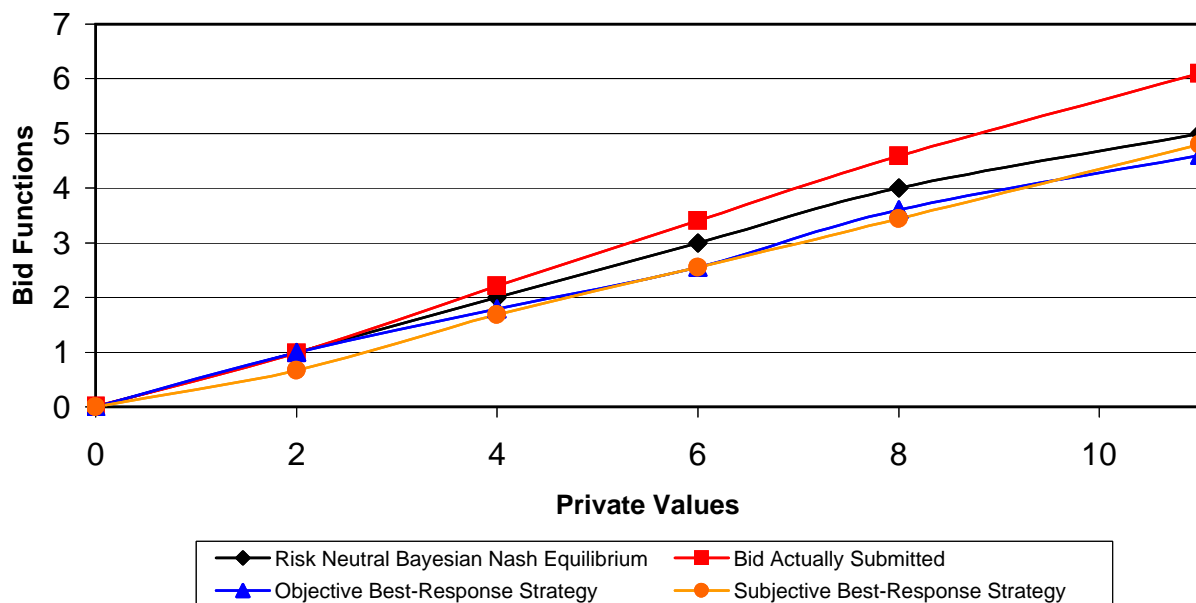**Probability of Winning for Each Possible Bid**
Treatment 1
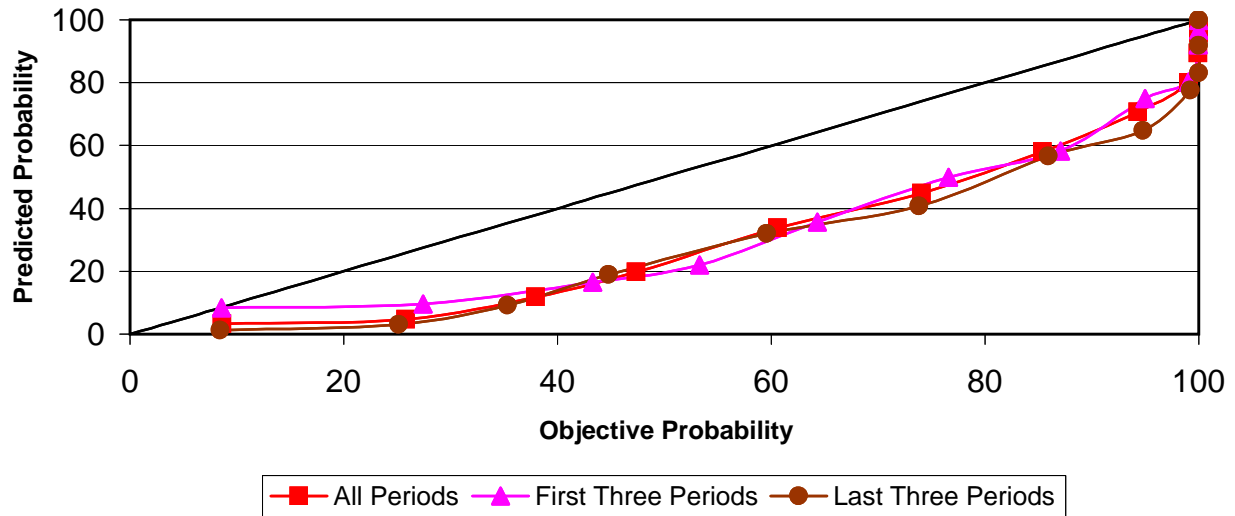
Figure 2
Subjects' Predictions of their
Probability of Winning for Each Possible Bid
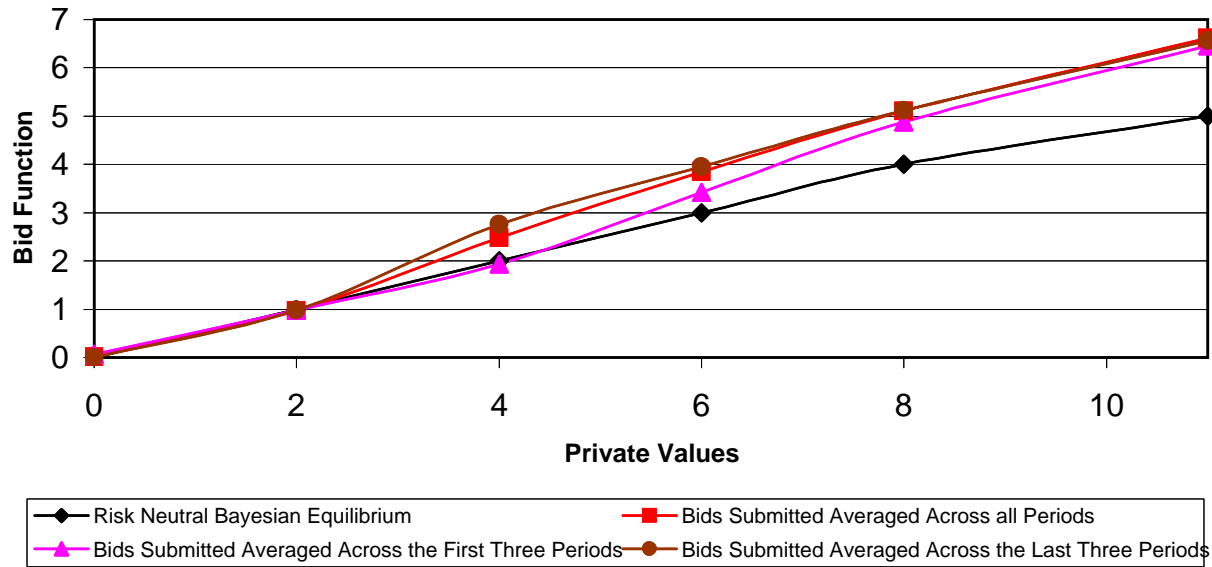**Treatment 2**
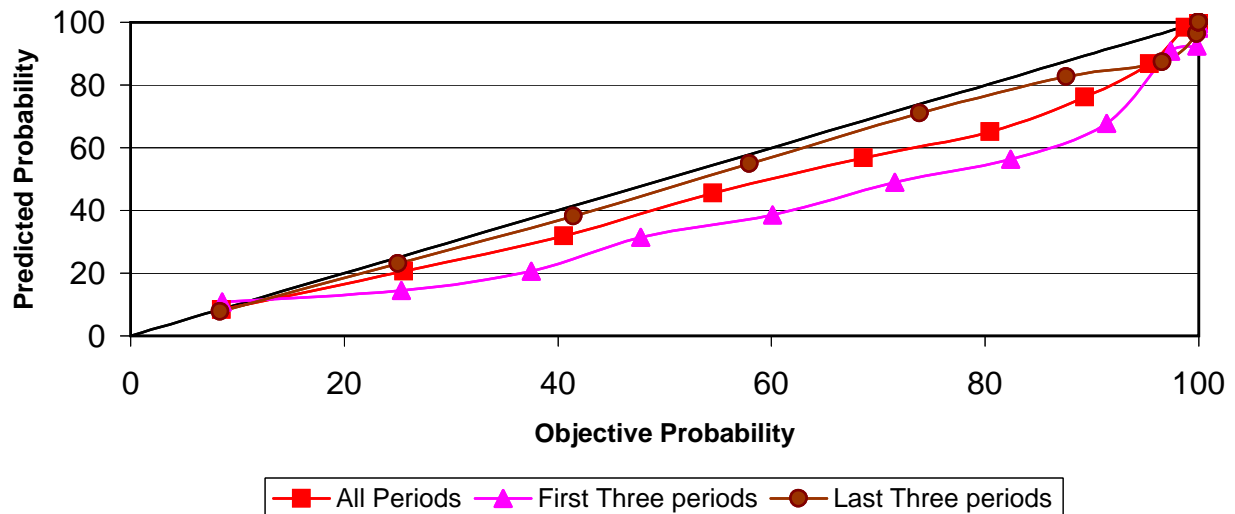
Figure 3
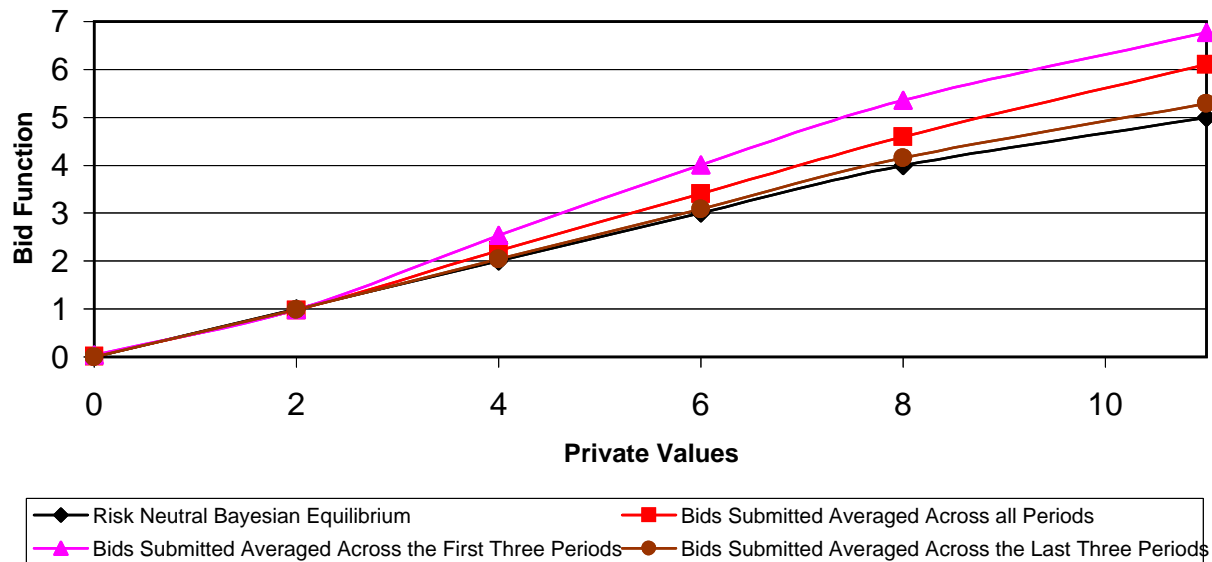**Comparison of Bid Functions**
Treatment 1

Figure 4
Comparison of Bid Functions
Treatment 2

Figure 5
Evolution of Bids Submitted
Treatment 1

Bid Function (y-axis) vs. *Private Values* (x-axis)

Legend:
- Risk Neutral Bayesian Equilibrium
- Bids Submitted Averaged Across all Periods
- Bids Submitted Averaged Across the First Three Periods
- Bids Submitted Averaged Across the Last Three Periods

Figure 6
Evolution of Bids Submitted
Treatment 2

Bid Function (y-axis) vs. *Private Values* (x-axis)

Legend:
- Risk Neutral Bayesian Equilibrium
- Bids Submitted Averaged Across all Periods
- Bids Submitted Averaged Across the First Three Periods
- Bids Submitted Averaged Across the Last Three Periods

Figure 7
A Star-Shaped Probability Weighting Function
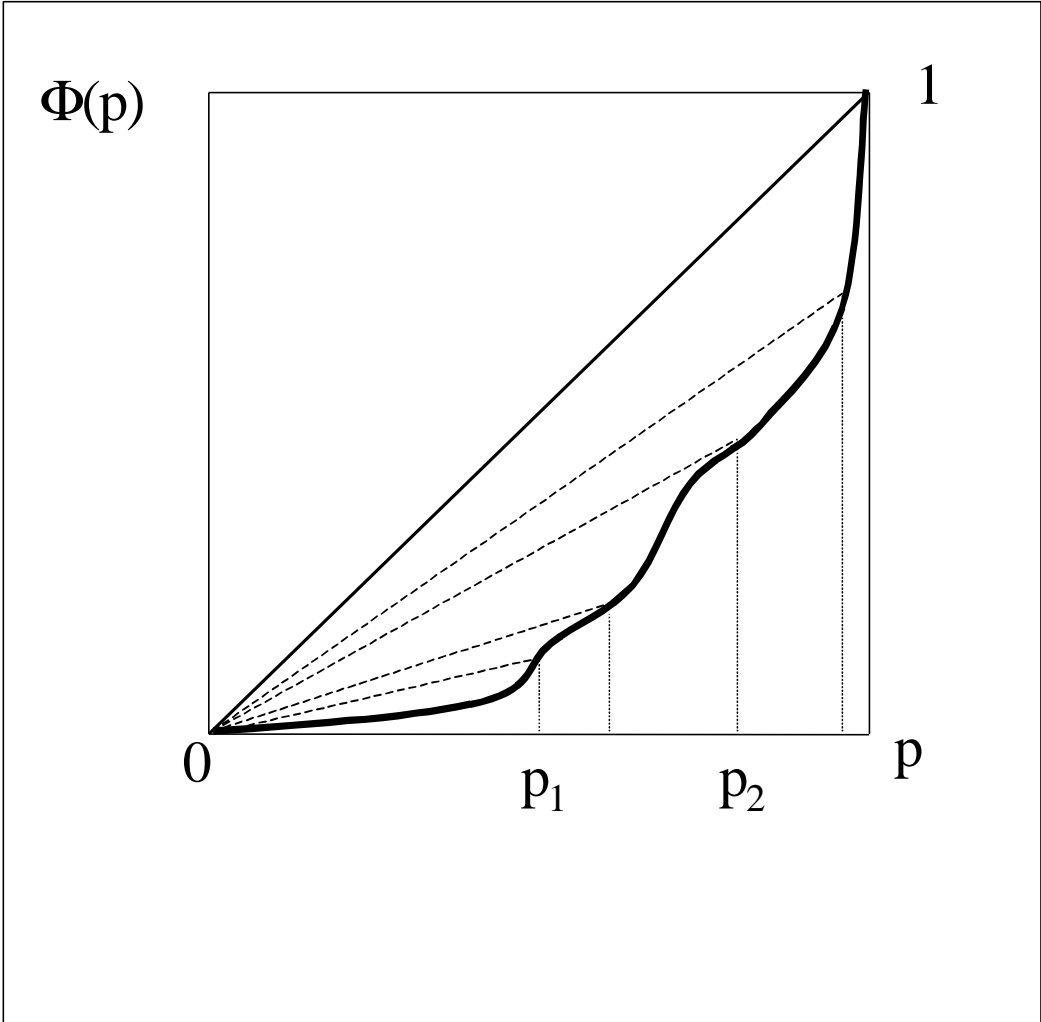
APPENDIX A

INSTRUCTIONS FOR THE EXPERIMENT

You are about to participate in an experimental study on decision-making. You can earn real money, which will be paid to you in cash privately at the end of the experiment. If you read and follow the instructions carefully, and if you make thoughtful decisions, you might earn a substantial amount of money. Please, do not hesitate to ask any question while we read the instructions.

Structure
The experiment consists of 15 rounds, and, as announced, it should last less than 1 hour and 30 minutes.
Each round is decomposed in three phases. Each phase is explained below.

Matching
In each round, you will be identified by a new ID number determined randomly by the computer. Your ID number is shown at the top left corner of your screen.
You will be paired in each round with another participant, using draws of numbered ping pong balls. Each ping pong ball is marked with one of the ID numbers of the 10 persons in the room. At the beginning of phase 3, we will draw two ping pong balls at a time from this bucket to determine who is matched with whom for this round.

The decision problem
In each round, you and the person you are matched with will participate in an auction. Before we explain in details what you will be asked to do, let us briefly describe how the auction works.

At each auction, you and the person you are matched with will bid for a prize. The value of the prize to you will be randomly determined, and equally likely to be $0, $2, $4, $6, $8, or $11. You will not know the value of the prize to the person you are matched with. All you know is that the value to him or her is equally likely to be $0, $2, $4, $6, $8, or $11. As explained below, you and the person you are matched with will be asked to make a monetary bid. The prize goes to the higher bidder. If you are the high bidder your payoff from the auction is equal to the difference between your own prize value and your bid. Otherwise, your payoff from the auction is $0.

Let us now explain in details what you will be asked to do. The decision problem consists in three phases.

- **Phase 1 : Auction Conditional Choices**

At this point, you do not know yet the value of the prize to you. It is only in phase 3 that we will determine randomly the value of the prize for each person in the room.

In phase 1, you will be asked to make a bid in an integer dollar amount for each of the possible prize value you might receive later on. More precisely, you will have to fill each cell of the following table:

Table 1: Bid for Each of the Possible Prize Value						
Your Possible Prize Value	$0	$2	$4	$6	$8	$11
Your Bid for that Prize Value	?	?	?	?	?	?

Note: your bids should be integer dollar amounts

When you fill each cell of this table you must considerer each time what you would do in this situation. For instance, when you fill the third cell, corresponding to a value of $4, you must ask yourself: "What bid should I make if the value of the prize to me is $4?".

After your prize value is randomly assigned to you in phase 3, we will look for the corresponding cell in Table 1 to determine your "effective bid". As further explained below, it is this "effective bid" that will be used to decide whether or not you win the auction. It is therefore very important that you fill each cell in Table 1 very carefully.

- **Phase 2: Predictions**

In phase 2, you will be given an opportunity to earn additional money by making some predictions about your chances of winning the auction for a given list of bids. First, we will flip a coin to decide for which bids you will make your predictions. Heads, and you must make a prediction for each of the following even bids: $0, $2, $4, $6, $8, and $10. Tails, and you must make a prediction for each of the following uneven bids: $1, $3, $5, $7, $9, and $11. For instance, if the coin toss selects "Heads", then, for each even bid in the following table, you must predict how many chances out of 100 does that bid has to be higher than the "effective bid" of a random person in this room.

Table 2: Chances of Winning the Auction for Each of the Following Bids						
Bid	$0	$2	$4	$6	$8	$10
Chances (out of 100) of Winning the Auction With this Bid	?	?	?	?	?	?

Note: in each cell you must enter a number between 0 and 100

If instead the coin toss selects "Tails", then you will have to fill the same table for the uneven bids.

When you fill each cell of table 2, remember that you must evaluate each time the chances this bid has to win the auction against the "effective bid" of a randomly selected person in this room. For instance, when you fill the second cell, corresponding to a bid of $2, we want you to answer the following question: "If I was to play this auction 100 times with someone different each time, how many times would I win the auction if I always bid $2?".

Note that the higher the bid, the higher the chances to win the auction. For instance, the chances to win the auction with a bid of $8 have to be at least as large as the chances to win with a bid of $6. Therefore, the numbers in two consecutive cells should not decrease.

It is important for you to understand that the experiment has been designed so that your choices and payoff in phase 2 are completely independent of your choices and payoff in phase 1, and vice versa. As we will see, your choices and your payoff in phase 2 are also completely independent of the prize value that will be revealed to you in phase 3. In fact, when filling Table 2, you should only concentrate on predicting the chances of winning the auction with each bid on the list, given what you think the other participants in this room will choose on average in phase 1.

Just like with table 1, you are advised to fill each cell of Table 2 very carefully, as one of them will determine your payoffs in phase 2.

- **Phase 3: Auction and Predictions Payoffs**

Auction Payoff

In phase 3, we will first draw ping pong balls to determine who is matched with whom. Then, we will come to your desk to roll a 6-sided die that will determine the value of the prize to you. A throw of 1 will determine a value of $0, a throw of 2 will determine a value of $2, a throw of 3 will determine a value of $4, a throw of 4 will determine a value of $6, a throw of 5 will determine a value of $8, and a throw of 6 will determine a value of $11. As you can see, your prize value is completely independent of the choices you made in phase 1 and phase 2.

Once your value is established, we will enter it in the computer. The computer will then look in Table 1 for the cell corresponding to your prize value in order to determine your "effective bid". The same procedure will be independently repeated with the person you are matched with in order to determine his or her prize value and "effective bid".

The prize will be awarded to you if your "effective bid" is higher than the "effective bid" of the person with whom you are matched. In the event of a tie, we will decide who wins with the flip of a coin (Heads and the person with the higher ID number wins, Tails and the person with the lower ID number wins). If you have the highest "effective bid" (or win the coin flip in the event of a tie), your payoff from the auction is equal to the difference between your own prize value and your "effective bid". If you are the low bidder, you earn nothing in this auction. To summarize,

Auction payoff = your own prize value - your "**effective bid**"
(If you have the highest "effective bid" or win the coin flip in case of a tie)

Auction payoff = 0
(If you have the lowest "effective bid" or lose the coin flip in case of a tie)

As you can see, it is very important that you fill each cell of Table 1 very carefully, as one of them will determine your "effective bid", which will directly influence your auction payoff.

Predictions Payoff

In addition to your auction payoff, we will also determine your payoff from your predictions in phase 2. To do so, we will first roll a 6-sided die. The outcome of this roll will determine which of the six bids in Table 2 will be used to measure the precision of the predictions for all the participants in this room.

Your prediction, and the prediction of the person with whom you are matched, will be compared with your true average probability of winning the auction for the bid randomly selected. This true probability is calculated precisely by the computer by looking at the decisions made in that round by all the participants in this room, other than you and the person with whom you are matched. The true probability is therefore the same for you and the person you are matched with, and it does not depend on your decisions in phase 1.

Your prediction payoff will be $4 if your prediction for the bid randomly selected is closer to the true average probability than the prediction of the person with whom you are matched. In the event of a tie, we will decide who wins the $4 with the flip of a coin (Heads and the person with the higher ID number wins, Tails and the person with the lower ID number wins). Otherwise, your prediction payoff is $0.
If you make thoughtful decisions in phases 1 and 2, your prediction payoff should be roughly the same as your auction payoff. It is therefore very important that you fill Table 1 and Table 2 with equal attention.

Note that you will not be immediately informed of your prediction payoff in each round. They will all be revealed to you at once, at the end of the experiment, that is after the 15 rounds.

Paragraph added in treatment 2.

Once the precision of your predictions has been evaluated, we will print on your screen (both in a table and on a graph) your predictions, as well as your true probabilities of winning the auction for each of the bids in Table 2. Note that these true probabilities are likely to change in the next round, since people in this room may take different decisions. However, we advise you to analyze this information carefully, as it may help you improve your predictions in future rounds.

Payment

You will receive $5 simply for showing up today.

In addition, we will add the auction payoffs and the predictions payoffs you accumulated during the 15 rounds, and we will pay you in cash half of this total.

We remind you to make your choices in each phase of each round very carefully as they will influence how much money you will receive today. In particular, note that if during a particular round, you win the auction with an "effective bid" higher than your prize value, then your auction payoff for that round will be negative and will be subtracted from your total earnings.

Questions

If you have any question, or if any part of these instructions was unclear, please raise your hand.

At anytime during the experiment, do not hesitate to raise your hand if you have a question. One of the instructors will come to you and answer your question privately. At no point should you ask a question aloud, or talk with another participant.

Before we actually start the experiment let's make sure you know how to operate the computer program. From now on you should follow the instructions on the screen, which will lead into the actual experiment.