



Universität St.Gallen

Eliciting Beliefs in Continuous-Choice Games: A Double Auction Experiment

Claudia Neri

December 2012 Discussion Paper no. 2012-7

Editor: Martina Flockerzi
University of St. Gallen
School of Economics and Political Science
Department of Economics
Varnbuelstrasse 19
CH-9000 St. Gallen
Phone +41 71 224 23 25
Fax +41 71 224 31 35
Email seps@unisg.ch

Publisher: School of Economics and Political Science
Department of Economics
University of St. Gallen
Varnbuelstrasse 19
CH-9000 St. Gallen
Phone +41 71 224 23 25
Fax +41 71 224 31 35

Electronic Publication: <http://www.seps.unisg.ch>

Eliciting Beliefs in Continuous-Choice Games: A Double Auction Experiment

Claudia Neri¹

Author's address:

Prof. Dr. Claudia Neri
Univerisity of St.Gallen
Department of Economics
Bodanstrasse 6
CH-9000 St. Gallen
Phone +41 71 224 2757
Fax +41 71 224 2894
Email claudia.neri@unisg.ch
Website <http://www.claudianeri.com>

¹I have benefitted from the opportunity to present earlier versions of this work at the Annual International Meeting of the Economic Science Association (University of Copenhagen), and at Northwestern University, University of St.Gallen, University of Zurich, LUISS, Universidade Nova de Lisboa, CSEF, Bank of Italy. I am grateful to seminar participants and to Charles Manski, Marco Ottaviani, Brian Rogers, Charles Bellemare, and Jacob Goeree for their comments. A previous version of this work circulated as 'Strategic Thinking and Subjective Expectations in a Double Auction Experiment'.

Abstract

This paper proposes a methodology to implement probabilistic belief elicitation in continuous-choice games. Representing subjective probabilistic beliefs about a continuous variable as a continuous subjective probability distribution, the methodology involves eliciting partial information about the subjective distribution and fitting a parametric distribution on the elicited data. As an illustration, the methodology is applied to a double auction experiment, where traders' beliefs about the bidding choices of other market participants are elicited. Elicited subjective beliefs are found to differ from proxies such as Bayesian Nash Equilibrium (BNE) beliefs and empirical beliefs, both in terms of the forecasts of other traders' bidding choices and in terms of the best-response bidding choices prescribed by beliefs. Elicited subjective beliefs help explain observed bidding choices better than BNE beliefs and empirical beliefs. By extending probabilistic belief elicitation beyond discrete-choice games to continuous-choice games, the proposed methodology enables to investigate the role of beliefs in a wider range of applications.

Keywords

Probabilistic beliefs, subjective expectations, private information, experiments, auctions

JEL Classification

E24, E25, E32, J10, J64

1 Introduction

Beliefs play an important role in game-theoretic models of strategic decision-making. In equilibrium models, such as the Nash Equilibrium, the outcome of a game is interpreted as a steady state where players hold correct beliefs about opponents' behavior and act rationally by best responding to those beliefs. In learning models, such as the belief learning model, which focus on if and how a steady state outcome is reached, players update their beliefs about opponents' actions by observing opponents' past actions and then make their choice by best responding to those beliefs.

When game-theoretic models are used to explain empirical or experimental choice data in the absence of beliefs data, proxies are used to represent beliefs. In the equilibrium framework, the proxy is provided by equilibrium beliefs defined in accordance with equilibrium strategies. In the belief learning framework, the proxy is provided by opponents' past behavior. Cheung and Friedman (1997) and Fudenberg and Levine (1998) introduce γ -weighted empirical (empirical, for short) beliefs as a weighted average of opponents' past actions, where the weights decline geometrically at rate γ .¹

In recent years a literature on the probabilistic elicitation of subjective beliefs in games has emerged, with the aim to verify (otherwise non-verifiable) model assumptions and to explore whether elicited beliefs lead to better predictions of choice behavior than alternative proxies do. The literature has focused on normal-form games with a discrete-choice variable, either in a finitely-repeated (Nyarko and Schotter 2002, Rutstrom and Wilcox 2009, Hyndman, Özbay, Schotter and Erhblatt 2012, Danz, Fehr, and Kübler 2012) or in a one-shot setting (Costa-Gomes and Weizsäcker 2008, Rey-Biel 2009).

This paper proposes a methodology to extend probabilistic belief elicitation to games with a continuous-choice variable, thus allowing for the use of beliefs data in a wider range of applications than the one covered by the existing literature. As an illustration, the methodology is applied to a double auction experiment.

In the experiment multiple buyers and sellers with independent private values and costs submit simultaneously bids and offers, respectively. Bids and offers are then aggregated into market demand and supply curves and the market is cleared at a uniform price.² As participants make strategic

¹ Specific cases of empirical beliefs include Cournot beliefs, defined as opponents' last-period actions, and fictitious-play beliefs, defined as the average of opponents' past actions. They are obtained by setting $\gamma = 0$ or $\gamma = 1$, respectively.

² In a static double auction (also called uniform-price two-sided sealed auction, call

bidding decisions in a continuous-choice setting, their probabilistic beliefs about other participants' bidding choices are elicited. The elicited beliefs data then allow for the investigation of how beliefs affect bidding choices.

Moving from a discrete- to a continuous-choice setting requires profound modification to the beliefs elicitation procedure. Since choices are defined over a continuous set, probabilistic beliefs are represented by a continuous subjective probability distribution. Partial information about a subjective probability distribution is elicited by partitioning the choice set into intervals and requiring each participant to report the probability that opponents' choices will fall within each interval. A parametric distribution is then fitted on the elicited beliefs data in order to recover beliefs in the form of a fitted continuous subjective distribution. The procedure builds on previous work on non-strategic decisions and a survey setting, such as the elicitation design used in the Survey of Economic Expectations (SEE³) and the fitting methods presented by Engelberg, Manski and Williams (2009), modifying them to be appropriate for strategic decisions and an experimental setting.

As mentioned above, equilibrium beliefs or empirical beliefs can be used as proxies for subjective beliefs. After implementing probabilistic belief elicitation within the experiment, I compare elicited subjective beliefs with its proxies. As the equilibrium benchmark for the double auction experiment consists of the Bayesian Nash Equilibrium (BNE), equilibrium beliefs are represented by BNE beliefs defined consistently with BNE strategies. I show that elicited subjective beliefs are appropriately proxied by neither BNE nor empirical beliefs and that elicited subjective beliefs help explain observed choices relatively better than BNE and empirical beliefs. Elicited subjective beliefs differ from BNE and empirical beliefs not only in terms of the forecasts of other agents' bidding choices but also in terms of the best-response bidding choices prescribed by beliefs. Moreover, differences in prescribed best responses are accompanied by differences in expected payoffs.

In order to investigate the role of subjective beliefs in explaining bidding behavior, I compare observed choices with the best-response actions

auction, or clearinghouse auction) all traders trade at the same time, when the market is 'called'. In a dynamic double auction (also called continuous double auctions) traders may instead trade anytime the market is open. Conducting a static double auction experiment offers two advantages: the static double auction is widely used as a trading mechanism by stock, bond and commodity exchanges and it can also be easily implemented as a laboratory experiment.

³The complete questionnaire of the Survey of Economic Expectations, conducted by the University of Wisconsin Survey Center, can be found at http://www.disc.wisc.edu/archive/econexpect/cbk_econexpect.txt.

prescribed by subjective, BNE or empirical beliefs. Assuming that subjects best respond, I find that they best respond more often to subjective beliefs than to BNE or empirical beliefs. Thus, a model of best response to the elicited subjective beliefs explains observed choices relatively better than a model of best response to BNE or empirical beliefs. Subjective beliefs, however, perform poorly in predicting other subjects' behavior. Analogous evidence of subjective beliefs' ability to explain observed choices despite low predictive accuracy is reported by Nyarko and Schotter (2002), who elicit beliefs in a 2x2 game and compare subjective beliefs with empirical beliefs.

This paper, while being closely related to previous experimental work in auction markets, differs from it in several respects. Cason and Friedman (1997) conduct a double auction experiment under the same trading rules used in this paper and report analogous evidence that bidding choices deviate from the risk-neutral BNE predictions.⁴ By collecting only bidding choices, they face an identification problem, which I solve, following Manski (2002, 2004), by collecting both choice and beliefs data. I then show that the deviation of observed choice from BNE predictions is related to the deviation of subjective beliefs from BNE beliefs. The result further cautions against proxing subjective beliefs by means of BNE beliefs, especially when trying to explain deviations of observed choices from BNE predictions.

Several papers elicit traders' non-probabilistic beliefs about future prices in a double auction experiment (Smith, Suchanek, and Williams (1988), Hommes, Sonnemans, Tuinstra, and van de Velden (2004, 2005), Haruvy, Lahav, and Noussair (2007)). The definition of beliefs as market price forecasts is certainly appropriate in the investigation of the effect of past prices on traders' beliefs about future price movements, which is the objective of the above-mentioned papers. However, such a definition is less appropriate in the investigation of how beliefs affect strategic decision-making, which is the objective of this paper. The market price is determined jointly by the choice made by a subject holding beliefs and the choices made by other subjects and it is therefore the *result* of strategic interaction.

Armantier and Treich (2009) and Kirchkamp and Reiß (2011) conduct a first-price independent-private-values auction experiment in which bidding strategies are recorded by means of the strategy method. Armantier and Treich (2009) elicit participants' probabilistic beliefs of winning the auction for a given list of bids and find that beliefs are inaccurate. Kirchkamp

⁴ In Cason and Friedman (1997), private valuations are drawn from the uniform distribution over $\$[0, 4.99]$ and participants play 30 trading rounds. In this paper, instead, private valuations are drawn on $\$[0, 9.99]$ and participants play 15 trading rounds.

and Reiß (2011) elicit participants’ non-probabilistic beliefs about the opponent’s bidding strategy and find that beliefs are fairly accurate.⁵ Despite the different auction format, the procedure and the evidence presented in my paper help reconcile the previous apparently-conflicting results. In fact, the belief of winning the auction for some bid depends on (i) the beliefs about other bidders’ *bidding strategies*, (ii) the beliefs about other bidders’ *private values*, (iii) how i-ii are combined into the beliefs about other bidders’ *bids*, and (iv) how the beliefs about other bidders’ bids are combined into the belief of winning. By eliciting probabilistic beliefs about other bidders’ bids, I show that inaccuracies arise as soon as in step *iii*.⁶

The remainder of the paper is organized as follows. Section 2 describes the experimental design and the procedures employed for the probabilistic elicitation and the parametric fitting of subjective beliefs. Section 3 contains descriptive results about subjective beliefs, bidding behavior and market performance. Section 4 presents the main discussion. Finally, Section 5 concludes.

2 Experimental Design and Procedures

2.1 The double auction

The experiment follows the rules of the double auction described by Rustichini, Satterthwaite and Williams (1994). In this section I detail the functioning of the double auction in a general setting. In the following section I outline aspects specific to the implementation of the double auction in a laboratory experiment.

The market is populated with n buyers and n sellers. Each buyer can buy at most one unit and each seller can sell at most one unit of an indivisible good. Each buyer has a private value and each seller has a private cost for the good. Private values and costs are independently drawn from the uniform distribution over $[0, 10]$. A subject’s private value or cost is her own private information, and the process by which private values and costs are drawn is common knowledge among subjects.

All buyers and sellers simultaneously submit their bids or offers. After

⁵ I refer to Manski (2004) for an overview of the limitations of eliciting beliefs without a probabilistic question format.

⁶ Kirchkamp and Reiß (2011) argue that ‘one reason for erroneous best replies is connected to the handling of probabilities, particularly when transforming the expected bidding strategies used by others (together with the underlying distribution of valuations) into the probability distribution of winning bids’.

sorting bids and offers in increasing order as $\psi_{(1)} \leq \psi_{(2)} \leq \dots \leq \psi_{(2n)}$, the market price is set at the midpoint between the middle two figures, at $p = 0.5\psi_{(n)} + 0.5\psi_{(n+1)}$. This procedure is equivalent to determining the intersection between the aggregated demand and supply schedules, as Figure 1 shows with an example. Buyers who have submitted a bid larger than or equal to p and sellers who have submitted an offer smaller than or equal to p trade. A buyer with private value v earns a payoff equal to $v - p$ if she buys the good at the market price, p , and earns zero otherwise. A seller with private cost c earns a payoff equal to $p - c$ if she sells the good at the market price, p , and earns zero otherwise. If excess demand or excess supply arises, priority is given to sellers whose offers are smallest and to buyers whose bids are largest. A fair lottery then determines who trades among the remaining subjects on the long side of the market.

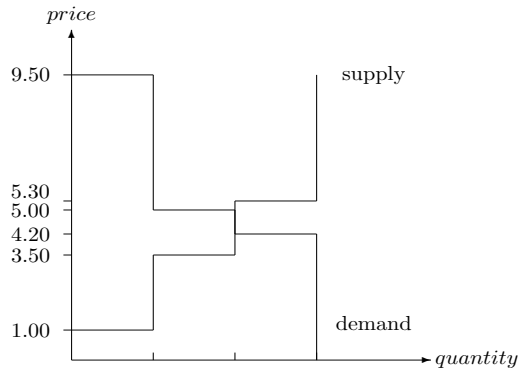


Figure 1: Example: a market with three buyers and three sellers.

Note: Buyers bid at 9.50, 5.00 and 4.20 and sellers offer at 1.00, 3.50 and 5.30. The market-clearing price is at the midpoint between the middle two figures, $p = 4.60$.

For buyer i with private value v submitting bid b , the relationship between her bid b and the bids and offers chosen by the other market participants will determine what her payoff will be, depending on whether she will trade and, if so, at what price. Buyer i will trade if her bid, b , is at least equal to the market price, p . Since the market price is set at $p = 0.5\psi_{(n)} + 0.5\psi_{(n+1)}$, a winning bid b must be at least equal to the $(n + 1)$ -th submitted choice, labeled $\psi_{(n+1)}$. However, from the point of view of buyer i , the uncertainty rests in the choices made by the other $2n - 1$ participants, excluding herself. Therefore, the condition about the winning bid can be formulated in terms of the choices submitted by the other market participants. If their sorted bids and offers are denoted with

$\zeta_{(1)} \leq \zeta_{(2)} \leq \dots \leq \zeta_{(2n-1)}$, then a winning bid b must be larger than the n -th submitted choice, labeled $\zeta_{(n)}$, among the $2n - 1$ choices submitted by the other agents.⁷

Buyer i 's beliefs can then be expressed as beliefs about the realizations of variables $\zeta_{(n)}$ and $\zeta_{(n+1)}$, since the relationship between b , $\zeta_{(n)}$ and $\zeta_{(n+1)}$ will determine whether buyer i trades and, if so, at what price. Assuming buyer i behaves as a risk-neutral subjective expected utility maximizer, she chooses bid b in order to maximize her expected payoff:

$$\begin{aligned} \pi(v, b) = & \int_b^{\bar{z}} \int_{\underline{z}}^b (v - [0.5s + 0.5b]) f_{(n),(n+1)}(s, t) ds dt + \\ & + \int_{\underline{z}}^b \int_{\underline{z}}^t (v - [0.5s + 0.5t]) f_{(n),(n+1)}(s, t) ds dt, \end{aligned} \quad (1)$$

where $f_{(n),(n+1)}(s, t)$ denotes the joint density of $\zeta_{(n)}$ and $\zeta_{(n+1)}$ evaluated at point (s, t) , and the set $[\underline{z}, \bar{z}]$ represents the set over which bids and offers are defined.⁸ In the double integrals, the inner one integrates over all possible values of $\zeta_{(n)}$ and the outer one integrates over all possible values of $\zeta_{(n+1)}$. The first term is the expected payoff that buyer i will receive if her bid, b , turns out to be $\zeta_{(n)} < b < \zeta_{(n+1)}$. In this case $\zeta_{(n)}$ can take a value between \underline{z} and b , while $\zeta_{(n+1)}$ can take a value between b and \bar{z} , and the price is $p = 0.5\zeta_{(n)} + 0.5b$. The second term is the expected payoff that buyer i will receive if her bid, b , turns out to be $b > \zeta_{(n+1)}$. In this case $\zeta_{(n)}$ can take a value between \underline{z} and $\zeta_{(n+1)}$, while $\zeta_{(n+1)}$ can take a value between \underline{z} and b , and the price is $p = 0.5\zeta_{(n)} + 0.5\zeta_{(n+1)}$.

Analogously, the beliefs of seller i , who has private cost c and who submits offer a , can be expressed as beliefs about the realizations of variables $\zeta_{(n-1)}$ and $\zeta_{(n)}$, and her expected payoff can be written as

$$\begin{aligned} \pi(c, a) = & \int_a^{\bar{z}} \int_{\underline{z}}^a (0.5a + 0.5t - c) f_{(n-1),(n)}(s, t) ds dt + \\ & + \int_s^{\bar{z}} \int_a^{\bar{z}} (0.5s + 0.5t - c) f_{(n-1),(n)}(s, t) ds dt. \end{aligned} \quad (2)$$

Therefore, a buyer's beliefs are represented by joint density $f_{(n),(n+1)}$ and a seller's beliefs are represented by joint density $f_{(n-1),(n)}$. Both densities are a function of the density and cumulative distribution functions of other buyers' bids (let us denote them as g_b and G_b , respectively), and the density and

⁷ $\zeta_{(n)}$ is the n -th order statistic among the $2n - 1$ choices submitted by the other agents.

⁸ The formulas for $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$ are reported in Appendix A.

cumulative distribution functions of other sellers' offers (let us denote them as g_s and G_s , respectively). Therefore, beliefs data could be collected either (i) by eliciting directly the joint densities, $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$, or (ii) by first eliciting g_b , G_b , g_s and G_s , and then computing the joint densities, $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$, using the order statistic formulas in Appendix A. In Section 2.3 I will argue in favor of the second method and I will present the details about its implementation.

2.2 Experimental design

A total of 66 subjects, recruited among Northwestern University undergraduate students, participated in a total of 11 auction markets designed as computer lab experiments using the z-Tree software. Four auctions were conducted with eight traders, four with four traders and three with six traders. At the beginning of the experiment participants were randomly matched into groups of traders and interacted within the same group over the course of 15 rounds.

In the first round each player was randomly assigned the role of either buyer or seller. If a player was initially assigned to the role of buyer (seller), then she was a buyer in rounds 1 through 5, a seller (buyer) in rounds 6 through 10, and a buyer (seller) again in rounds 11 through 15. The information about each participant's current role always appeared on the computer screen. In each round, each buyer had \$10 and each seller had one unit of the commodity. Each buyer could purchase a single unit of the commodity from any seller, and each seller could sell a single unit of the commodity to any buyer.

Before the start of each round, each buyer was informed about her personal value, v , and each seller was informed about her personal cost, c . The personal values and costs of other buyers and sellers were never revealed. Each buyer's personal value and seller's personal cost were assigned randomly at the beginning of each round: a computer determined each of them as an independent random draw of the 1,000 numbers (to the nearest penny) between \$0.00 and \$9.99. Thus, in each round I collected a new set of observations of the subjects' decisions and of the resulting market price and trades. After being informed about their personal values and costs, buyers and sellers submitted their bids and offers. Any non-negative dollar amount was allowed. Amounts could be specified up to a precision of cents of a dollar. After submitting their choices, players were not allowed to modify them.

At the end of each round, each participant received detailed feedback

about the results for that trading period. Information included the following: a graph representing the bids and offers submitted by all market participants (i.e., the demand and supply curves), the market price, the participant’s trading status (indicating whether she traded), her profit (with a reminder of the formula for profit), the accuracy of her beliefs about the bidding choices of the other buyers and sellers, and the monetary reward earned for holding accurate beliefs. If a player did not trade, her initial cash (if a buyer) or the unsold unit of the commodity (if a seller) did not count towards her profits.

When the last round ended, the computer randomly drew one of the rounds and the participants were paid according to their performance in that round only.⁹ Payments included \$5 for attending the session, \$5 for answering correctly all questions in a comprehension quiz administered before the beginning of the auction, and the amount earned in the auction in the randomly selected round. Any negative trading profit was deducted from the total payments.¹⁰ All payments were announced and paid out in US dollars. On average, subjects earned \$10.26.¹¹

Appendix C contains additional information about the experimental procedures, including the complete instructions, an overview of participants’ performance in the comprehension quiz, and a description of the subjects pool.

2.3 Eliciting subjective beliefs

Section 2.1 has shown that a buyer’s and a seller’s beliefs are represented by joint density $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$, respectively. Eliciting beliefs in the form of $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$ is likely to prove very cumbersome. First, the subjects’ assessment of the joint probability of the simultaneous occurrence of events is usually not coherent.¹² Second, $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$ are joint densities of order statistics, and the concept of order statistic is likely to produce confusion and misunderstanding. Section 2.1 has also shown that the beliefs about an order statistic of other subjects’ bids and offers is a

⁹ Cason and Friedman (1997) instead allowed profits to accumulate throughout the rounds and paid participants the accumulated amount.

¹⁰ No participant received a negative payoff at the end of the session.

¹¹ On average, \$2.24 for forecasting, \$0.82 for trading, and \$2.20 for completing the comprehension quiz. Everyone received \$5 for attendance. The average profit for trading was \$2.08 for players who concluded a trade.

¹² For example, the assessment can be affected by the so-called conjunction fallacy: the assessment of the joint probability is higher than the probabilities of the constituent events.

function of (i) the beliefs about other buyers' bids and (ii) the beliefs about other sellers' offers. Based on this feature, the approach explored in this paper consists of measuring separately the cumulative distribution functions of other buyers' bids (G_b) and the cumulative distribution functions of other sellers' offers (G_s), and then computing the joint densities $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$.¹³ Thus, experiment participants are required to report separately beliefs about the bids and the offers of other market participants, but are not required to refer to a relative ordering among them.

Since in the experiment the only restriction on bidding choices is that they must be non-negative and must be expressed in dollars and cents, the action space for buyers and sellers consists of all numbers from 0 to ∞ rounded up to the second decimal. Thus, the action space is continuous and unbounded. The action space being continuous, probabilistic beliefs about other subjects' bidding choices need to be represented by a continuous probability distribution. The action space being unbounded, the continuous probability distribution representing beliefs needs to be defined over a potentially unbounded support.¹⁴

The set $[0, \infty)$ is divided into six intervals: $[0, 2]$, $[2.01, 4]$, $[4.01, 6]$, $[6.01, 8]$, $[8.01, 10]$ and $[10.01, \infty)$. I place the cutoff at \$10 because realizations of private values and costs are not higher than \$10. I choose a width equal to \$2 for the intervals and, therefore, a total number of six intervals to allow for a balance between the goal of eliciting multiple pieces of information regarding subjective beliefs and the goal of keeping the task manageable and fairly quick for the respondents. The respondents are required to assign to each interval a probability that represents their beliefs that the variables of interest (bids or offers) will fall within each interval. The task is presented in each round, after each subject has submitted her bid or offer.¹⁵ The exact wording is the following.

Think about the other buyers and sellers in the market. Their personal values and costs, as your own value¹⁶, are also determined by the computer as a random draw of the numbers between 0.00 and 9.99. What do you think they will choose? Please answer the following questions.

¹³ See the order statistic formulas in Appendix A.

¹⁴ By not restricting the support to be bounded, respondents are not misled to believe that other subjects' bidding choices are restricted to be within a certain upper bound.

¹⁵ The elicitation method that I adopt builds on the procedures used in survey studies such as the Survey of Economic Expectations. See Dominitz and Manski (1997).

¹⁶ 'Value' is displayed if the player is a buyer, 'cost' is displayed if the player is a seller.

What do you think the percent chance is that another¹⁷ BUYER in the market will choose a bid in each of the following intervals? Assign a percent chance to each interval. Remember: a percent chance is a number between 0 and 100. Also, percent chances should add up to 100.

between \$0 and \$2	between \$2.01 and \$4	between \$4.01 and \$6	between \$6.01 and \$8	between \$8.01 and \$10	larger than \$10

What do you think the percent chance is that a¹⁸ SELLER in the market will choose an offer in each of the following intervals? Assign a percent chance to each interval. Remember: a percent chance is a number between 0 and 100. Also, percent chances should add up to 100.

between \$0 and \$2	between \$2.01 and \$4	between \$4.01 and \$6	between \$6.01 and \$8	between \$8.01 and \$10	larger than \$10

Subjects are rewarded separately for the accuracy of their beliefs about the choices made by other buyers and for the accuracy of their beliefs about the choices made by other sellers. Beliefs are accurate when they correctly predict the fraction of other subjects in the market submitting a choice within each interval. I use a quadratic scoring rule to determine the reward amount. The quadratic scoring rule is an incentive-compatible mechanism, which provides subjects with an incentive to report their beliefs truthfully, provided that subjects are risk neutral and do not distort probabilities.¹⁹ Subjects are provided with a simplified explanation of the quadratic scoring rule in the instructions.²⁰

¹⁷ ‘Another’ is displayed if the player is a buyer, ‘a’ is displayed if the player is a seller.

¹⁸ ‘A’ is displayed if the player is a buyer, ‘another’ is displayed if the player is a seller.

¹⁹ Rewarding subjects for the accuracy of their beliefs and using the quadratic scoring rule to determine the reward may have several limitations. Rutström and Wilcox (2009) and Palfrey and Wang (2009) report evidence on the effects that the elicitation of beliefs may have on choice behavior, including the possibility of more strategic behavior, lower risk aversion and overconfidence. Blanco, Engelmann, Koch and Normann (2010) investigate whether subjects employ the belief elicitation task in order to hedge against adverse outcomes in the choice task. They find no evidence of hedging being a major problem in belief elicitation.

²⁰ See Appendix C for the complete text of the instructions.

Suppose that a subject, when asked to report her beliefs about the bids made by N buyers, assigns probabilities $y = (y_1, \dots, y_6)$ to the six intervals, $j = 1, 2, \dots, 6$. Suppose also that (N_1, \dots, N_6) are the actual numbers of bids falling within each interval. The quadratic scoring rule then determines that the subject's reward is equal to $\$2 - \frac{1}{N} \sum_{j=1}^6 N_j [(y_j - 1)^2 + \sum_{h \neq j} y_h^2]$. Such a scoring rule implies that the reward is bounded between \$0 and \$2.²¹

2.4 Parametric analysis of subjective beliefs

Each observation in the data consists of the six probabilities, which a subject assigned to each one of the six intervals $\$[0,2]$, $[2.01,4]$, $[4.01,6]$, $[6.01,8]$, $[8.01,10]$ and $[10.01,\infty)$. From knowledge of the probabilities assigned to each interval, the value of the subjective cumulative distribution function at the right endpoints of the six intervals can be easily computed. I denote the values of the cumulative distribution function at these points as $G(r_1), \dots, G(r_6)$, where r_1, \dots, r_6 denote the right endpoints of the six intervals. Fitting a parametric distribution over the data $G(r_1), \dots, G(r_6)$ allows to obtain the fitted density and cumulative distribution functions of other subjects' bids, \hat{g}_b and \hat{G}_b , and the fitted density and cumulative distribution functions of other subjects' offers, \hat{g}_s and \hat{G}_s .

Depending on (a) the number of intervals where positive probability is placed, and (b) whether the intervals, where positive probability is placed, are adjacent to each other, I fit subjective beliefs using a unimodal Beta distribution, a triangular distribution, the union of two or three triangular distributions, or the union of a unimodal Beta and a triangular distribution. Empirically, the most relevant case (approximately 90% of observations) is the one in which beliefs are fitted with a unimodal Beta distribution. Appendix B contains an overview of the fitting methods, which build on procedures introduced by Engelberg, Manski and Williams (2009). After obtaining \hat{g}_b , \hat{G}_b , \hat{g}_s and \hat{G}_s , the fitted joint densities $f_{(n),(n+1)}$ and $f_{(n-1),(n)}$ can be computed using the formulas in Appendix A.

²¹ In a two-person two-action (action A and action B) experiment this quadratic scoring rule would simplify to $\$2 - \{[1_A - p_A]^2 + [(1 - 1_A) - (1 - p_A)]^2\}$, where 1_A is an indicator functions that equals 1 if the opponent chooses action A and 0 if the opponent chooses B and p_A is the subjective probability that the opponent chooses action A .

3 Descriptive Results

3.1 Descriptive results about subjective beliefs

The elicited beliefs reveal substantial heterogeneity across subjects. Among the beliefs about other buyers' choices as well as among the beliefs about other sellers' choices, the subjective probabilities assigned to each interval range from 0% to 100% for most intervals.²²

The only restriction imposed on the bidding choices that subjects submit and on the beliefs that they report is that bidding choices must be non-negative and that positive subjective probabilities can be assigned only to non-negative bidding choices. This implies that subjects can submit a bidding choice above \$10 and can assign positive probability to bidding choices above \$10 (i.e., above the highest possible private valuation). While an actual bid or offer higher than \$10 is extremely rare, beliefs assigning a positive probability to bidding choices falling in the interval $[\$10.01, \infty)$ are not rare (19% among beliefs about bids and among beliefs about offers).²³

3.2 Descriptive results about bidding behavior

I now illustrate how the observed bidding choices compare with the bidding behavior prescribed by the risk-neutral BNE model, as presented by Rustichini, Satterthwaite and Williams (1994).²⁴

Since the double auction is a strategic environment with private information, strategic misrepresentation of private information is a key feature within the BNE model. As a measure of how much subjects reveal of the private information they hold, I use the ratios also employed by Cason and Friedman (1997). The value underrevelation ratio $VUR(v, b)$ is the percentage of the buyer's private value, v , revealed in the chosen bid, b , and the cost underrevelation ratio $CUR(c, a)$ is the percentage of the seller's private cost, c , revealed in the chosen offer, a . Thus, $VUR(v, b) = (v - b)/v$ and

²² Tables are omitted for brevity.

²³ There are only 4 observations (out of 990) of a bid or offer above \$10. Buyer $i = 49$ in round 12 bids at \$10.60. Seller $i = 21$ in round 9 offers at \$11, seller $i = 7$ in round 2 offers at \$15, and seller $i = 29$ in round 2 offers at \$18.50.

²⁴ Rustichini, Satterthwaite and Williams (1994) show that there exists a family of asymmetric smooth equilibria, which can be computed numerically. To simplify the analysis, I consider approximate symmetric bid and offer functions. The approximation should not affect the analysis since the family of asymmetric smooth equilibria is contained in a small neighborhood of the symmetric bid and offer functions. Also Cason and Friedman (1997) use approximate symmetric bid and offer functions in their experimental implementation of the Rustichini, Satterthwaite and Williams (1994) double auction.

$CUR(c, a) = (a - c)/(9.99 - c)$.²⁵ A positive ratio corresponds to underrevelation and a negative ratio corresponds to overrevelation.²⁶

Table 1: Underrevelation of private information and deviation of the BNE best response from the observed choice. Median, 1st quartile (Q1), and 3rd quartile (Q3).

		(i) all %	(ii) bid ≤ value or offer ≥ cost %	(iii) bid > value or offer < cost %
underrevelation of private information				
$VUR(v, b) = (v - b)/v$	median	3***	7***	-34***
	[Q1, Q3]	[0, 20]	[0, 23]	[-203, -13]
$CUR(c, a) = (a - c)/(9.99 - c)$	median	1***	4***	-36***
	[Q1, Q3]	[0, 13]	[0, 20]	[-141, -10]
deviation of BNE best response from choice				
$D(b_{BNE}, b) = (b_{BNE} - b)/b_{BNE}$	median	-4***	0	-46***
	[Q1, Q3]	[-12, 10]	[-10, 14]	[-208, -28]
$D(a_{BNE}, a) = (a_{BNE} - a)/a_{BNE}$	median	6***	2***	42***
	[Q1, Q3]	[-2, 32]	[-4, 12]	[25, 59]
obs.		990	802	188

Note: Wilcoxon signed-rank test of the null hypothesis that the median equals zero. *** denotes significantly different from zero at 0.1 percent, ** at 1 percent, * at 5 percent.

The upper panel of Table 1 reports the median value and cost underrevelation ratios. Results are reported for (i) the entire sample, and separately for (ii) the subsample in which underrevelation occurs and (iii) the subsample in which overrevelation occurs. Within subsample (ii) the median

²⁵Recall that the highest possible private cost in the experiment is \$9.99.

²⁶Table 1 reveals that 19% of the observations (188 out of 990) consist of either a buyer bidding above value or a seller offering below cost. While a percentage of 19% is surprisingly high, several qualifications need to be made. Differences exist across subjects: some subjects incur in overrevelation more often than others. Among participants, 70% of them overreveal their private information in 3 or fewer of the 15 total rounds. Most importantly, the expected monetary losses generated by overrevelation are small. I compute, based on the elicited subjective beliefs, the expected payoffs generated by the actual bidding choices above value or below cost and I compare them with the expected payoffs generated by bidding at value or at cost. The percentage of observations, for which the expected monetary loss generated by overrevelation is greater than or equal to 10 cents (50 cents) is equal to 10% (5%).

$VUR(v, b)$ and $CUR(c, a)$ are 7 and 4 percent, respectively.²⁷

The lower panel of Table 1 reports the magnitude of the percent deviation of the BNE best response from the observed choice, defined as $D(b_{BNE}, b) = (b_{BNE} - b)/b_{BNE}$ for buyers and $D(a_{BNE}, a) = (a_{BNE} - a)/a_{BNE}$ for sellers. Within subsample (ii), the median deviation is approximately 0 for buyers and 2 percent for sellers.²⁸

3.3 Descriptive results about market performance

Table 2 reports a description of market performance across the experimental auctions. The number of trades taking place in each round ranges between zero and three, with one or two trades per round being the most common outcome. Over all observations, including those when no trade occurs, mean profits are \$0.88. In a competitive equilibrium (CE) buyers and sellers submit respectively their private values and private costs without engaging in any strategic misrepresentation of their private information. Across all experimental auctions, trading efficiency, defined as the percentage of the gains from exchange realized by traders in comparison with the CE, is 78% and prices are within the CE price interval in 50% of the rounds.²⁹

4 Bidding Behavior and Subjective Beliefs

In this section, I will discuss the main results and describe the role of subjective beliefs in explaining observed choice behavior. The first question, which I address, is what alternatives would be available if subjective beliefs were

²⁷ I compare $VUR(v, b)$ and $CUR(c, a)$, which are computed with respect to the observed bids and offers b and a , with the underrevelation ratios computed with respect to the bids and offers prescribed by the BNE strategies b_{BNE} and a_{BNE} , defined as $VUR(v, b_{BNE}) = (v - b_{BNE})/v$ and $CUR(c, a_{BNE}) = (a_{BNE} - c)/(9.99 - c)$. The medians of $VUR(v, b_{BNE})$ and $CUR(c, a_{BNE})$ are equal and approximately 10 percent, indicating that experiment participants reveal more private information than they would according to the BNE.

²⁸ Cason and Friedman (1997) report qualitatively comparable results, in spite of several differences in the experimental design. First, buyers choose a bid above what the BNE would prescribe and sellers choose an offer below what BNE would prescribe. Second, subjects reveal more of their private information than what they would according to BNE strategies: for inexperienced subjects the median value underrevelation ratio is 6.5 percent and the median cost underrevelation ratio is 9.1 percent (for experienced subjects the ratios are 2.5 and 2.4 percent, respectively).

²⁹ A comparison with Cason and Friedman (1997) is possible for the markets with eight subjects. They find a higher trading efficiency and a slightly higher fraction of prices within the CE price interval. Recall that in their experiment there are more trading rounds (30 instead of 15).

Table 2: Market Performance

	2 buyers, 2 sellers	3 buyers, 3 sellers	4 buyers, 4 sellers	all markets
number of trades				
0	20%			7%
1	67%	53%	27%	48%
2	13%	44%	55%	37%
3		2%	18%	7%
mean profits	\$0.94	\$0.78	\$0.91	\$0.88
trading efficiency	84.8	66.4	79.8	78
<i>Cason & Friedman 1997</i>			87.3	
within CE				
price interval	68.3%	37.8%	41.7%	50%
<i>Cason & Friedman 1997</i>			45%	

not elicited. If that were the case, choice behavior could be studied assuming that the beliefs that subjects hold can be proxied with either BNE beliefs or empirical beliefs. Therefore, I now define BNE beliefs and empirical beliefs.

Bayesian Nash Equilibrium beliefs. BNE beliefs can be represented by means of density and cumulative distribution functions, analogously to the way subjective beliefs are represented. The bid density and the offer density can be computed using the knowledge of the BNE strategies. Denoting with $b_{BNE}(v)$ and $a_{BNE}(c)$ the BNE bid and offer functions, respectively, the bid density evaluated at a bid equal to b is $g_{b,BNE}(b) = f_b(v) \frac{1}{b_{BNE}(v)}$, where $b_{BNE}(v) = b$ and $f_b(v)$ is the density of private values at v . Analogously, the offer density evaluated at an offer equal to a is $g_{s,BNE}(a) = f_s(c) \frac{1}{a_{BNE}(c)}$, where $a_{BNE}(c) = a$ and $f_s(c)$ is the density of private costs at c . Recalling that private values and costs are distributed uniformly over $[0, 10]$, we can substitute $f_b(v) = f_s(c) = \frac{1}{10}$.

Empirical beliefs. Following Cheung and Friedman (1997), I define subject i 's γ -weighted empirical beliefs about the likelihood that another subject will choose a bid (offer) in the j -th interval in round $t + 1$ as $\eta_{i,t+1,j} = (N_{t,j} + \sum_{u=1}^{t-1} \gamma_i^u N_{t-u,j}) / (N_t + \sum_{u=1}^{t-1} \gamma_i^u N_{t-u})$, where $N_{t,j}$ is the number of other players who chose a bid (offer) in the j -th interval in period t , γ_i^u is the weight given to the observation of a bid (offer) in the j -th interval being chosen by another buyer (seller) in round $t - u$, and N_t is the total number of buyers (sellers) whose behavior is observed.³⁰

³⁰ Note that, since in the experimental design a subject changes roles in the sixth and

I define γ^* -weighted empirical beliefs as the γ -weighted empirical beliefs such that γ_i minimizes the absolute squared difference between subjective beliefs and γ -weighted empirical beliefs.³¹ Thus, for each subject i in the data, I find the γ_i that solves $\sum_{t=1}^{15} \sum_{j=1}^6 |y_{i,t,j} - \eta_{i,t,j}|^2$, where $y_{i,t,j}$ denotes the subjective probability that subject i assigned to another buyer (seller) choosing a bid (offer) in the j -th interval in round t and $\eta_{i,t,j}$ is the corresponding probability according to γ -weighted empirical beliefs.³² In what follows, I will refer to empirical beliefs according to the definition of γ^* -weighted empirical beliefs. After computing in such a way empirical beliefs, I apply to them the same fitting methods applied to subjective beliefs.³³

I now turn to presenting the main results. I organize my investigation according to the following questions. *Question 1:* Can subjective beliefs be modeled by BNE or γ^* -empirical beliefs? *Question 2:* If subjects best respond, which beliefs do they best respond to? *Question 3:* How accurate are subjective beliefs in predicting other subjects' choices? *Question 4:* Is there evidence of convergence of beliefs and/or choices to the BNE?

4.1 Can subjective beliefs be modeled by BNE or γ^* -empirical beliefs?

This question is answered separately for beliefs about bids and for beliefs about offers. Given that results for both are similar, for brevity I will report only results regarding beliefs about bids.

I compute for each participant i and for each round t the average absolute differences (AAD) between the subjective beliefs $y_{i,t}$ and the γ^* -empirical beliefs $\eta_{i,t}$, and between the subjective beliefs $y_{i,t}$ and the BNE beliefs $\nu_{i,t,j}$, which I define respectively $AAD_{i,t}(y, \eta) = \frac{\sum_{j=1}^6 |y_{i,t,j} - \eta_{i,t,j}|}{6}$ and $AAD_{i,t}(y, \nu) = \frac{\sum_{j=1}^6 |y_{i,t,j} - \nu_{i,t,j}|}{6}$. I then compute for each subject i the average AADs over three subperiods. Since participants change roles (from buyer

again in the eleventh rounds, the total number of other buyers (sellers) whose choices she predicts is not constant during the entire experimental session.

³¹ For previous work employing γ^* -weighted empirical beliefs see Nyarko and Schotter (2002).

³² The median estimated γ^* is equal to 1.02 for beliefs about bids and equal to 0.99 for beliefs about offers. For round $t = 1$, I assume that γ -weighted empirical beliefs prescribe $\eta_{i,1,j} = 1/6$ for $j = 1, \dots, 6$.

³³ Thus, for empirical beliefs, I also first obtain the fitted density and cumulative distribution functions of other subjects' bids, \hat{g}_b and \hat{G}_b , and the fitted density and cumulative distribution functions of other subjects' offers, \hat{g}_s and \hat{G}_s , and then use \hat{g}_b , \hat{G}_b , \hat{g}_s and \hat{G}_s to compute the fitted joint densities $\hat{f}_{(n),(n+1)}$ and $\hat{f}_{(n-1),(n)}$.

to seller, or vice versa) in the 6th and 11th rounds, I consider separately rounds 1-5, 6-10 and 11-15. For a generic subperiod with $\underline{t} \leq t \leq \bar{t}$, I define $\overline{AAD}_{i,t \leq \underline{t} \leq \bar{t}}(y, \eta) = \frac{\sum_{t=\underline{t}}^{\bar{t}} AAD_{i,t}(y, \eta)}{\bar{t} - \underline{t} + 1}$ and $\overline{AAD}_{i,t \leq \underline{t} \leq \bar{t}}(y, \nu) = \frac{\sum_{t=\underline{t}}^{\bar{t}} AAD_{i,t}(y, \nu)}{\bar{t} - \underline{t} + 1}$, respectively. Figure 2 presents the sample distributions across subjects.

Subjective beliefs differ from γ^* -empirical and BNE beliefs, since the distributions in Figure 2 are spread away from 0. In all subperiods the mean and the median of the distribution of the average AADs between subjective and γ^* -empirical beliefs are between 9 and 17 percentage points. Analogously, the mean and the median of the distribution of the average AADs between subjective and BNE beliefs are between 8 and 9 percentage points.³⁴ Figure 2 also shows that the distributions of average AADs between subjective and γ^* -empirical beliefs and between subjective and BNE beliefs changes in later periods.³⁵ In order to investigate what changes occur over time, Figure 3 plots for each round the median AAD between subjective and BNE beliefs and the median AAD between subjective and γ^* -empirical beliefs, and shows that the first remains stable through time at around 8 percentage points, while the second declines from 20 to 8 percentage points.

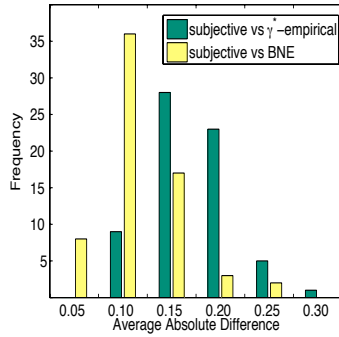
I draw two conclusions. First, neither BNE beliefs nor γ^* -empirical beliefs provide an appropriate proxy for subjective beliefs, since the differences between beliefs are not zero. Second, if we wish to use a proxy for subjective beliefs, γ^* -empirical beliefs are a poorer proxy than BNE beliefs in earlier rounds, and appear to become as good as BNE beliefs in later rounds, as information about a larger number of previous rounds becomes available.³⁶

Even if, as shown above, there are non-negligible differences between subjective beliefs and BNE beliefs or γ^* -empirical beliefs, either BNE beliefs or γ^* -empirical beliefs or both may still prescribe the same best response as subjective beliefs. If that were the case, eliciting subjective beliefs to determine which best response they prescribe, in order to later investigate


³⁴ A two-sided sign test test rejects at the 1% significance level, for all subperiods and for both BNE and γ^* -empirical beliefs, the null hypothesis that the median of the distribution of average AADs is zero.

³⁵ A Kolmogorov-Smirnov test rejects at the 1% level the hypothesis that the distribution of $\overline{AAD}_{1 \leq t \leq 5}(y, \eta)$ is the same as the one of $\overline{AAD}_{11 \leq t \leq 15}(y, \eta)$ and at the 5% level the hypothesis that the distribution of $\overline{AAD}_{1 \leq t \leq 5}(y, \nu)$ is the same as the one of $\overline{AAD}_{11 \leq t \leq 15}(y, \nu)$.

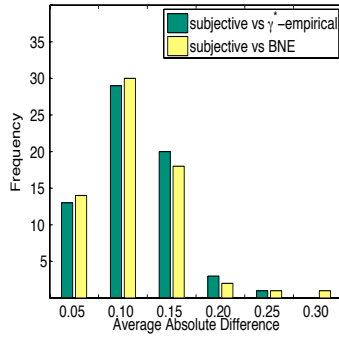
³⁶ Figure 3 also reveals that the difference between subjective and γ^* -empirical beliefs is smaller in the sessions with eight players than in ones with six or four players. However, the previous two conclusions still hold: neither γ^* -empirical beliefs nor BNE beliefs provide a good proxy for subjective beliefs, and γ^* -empirical beliefs are a poorer proxy than BNE beliefs in earlier rounds.



(a) Rounds 1-5




(b) Rounds 6-10




(c) Rounds 11-15

Figure 2: Absolute average difference (AAD) between subjective and γ^* -empirical and between subjective and BNE beliefs. Distribution across subjects of each subperiod average AADs: rounds 1-5, 6-10, 11-15.


choice data, could be avoided and a simple computation of BNE beliefs or γ^* -empirical beliefs would suffice. Therefore, the next question is whether subjective beliefs prescribe the same best response as BNE or γ^* -empirical




(a) All Markets



(b) Markets with 4 buyers and 4 sellers



(c) Markets with 3 buyers and 3 sellers



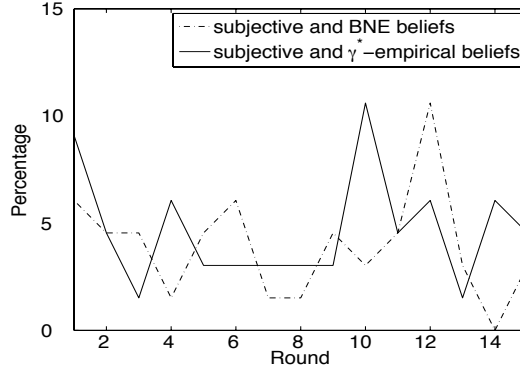
(d) Markets with 2 buyers and 2 sellers

Figure 3: Absolute average difference (AAD) between subjective, γ^* -empirical and BNE beliefs. Median AAD for each round.

beliefs.

Figure 4 depicts for each round the percentage of subjects for whom sub-

Figure 4: The percentage of subjects for whom subjective beliefs prescribe the same best response as BNE beliefs or the same best response as γ^* -empirical beliefs.



jective beliefs prescribe the same best response as BNE beliefs, or the same best response as γ^* -empirical beliefs.³⁷ Subjective beliefs usually prescribe a different best response from the one prescribed by BNE or γ^* -empirical beliefs: on average across rounds, the best response to subjective beliefs coincides with the one to γ^* -empirical beliefs in 5 percent of the cases and with the one to BNE beliefs in 4 percent of the cases.³⁸ When the best response to subjective beliefs and the best response to BNE beliefs differ, the absolute difference between them ranges from 1 cent to 5 dollars, with an average of 30 cents. Similarly, when the best response to subjective beliefs and the best response to γ^* -empirical beliefs differ, the absolute difference between them ranges from 1 cent to 5.57 dollars, with an average of 47 cents. Moreover, differences in best responses imply differences in expected payoffs.³⁹ Payoff differences range from \$0 to \$3.61 (average \$0.36) when the best response to subjective beliefs differs from the one to BNE beliefs, and from \$0 to \$4.89 (average \$0.45) when the best response to subjective beliefs differs from the

³⁷ Since in the experimental design bids and offers are expressed up to cents of a dollar, I round the best responses to cents before comparing them to each other.

³⁸ The percentages range from 2 to 11 percent without a tendency to increase.

³⁹ I compute the expected payoffs that would be generated by the best responses prescribed by each type of belief. The results are reported for cases in which subjective beliefs prescribe a best response different from the best response prescribed by BNE beliefs or γ^* -empirical beliefs. However, it is possible that even when different types of beliefs prescribe the same best response, the expected payoffs computed at the best response may differ. Considering all cases, the average payoff differences are \$0.35 when comparing subjective beliefs to BNE beliefs, and \$0.44 when comparing subjective beliefs to γ^* -empirical beliefs.

one to γ^* -empirical beliefs.

4.2 If subjects best respond, which beliefs do they best respond to?

I now turn to describing which bidding choices are actually made by the experiment participants and which beliefs can better explain those choices. More precisely, if subjects best respond, which beliefs do they best respond to? In order to answer this question, I will determine when the observed bid or offer is consistent with the best response prescribed by the different types of beliefs. The following argument illustrates the criterion used to verify consistency of observed choice with best-response behavior. Suppose subject i submits choice b . Let us denote the best responses prescribed by her subjective beliefs, γ^* -empirical beliefs, and BNE beliefs with $b_{BR,y}$, $b_{BR,\eta}$ and $b_{BR,\nu}$, respectively. It would be overly restrictive to simply verify whether $b = b_{BR,y}$, $b = b_{BR,\eta}$ and $b = b_{BR,\nu}$. In fact, it may be the case that $b \neq b_{BR,y}$ but the expected payoffs generated by b and $b_{BR,y}$, $\pi(b|y)$ and $\pi(b_{BR,y}|y)$, are equal. Then, we would still want to consider the observed choice as consistent with the best response to the subjective beliefs.

Thus, I implement the following procedure. First, I compute the payoffs that subject i would expect to receive by choosing b given her subjective beliefs y , her γ^* -empirical beliefs η , or the BNE beliefs ν , denoted with $\pi(b|y)$, $\pi(b|\eta)$ and $\pi(b|\nu)$, respectively. Second, I compute the payoff that subject i would expect to receive by choosing the best response prescribed by her subjective beliefs, $\pi(b_{BR,y}|y)$, the best response prescribed by her γ^* -empirical beliefs, $\pi(b_{BR,\eta}|\eta)$, or the best response prescribed by the BNE beliefs, $\pi(b_{BR,\nu}|\nu)$. I then compare $\pi(b|y)$ with $\pi(b_{BR,y}|y)$, $\pi(b|\eta)$ with $\pi(b_{BR,\eta}|\eta)$, and $\pi(b|\nu)$ with $\pi(b_{BR,\nu}|\nu)$.⁴⁰

Since the criterion is based on the expected payoff, it is useful to assess how the sensitivity of the expected payoff to the bidding choice depends on the private valuation. On the one hand, buyers with a low value and sellers with a high cost face a very small probability of trading and, provided they bid at or below value or offer at or above cost, the largest payoff they can expect to receive is most likely approximately zero. On the other hand, buyers with a high value and sellers with a low cost are likely to trade at a price determined by the bids and offers submitted by other subjects. Thus, the expected payoffs for subjects with an extreme private valuation are less sensitive to their choices compared to the expected payoffs of subjects

⁴⁰The comparison is done after rounding expected payoffs to the nearest cent.

with a mid-range valuation.⁴¹ Therefore, I will determine how often choice is consistent with best response to subjective, γ^* -empirical or BNE beliefs, and verify how the result changes depending on private valuations. Evidence that subjective beliefs perform relatively better than γ^* -empirical or BNE beliefs in explaining observed choice will be stronger if found when expected payoff is more sensitive to choice (i.e., for observations with mid-range valuation).

The two Venn diagrams in Figure 5 visualizes as proportional areas the number of observations for which observed choice is consistent with the best response to each type of belief. The intersection of two different sets represents the number of observations for which choice is consistent with the best responses prescribed by two different types of beliefs. The intersection of all three sets represents the number of observations for which choice is consistent with the best responses prescribed by all three types of beliefs. Both diagrams visualize how often choice is consistent (*i*) uniquely with the best response to subjective beliefs; (*ii*) with both best responses to subjective and γ^* -empirical beliefs; (*iii*) with both best responses to subjective and BNE beliefs; (*iv*) simultaneously with the best responses to all three types of beliefs. The occurrence of (*i*)–(*iv*) is verified in diagram (*a*) over the entire sample and in diagram (*b*) over the subsample with mid-range valuations. Valuations between 3 and 7 dollars are used as mid-range valuations.⁴²

In diagram (*a*), out of 990 observations, cases (*i*) – (*iv*) occur in 57, 29, 55 and 212 observations, respectively.⁴³ Therefore, choice is consistent with the best response to subjective beliefs, γ^* -empirical beliefs or BNE beliefs in 36%, 32% and 34% of the observations, respectively.⁴⁴ In order to

⁴¹ A low sensitivity of expected payoffs to bidding choices can be visually illustrated by the width of the interval of choices generating the same expected payoff (rounded to the nearest cent) as the one generated by the best response to subjective, γ^* -empirical or BNE beliefs. Irrespective of the type of beliefs used to compute best responses and expected payoffs, both the median width for observations corresponding to a low value or a high cost and the median width for observations corresponding to a high value or a low cost are larger than the median width for observations corresponding to a mid-range valuation. Specifically, the median width for observations corresponding to a low value or a high cost is about four times larger.

⁴² In a setting in which valuations are drawn independently from a uniform distribution over $\$[0,10]$, the range $\$[3,7]$ exclude those valuations for which the BNE model predicts that buyers with a low value and sellers with a high cost will have a zero probability of trading. See Rustichini, Satterthwaite and Williams (1994).

⁴³ Choice is not consistent with the best response to any of the beliefs in 53% of the observations.

⁴⁴ The percentage of observations in which choice is consistent with the best response to subjective beliefs appears to increase in later rounds, irrespective of the number of market participants. The same cannot be said for γ^* -empirical or BNE beliefs. Also,

assess whether the same results could be obtained by simply assuming that subjects choose bids and offers randomly, I compute, for each observation, the probability that a choice consistent with the best response to subjective beliefs, γ^* -empirical or BNE beliefs would be picked simply by chance.⁴⁵ Under this assumption, the probability that choice is consistent with the best response to subjective beliefs, γ^* -empirical or BNE beliefs would be 12.6%, 10% or 9%, respectively.

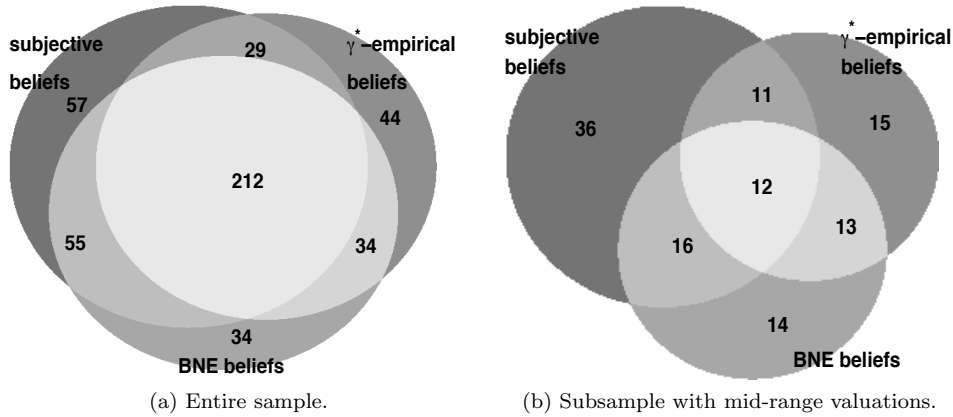


Figure 5: The number of observations for which choice is consistent with the best response to subjective, γ^* -empirical and/or BNE beliefs.

Restricting to the subsample with mid-range valuations (372 out of 990 observations), diagram (b) shows that choice is consistent with the best response to subjective beliefs, γ^* -empirical beliefs or BNE beliefs in 20.2%, 13.7% and 14.8% of the observations, respectively.⁴⁶ Therefore, while these percentages are smaller compared to the ones computed over the entire sample, subjective beliefs perform relatively better than γ^* -empirical and BNE

the percentage of observations in which choice is consistent with the best response to subjective beliefs appears to be higher in markets with more participants. Tables are omitted for brevity.

⁴⁵ The probability is defined as the ratio between the number of the choices that would generate an expected payoff equal to the one generated by the best response and the number of all possible alternatives from which a subject could randomly choose. Since in the experiment choices and payoffs are defined to the nearest cent, numbers are expressed in cents also when computing the probability. The possible alternatives from which to randomly choose are all values between \$0 and \$10.

⁴⁶ Choice is not consistent with the best response to any of the beliefs in 69% of the observations.

beliefs in explaining observed choice.⁴⁷

In the above discussion I argued that, in order to compare the ability of subjective, γ^* -empirical and BNE beliefs to explain observed choice, it is appropriate to restrict attention to the subsample of observations for which the expected payoff is more sensitive to choice itself. Within such a sample, subjective beliefs perform better than γ^* -empirical or BNE beliefs in explaining observed choice.

4.3 How accurate are subjective beliefs in predicting other subjects' choices?

The same criterion, according to which subjects are rewarded in the experiment (i.e., the quadratic scoring rule), can be used to compare subjective beliefs to γ^* -empirical and BNE beliefs in terms of the accuracy in the prediction of other subjects' choices.⁴⁸ For each participant I compute the average reward earned over all rounds and, subsequently, the median over all participants. Since subjects report separately their beliefs about the choices made by other buyers and by other sellers, evidence could be presented either separately or jointly for both beliefs. In the analysis that follows, I consider the joint accuracy of beliefs.

On average subjective beliefs are slightly more accurate than γ^* -empirical beliefs, but less accurate than BNE beliefs. The median average reward over all rounds is \$2.27 if generated by subjective beliefs, \$2.29 if generated by γ^* -empirical beliefs, and \$2.40 if generated by BNE beliefs. While the difference in accuracy between subjective beliefs and γ^* -empirical beliefs is not statistically significant, the one between subjective beliefs and BNE beliefs is statistically significant. Also, BNE beliefs are more accurate irrespective of the number of market participants. Table 3 shows the comparison between

⁴⁷ How do the results for observations with a mid-range valuation compare with the results for observations with an extreme valuation? Among observations corresponding to a low value or a high cost, the percentage of observations for which choice is consistent with the best response to subjective beliefs, γ^* -empirical beliefs or BNE beliefs is 69.5%, 69.5% and 71%, respectively. Within this subsample, the percentage of observations for which choice is consistent with the best response to either type of belief artificially increases due to the extremely low sensitivity of the expected payoff to choice. Among observations corresponding to a high value or a low cost, the percentage of observations for which choice is consistent with the best response to subjective beliefs, γ^* -empirical beliefs or BNE beliefs is 19.5%, 16.2% and 18.5%, respectively. For both subsamples (low value or high cost, high value or low cost) there is no evidence of a better performance of subjective beliefs, compared to γ^* -empirical or BNE beliefs, in explaining observed choice.

⁴⁸ Recall from Section 2.3 that subjects are rewarded based on the comparison of their subjective beliefs to the bidding behavior of other participants in the same auction market.

subjective and BNE beliefs.

Table 3: Accuracy of subjective beliefs versus accuracy of BNE beliefs.

	all markets	with 4 players	with 6 players	with 8 players
quadratic scoring rule	***	**	***	***
subjective beliefs	\$2.27	\$2.23	\$2.24	\$2.28
BNE beliefs	\$2.40	\$2.37	\$2.35	\$2.41
linear scoring rule				
subjective beliefs	\$0.82	\$0.83	\$0.83	\$0.82
BNE beliefs	\$0.83	\$0.84	\$0.78	\$0.83
average absolute difference	*			***
subjective beliefs	0.39	0.48	0.40	0.33
BNE beliefs	0.37	0.47	0.39	0.29

Note: *** significantly different from zero at 0.1 percent, ** significantly different from zero at 1 percent, * significantly different from zero at 5 percent.

The result that BNE beliefs would obtain a more accurate prediction than subjective beliefs could depend on the use of the quadratic scoring rule to measure accuracy. As alternative measures, I use the linear scoring rule and the average absolute difference. Evaluating accuracy with a linear scoring rule, BNE beliefs are more accurate than subjective beliefs in all markets except for the ones with six players, but the difference in accuracy is not statistically different from zero in any market. Measuring (in)accuracy by the average absolute difference (*AAD*), BNE beliefs are more accurate than subjective beliefs in all markets, but the difference in accuracy is statistically different from zero only for markets with eight players.⁴⁹

Despite being less accurate than BNE beliefs, do subjective beliefs become more accurate in later rounds? The median average accuracy of subjective beliefs does not appear to increase over time. Independently of the number of market participants, the median average accuracy of subjective beliefs in rounds 1-5 and in rounds 11-15 are not significantly different. Results are omitted for brevity.

⁴⁹ The average absolute difference (*AAD*) is defined as $\frac{\sum_{j=1}^6 |y_j - n_j|}{6}$, where y_j is the probability assigned to an opponent submitting a choice within interval j , and n_j is the percent of opponents submitting a choice within interval j .

4.4 Is there evidence of convergence of beliefs and/or choices to the BNE?

Given the definition of the BNE, finding evidence of choices and beliefs satisfying its definition requires finding evidence of convergence of both choices and beliefs to the BNE. Convergence in choices to the BNE requires choices to be consistent with the best response prescribed by BNE beliefs from a certain point in time onwards without deviation. Similarly, convergence in beliefs to the BNE requires subjective beliefs to prescribe a best response consistent with the best response prescribed by BNE beliefs from a certain point in time onwards without deviation.

In the collected experimental data there is no evidence of convergence to the BNE in either choices or beliefs. Thus, it is not possible to perform a comparison between convergent and non-convergent cases, which could provide insight into the conditions leading to convergence itself. However, it is possible to simply report the frequency of observations in which choice is consistent with the BNE best response and/or the best response to subjective beliefs is consistent with the BNE best response, as Table 4 does.

Measuring the frequency with which choice is consistent with the BNE best response depending on whether the best response to subjective beliefs is consistent or not with the BNE best response, and vice versa measuring the frequency with which the best response to subjective beliefs is consistent with the BNE best response depending on whether choice is consistent or not with the BNE best response, can provide an insight into assessing the sufficiency and/or necessity of the two conditions required for convergence to the BNE: convergence in beliefs and convergence in choices.

Table 4: Frequency of choice and/or best response to subjective beliefs consistent with the best response to BNE beliefs

choice and best response to BNE beliefs	best response to subjective beliefs and best response to BNE beliefs			
	consistent	not consistent	all	
consistent	286 44%	85% 49	15% 14%	335
not consistent	363 46%	55% 292	45% 86%	655
all	649	341		

The frequency with which choice is consistent with the BNE best re-

sponse changes from 14% to 44% when the best response to subjective beliefs changes from being non-consistent to being consistent with the BNE best response. The frequency with which the best response to subjective beliefs is consistent with the BNE best response changes from 55% to 85% when choice changes from being non-consistent to being consistent with the BNE best response. This is suggesting evidence that (i) the best response to subjective beliefs being consistent with the BNE best response is a necessary but not a sufficient condition in order for choice to be consistent with the BNE best response, and that (ii) choice being consistent with the BNE best response is a sufficient but not a necessary condition in order for the best response to subjective beliefs to be consistent with the BNE best response.

5 Conclusion

This paper has proposed a methodology to implement probabilistic belief elicitation in continuous-choice games. Representing subjective probabilistic beliefs about a continuous variable as a continuous subjective probability distribution, the methodology involves eliciting partial information about the subjective distribution and fitting a parametric distribution on the elicited data. As an illustration, the methodology was applied to a double auction experiment, where traders' beliefs about the bidding choices of other market participants are elicited. Elicited subjective beliefs are found to differ from proxies such as BNE beliefs and empirical beliefs, both in terms of the forecasts of other traders' bidding choices and in terms of the best-response bidding choices prescribed by beliefs. Moreover, elicited subjective beliefs are found to explain observed bidding choices better than BNE beliefs and empirical beliefs. By extending probabilistic belief elicitation beyond discrete-choice games to continuous-choice games, the proposed methodology enables to investigate the role of beliefs in a wider range of applications. Thus, the use of beliefs data can be applied to investigations in which traditionally only choice data are studied.

References

- [1] O. Armantier and N. Treich. Subjective probabilities in games: an application to the overbidding puzzle. *International Economic Review*, 50(4):1079–1102, 2009.
- [2] T.N. Cason and D. Friedman. Price formation in single call markets. *Econometrica*, 65(2):311–345, 1997.
- [3] J. Dominitz and C.F. Manski. Using expectations data to study subjective income expectations. *Journal of the American Statistical Association*, 92(439):855–867, 1997.
- [4] J. Engelberg, C.F. Manski, and J. Williams. Comparing the point predictions and subjective probability distributions of professional forecasters. *Journal of Business and Economic Statistics*, 27(1):30–41, 2009.
- [5] U. Fischbacher. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.
- [6] E. Haruvy, Y. Lahav, and C. Noussair. Traders’ expectations in asset markets: experimental evidence. *The American Economic Review*, 97(5):1901–1920, 2007.
- [7] C. Hommes, J. Sonnemans, J. Tuinstra, and H. Van De Velden. Coordination of expectations in asset pricing experiments. *Review of Financial Studies*, 18(3):955–980, 2005.
- [8] O. Kirchkamp and J.P. Reiß. Out-of-equilibrium bids in first-price auctions: wrong expectations or wrong bids. *The Economic Journal*, 121(557):1361–1397, 2011.
- [9] C.F. Manski. Identification of decision rules in experiments on simple games of proposal and response. *European Economic Review*, 46(4-5):880–891, 2002.
- [10] C.F. Manski. Measuring Expectations. *Econometrica*, 72(5):1329–1376, 2004.
- [11] Y. Nyarko and A. Schotter. An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3):971–1005, 2002.

- [12] T.R. Palfrey and S.W. Wang. On eliciting beliefs in strategic games. *Journal of Economic Behavior and Organization*, 71(2):98–109, 2009.
- [13] A. Rustichini, M. Satterthwaite, and S. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–1063, 1994.
- [14] E.E. Rutström and N.T. Wilcox. Stated beliefs versus inferred beliefs: A methodological inquiry and experimental test. *Games and Economic Behavior*, 67(2):616–632, 2009.
- [15] V.L. Smith, G.L. Suchanek, and A.W. Williams. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5):1119–1151, 1988.
- [16] J. Sonnemans, C. Hommes, J. Tuinstra, and H. Van De Velden. The instability of a heterogeneous cobweb economy: a strategy experiment on expectation formation. *Journal of Economic Behavior and Organization*, 54(4):453–481, 2004.

A Formulas for the Joint Densities of Order Statistics

In a market with n buyers and n sellers, the equation for the expected payoff of the buyer with value v and bid b employs the formula of the joint density of the n th and $(n + 1)$ th order statistics of the $2n - 1$ bids and asks. Denote the density functions of the bids and the asks with g_b and g_s respectively, and the cumulative density functions with G_b and G_s respectively. Consider a general formulation with m buyers and n sellers. Denote $k = m$. Denote with $f_{(k),(k+1)}$ the joint density of the m th and $(m + 1)$ th order statistics of the $m + n - 1$ bids and asks. Then:

$$\begin{aligned}
 f_{(k),(k+1)}(x, y) = & \\
 & n(n-1)g_s(x)g_s(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-1 \\ 0 \leq j \leq n-2}} \binom{m-1}{i} \binom{n-2}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-1-i} (1-G_s(y))^{n-2-j} + \\
 & +n(m-1)g_s(x)g_b(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-2 \\ 0 \leq j \leq n-1}} \binom{m-2}{i} \binom{n-1}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-2-i} (1-G_s(y))^{n-1-j} + \\
 & +(m-1)ng_b(x)g_s(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-2 \\ 0 \leq j \leq n-1}} \binom{m-2}{i} \binom{n-1}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-2-i} (1-G_s(y))^{n-1-j} + \\
 & +(m-1)(m-2)g_b(x)g_b(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-3 \\ 0 \leq j \leq n}} \binom{m-3}{i} \binom{n}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-3-i} (1-G_s(y))^{n-j}
 \end{aligned}$$

In the equation for the expected payoff of the seller with cost c and ask a , I employ the formula of the joint density of the $(n - 1)$ th and n th order statistics of the $2n - 1$ bids and asks. Denote the density functions of the bids and the asks with g_b and g_s respectively, and the cumulative density functions with G_b and G_s respectively. Consider a general formulation with m buyers and n sellers. Denote $k = m - 1$. Denote with $f_{(k),(k+1)}$ the joint density of the $(m - 1)$ th and m th order statistics of the $m + n - 1$ bids and asks. Then:

$$\begin{aligned}
f_{(k),(k+1)}(x, y) = & \\
& (n-1)(n-2)g_s(x)g_s(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m \\ 0 \leq j \leq n-3}} \binom{m}{i} \binom{n-3}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-i} (1-G_s(y))^{n-3-j} + \\
& +(n-1)mg_s(x)g_b(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-1 \\ 0 \leq j \leq n-2}} \binom{m-1}{i} \binom{n-2}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-1-i} (1-G_s(y))^{n-2-j} + \\
& +m(n-1)g_b(x)g_s(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-1 \\ 0 \leq j \leq n-2}} \binom{m-1}{i} \binom{n-2}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-1-i} (1-G_s(y))^{n-2-j} + \\
& +m(m-1)g_b(x)g_b(y) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-2 \\ 0 \leq j \leq n-1}} \binom{m-2}{i} \binom{n-1}{j} G_b(x)^i G_s(x)^j (1-G_b(y))^{m-2-i} (1-G_s(y))^{n-1-j}
\end{aligned}$$

I also include formulas for the marginal densities of order statistics. Denote the density functions of the bids and the asks with g_b and g_s respectively, and the cumulative density functions with G_b and G_s respectively. Denote $k = m$.

For the buyer:

$$\begin{aligned}
f_{(k)}(x) = & \\
& ng_s(x) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-1 \\ 0 \leq j \leq n-1}} \binom{m-1}{i} \binom{n-1}{j} G_b(x)^i G_s(x)^j (1-G_b(x))^{m-1-i} (1-G_s(x))^{n-1-j} + \\
& +(m-1)g_b(x) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-2 \\ 0 \leq j \leq n}} \binom{m-2}{i} \binom{n}{j} G_b(x)^i G_s(x)^j (1-G_b(x))^{m-2-i} (1-G_s(x))^{n-j}
\end{aligned}$$

For the seller:

$$\begin{aligned}
f_{(k)}(x) = & \\
& (n-1)g_s(x) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m \\ 0 \leq j \leq n-2}} \binom{m}{i} \binom{n-2}{j} G_b(x)^i G_s(x)^j (1-G_b(x))^{m-i} (1-G_s(x))^{n-2-j} + \\
& +mg_b(x) \sum_{\substack{i+j=k-1 \\ 0 \leq i \leq m-1 \\ 0 \leq j \leq n-1}} \binom{m-1}{i} \binom{n-1}{j} G_b(x)^i G_s(x)^j (1-G_b(x))^{m-1-i} (1-G_s(x))^{n-1-j}
\end{aligned}$$

B Parametric Fitting of Beliefs

Depending on (a) the number of intervals where positive probability is placed, and (b) whether the intervals, where positive probability is placed, are adjacent to each other, I fit:

1. a triangular distribution: when positive probability is placed (i) over a single interval or (ii) over two adjacent intervals,
2. the union of two triangular distributions: when positive probability is placed (i) over two non-adjacent intervals or (ii) over three intervals of which two but not all three are adjacent (iii) over four intervals, consisting of two disjoint pairs of adjacent intervals,
3. the union of three triangular distributions: when positive probability is placed over three interval none of which is adjacent to any other,
4. a unimodal beta distribution: when positive probability is placed over three or more intervals all adjacent to each other,
5. the union of a beta distribution and a triangular distribution: when positive probability is placed over more than three intervals, of which at least three, but not all, are adjacent to each other.

Here below I describe how the fitting is performed in each case. Figure 6 shows an example for each case, the most common being the case in which beliefs are fitted with a unimodal Beta distribution (approximately 90% of observations), followed by the case in which beliefs are fitted with a triangular distribution (approximately 7% of observations).

Triangular distribution. If positive probability is assigned to only one interval, $[l, r]$, then I assume that the support of the subjective distribution is $[l, r]$ and I fit a triangular distribution over it. The fitted isosceles triangle has base $r - l$ and height $\frac{2}{r-l}$. If positive probability is assigned to two adjacent intervals and if equal probability is assigned to each interval, then I assume that the support of the subjective distribution is the union of the two intervals and the fitted isosceles triangle has base 4 and height $\frac{1}{2}$.

If positive probability is assigned to two adjacent intervals and a higher probability is assigned to one interval than to the other, then I assume that the subjective distribution has the shape of an isosceles triangle and that its support contains entirely the interval that was assigned a higher probability and partly the other interval. If the subject assigns probability α and $1 - \alpha$ to the intervals $[y, y + 2]$ and $(y + 2, y + 4]$, respectively, where $\alpha < 0.5$, then,

the fitted isosceles triangle has base with endpoints $y + 2 - t$ and $y + 4$ and height $h = \frac{2}{t+2}$, with $t = \frac{2\sqrt{\frac{\alpha}{2}}}{1-\sqrt{\frac{\alpha}{2}}}$.⁵⁰

Union of two triangular distributions. If positive probability is assigned to two non-adjacent intervals, then I assume that the support of the subjective distribution is the union of the two intervals and I fit a triangular distribution over each interval. For example, probability α is assigned to the interval $[l_1, r_1]$ and probability $1 - \alpha$ to the interval $[l_2, r_2]$, then I assume that the support of the distribution is the union of $[l_1, r_1]$ and $[l_2, r_2]$. The isosceles triangle fitted over $[l_1, r_1]$ has base $r_1 - l_1$ and height $\frac{2\alpha}{r_1 - l_1}$. The isosceles triangle fitted over $[l_2, r_2]$ has base $r_2 - l_2$ and height $\frac{2(1-\alpha)}{r_2 - l_2}$.

If positive probability is assigned to three intervals of which two but not all three are adjacent, then I assume that the support of the subjective distribution is the union of the intervals and I fit two triangular distributions, one over the two adjacent intervals and another over the non-adjacent interval. For example, suppose that the intervals $[l_1, r_1]$, $[l_2, r_2]$ and $[l_3, r_3]$ are assigned probability α , β and $1 - \alpha - \beta$, respectively, and that $[l_1, r_1]$ and $[l_2, r_2]$ are adjacent to each other (with $l_2 > l_1$), while $[l_3, r_3]$ is not adjacent to any of the other intervals. Then one triangle is fitted over the union of $[l_1, r_1]$ and $[l_2, r_2]$ and one triangle is fitted over $[l_3, r_3]$, following the procedures already described for fitting one triangular distribution.⁵¹

If positive probability is assigned to four intervals, consisting of two disjoint pairs of adjacent intervals, then I also fit two triangular distributions. Each triangular distribution has support over a pair of adjacent intervals and the fitting is done following the procedure already described for fitting one triangular distribution.

Union of three triangular distributions. If positive probability is assigned to three intervals none of which is adjacent to any other, then I assume that the support of the subjective distribution is the union of the intervals and

⁵⁰ Analogously, if the subject assigns probability α and $1 - \alpha$ to the intervals $[y, y + 2]$ and $(y + 2, y + 4]$ respectively, where $\alpha > 0.5$, then I let $t = \frac{2\sqrt{\frac{1-\alpha}{2}}}{1-\sqrt{\frac{1-\alpha}{2}}}$. Then the subjective probability density function takes the form of a triangle with a base with endpoints y and $y + 2 + t$ and with a height $h = \frac{2}{t+2}$.

⁵¹ Therefore, if $\alpha < \beta$ then the first triangle has base with endpoints $l_1 + 2 - t$ and $l_1 + 4$ and height $h = \frac{2(\alpha+\beta)}{t+2}$ with $t = \frac{2\sqrt{\frac{\alpha}{2(\alpha+\beta)}}}{1-\sqrt{\frac{\alpha}{2(\alpha+\beta)}}}$. If $\alpha > \beta$, then the first triangle has base with endpoints l_1 and $l_1 + 2 + t$ and height $h = \frac{2(\alpha+\beta)}{t+2}$ with $t = \frac{2\sqrt{\frac{\beta}{2(\alpha+\beta)}}}{1-\sqrt{\frac{\beta}{2(\alpha+\beta)}}}$. The second triangle has base with endpoints l_3 and r_3 and height $\frac{2(1-\alpha-\beta)}{r_3 - l_3}$.

I fit a triangular distribution over each intervals, following the procedure already described for fitting one triangular distribution.

Unimodal Beta distribution. If positive probability is assigned to three or more intervals all adjacent to each other, then I fit a generalized unimodal Beta distribution over the intervals. The cumulative distribution function for a unimodal Beta distribution evaluated at x is denoted $Beta(x, \alpha, \beta, l, r)$, where α and β are shape parameters and l and r are location parameters determining the support for the distribution over the range $[l, r]$. If a subject does not assign positive probability to the right tail interval $[10.01, \infty)$, then the lower bound l of the support for the fitted Beta distribution will coincide with the left endpoint of the leftmost interval with positive probability and the upper bound r of the support will coincide with the right endpoint of the rightmost interval with positive probability and, therefore, the parameters l and r will be fixed. Thus, fitting the data with a Beta distribution requires solving the problem $\min_{\alpha, \beta} \sum_{j=1}^6 [Beta(r_j, \alpha, \beta, l, r) - G(r_j)]^2$, where $G(r_j)$ is the sum of the subjective probabilities assigned up to the interval with right endpoint r_j , inclusive.

If instead a subject assigns positive probability to the upper unbounded interval $[10.01, \infty)$, I let the location parameter r be a free parameter in the minimization of the least squares problem. I restrict r to lie within the most extreme value recorded in the data, which is 18.50. Thus, the problem becomes $\min_{\alpha, \beta, r < 18.50} \sum_{j=1}^6 [Beta(r_j, \alpha, \beta, l, r) - G(r_j)]^2$.

Union of a Beta distribution and a triangular distributions If positive probability is assigned to more than three intervals, of which at least three but not all are adjacent to each other, then I fit a unimodal beta distribution over the three or more adjacent intervals and a triangular distribution over the remaining one or two intervals. I follow the procedures already described for fitting a triangular distribution and a unimodal Beta distribution.

Goodness of fit is assessed by the average absolute deviation between the fitted and the elicited beliefs. In most cases the average absolute deviation is below 0.01, and in all cases below 0.06.

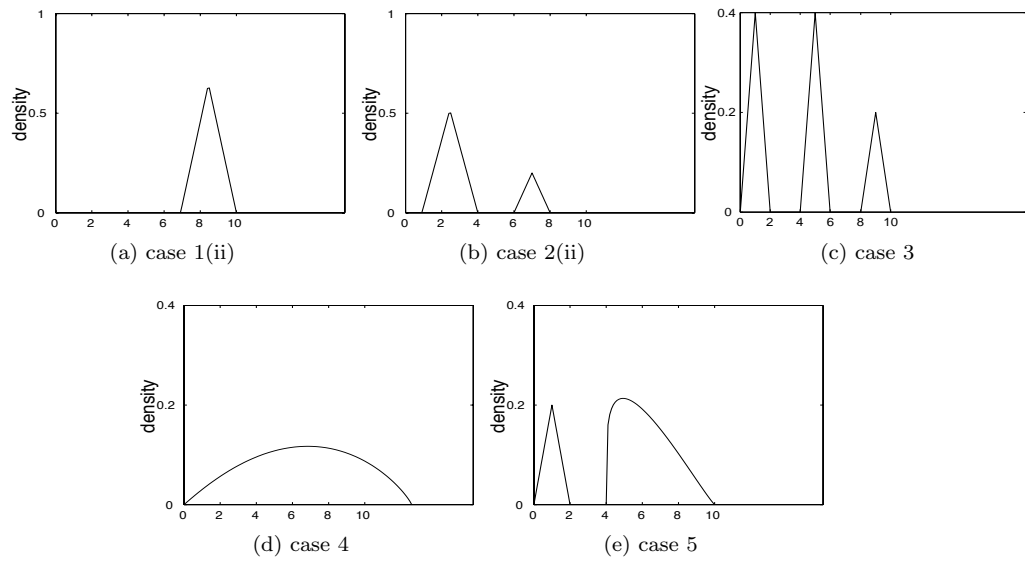


Figure 6: A selection of the fitting methods 1-5.

C Experimental Procedures

This Appendix contains additional information about the experimental procedures, including instructions, comprehension quiz, and subjects pool.

Eleven auction laboratory experiments were conducted between May 26th and June 5th 2009 in the Computer Laboratory of the Main Library at Northwestern University using the z-Tree experimental software. A total of 66 subjects, recruited by email and via campus posters among Northwestern University undergraduate students, participated. Each session lasted for approximately one hour, including the time for reviewing the instructions, and was identically administered by me personally. First, a welcoming speech was given, describing the structure and timing of the experiment. Then, a four-page copy of the instructions was distributed to all of the participants, who then had 10 minutes to read the instructions and ask questions. Students who wished to ask questions would raise their hand and I would answer their questions privately. Students were allowed to keep a copy of the instructions during the entire session.

After reading the instructions, the subjects completed a comprehension quiz⁵², which helped them better understand the rules of the game and test their comprehension. Being the quiz a self-test tool, all subject, irrespective of their performance on the quiz, were allowed to proceed in the experiment. Section C.2 includes descriptive results about the subjects' performance on the quiz, providing evidence that most of the participants understood well the rules of the game. Once the experiment was over, subjects filled in a questionnaire, consisting of questions about gender, age, major, year of graduation, familiarity with the game, and number of classes taken in several fields.⁵³ Subjects were paid individually in a sealed envelope.

⁵² The comprehension quiz was referred to as a 'guided example'.

⁵³ The participants were also given the option to leave specific comments about the way they played the game and/or general comments about the experiment.

C.1 Instructions (Market with 8 players)

Welcome! This is an experiment in decision-making. Please follow these instructions carefully. In this experiment we are going to simulate a market of a commodity in which some of you will be buyers and some of you will be sellers. You will use the computer to buy or sell. To ensure the best results for yourself, do not talk with other people in the room, and do not discuss your information with others at any point during the experiment.

A market consists of 8 participants, including you. Of the 8 participants, 4 are buyers and 4 are sellers.

You will play the game for 15 rounds. The participants in the market are randomly matched at the beginning of the experiment and don't change across rounds.

At the end of the session, the computer will randomly select one of the rounds, and you will be paid according to your performance in that round only.

In the first round the computer will randomly select your role as either a buyer or a seller. If the computer selects the role of buyer for you, then you will be a buyer in rounds 1 to 5, a seller in rounds 6 to 10, and a buyer again in rounds 11 to 15. Similarly, if the computer selects the role of seller for you, then you will be a seller in rounds 1 to 5, a buyer in rounds 6 to 10, and a seller again in rounds 11 to 15. You will see the information on the screen.

If you are a buyer, you will have \$10, and if you are a seller, you will have 1 unit of the commodity.

There are two ways to earn profits: trading, and forecasting.

HOW TO TRADE

Each buyer can purchase a single unit of the commodity from any seller, and each seller can sell a single unit of the commodity to any buyer.

How buyers and sellers submit their choices

Each buyer will see on the screen his or her **personal value V** of the unit, and each seller will see his or her **personal cost C** . You will not know the personal values and costs of other buyers and sellers.

Each buyer's personal value and seller's personal cost are assigned randomly: the computer determines them by a random draw of the 1000 numbers (to the nearest penny) between \$0.00 and \$9.99. Each number is equally likely. Each draw is independent from any other draw: the value for each buyer and the cost for each seller do not depend on the numbers drawn for other buyers or other sellers.

If you are a **buyer**, after being informed about your personal value, you'll be asked to **submit your bid B** . Your bid is the highest price, which you are willing to pay to purchase the commodity.

If you are a **seller**, after being informed about your personal cost, you'll be asked to **submit your offer O** . Your offer is the lowest price, which you are willing to accept to sell the commodity.

Type in the your choice and click on the OK button to submit it. Once you have clicked on OK you can't change your choice.

How the market price is determined

All trades in the market occur at a unique price: the market price P. What the market price turns out to be depends on the specific bids and offers submitted by you and by the other three participants in the market.

The computer gathers the submitted bids and offers, and it computes the market price in such a way that the numbers of units sold and purchased is the same, i.e., to equalize demand and supply. Later, we explain, using an example, how the bids and offers jointly determine the market price.

Who trades at the market price

If you are a **buyer**, whether you buy or not depends on the relation between your bid and the market price.

IF your bid \geq market price \rightarrow THEN you buy and pay the market price

Otherwise, you will not buy.

If you are a **seller**, whether you sell or not depends on the relation between your offer and the market price.

IF your offer \leq market price \rightarrow THEN you sell and receive the market price

Otherwise, you will not sell.

How trading profits are computed

Trading profits can be either positive, negative, or zero. For a buyer or a seller who does not trade, trading profits are zero. For a buyer or a seller who trades, trading profits depend on the market price P and on the personal value (if a buyer) or the personal cost (if a seller).

A **buyer** with a personal value V who buys a unit at market price P will earn:

$$\text{buyer's profit} = \text{personal value} - \text{market price} = V - P.$$

For example, a buyer with a personal value of \$3.08 who buys at market price \$2.58 earns a profit of $\$3.08 - \$2.58 = \$0.50$. A buyer will earn a negative profit (lose money) if she buys at a market price above her personal value. For example, a buyer with a personal value of \$3.08 who buys at market price \$3.20 earns a negative profit of $\$3.08 - \$3.20 = -\$0.12$, i.e., loses \$0.12.

A **seller** with a personal cost C who sells a unit at market price P will earn:

$$\text{seller's profit} = \text{market price} - \text{personal cost} = P - C.$$

For example, a seller with a personal cost of \$0.63 who sells at market price \$2.70 earns a profit of $\$2.70 - \$0.63 = \$2.07$. A seller will earn a negative profit (lose money) if she sells at a market price below her personal cost. For example, a seller with a personal cost of \$0.63 who sells at market price \$0.60 earns a negative profit of $\$0.60 - \$0.63 = -\$0.03$, i.e., loses \$0.03.

Whatever profit you make by trading is yours to keep. However, a buyer's initial cash and a seller's unsold unit do not count towards profits. If you earn a negative profit, such amount will be deducted from your show-up fee.

HOW TO FORECAST

Besides having the opportunity to trade, you will also be given the opportunity to make predictions. The questions will appear on your screen automatically.

You will be asked to forecast what the percent chance is that other participants in the market will make particular choices. Specifically, you will **assign a percent chance to each possible outcome**. A percent chance is a number between 0 and 100 percent, where 100 percent chance assigned to an outcome means that you are certain that such outcome is going to be the correct one, and 0 percent chance means that you are certain that such outcome is *not* going to be the correct one.

You will be paid based on the accuracy of your forecasts. Specifically, we will give you \$2 from which we will subtract an amount which depends on how inaccurate your prediction was.

Suppose that you are a seller and that you have to forecast the percent chance that any of the 4 other buyers in the market will choose one of two possible alternatives X and Y . Suppose that your forecasts are p_X and p_Y respectively.


For every other buyer who chooses X , the amount $\frac{1}{4}\{(\frac{p_X}{100} - 1)^2 + (\frac{p_Y}{100})^2\}$ is subtracted from the initial \$2. For every other buyer who chooses Y , the amount $\frac{1}{4}\{(\frac{p_X}{100})^2 + (\frac{p_Y}{100} - 1)^2\}$ is subtracted from the initial \$2.

The worst you can do is to assign a 100 percent chance to X when all 4 buyers choose Y instead. In this case your payoff is $\$2 - \frac{1}{4}\{0 + 4[(1)^2 + (0-1)^2]\} = \0 .

The best you can do is instead to assign a 100 percent chance to X when in fact all 4 buyers choose X . In this case your payoff is $\$2 - \frac{1}{4}\{4[(1-1)^2 + (0)^2] + 0\} = \2 .

Since your predictions are made when you don't know what other participants have chosen, **the best thing you can do to maximize the expected size of your payoff is to simply state what you think**.

Example Suppose that a market consists of 8 participants. Of the 8 participants, 4 are buyers and 4 are sellers. Suppose that the 4 buyers submitted bids of \$2.80, \$4.20, \$5.00, \$9.50, and the 4 sellers submitted offers of \$1.00, \$3.50, \$5.30, \$8.50. The bids and offers are represented in this figure.



The red line is the demand curve. The steps of the curve represent the bids submitted by the buyers: one buyer is willing to pay at most \$9.50, another buyer at most \$5.00, another buyer at most \$4.20, and another buyer at most \$2.80.

The blue line is the supply curve. The steps of the curve represent the offers submitted by the sellers: one seller is willing to accept no less than \$1.00, another seller no less than \$3.50, another seller no less than \$5.30, and another seller no less than \$8.50.

The two curves intersect within the range of prices between \$4.20 and \$5.00. The market price is set at the midpoint of this range, at $\frac{\$4.20 + \$5.00}{2} = \$4.60$.

Who will trade at the market price? Those buyers willing to pay a price higher than or equal to \$4.60, and those sellers willing to accept a price lower than or equal to \$4.60. Therefore, the buyers who submitted bids of \$5.00 and \$9.50, and the sellers who submitted offers of \$1.00 and \$3.50 will trade. On the other hand, the buyers who submitted bids of \$2.80 and \$4.20 will not buy (for them the market price is too high!) and the sellers who submitted offers of \$5.30 and \$8.50 will not sell (for them the market price is too low!). Demand equals supply at the market price: two buyers buy and two sellers sell, i.e., two units are bought and two units are sold.

What are the profits? The buyers and the sellers who don't trade earn zero profit. The buyers who trade earn a profit equal to the difference between their value and the market price. For example, if the buyer with bid of \$5.00 had a value of \$5.64, then he or she would earn $\$5.64 - \$4.60 = \$1.04$. The sellers who trade earn a profit equal to the difference between the market price and their cost. For example, if the seller with offer of \$3.50 had a cost of \$2.98, then he or she would earn $\$4.60 - \$2.98 = \$1.62$.

The market price can also be quickly determined in the following way. Let's sort all 8 bids and offers in increasing order, irrespective of whether it's a bid or an offer. We have: \$1.00, \$2.80, \$3.50, \$4.20, \$5.00, \$5.30, \$8.50, and \$9.50. The middle two (i.e. the 4th and 5th highest numbers) are \$4.20 and \$5.00. *The market price is set at the midpoint between the middle two bids and offers*, i.e. $\frac{\$4.20 + \$5.00}{2} = \$4.60$.

Priority for trading is always given to sellers whose offers are smallest and to buyers whose bids are largest. If this criterion does not make demand and supply equal, then a coin toss determines who trades among the remaining participants.

C.2 Performance in the Comprehension Quiz

After reading the instructions, subjects had 10 minutes to complete a comprehension quiz, which consisted of an example of an auction followed by four questions. Each question allowed for two attempts. After an answer was submitted, the computer announced whether the answer was correct. If incorrect, a short explanation and a hint appeared on the screen, and the subject was invited to answer the question again. After the second attempt, the computer announced once again whether the answer was correct. If incorrect, a longer explanation appeared on the screen. Subjects received \$5 if they answered correctly all four questions, regardless of whether the answers were correct the first or second time around.

I introduced the comprehension quiz and rewarded correct answers in order to provide the subjects with a stronger incentive to concentrate on reading and understanding the rules of the game. For the same reason, each question allowed for two attempts and correct answers were remunerated whether they were correct the first or second time around. Additionally, the subjects were allowed to consult the printed instructions while answering the questions. The data collected from the comprehension quiz provide evidence that most of the participants understood the rules of the game.

Table 5 shows details about the participants' performance in the comprehension quiz. The first question required sorting in increasing order the bids and offers in the auction example. The second question required determining the market price. The third question required determining whether each trader in the example would trade or not at the market price. A subject's answer was correct if she correctly stated whether each trader would trade or not. The fourth question required determining the profit of a selection of traders in the market. For each trader three alternatives were shown and only one of them was correct. A subject's answer was correct if she correctly stated the profit of each trader.

C.3 Subject Pool

The sample contains similar proportions of freshmen, sophomore, junior and senior students. Among participants, 9%, 18%, 49%, and 24% have a major in the Humanities, the Sciences, the Social Sciences (excluding Economics), and Economics, respectively. According to the most recent available statistics provided by the Office of the Registrar, in the 2007 graduating class the proportions were 18%, 25%, 47% and 10%, respectively. If the 2007 graduating class is considered as the population of Northwestern undergraduates,

Table 5: Performance in the comprehension quiz.

topic		1st attempt		2nd attempt*	
		No.	%	No.	%
ordering	<i>correct</i>	66	100		
price	<i>correct</i>	57	86	8	89
	<i>incorrect</i>	9	14	1	11
trades	<i>correct</i>	56	85	7	70
	<i>incorrect</i>	10	15	3	30
profits	<i>correct</i>	16	24	20	42
	<i>incorrect</i>	48	73	25	52
	<i>no answer</i>	2	3	3	6

* if 1st attempt is incorrect

the following differences between the sample and the population stand out: (1) students with a major in the Social Sciences are over-represented, (2) students with a major in the Humanities are under-represented, (3) the over-representation of students with a major in the Social Sciences coincides with an over-representation of students with a major in Economics. Finally, female students are over-represented in the sample compared to the population (71% versus 53%).