EXPERIMENTS ON DECISION MAKING AND AUCTIONS

A Dissertation

by

ELIZBETH ANN WATSON

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

August 2007

Major Subject: Economics

EXPERIMENTS ON DECISION MAKING AND AUCTIONS

A Dissertation

by

ELIZBETH ANN WATSON

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Approved by:

Chair of Committee,	Rajiv Sarin
Committee Members,	Brit Grosskopf
	Theodore Turocy
	Nehemia Geva
Head of Department,	Amy Glass

August 2007

Major Subject: Economics

ABSTRACT

Experiments on Decision Making and Auctions. (August 2007)

Elizabeth Ann Watson, B.G.S., University of Louisiana at Lafayette

Chair of Advisory Committee: Dr. Rajiv Sarin

Experimental economics is often called upon to inform theory and aid in explaining real world behavior. As such it is important to carefully design laboratory experiments to test the validity of new theories and to re-examine results that demonstrate robust anomalies of classic theory. My first aim is to design an experiment that will allow me to test the propensity of subjects to use Case-Based Decision Theory (from now on referred to as CBDT). I carefully design a setting in which the predicted choices of CBDT are deterministic and unique to CBDT (i.e. different from the predicted choices of other relevant decision making processes). I examine how well CBDT organizes subject choices when subjects are asked to make thirty independent decisions each with a fixed and given history. I find that some subjects do appear to be using the information given to them in the form of Case-Based Decision making.

My second goal is to revisit traditional first-price private values auction experiments with the idea of making the price-probability trade-off, the central consideration in auctions, more salient to subjects. I approach this in two different ways. First, I use a custom-designed graphical interface which displays all results both visually and numerically. In this treatment I find that subjects bid more aggressively than predicted by risk-neutral Bayes-Nash equilibrium. Second, I restructure the presentation of the idiosyncratic reservation value. Subjects are now engaging in some economic behavior and earning their total consumer surplus each period. This differs from traditional first-price private values auction experiments in which subjects only earn a payoff if they win the auction. Here I observe that market prices in a sealed-bid

implementation are significantly lower than those reported using the standard auction set-up.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER I
INTRODUCTION

Experimental economics is often called upon to inform theory and aid in explaining real world behavior. As such it is important to carefully design laboratory experiments to test the validity of new theories and to re-examine results that demonstrate robust anomalies of classic theory. Our first aim is to design an experiment that will allow us to test the propensity of subjects to use Case-Based Decision Theory (from now on referred to as CBDT). This theory remains, as yet, unexplored in experimental economics. We carefully design a setting in which the predicted choices of CBDT are deterministic and unique to CBDT (i.e. different from the predicted choices of other relevant decision making processes). We examine how well CBDT organizes subject choices when subjects are asked to make thirty independent decisions each with a fixed and given history. Our second goal is to revisit traditional first-price private values auction experiments with the idea of making the price-probability trade-off, the central consideration in auctions, more salient to subjects. We approach this in two different ways. First, we use a custom-designed graphical interface which displays all results both visually and numerically. Second, we restructure the presentation of the idiosyncratic reservation value. Subjects are now engaging in some economic behavior and earning their total consumer surplus each period. This differs from traditional first-price private values auction experiments in which subjects only earn a payoff if they win the auction.

In Chapter II, "An Experiment on Case-Based Decision Making", we design an experiment to investigate the disposition of subjects to use case based reasoning as predicted by Case-Based Decision Theory. For many decades the driving theory behind decision making has been that decision makers attempt to maximize their expected utility. Even though this theory is widely known and commonly used, both in theoretical

---

This dissertation follows the style of *Experimental Economics*.

and applied work, not all decision problems under uncertainty easily lend themselves to analysis with expected utility theory. CBDT was developed to describe behavior in limited information environments where expected utility theory is not a viable process. CBDT assumes that decision makers use their past experience to inform current choices and that they weight the outcomes of those past experiences using a similarity comparison between their current situation and the ones from the past. In order to obtain predictions and make comparisons to actual subject choices we use a feature-based similarity function to assess the similarity between current situations and the ones from the past. The theory is implemented by having the subjects act as monopolists making production decisions for thirty independent markets. Subjects receive a different "history" for each of the 30 markets. Inducing the "history" allows us to control what the subjects know and enables us to predict the exact choices the subjects would make if they were indeed Case-Based Decision Makers. The results reveal that some subjects do engage in this form of Case-Based decision making.

In Chapter III, "Framing the First-Price Auction", we investigate the hypothesis that the non-isomorphism between sealed-bid and Dutch laboratory auctions arises from framing and presentation effects. We carefully construct a graphical interface that is identical across the two treatments, except for the differences necessary for subjects to be able to submit a bid in the sealed-bid format and to allow a clock to tick down in the Dutch. In addition, we present the results of the auction both visually (as rectangles that increase in size as the payoff increases) and numerically (in tabular form with a list of actual payoffs) to help emphasize the price-probability trade-off that is central to auction theory. The graphical interface uses the standard method for presenting idiosyncratic reservation values: each subject is given a "resale value" at which he can sell the object back to the experimenter upon winning the auction. A subject's earnings are then computed as resale value minus the price paid in the auction if he wins the auction, and his earnings are zero if he does not win the auction. We find, however, the difference in bidding behavior across the two treatments persists, even when subjects participate in sessions lasting 60 periods. Furthermore, subjects consistently bid significantly more aggressively than is predicted by risk-neutral Bayes-Nash equilibrium. This result indicates that employing a graphical interface to present prices and payoffs does not make the price-probability trade-off more salient to subjects.

In Chapter IV, “Reservation Values in Laboratory Auctions: Context and Bidding Behavior”, we investigate the hypothesis that the standard presentation of reservation values distracts from the price-probability trade-off and results in bidding behavior that is more aggressive than that predicted by risk-neutral Bayes-Nash equilibrium. Using the same graphical interface we introduced in Chapter III, we present the reservation value as an “outside price” rather than a “resale value”. The “outside price” is a price at which the bidder can purchase a close substitute outside the auction market. Subjects are instructed that they will purchase one unit of commodity each period either in the auction at the price they bid or outside the auction at their outside price. Subjects’ earnings are calculated as their total consumer surplus, regardless of whether they win the auction. This differs from the typical earnings calculation in laboratory auctions as described above. Discussions following classroom experiments in which students participated in “traditional” auction experiments suggest that a central consideration in their decision making was to “try to win” the auction, because that is the only way to receive a positive payoff. We hypothesize that subjects earning their total consumer surplus each period lessens the focus on the “just win” paradigm and allows the price-probability trade-off to become more salient. We find that bidding behavior in the outside price frame is significantly less aggressive than bidding behavior in the resale value frame for first-price sealed-bid auctions, even though the risk neutral Bayes-Nash equilibrium remains unchanged across the two treatments.

CHAPTER II
AN EXPERIMENT ON CASE-BASED DECISION MAKING

INTRODUCTION

For many decades the driving theory behind decision making has been that decision makers attempt to maximize their expected utility. Even though this theory is widely known and commonly used, both in theoretical and applied work, not all decision problems under uncertainty easily lend themselves to analysis with expected utility theory. In order to apply expected utility theory a decision maker (DM) needs to know all possible states of the world and the outcomes associated with them. In many decision problems under uncertainty, states of the world are neither naturally given nor can they be simply formulated. Oftentimes, even a comprehensive list of all possible outcomes is neither readily available nor easily imagined. How should a DM choose in such situations? One alternative model of the decision making process that has been formulated is Case-Based Decision Theory (Gilboa and Schmeidler, 1995). The basic premise behind Case-Based Decision Theory (henceforth CBDT) is that a DM uses her past experience to help evaluate current choices, rather than constructing beliefs about certain states of the world occurring. CBDT suggests that agents believe that taking similar actions in similar situations will result in similar outcomes.

While CBDT has been the focus of some applied work in economics (e.g. in financial markets, Guerdjikova, 2002 and 2003; housing markets, Gayer et al., 2004; and capacity planning, Jahnke et al., 2005), little experimental evidence has been gathered by economists to test the theory. The purpose of this paper is to design an experiment to test the predictive power of CBDT for situations where a DM knows very little about the underlying environment and would have difficulty imagining all possible states or outcomes relevant for the decision. Although similarity is central to the decision making process of a Case-Based decision maker, CBDT itself is fairly silent about the details of the similarity function. CBDT shows that similarity is derived from preferences and therefore could be unique to each individual. However, much of the theoretical work takes similarity as given and the same across individuals and applied work simplifies the analyses further by assuming that similarity between situations is always equal to one or by treating similarity as a binary variable in which all situations are either identical (similarity equal to one) or completely different (similarity equal to zero).

In this paper, we assume a specific form of the similarity function that is widely used in the psychology literature: Feature based similarity (Tversky, 1977). We argue that this specific similarity function is plausible given the description of the environment the subjects in our experiment face. In the experiment, the similarity function can take on values between, and including, 0 and 1. With this implementation we are able to predict the choices that a Case-Based DM would make and can compare them with actual behavior of participants in the experiment. This allows us to test whether our subjects employ case based reasoning (coupled with feature based similarity) in their decision making. Note that our paper is not aimed at deriving individual similarity functions given the choices people make. Such an objective would require a completely different experimental design.

In order to test the robustness of the predictions of CBDT we vary two things. First, we contrast behavior in situations when CBDT can be used to situations when it cannot be used. This is achieved by comparing choices in situations in which the DM has some information about the current problem to when she does not. In the latter the agent has no information on which to base a similarity comparison. Second, we analyze if economic framing of the decision task makes the subject more prone to use CBDT. We do this by framing the decision task either as a choice of output problem for a monopolist, or by framing it as an abstract decision problem. Our results reveal some support for CBDT. Predicted choices coincide more often with actual choices when participants are given information about the current situation alongside a "history." Framing does not appear to matter for the ability of CBDT to explain subjects' choices. However, we also find support for a simple heuristic being used in the limited information environments we analyze. In particular, equally many participants seem to be guided by CBDT as by the heuristic which assumes that the action with the highest payoff in memory will give the highest current payoff. The remainder of the paper is organized as follows. Section 2 introduces CBDT and the notion of similarity. It gives a motivation for the specific functional form of similarity that we are using. Section 3 explains the experimental design. Section 4 states our hypotheses. Section 5 discusses the experimental results and Section 6 provides concluding remarks. Appendix A contains a sample set of instructions.

BACKGROUND

Case-Based Decision Theory

Often times, decisions have to be made with very little information about the underlying environment. Consider the following example where only limited information is available. The manager of a firm is looking to hire a technician for her IT department. Her choice set is given by the applicants for the job. She knows that she is looking for a technician who is highly skilled at computer networking, fluent in visual basic, and can lead and motivate the rest of the IT team. She, however, does not know how each candidate would perform if hired. For example, a candidate may be highly skilled in all of the requisite areas, but it may turn out that he is going through a painful divorce and is continuously late for work and often depressed. Or a candidate may display great leadership and organizational skills but may turn out to be very poorly skilled at computer networking. The more she thinks about it the more she realizes that other problems may also occur and that she has no way of knowing what they might be or how they might affect the company. The manager is facing uncertainty, ambiguity, and a lack of information on several measures.

There are several difficulties with fitting this problem into the framework of expected utility. First, the states of the world do not naturally suggest themselves. Second, imagining all of the possible outcomes for each action is not a trivial task. This would amount to imagining every possible thing that could happen once an employee is hired and imagining all of those things for every possible employee. Lastly, even after an action has been taken, the outcome may not reveal the realized state of the world or whether the action chosen was optimal or not. For situations like this, when DMs cannot be guided by expected utility theory, CBDT has been suggested as an alternative.

The basic premise behind CBDT is that a DM uses her past experiences (or the experiences of others) to help evaluate current choices, rather than relying on beliefs about certain states of the world occurring. In the above example, if each job candidate provided references, then the manager could use the candidate's previous performance to help assess how each candidate would perform if she was hired. In order to help evaluate past outcomes, an agent possesses a similarity function that measures how similar the current situation is to a past situation. The agent is assumed to compare the current situation to all available past situations. The more similar the current situation is to a past situation the more heavily the agent will weight the outcome of that past situation. The agent is then assumed to choose the action that maximizes the sum of the similarity weighted outcomes of all past situations.

Formally, a Case-Based DM is assumed to have a memory, *M*, consisting of a set of cases. A case is formed by a problem or situation, *q*, the action chosen in that situation, *a*, and the utility, *u*(*r*), gained from choosing action *a* in situation *q* and receiving the result *r*. The Case-Based DM is shown to possess a similarity function, *s*, which evaluates the similarity between the current situation and any past situation. When confronted with a problem, *p*, the Case-Based DM chooses the action *a* that maximizes the following expression

$$U(a) = U_{p,M}(a) = \sum_{(q,a,r)\in M} s(p,q)u(r) \tag{2.1}$$

where *s*(*p*, *q*) measures the similarity between the current situation, *p*, and some past situation, *q*. When considering any action *a*, the Case-Based DM only concerns herself with past situations in which that particular action was chosen. If in a past situation action *a* was not chosen, then the result and the subsequent utility obtained in that situation are ignored[1].

[1] For a version of case-based decision theory that allows the agent to use such information, see Gilboa and Schmeidler (1997).

In other words, a case-based DM adds up, over all cases in her memory, the similarity weighted utility that each action has received[2]. Whichever action has the largest sum is the action that is predicted to be chosen in the current situation. Note that this means that any action that has never been chosen in the past will never be chosen in the current situation unless all actions chosen in the past have resulted in negative utility[3].

Similarity

Let's reconsider the example of the IT manager. Assume the manager received equally outstanding references for two candidates (Bob and Betty) and she must decide between them. Suppose Bob's reference was from a previous job in which he designed and maintained web pages. However, Betty's reference was from a previous job in which she was the head of a large corporation's IT department and was responsible for maintaining all networking. It seems obvious that the similarity between the past situation in which Betty was hired and the manager's current one is greater than the similarity between the past situation in which Bob was hired and the current one. Therefore Betty's outstanding recommendation will receive more weight than Bob's and the manager will choose to hire Betty.

While the above example seems intuitive, we have to consider a specific form for the similarity function in order to obtain actual choice predictions from CBDT. While the notion of similarity has not been widely studied in the economics literature, it has been the subject of much discourse in the psychology literature (see Goldstone and Son, 2005 for an overview)[4]. Most of the models of similarity can be divided into two groups: geometric models and feature-matching models.

Geometric models assume that the objects that are being evaluated can be represented in some *n*-dimensional space. The (dis)similarity between two objects is then calculated as some measure of distance between the two objects. The most typical

[2] CBDT does not make any distinction between an action that resulted in zero utility and one that simple was not chosen, since zero utility is taken as the default aspiration level.
[3] In such a scenario a case-based DM is assumed to randomly choose an action from the set of available actions that have not yet been chosen.
[4] See Rubenstein (1988) and Sarin and Vahid (2004) for previous applications in economics.

measures are the Euclidean distance and the City-Block distance. While these models have desirable mathematical properties, experimental studies have shown that they do not do well in representing how subjects actually perceive similarity (see Goldstone and Son, 2005 and Tversky, 1977 for overviews)[5].

In response to some of these findings, Tversky (1977) developed a model of similarity that assumes that objects can be described by a set of features and that similarity is defined over the features that two objects have in common and those that they do not have in common. This allows an agent much more flexibility in measuring the similarity between two objects, and also allows similarity to be measured among objects that do not naturally lend themselves to placement in some n-dimensional space. Specifically, Tversky's model says that similarity is calculated in the following manner. Let A be the set of features associated with object a and let B be the set of features associated with object b. The measure of how similar a is to b is given by $s(a, b) = \theta f(A \cap B) - \beta f(A - B) - \gamma f(B - A)$, where θ, β and γ are positive constants. Thus, the similarity between two objects, a and b, is a function of the set of features the two have in common, those that a has but b does not, and those that b has but a does not. This allows the measure of similarity between a and b to be positive or negative and it allows the similarity between a and b to differ from the similarity between b and a.

In order to remain within the bounds of Case-Based Decision Theory, we use a simplified version of Tversky's feature based similarity function when calculating its predictions in our setup. The underlying axioms of CBDT imply that the similarity function a Case-Based DM uses can only take on values between 0 and 1. To achieve this we choose $\beta = \gamma = 0$ to prevent the similarity function from taking negative values. We assume that all features are given the same weight when calculating similarity, i.e. it

[5] In particular, several of the properties of the geometric models are consistently violated by experimental subjects. First, it has been shown that the identity property does not hold, i.e. subjects do not always perceive an object as identical to itself (see Podgorny and Garner, 1979). Second, actual similarity evaluations are not always symmetric (see Holyoak and Gordon, 1983; and Ortony et al., 1985). For instance, a subject reporting that domestic cats are very similar to tigers does not necessarily indicate that the same subject will report that tigers are very similar to domestic cats. Lastly, the triangle inequality often does not hold, nor does transitivity (Tverksy and Gati, 1982). Finding objects A and B very similar and objects B and C very similar does not necessarily indicate that the subject will find objects A and C very similar.

is not more important to have feature 1 in common than it is to have, say, feature 2 in common. We let f count the number of features two objects have in common. CBDT assumes that if two objects are identical then the similarity between them is equal to 1 and that it is equal to 0 if they have no features in common. To ensure this we set the parameter θ equal to the reciprocal of the maximum number of features two objects could possibly have in common. With the additional assumption of $u(r) = r$ the decision problem of a Case-Based DM can be formulated as,

$$\max_a U(a) = U_{p,M}(a) = \sum_{(q,a,r)\in M} \theta f(p \cap q) r \tag{2.2}$$

EXPERIMENTAL DESIGN

Given that our experiment is the first to investigate Case-Based decision making in economics, there are several aspects of our design that we need to discuss in detail. CBDT aims to guide behavior in situations where a DM never faces the same problem twice. As such we decided to confront subjects with 30 different situations. In the instructions, and in the questionnaire following the instructions, we made sure that the subjects understood that the situations were independent from one another and that decisions made for one situation would not affect any other situation. It was also emphasized that the decisions and resulting outcomes are independent across participants.

Central to CBDT is the idea that a current case is compared to the cases in the decision maker's "memory." Therefore, control over the memory is important in order to make accurate predictions about the behavior of a true Case-Based DM. We did not want subjects to "build" a memory over the 30 different situations they faced during the experiment. We therefore decided to induce a separate "history" for each of the 30 situations, i.e., a different memory for each situation. For each situation, we displayed four "scenarios" or cases. The scenarios that were displayed were different for each situation. In each of the four scenarios, a different choice was made and that choice and the associated payoff that would have been earned were displayed. Thus, the memory of each subject, for each situation, consisted of 4 cases, their descriptions, actions chosen and resulting payoffs[6].

We ensured that in each situation each action was a possible choice for a Case-Based decision maker by having each memory contain one observation of each possible action. A Case-Based decision maker would calculate the similarity between the current situation and the four scenarios and use the similarity measure to weight the outcome that was “received” in that scenario. The choice with the highest similarity-weighted payoff is the one that a Case-Based decision maker would choose. While the scenarios and choices were randomly generated, the final display was not random. We displayed four scenarios and their corresponding choices such that predictions for CBDT were different from those of expected utility theory[7]. It was also the case that the predictions for CBDT differed from the optimal (profit maximizing) choice in all but 5 of the 30 markets. This was done so that there would be no confusion as to whether subjects were really using CBDT to make their choices or if they were simply maximizing their earnings. As such, in this design, any subject who does use CBDT to make their choices will be (unconsciously) doing so at a cost. This could be seen as making the case for CBDT stronger.

Since there is no standard practice in implementing CBDT in a laboratory environment, we test the robustness of our findings by varying our experimental implementation along two dimensions. First, we vary whether the situations are framed in an economic context or not. Second, we vary whether subjects are given information about the current situation in addition to the four scenarios (cases) in the induced memory or not. Table 2.1 summarizes our 2×2 between-subject design.

[6] See Appendix A for a full set of the instructions. We emphasized in the instructions that these choices were not chosen to maximize earnings, but were selected randomly with the constraint that each possible choice was chosen exactly once. We further emphasized that these were hypothetical scenarios and that subjects would not be paid for these choices. They would only be paid for choices they made during the experiment.

[7] This was done so that it was possible to distinguish between subjects behaving as if they were case-based DMs and those behaving as if they were maximizing expected utility. Of course, the latter was not calculable in our setting.

Table 2.1 The 2x2 experimental design.

	Information	
Frame	Past + Current	Past Only
Monopoly	39 subjects	30 subjects
Abstract	31 subjects	32 subjects

The framing treatments were done to investigate whether certain frames trigger subjects to think more carefully about the similarity between their current situation and the ones from their "memory" or whether they hinder such comparisons. For the treatments with a frame we were looking for an individual decision making task with an economic content. We chose a monopoly situation.

The information treatments were done in order to make sure that subjects were not just "accidentally" making a choice in line with CBDT. Given that there are only four available choices even a subject acting randomly would appear to choose the action predicted by CBDT approximately 25% of the time. Furthermore, if we see subjects' choices aligned with CBDT predictions in roughly the same proportion over both information treatments then we should conclude that CBDT is not actually explaining their choices but rather that this is happening just by coincidence. However, if we find that the proportion of subject choices explained by CBDT is significantly higher in the treatments with the current information than without, then we can conclude that subjects are making use of this extra information and that they appear to be doing so in the manner suggested by Case-Based decision theory[8].

In the treatments with an economic frame, subjects were told that they were a monopoly firm that had to make production choices for 30 different markets[9]. The production choice could be one of four values, 50, 100, 150, or 200 units. In the treatments in which information about the current situation was given, this information was presented as a marketing report that included information on a set of market conditions. Subjects were informed that their profits would depend on their production

[8] Note that we calculate the predictions of CBDT in the treatments without current information *as if* this information was available.

[9] Participants were told that they could think of making production decisions for 30 different islands.

choice and may depend on some of these market conditions. The four scenarios in the "memory" were presented as four different marketing reports, with information on the same set of market conditions as the current report. One of the four production values was then chosen at random to be displayed with each report and the profits that would have resulted from that production value given the market conditions in the report were displayed alongside the production choice. After making a production choice, subjects were informed about their profits from that market and then moved on to the next situation (i.e. market).

We used a total of 12 different market conditions. However, subjects did not know this. In fact, in each market subjects were only given information on a random subset of 5 conditions[10] [11]. In order not to overload our monopoly frame, we decided to name our market conditions as neutrally as possible. In particular, we did not want to give the names too much economic meaning and have subjects trying to guess which ones might influence their sales and therefore their profits. We came up with Tourist Population, Wind, UV Factor, Chance of Rain, % of Population Female, Humidity, Traffic Conditions, Temperature, Literacy Rate, Median Age, # of Potential Buyers, and Gas Price[12]. For each market each condition was randomly chosen to have a value of either 1, 2, or 3. In order to implement a feature-based similarity approach we displayed all values as symbols[13]. Subjects were told that the marketing report in each market was transmitted by their marketing department with an error that resulted in all numbers/levels being erased and only symbols being reported. Subjects were informed that the error was consistent, i.e. the same symbol for the same market condition always meant the same thing, while it could mean different things for different market conditions.

---

[10] By displaying different market conditions for each market we wanted to emphasize that each of these markets/decisions was in fact independent.

[11] The market conditions displayed in the hypothetical scenarios were the same as those given in the marketing report so that similarity comparisons could be made between them.

[12] Subjects never saw the last two conditions: # of potential buyers and gas price. Subjects did not even know that these conditions existed.

[13] Besides eliminating geometric similarity considerations this also got rid of potential home made priors of the following type: "The temperature is very high today so I shouldn't produce very much because it's too hot for buyers to want to be out shopping."

CBDT was conceived under the premise that in many decision problems states of the world are neither naturally given nor can they be simply formulated. Furthermore, it assumes that often even a comprehensive list of all possible outcomes is neither readily available nor easily imagined. We implement this lack of information in a potentially complex decision environment by designing different payoff functions for each market. We generated different payoff functions and used each one twice. The payoff functions varied according to which choice maximized expected payoff and the "penalty" for making a non-optimal choice[14]. Of the 12 market conditions a randomly selected set of 4 would enter the payoff function in each market. Out of the 5 conditions that the subjects saw in each market, 3 actually mattered (i.e. were payoff relevant) and the other 2 did not. We reserved two market conditions that were never reported to the subjects. One of these was chosen at random to enter the payoff function in each market. We did this in order to create a complex environment, in which it is impossible for subjects to figure out potential states of the world or associated outcomes.

In the abstract frame the basic setup remained the same as in the monopoly frame, however, all economic and market terminology was removed from the instructions and the computer screens. Subjects were told that they were making decisions for 30 different, independent situations. The available choices were now 1, 2, 3, or 4 and the labels for the conditions in each situation were changed to be A, B, C, D, E, F, G, H, I, or J. Also, the payoffs in the abstract frame were represented as Earnings instead of Profits.

The experiments were conducted at the Economic Research Laboratory at Texas A&M University between March and October 2006. Elicitations to participate in the experiments were randomly sent out to all undergraduate students in our database of about 1,200 from a diverse background of majors. The experimental interface was programmed in zTree (Fischbacher, 1999). On average the experimental sessions lasted about 90 minutes and average earnings were $21.32 with a minimum of $15.37 and a

[14] To ensure comparability across treatments, the subjects faced the same payoff functions in the same order.

maximum of $24.56 (this includes a $5 show-up fee)[15]. We tested the participants' comprehension in a questionnaire after reading the instructions and subjects seemed to understand since they answered the majority of all questions correctly. We also collected some demographic data in an ex-post questionnaire.

HYPOTHESES

We analyze the predictive power of CBDT in the different treatments. Since EUT is not a reasonable alternative decision making procedure in the environments we consider, we looked for simple decision making principles that deliver predictions in our environment and that can pose as an alternative to CBDT. Heuristics, or rules of thumb, referring to useful and indispensable cognitive processes for solving problems that cannot be handled by logic and probability theory (e.g. Polya, 1954), suggest themselves. Heuristics are strategies that guide information search and modify representations to facilitate solutions (e.g. Simon, 1955).

Gigerenzer and Goldstein (1996) introduced such a fast and frugal algorithm, called "Take the Best" for a search problem. While we use the same basic principle as "Take the Best," the implementation of this heuristic is quite different (and much simpler) in our environment. As such we call this heuristic the "Max-Heuristic" in our setting. Given that our subjects are given 4 scenarios (i.e., cases) for each of the decisions that they are facing (with each possible choice being chosen once), the Max-Heuristic would predict that a DM choose the choice that returned the highest payoff in those cases. Note that in the treatment where information about the current situation is given this would mean that DMs ignore that information. In the treatments when such information is not given, the Max-Heuristic seems a natural choice.

In order to distinguish whether our subjects used Case-Based reasoning or the Max-Heuristic, we calibrated the predicted choices such that the choices predicted by CBDT were different from those of the Max-Heuristic in all but 7 situations. Furthermore, the earnings from using the Max-Heuristic were chosen to be higher than

[15] Sessions varied in size from 4 to 18 participants.

the earnings from using CBDT. This was done in order to discriminate against CBDT making any findings of subjects behaving like a Case-Based DM more credible. We formulate the following two hypotheses. Hypothesis 1 refers to comparisons across treatments and Hypothesis 2 refers to a comparison between CBDT and the Max-Heuristic.

H1. CBDT is a better predictor of choices in treatments when the current information is available than when it is not.

H2. Max-Heuristic is a better predictor of choices than CBDT in treatments when the current information is not available.

We do not have any hypotheses regarding the framing of the situations. Given our choice of abstract symbols to represent features, one might argue that the abstract frame goes better with such a representation, and hence favors CBDT. But one could also argue that similarity comparisons are more difficult when there is no frame of reference to indicate that it is preferable or informative for two conditions to have the same symbol, hence making the use of CBDT less likely.

RESULTS

Behavior across Treatments

We first analyze how well the predictions of CBDT and the Max-Heuristic match observed behavior[16]. Our interest lies in predicting individual behavior rather than the average behavior of participants. We therefore calculate the mean squared deviations (MSDs) of the theoretical prediction (either CBDT or Max-Heuristic) from the observed choice for all 30 decisions a subject faced. We did this for each subject and calculated the mean for each subject over the 30 situations. If every period's choice coincides with its prediction a subject would show a MSD of 0, and if the subject never selected as predicted by theory her MSD would equal 2.

[16] The predictions of CBDT in the treatments without current information were calculated as if the current information was available. All subjects were given the same order of underlying payoff functions for comparability across treatments.

For a first test of Hypothesis 1 we compare the median of the individual MSDs across the different information treatments. We find that the median of subjects' MSDs is significantly higher in the treatments where the current information is not available, indicating worse performance of the CBDT predictions in that case and supporting Hypothesis 1. Results of Robust Rank Order Tests are given in Table 2.2. These tests also show that there is no framing effect. CBDT predicts equally well in the abstract and in the monopoly frame when current information is provided and equally poorly when such information is not given.

Table 2.2 Tests for differences in Case-Based MSDs.

Robust Rank Order Test	z-statistic
Monopoly: w/o Current vs w/ Current	2.56**
Abstract: w/o Current vs w/ Current	2.73**
w/ Current: Monopoly vs Abstract	-0.08
w/o Current: Monopoly vs Abstract	-1.14

** indicates significance at the 5% level

Table 2.3 Tests for differences in Max-Heuristic MSDs.

Robust Rank Order Test	z-statistic
Monopoly: w/ Current vs w/o Current	2.28**
Abstract: w/ Current vs w/o Current	-0.51
w/ Current: Monopoly vs Abstract	0.13
w/o Current: Monopoly vs Abstract	-1.31

** indicates significance at the 5% level

Table 2.3 indicates that the Max-Heuristic predicts actual choices significantly better in the monopoly frame when no current information is available as compared to the same frame when such information is available. However, it seems as if there is a slight framing effect, since this difference does not hold up in the abstract frame. This could possibly be because subjects in the abstract frame seem to be guided by the Max-Heuristic even when information on the current situation is available.

Behavior within Treatments

Table 2.4 reveals support for Hypothesis 2. It shows that when information about the current situation is not given, the Max-Heuristic predicts choices better than CBDT (i.e. individual MSDs of observed choices from predicted choices are lower). Table 2.4, however, also demonstrates that CBDT does not predict a larger proportion of choices when compared to the Max-Heuristic, indicating that equally many participants are guided by CBDT as are guided by the Max-Heuristic when both could provide decision making guidance.

Table 2.4 Case-Based versus Max-Heuristic.

Robust Rank Order Test	z-statistic
Monopoly w/ Current	-0.46
Monopoly w/o Current	4.83**
Abstract w/ Current	0.48
Abstract w/o Current	2.91**

** indicates significance at the 5% level

In a next step we would like to not just compare levels of MSDs across the treatments, we would like to get a better understanding of what these levels mean. For example, we are interested in whether a pure random choice model would do better than either, CBDT or Max-Heuristic. Furthermore, we would like to understand what percentage of an individual's choices are predicted correctly in each treatment.

If a subject was choosing randomly, one could interpret this as meaning that she would coincide with the predictions 25% of the time (since there are four choices to choose from). Given that CBDT as well as the Max-Heuristic are deterministic models, it seems unfair to compare them to a probabilistic one where each entry in the prediction vector is 0.25[17]. We therefore establish the benchmark of random choice as being correct

[17] Such a calculation would lead to an MSD of 0.75. In general, the calculation of MSDs favors probabilistic models over point predictions. See Selten (1998) for an axiomatization of quadratic scoring rules.

25% of the time and obtaining a MSD of 0 and being incorrect 75% of the time realizing a MSD of 2. The average over all 30 decisions is then 1.5.

We also establish two other benchmarks, predicting correctly at least half of the time (15 out of 30 situations) and predicting correctly at least two-thirds of the time (20 out of 30 situations). Table 2.5 shows the percentage of people for whom CBDT predicts at least as well as the given category, i.e. returns a lower MSD. Table 2.6 illustrates the same for the Max-Heuristic.

Table 2.5 Observed frequencies of individual MSDs when comparing choices to CBDT predictions.

	CBDT predicts correctly more often than:		
	Random (MSD<1.5)	Half MSD(<1)	Two-Thirds MSD(<0.67)
Monopoly w/ Current	79.5% (31/39)	25.6% (10/39)	15.4% (6/39)
Monopoly w/o Current	70.0% (21/30)	3.3% (1/30)	0% (0/30)
Abstract w/ Current	80.6% (25/31)	22.6% (7/31)	16.1% (5/31)
Abstract w/o Current	43.8% (14/32)	0% (0/32)	0% (0/32)

Table 2.6 Observed frequencies of individual MSDs when comparing choices to Max-Heuristic predictions

	Max-Heuristic predicts correctly more often than:		
	Random (MSD<1.5)	Half MSD(<1)	Two-Thirds MSD(<0.67)
Monopoly w/ Current	87.2% (34/39)	20.5% (8/39)	10.3% (4/39)
Monopoly w/o Current	93.3% (28/30)	33.3% (10/30)	26.7% (8/30)
Abstract w/ Current	71.0% (22/31)	45.2% (14/31)	32.3% (10/31)
Abstract w/o Current	78.1% (25/32)	28.1% (9/32)	21.9% (7/32)

By comparing the last column of Table 2.5 with the last column of Table 2.6, we find additional support for hypotheses 1 and 2. CBDT predicts choices correctly at least two-thirds of the time for more subjects in treatments when current information is provided compared to when such information is not available (Test of Equality of Proportions: monopoly frame, z-value= 2.25, one-tailed p = 0.0122; abstract frame, z-value= 2.37, one- tailed p = 0.0089)[18]. In fact, CBDT cannot be applied in the treatments without the current information and observing that for no subject CBDT predicts more than two-thirds of their choices is reassuring. In those treatments, we find that the alternative Max-Heuristic predicts more than two-thirds of the choices correctly for more subjects than CBDT (Test of Equality of Proportions: monopoly frame, z-value= 3.04, one-tailed p = 0.0012; abstract frame, z-value= 2.80, one-tailed p = 0.0026), again supporting Hypothesis 2[19].

Learning

There are different types of learning that could take place in our experiment. First, it seems natural to ask whether our subjects learn to become case- based decision makers. If this was the case, we should observe the individual MSDs (calculated with respect to the CBDT predictions) to get smaller over time. Second, our subjects could learn across situations. For the purpose of our paper it is important that there be no such learning. However, since we gave subjects feedback about their performance after each decision, it could be argued that some choices get positively reinforced.

In order to address the first type of learning we calculate the average of an individual's MSDs over the first half of the experiment and compare this with the average of an individual's MSDs over the second half. We find that 44% (17/39) show a smaller MSD in the second half of the experiment in the monopoly frame when current information is provided. This holds for 42% (13/31) in the abstract frame when current

---

[18] The specific test statistic is $z = (p_1 - p_2)/S_{p_c}$, where p_i is the proportion in sub-sample i, and

$$S_{p_c} = \sqrt{p_c(1-p_c)(\frac{1}{N_1}+\frac{1}{N_2})}$$ is an estimate of the standard error of the difference in proportions, $p_1 - p_2$. p_c is an estimate of the population proportion under the null hypothesis of equal proportions, $p_c = (p_1 N_1 + p_2 N_2)/(N_1 + N_2)$ where N_i is the total number of subjects in sub-sample i (see Glasnapp and Poggio, 1985).

[19] Doing the same exercise for the Max-Heuristic, we find that the Max-Heuristic predicts better (i.e. a higher proportion of people for whom it predicts correctly at least 2/3 of the time) in the monopoly frame when current information is not provided than when it is. This does not hold in the abstract frame (Test of Equality of Proportions: monopoly frame, z-value = 1.78 one-tailed p-value = 0.0375; abstract frame, z-value = -0.92, one-tailed p-value = 0.1788).

information is provided. Both of these percentages are not different from those that show an increase in their MSDs (Test of Equality of Proportions, monopoly frame, z-value= −1.13, one-tailed p = 0.1292; abstract frame, z-value= −1.27, one-tailed p = 0.1020). We can therefore conclude that subjects do not tend to use CBDT more often over time, i.e. there is no learning towards becoming a Case-Based DM.

In order to address the second type of learning, we simulated two different learning models, a reinforcement learning model (Roth and Erev, 1995) and a payoff assessment learning model (Sarin and Vahid, 1999)[20 21].

Since the reinforcement learning model (RL) is a probabilistic learning model and the payoff assessment (PA) model is a deterministic model, we evaluate each against a different "calibration" of a random model. The MSD of a random model is 0.75 when a predicted probability vector is compared to the actual choice. We use this as a benchmark for the RL model. However, as we have noted before, such a comparison would bias against any deterministic prediction. We therefore use a benchmark of 1.5 for the PA model. Table 2.7 reports the results of these comparisons.

While the PA model seems to do better than the RL model when each is compared to a random model, it is worthwhile to note that the PA model could not predict more than 50% of the choices any subject made (i.e. no individual MSD was lower than 1). For the RL model we find very little variance in the MSDs, with a minimum of 0.7249 and a maximum of 0.9343 indicating close to random performance.

While it seems easy to reject the notion that our subjects learned across situations in a reinforcement type of way, we find that behavior of some subjects, at least for some situations (although never for the majority of them), can be modeled by a payoff assessment type of learning. This raises the question of whether we would have observed

[20] Both learning models were simulated 100 times with the weight of the current payoff being 0.9. Note that this is a modification of the RL model used by Roth and Erev (1995) in order to "equalize" treatment of obtained payoffs across the two different models. With this modification, if player j chooses k at time t and receives a payoff of x, then the propensity to choose k at time $t+1$ is $q_{jk}(t+1) = 0.1\ q_{jk}(t) + 0.9x$, while for all other possible choices n, $q_{jn}(t+1) = q_{jn}(t)$.

[21] Payoffs for the reinforcement learning model were transformed such that they were all positive and initial propensities were set at the mean of the newly transformed payoffs in the first round. For the payoff assessment model, initial assessments were set equal to the ones seen in the first four scenarios in the first situation of the experiment.

stronger evidence for CBDT if we had not given subjects information about their earnings after each situation.

Table 2.7 Observed frequencies of individual MSDs when comparing choices to PA and RL predictions.

	PA (MSD<1.5)	RL MSD(<0.75)
Monopoly w/ Current	69.2% (27/39)	25.6% (10.39)
Monopoly w/o Current	73.3% (22/30)	20.0% (6/30)
Abstract w/ Current	54.8% (17/31)	38.7% (12/31)
Abstract w/o Current	71..9% (23/32)	28.1% (9/32)

Decision Time and Earnings

Figure 2.1 shows the average time in seconds that it took participants to confirm their choices in the treatments when current information was provided. Participants are put into three distinct groups. The "CBDT" group includes participants for whom CBDT predicts choices correctly more than two-thirds of the time. The "Max-Heuristic" group includes participants for whom the Max-Heuristic predicts choices correctly more than two-thirds of the time. Everybody else is put into the group called "Rest." Participants who seem to be guided predominantly by CBDT take equally long in either framework when the current information is provided (Robust Rank Order Test, U = 0.1631, n.s.). While it appears that subjects for whom the Max-Heuristic predicts choices correctly more than two-thirds of the time take longer in the abstract frame than in the monopoly frame when the current information is provided, this difference is not significant (Robust Rank Order Test, U = 0.6159, n.s.). Subjects who can neither be "classified" as users of CBDT nor as users of the Max-Heuristic, take significantly longer in the monopoly frame than in the abstract frame when current information is provided (Robust Rank Order Test, U = −3.1930, p = 0.0007). One explanation could be that subjects in the monopoly frame who do not use a "rule" to make their decisions try to figure out what certain symbols meant and how these are related to their profits. Such “calculations” are impossible in the abstract frame and random choices are much quicker to make.

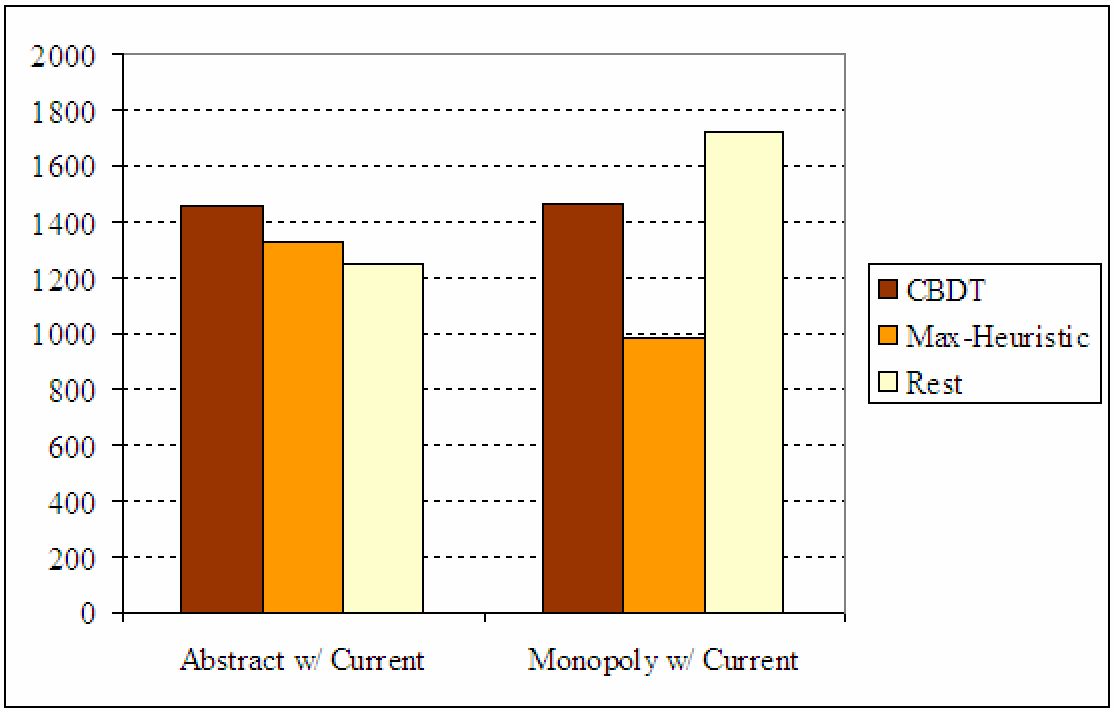

Figure 2.1 Time in seconds before participants confirmed their choices.

Figure 2.2 shows average earnings of the same set of people “categorized” in Figure 2.1. As expected (and in line with the calibration of the experiment) subjects guided by CBDT for more than two-thirds of their decisions earn similar amounts in either frame (Robust Rank Order Test, U = 0.9129, n.s.). This is also true for subjects who seem to employ the Max-Heuristic more than two-thirds of the time (Robust Rank Order Test, U = 0, n.s.). However, for the subjects “without” a rule, we find that they earn significantly more in the monopoly frame, when they also took more time (Robust Rank Order Test, U = −3.1978, p = 0.0007).

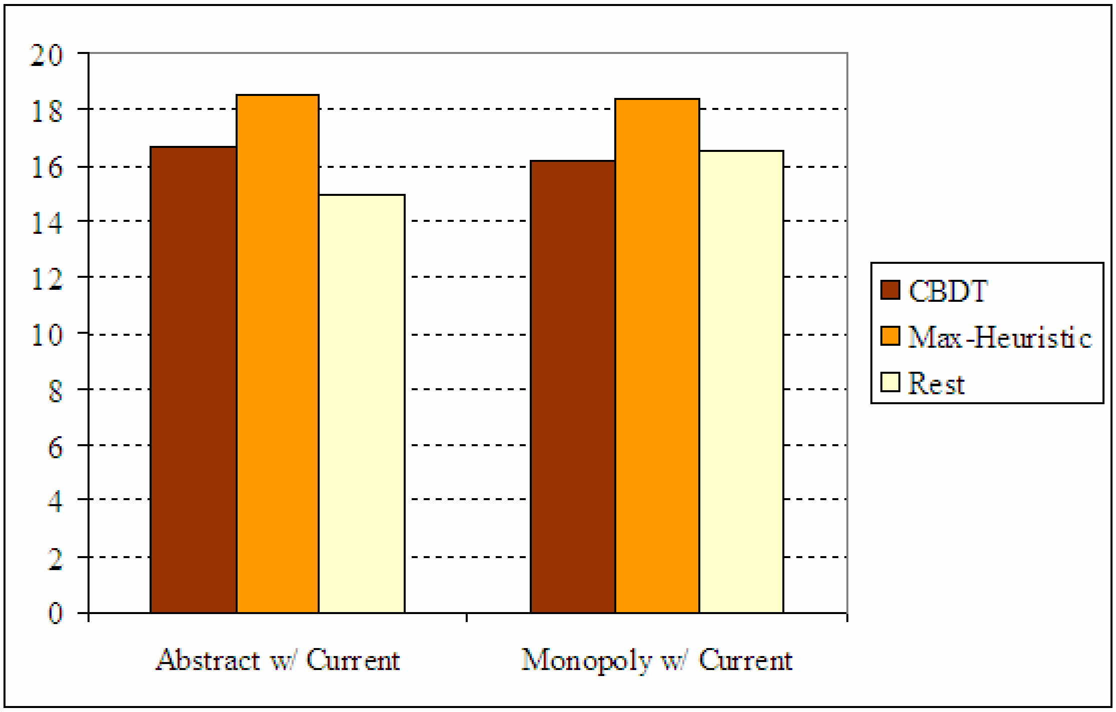

Figure 2.2 Earnings in dollars (excluding show up fees).

DISCUSSION

We designed an experiment to test whether subjects use case based reasoning in an individual decision making environment. Our results provide some support for Case-Based Decision Theory. This support did not depend on the framing of the problem. We also found support for a heuristic that was much simpler to use (and resulted in better payoffs). Very little is known about decision making procedures used in complex, limited information environments. Further research in this area is needed.

CHAPTER III

FRAMING THE FIRST-PRICE AUCTION

INTRODUCTION

Despite being one of the simplest auction environments, the single-unit independent private values first-price auction has exhibited two intriguing stylized facts in laboratory studies. First, subjects typically bid more aggressively than predicted by the risk-neutral Bayes-Nash equilibrium (RNNE for short), resulting in significant reductions in their overall monetary earnings. In addition, the first-price auction may be implemented in two ways, as a sealed-bid or as a Dutch descending-clock auction. Bayes-Nash equilibrium theory suggests these should be isomorphic, giving identical results. Yet, bidding behavior in the Dutch implementation is significantly less aggressive[22].

The seminal papers in this area include those of Coppinger et al. (1980), Cox et al. (1982), and Cox et al. (1988), with a substantial subsequent literature. Kagel (1995) provides a good survey and summary of the results and the theories put forward to explain them. Instead of testing these theories specifically, we return to these stylized facts from a new perspective, in which more careful consideration of the presentation and framing of these auctions plays a central part. Our design considers two aspects of framing in the first-price auction.

We investigate the result that both auction implementations produce market prices in excess of the RNNE by presenting the games in the context of a custom-designed graphical interface. This interface presents the information, action, and payoff spaces within a unified rectangular area, visually organizing the interrelations among private values, bids, and earnings. Our subjects make choices and receive feedback in the same frame, in keeping with principles of interaction design (e.g., Cooper and

[22] We will use the term "first-price" auction to refer generically to the institution, with the term "implementation" referring to one of the ways in which the theoretical structure of the auction is realized.

Reimann, 2003). The graphical display is identical across implementations, except for the minimal changes necessary due to the rules of the games[23].

As a further control to keep the environments as similar as possible, and in view of the observations of Katok and Kwasnica (2003), we attempt to control for the opportunity cost of subjects' time by choosing the speed of the clock in the Dutch auction in such a way that the typical period and session length for all implementations are approximately the same. We view the results of Katok and Kwasnica as an illustration of the importance of maintaining the dominance of the payoff implications of actions within the session (Smith, 1982). In a pilot session not reported in this paper, in which we did not institute this control, our subjects completed 80 sealed-bid auctions in about 30 minutes. In subsequent sealed-bid sessions, a common question was whether the session would end sooner, or if more periods would be conducted, if they "bid faster." Thus, we view this control as a significant design feature for the purpose of testing isomorphism.

The sessions reported in this paper consist of 60 auction periods, in which subjects participate in the same implementation throughout. This duration is announced to the subjects during the instructions for the session. This design feature has two objectives. First, with a large number of auction periods, the cumulative effects of suboptimal bidding, in terms of foregone earnings, will be more significant. Second, if the behavioral features of these auction environments arise largely from an initial misperception of the tradeoffs made in setting bids, a larger number of periods offers time for subjects to adapt.

Our second framing treatment addresses more directly the difference in the extensive forms of the sealed and Dutch implementations. We contrast the way the uncertainty about other bidders' values and bidding behavior is presented in these two games. Dorsey and Razzolini (2003) observe that the choice of a bid in the sealed implementation is similar to the choice of a lottery from a menu, where each lottery i in the menu has some probability p_i of winning some prize q_i, with a prize of zero

[23] Specifically, subjects submit bids by clicking on the appropriate price in the sealed implementation, whereas they bid by clicking on a "Purchase" button in the Dutch.

otherwise, such that the p_i and q_i have an inverse relationship. In the Dutch implementation, those same lotteries are presented in a sequential format. At each clock price, the bidder has two choices. He may purchase at the current price and earn an amount with certainty. (For simplicity, we neglect the possibility of a tie.) Alternatively, he may not purchase now; in this case, there is a high probability he will have another opportunity to purchase at a slightly better price at the next clock increment. We conjecture that the clock-based presentation helps bidders to recognize the tradeoff between probability of winning and the amount won, which is the essential tradeoff in the first-price auction.

To investigate this, we introduce a synthetic implementation in which a clock counts down as in the Dutch implementation, but in which the outcome of the auction is not revealed until the clock reaches the lowest price. We will refer to this intermediate implementation as the “silent” implementation. It is functionally similar to the sealed implementation in terms of feedback, in that the results of the auction are not known until all choices are made. At the same time, it shares the property of the Dutch implementation in that it may cue the bidder to make assessments at the margin about the choice of bids. The silent implementation allows these to be separated to some degree.

We find that we replicate the standard empirical regularities. The subjects bid significantly more than RNNE in all implementations, and subjects in the sealed implementation bid more aggressively than in the Dutch. We also find that both of these results persist over the session. The results of the silent implementation fall in between the sealed and Dutch: market prices typically exceed those in the Dutch, but are less than those in the sealed implementation.

The paper is organized as follows. Section 2 describes the design of the experimental sessions. Section 3 outlines the experimental results, both at the market and individual levels. Section 4 concludes with a discussion and future directions.

DESIGN

Each experimental session was conducted with 18 subjects recruited from the undergraduate student body at Texas A&M University. The 18 subjects were randomly grouped into two cohorts of 9 subjects each, labeled *a* and *b*. Within each cohort, subjects were randomly assigned in each period to one of three markets, each consisting of three subjects. The matching was done anonymously, and no subject ID numbers or other information about which subjects were participating in which markets in which periods was known to the subjects. All interaction among the subjects was mediated via computer in the Economic Research Laboratory at Texas A&M.

The subjects participated in a uniform independent private values auction in each period. Subjects received a resale value for a single unit of a fictional commodity drawn from the range $0.15 to $6.00 in increments of $0.15; therefore, there were 40 possible resale values. The resale values were equally likely and drawn independently across periods and subjects. The subject who purchased the object earned the difference between her resale value and the market price; subjects who did not purchase the unit earned zero for the period. Ties were broken at random. Two sequences of private values and market assignments were used in all sessions, facilitating comparisons across implementations. The same sequence was assigned to cohort *a* in all sessions, and the other to cohort *b*.

Bids were constrained to $0.10 increments, starting at $0.10; the maximum permitted bid was $6.20. The bid and value spaces were chosen so that in our environment it is a symmetric Bayes-Nash equilibrium to submit a bid equal to two-thirds of the private value. This is the analog of the unique symmetric Bayes-Nash equilibrium of the first-price auction with risk-neutral bidders when values are distributed uniformly and the permitted bids are the nonnegative real numbers.

Figure 3.1 shows a screenshot of the subject interface presenting the results of a period. Subjects used the rectangle on the left of the screen to interact with the market. All decisions were made in this area, and feedback from the results of the period was presented in the same area. In the sealed implementation, subjects observed the realization of their private value, and submitted a bid by clicking on the corresponding area of the market rectangle. In the Dutch implementation, a clock price was displayed, which started at $6.20 and decreased by $0.10 each second; once a subject clicked the button to purchase the object, the clock stopped for that market and the results were displayed to all participants in that market. The silent implementation operates as the Dutch does, except the clock price decreases to zero in every period, with no feedback regarding the outcome of the market given until that time. Subjects were permitted to submit bids, or click to purchase, above their resale value in all implementations.

At the right of the screen is a record sheet, that summarized the subject's history, including their private values and bids, the market prices in the markets in which they participated, and their history of earnings. The subjects did not observe the results of other markets in which they did not participate. Since the design investigates the isomorphism between the sealed-bid and Dutch implementations, in order to keep information constant across implementations, no information regarding how others bid was presented, since that is not available in the Dutch implementation. The record sheet by default displayed the results of the last 25 periods, with scroll buttons available for subjects to review the results from earlier in the session.

The clock speed in the Dutch and silent implementations was chosen with two goals in mind. First, the clock ticked slowly enough that subjects could click to purchase while the clock was at their desired price with a high degree of accuracy[24]. Second, the clock speed was such that the overall length of the clock-based sessions would be roughly comparable to the sealed sessions.

The subjects participated in 60 periods, and the length of the session was announced during the instructions. Three sessions of each implementation were conducted. No subject participated in more than one session, and none had any prior experience in an auction environment at Texas A&M. The instructions were read aloud from a projector screen while subjects followed along on the screen at their stations.

[24] We implemented "practice" round during the instructions during which subjects were asked to "stop" the clock at a given price. No subject experienced any difficulty doing so. Thus we believe, at least in the vast majority of cases, that the subjects were able to register their desire to stop the clock and buy the object at the price they intended.

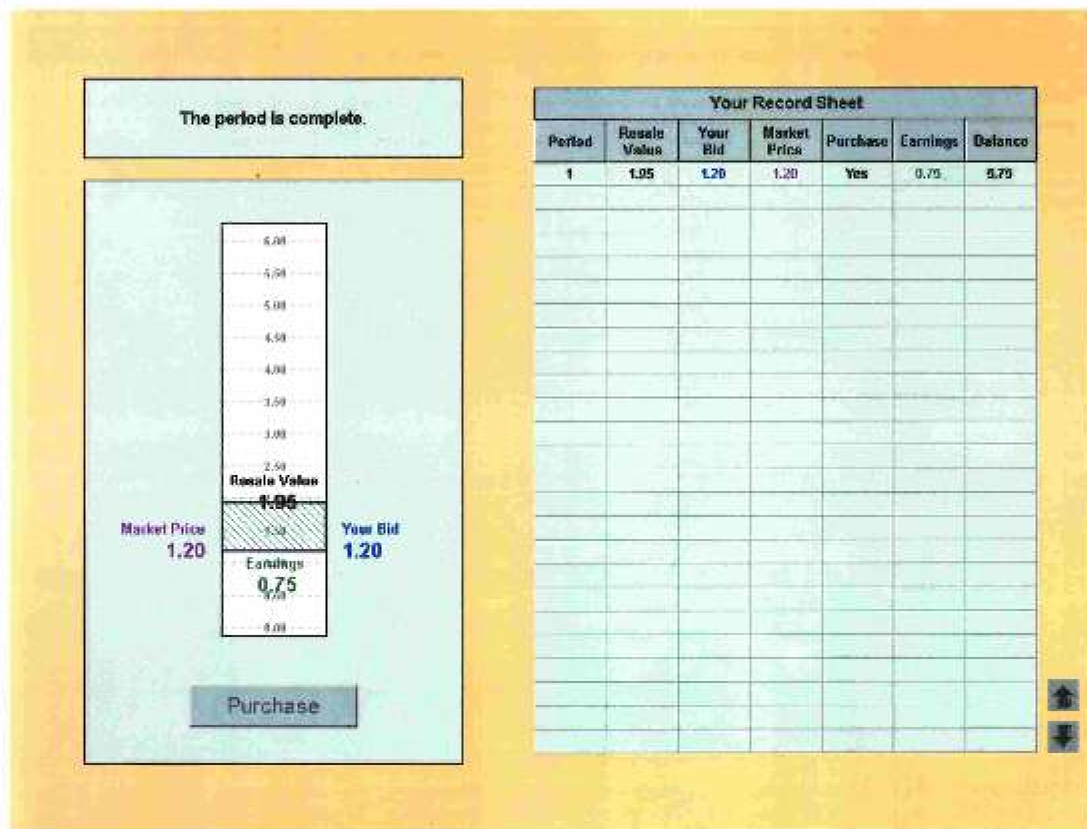

Figure 3.1 Screenshot of the subject interface, displaying the results of a period.

After the instructions were completed, subjects answered a questionnaire, which was checked for correctness by an experimenter. The answers were then reviewed aloud, referring to the instructions as appropriate. After an interval for clarifying questions, the session began.

RESULTS

Market performance

Table 3.1 presents summary statistics for overall market performance in each of the 18 cohorts. Each cohort contains 180 market observations (60 periods × 3 markets per period). The analysis of market performance takes the cohort as the smallest unit of independent observation, and we characterize performance based upon statistics aggregated at the cohort level.

Table 3.1 Statistics on market performance for all 18 cohorts.

		(a) Market/Theory			(b) Market vs Theory			(c) Efficiency	
Type	Cohort	1-10	6-15	51-60	Above	Equal	Below	Freq	Percent
Sealed	1a	1.295	1.299	1.296	168	3	9	86.1	98.1
	1b	1.390	1.311	1.264	175	2	3	91.7	98.6
	2a	1.356	1.350	1.238	179	2	0	92.8	99.4
	2b	1.269	1.290	1.251	169	2	9	91.1	98.7
	3a	1.301	1.261	1.257	175	2	3	90.0	98.8
	3b	1.310	1.300	1.289	178	0	2	92.2	99.0
	Mean	1.320	1.302	1.266	173.8	1.8	4.3	90.7	98.8
Dutch	1a	1.173	1.185	1.194	160	7	13	86..1	97.7
	1b	1.257	1.244	1.237	174	3	3	88.9	98.3
	2a	1.120	1.144	1.123	140	7	33	82.2	96.6
	2b	1.263	1.150	1.188	150	8	22	88.3	98.5
	3a	1.195	1.094	1.041	136	5	39	81.1	96.7
	3b	1.222	1.209	1.230	170	3	7	87.8	98.5
	Mean	1.205	1.171	1.169	155	5.5	19.5	85.7	97.7
Silent	1a	1.249	1.271	1.248	166	3	11	92.8	99.1
	1b	1.294	1.282	1.307	164	6	10	91.7	99.2
	2a	1.325	1.193	1.170	178	0	2	87.2	98.5
	2b	1.190	1.192	1.181	158	12	10	91.1	98.8
	3a	1.423	1.287	1.267	167	2	11	91.1	98.2
	3b	1.231	1.220	1.182	176	3	1	88.3	98.3
	Mean	1.285	1.241	1.226	168.2	4.3	7.5	90.4	98.7

We relate observed market prices to the RNNE prediction in two ways. To provide an overall indicator of how market prices typically compare to theory, for each market observation we construct the ratio of the market price to the RNNE theory price. We then report the average of this ratio for 10-period increments in order to summarize its evolution. This evolution is plotted for each of the 18 cohorts in Figure 3.2, and values for selected groups of periods are presented in part (a) in Table 3.1. Because we are interested in changes in market price over time, ratios for the first ten and last ten periods are shown. Since there were no practice periods, there may have been some early-period confusion about the rules of the game. To focus on how the markets evolve once a common understanding of the game has been established, we also present the ratios for periods 6 through 15.

Since average rates may be influenced by particular markets with atypical results, we also categorize each market by whether the observed market price exceeds, equals, or is less than the RNNE prediction. These counts are presented in part (b) of Table 3.1. By all the measures, the averages across cohorts order the implementations the same way, with sealed having the highest prices, Dutch the lowest, and silent in between. Also, the average market to theory ratios drop over time in all three implementations. We now formalize and test the significance of these observations.

Result 1. Prices exceed the risk-neutral prediction in a significant majority of markets in all implementations.

Clearly, the RNNE point prediction is not a tenable hypothesis given the data. A more generous hypothesis is that market prices are equally likely to fall above or below RNNE. We form a test statistic by computing for each cohort the number of markets exceeding RNNE minus the number in which the market price is less than RNNE. We reject the null hypothesis that the mean of this statistic is zero, against the two-sided alternative, for all implementations, with *p*-values 10−8 for sealed, 10−4 for Dutch, and 10−7 for silent.

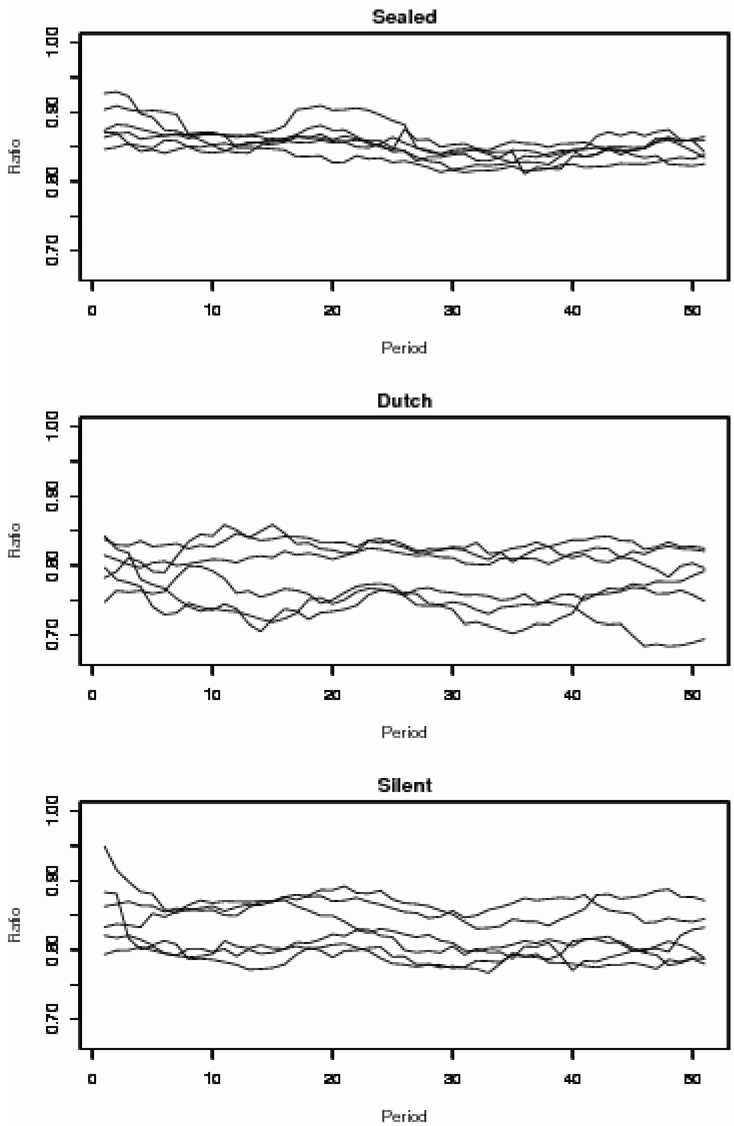

Figure 3.2 Ten period moving averages of market-to-theory ratios. Each line repressents one cohort.

Result 2. Prices are significantly higher in the sealed than in the Dutch. The difference in prices is persistent throughout all periods.

The phenomenon of higher prices in the sealed implementation is persistent in this environment, even when subjects are given the opportunity to participate in 60 market periods. We ask whether the market-to-theory ratios are significantly different across implementations at various points in the session. We are interested both in initial differences, as well as whether differences tend to diminish over time.

For each implementation, our statistic is the mean across cohorts of the ratios reported in part (a) of Table 3.1. We test the null hypothesis of equality against the two-sided alternative, using a two-sample t-test. The p-values for these tests are shown in Table 3.2. Cohorts using the Dutch implementation give significantly lower market-to-theory ratios throughout the session when compared to the sealed implementation. While in point estimate terms, the Dutch cohorts also exhibit lower ratios than the cohorts using the silent implementation, the differences are only significant at the .05 level in periods 6-15.

Table 3.2 p-values for two-sample t-test on market to theory ratios (across implementations, over three selected groups of periods).

	Periods		
Comparing	1-10	6-15	51-60
Sealed - Dutch	0.0026	0.00081	0.023
Sealed - Silent	0.390	0.0213	0.153
Silent - Dutch	0.0786	0.0336	0.168

Result 3. Market price adjustment within a session is minimal.

The across-cohort tests of Result 2 could mask significant trends within cohorts. To focus on within-cohort changes, we construct for each cohort the difference between the ratio in periods 1-10 and periods 51-60. We test the null hypothesis that the mean of these ratios for an implementation is zero, against the two-sided alternative. We repeat the process, using periods 6-15 instead of periods 1-10. Table 3.3 presents the *p*-values for these *t*-tests.

Table 3.3 *p*-values for paired *t*-test of early versus late price ratios (with null hypothesis of equal ratios against two-sided alternative).

Implementation	Periods	*p*-value
Sealed	1-10 vs 51-60	0.058
	6-15 vs 51-60	0.088
Dutch	1-10 vs 51-60	0.243
	6-15 vs 51-60	0.876
Silent	1-10 vs 51-60	0.118
	6-15 vs 51-60	0.147

Only the sealed implementation shows statistically significant evidence of decline in the ratio of market to RNNE prices. Even though overall prices decrease on average in all implementations, there is much more variation across cohorts in the clock-based implementations, leading to the insignificance of the tests. In fact, a striking feature of Figure 3.2 is that there is much more dispersion in the levels of these ratios across cohorts in the clock-based implementations than in the sealed implementation.

Result 4. All implementations result in high levels of efficiency, significantly exceeding the efficiency predictions of a zero-intelligence random bidding model. The Dutch implementation is less efficient than the other two.

The efficiency of these markets can be measured in two ways: the percentage of gains from exchange realized, and the frequency with which the highest-value bidder purchased the object. Part (c) of Table 3.1 summarizes the two efficiency measures for each cohort. To obtain a useful baseline for interpreting these levels, we consider a model of “zero-intelligence” (ZI) random bidding patterned after that of Gode and Sunder (1993) for continuous-time double-auction markets, and Cason and Friedman (1997) for call markets. In our environment, we operationalize the model by assuming bidders choose all individually-rational bids with equal probability. Simulation results for this zero-intelligence model give a percentage efficiency of 89.4% and a frequency of efficient allocation of 64.3%.

Formally, we construct the average frequency of efficient allocation, and average percentage efficiency, for each implementation. We then perform two- sample t-tests for each pair of implementations, testing the null of equality against the two-sided alternative, for both measures. We also perform *t*-tests of these measures against the ZI predictions, again with the null of equality against the two-sided alternative. *p*-values for all tests are reported in Table 3.4.

Table 3.4 *p*-values for *t*-tests on efficiency measures.

	Frequency			Percentage		
	Silent	Dutch	ZI	Silent	Dutch	ZI
Sealed	0.892	0.016	10^{-6}	0.741	0.033	10^{-8}
Silent		0.019	10^{-7}		0.044	10^{-6}
Dutch			10^{-5}			10^{-6}

Thus, observed efficiencies do significantly exceed those generated by random bidding behavior. The sealed and silent implementations are statistically indistinguishable, but both are significantly more efficient than the Dutch by both measures.

Even though overall the silent and sealed implementations result in almost exactly the same number of inefficient allocations, the markets in which misallocations occur are not the same. Figure 3.3 presents a boxplot of the distribution of the highest value in the market among the markets in which misallocation occurred. For comparison, "All Markets" plots the distribution of the highest value for all markets in the two cohorts. Inefficient allocation happens more often in the sealed implementation when the highest value in the market is closer to the top of the interval of possible private values. In contrast, the distribution for the silent and Dutch look much more like random draws from the set of markets, at least by this way of categorizing markets.

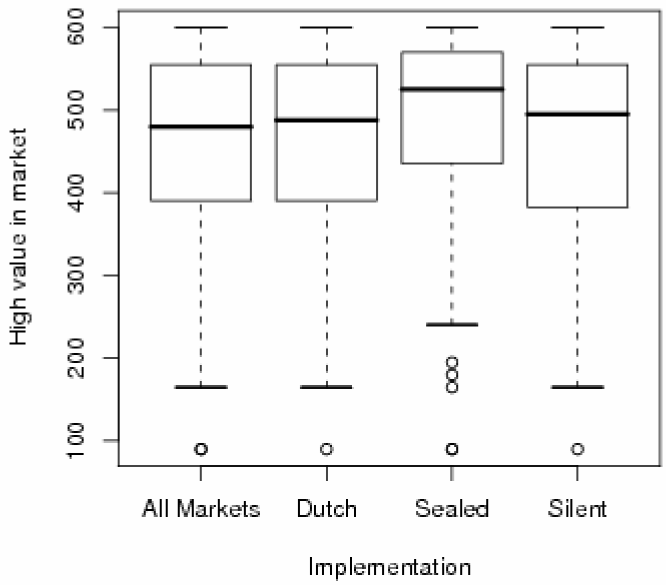

Figure 3.3 Boxplot distribution of the highest market value, among markets in which inefficient allocation occured. "All Markets" plots the distribution of the highest value for all markets in the parameter set.

Individual performance

Result 5. In all implementations, bidders leave significant amounts of earnings "on the table" due to bidding more aggressively than risk-neutral.

Summary statistics for subject earnings are presented in Table 3.5, along with the predictions assuming all subjects bid according to the RNNE strategy[25]. For comparison, we take the RNNE bidding function as an alternative heuristic which any bidder might have unilaterally chosen. In the sealed implementation, the mean earnings loss over the session relative to this benchmark was $7.93; for the silent implementation, the mean

[25] The earnings totals presented are for the contingent portion of the experiment only, and do not include the $10.00 participation fee

earnings loss was $7.38[26]. Thus, most bidders would have been significantly better off unilaterally using this simple, less aggressive heuristic, even holding constant aggressive bidding by other subjects.

Table 3.5 Summary statistics for distribution of subject earnings, by implementation. The theory column refers to the predictions for the risk-neutral Nash equilibrium.

	Sealed	Silent	Dutch	Theory
Mean	$12.86	$14.96	$17.65	$30.34
Median	$12.45	$15.03	$17.90	$29.81
Minimum	$7.25	$3.04	$6.70	$23.10
Maximum	$29.25	$30.00	$33.15	$42.10
Subjects	54	54	54	

Result 6. After winning an auction, bidders frequently, though not always, change bids in accordance with directional learning (Selten and Buchta, 1998). Directional learning fits the data better in the clock-based implementations.

Directional learning implies that a bidder who wins the object in one period would adjust his bid function at his realized private value downward, since he almost certainly could have won the object at a lower price. Although we do not directly observe bid functions, there are 26 instances in each session in which a bidder receives the same resale value in two consecutive periods. For these instances, we examine whether the bidder's behavior between those periods is consistent with the directional learning hypothesis. Because independence assumptions across observations in the same session may not be appropriate, we summarize our observations without performing formal statistical tests.

Table 3.6 summarizes the data on how bidders set their bids when receiving the same value as in the previous period. In an environment using the strategy method to elicit entire bid functions, Selten and Buchta (1998) report substantial inertia; more than

[26] Because only winning bids are observed in the Dutch, we cannot construct this counterfactual.

80 percent of the time their subjects did not adjust their bid function in response to the outcome of the previous period. We observe the opposite; with a similar frequency our subjects do change their bidding behavior[27]. Along with Selten and Buchta, we find that when the subjects do change their bids, they generally do so in the direction that directional learning predicts. Our results more closely parallel Cason and Friedman (1997), who also find evidence in favor of directional learning in a double-auction call market.

Table 3.6 Bidder reaction to winning an auction and having the same value draw the next period. Numbers in parentheses are percentages.

	Increase	Same	Decrease
Sealed	6 (16.7%)	10 (27.8%)	20 (55.6%)
Silent	4 (11.4%)	5 (14.3%)	26 (74.3%)
Dutch	0 (0.0%)	6 (24.0%)	19 (76.0%)
Dutch (worst case)	7 (21.9%)	6 (18.8%)	19 (59.4%)

The implementation may have some effect on the frequency with which subjects change their bids. In the silent implementation we observe less inertia and a greater frequency of decreasing bids after having won the previous period. A similar pattern appears in the Dutch implementation. There is a censoring bias in the Dutch implementation, since it is possible to win the auction in one period, and then lose the next period because the market price increases due to another bidder purchasing at a higher price. In these cases, we do not observe what the subject would have done. The worst case for directional learning is that in each of these instances, the bidder was in fact intending to increase his bid. The row in Table 3.6 labeled "Dutch (worst case)" presents the data assuming the bidder would have increased his bid in all 7 of these

[27] We conjecture that this result is driven by a bias in the design of Selten and Buchta's interface. In their design, subjects submit the same bid function as the previous period with a single mouse click, but had to redraw the entire bid function to modify any part of it. In our implementation, the amount of work to set the bid in the subsequent period is the same regardless of whether or not the subject changes their behavior relative to the previous period.

cases. Even with this extreme assumption, the data generally favor the directional learning hypothesis.

Result 7. There is no apparent trend, either individually or in aggregate, in how bidders change bidding behavior over time. There is weak evidence that bidders exhibit less inertia in the silent implementation than the sealed.

We consider cases in which the same subject received the same resale value in two different periods, within a span of 25 periods. Additionally, we require that the resale value was at least $3.00, to restrict attention to behavior in cases where the bidder likely believes he has a realistic chance of purchasing the object. Since we are interested in the overall bidding trends, we do not distinguish whether the subject won the auction in the earlier period of the pair; we simply ask whether the bidder chose a higher, lower, or the same bid in the later period than in the earlier one.

We use the silent implementation as a proxy for how subjects react to the clock-based presentation, since we do not observe behavior in the Dutch for bidders who do not win the auction. The data for the 54 subjects in the sealed and the 54 in the silent implementation are represented in the simplices in Figure 3.4. A subject's location on these barycentric plots is determined by the frequency with which he raised, lowered, or did not change his bid in the period pairs. Graphically, points closer to the apex of the simplex represent bidders who exhibited more inertia, that is, who more often submitted the same bid in both periods of the pair. Points to the left of the vertical line are bidders who, conditional on having changed their bid, increased their bid more often than decreased it[28].

The points cluster lower in the simplex for the silent implementation, indicating that typically bidders tended to change their bids more often in the silent implementation than in the sealed implementation. In terms of the representation in Figure 3.4, for any

[28] Of the 18 value sequences in the parameter set, the number of period pairs satisfying our restriction ranges from 6 to 14, with a mean of 10.8 and a median of 11.5. In the sealed implementation, we observe 23 bidders who increased their bid in the second period of the pair more often than decreased it, and 26 bidders who decreased more often than increased, with 5 bidders equally likely to go either way conditional on making a change. In the silent treatment, 22 bidders increase more often than decreased, and 26 decreased more often than increased, with 5 bidders equally likely to adjust in either direction.

line drawn horizontally across the two simplices at the same height, there are more bidders whose points fall on or below that line in the silent implementation graph than in the sealed implementation graph.

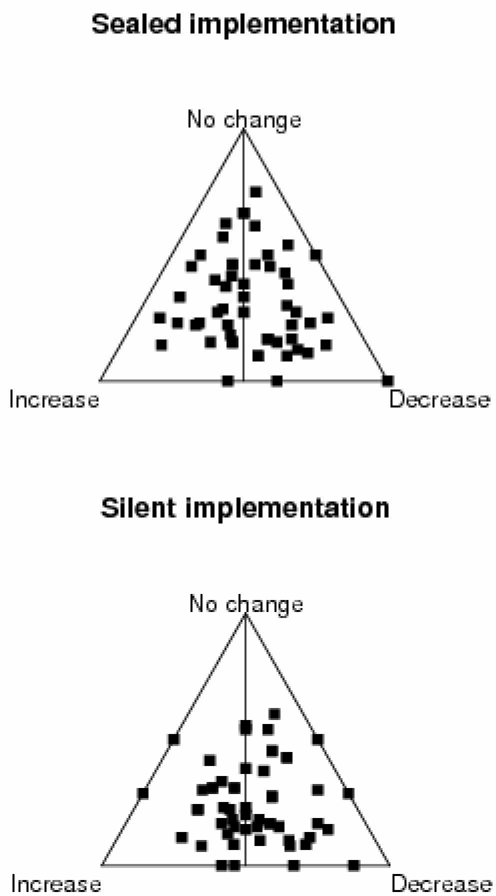

Figure 3.4 Frequency with which subjects increase, decrease, or do not change bids. Each point is one subject

DISCUSSION

We replicate two empirical regularities in laboratory first-price auctions with private values. A sealed implementation generates higher market prices than a Dutch implementation, even though the two are considered isomorphic in theory. Furthermore, both implementations give market prices in excess of those predicted by risk-neutral Bayes-Nash equilibrium.

Our design extends these replications in some new directions. These differences are persistent over a session of 60 market periods. Thus, advance knowledge of the relatively large number of periods does not seem to encourage subjects to recognize that bidding less aggressively is in fact (expected) earnings-enhancing, nor does repetition of the auctions lead to this realization. In addition, an interface in which the magnitude of earnings is prominently displayed in the feedback space does not encourage subjects to bid less aggressively. Finally, we observe that there is more heterogeneity among price levels in cohorts using the Dutch relative to those using the sealed implementation.

To investigate the differences further, we introduce an intermediate synthetic implementation in which the market operates like the Dutch auction with a descending clock, but in which the choices of the bidders are not revealed until the clock reaches the lowest price, simulating a sealed bid. We find that market prices in this implementation generally lie between those generated by the Dutch and the sealed, while maintaining the heterogeneous across-cohort behavior observed in the Dutch implementation.

We use the fact that we observe all bids in the silent implementation to investigate how subjects respond and adapt over the course of the session. Subjects tended to make changes in their bidding behavior more often in the silent implementation, compared to the sealed[29]. This could occur because subjects are cued to think more carefully about the consequences at the margin of purchasing at the current price, versus risking another tick of the clock. However, there is no indication of

[29] While we did not make systematic records, our observations during the session were that the body language and reactions of subjects in the clock-based sessions were different than in the sealed-bid. Subjects were generally more likely to be leaning forward and studying the screen intently during the clock treatments, and to react to the results visibly, while subjects in the sealed implementation often sprawled or reclined in their chairs.

systematic learning in the sense of decreasing bids. Even though recruiting procedures were identical, the two clock-based implementations exhibit more heterogeneity across cohorts in both the level of prices, and the way those prices adapt over time, than the sealed implementation. This suggests that, in addition to the institutional effects on the level of prices reported in previous work, there is more heterogeneity, at least in our subject population, in the perception of how to bid in the clock-based implementations.

Finally, we conclude with a few casual observations. When we distributed screenshots of the instructions to the subjects, we instructed them that they were "free to mark up those pages in whatever way they might find helpful." While most subjects made no marks or doodled, a few chose to attempt certain forms of market analysis. A few subjects chose to track their earnings per minute, which we interpreted as further evidence that controlling the length of the session is a significant aspect of testing the isomorphism. No subjects tracked the most relevant datum - the distribution of the market prices they observed - though several attempted to approximate the average market price. However, among those subjects, none attempted to distinguish periods in which they purchased the object from those in which they did not. Note that those market statistics, then, are generated by the maximum bid out of three bids, whereas in computing an optimal bid in a three-bidder market, the subject is interested in the assessed distribution of the maximum bid out of the two other bids in the market. While we do not assert a direct link between these statistics and the way those bidders formulated their bids, we note that failure to correct for one's own effect on the history of market prices would also lead to more aggressive bidding.

Dorsey and Razzolini (2003) investigate bidding behavior against robots programmed to bid according to the risk-neutral Nash equilibrium in four bidder sealed-bid auctions. They consider treatments in which the interface does compute the probability of winning with any given bid, and those in which it does not. For high realizations of the private value, the range for which the choice of bid is most important, they find that giving the subjects the probability of winning makes bidding less aggressive. Their results are consistent with the hypothesis that when subjects generate impressions of the pattern of market prices, they fail to correct for the fact they themselves are part of the process that generates those market prices.

CHAPTER IV
RESERVATION VALUES IN LABORATORY AUCTIONS: CONTEXT AND BIDDING BEHAVIOR

INTRODUCTION

Laboratory experiments in economics intermediate between pure theory and empirical observations in the field. In the lab, experimenters observe the decisions of real, human agents, while being able to control at least some environmental variables. Whether implicitly or explicitly, experimenters deal in two kinds of mappings: the mapping between a theoretical model and the laboratory environment, and the mapping between a field environment and its laboratory counterpart. The usefulness of laboratory results in improving theoretical models and in understanding field behavior depends on the validity of these mappings. In the terminology of the recent survey of Schram (2005), these can be thought of as the "internal validity" and the "external validity," respectively, of a design. The interplay among theory, lab, and field is particularly evident in auction markets. Results of laboratory auctions have been used to refine auction models and theories of bidding (Cox et al., 1988), and have informed the design of mechanisms in the world at large (for example, Roth, 2002).

The strategic consideration faced by a bidder in a private-values first-price auction is a price-probability tradeoff. A higher bid increases the probability of winning the auction, but decreases the surplus the bidder gains when he wins, because he pays a higher price. Theoretical models assume agents reason about this price-probability tradeoff. In the field, it is taken as a given that they do. The validity of a private-values, first-price auction experiment depends on the salience of this tradeoff to the subjects.

In the lab, a robust finding is that subjects bid significantly more aggressively than predicted by the risk-neutral Bayes-Nash equilibrium. (See the survey of Kagel, 1995 for cites and discussion.) We argue that this finding is an artifact of the standard method for inducing incentives in these experiments. Using an alternate presentation of the incentives, in which subjects are paid according to the total consumer surplus they generate, we find that bidding is significantly less aggressive, even though the risk-neutral Nash equilibrium prediction is unchanged.

This result is most closely related to the findings reported by Isaac and James (2000). In a within-subjects design, they show that the level of risk aversion implied by a subject's choice in the BDM mechanism is often quite different from the level implied by that subject's bidding behavior in a private-values first-price auction against simulated bidders. Furthermore, the ranking of subjects from most to least risk-averse differs substantially between the BDM and first-price auction tasks. These tasks, while isomorphic in the eyes of the decision theorist, are framed sufficiently differently that subjects may bring different heuristics to bear on the two tasks. Our results obtain within the same institutional space, in that subjects participate in a first-price auction in both treatments.

In evaluating an institution's performance, an objective of the experimental method is to separate regularities which are inherent to the institution from observations which are artifacts of experimental design. In the language of Smith (2002), interpreting laboratory results assumes that a set of auxiliary hypotheses relating to the implementation of the experiment hold. The validity of these auxiliary hypotheses cannot be directly observed, but their plausibility can be assessed in part by considering modifications to an experimental protocol.

In the context of understanding individual preferences and choice behavior, Plott and Zeiler (2005) investigate the “willingness to pay/willingness to accept gap,” the claim that there is a systematic difference between the amount a subject is willing to pay for an object and the amount for which he is willing to sell the same object. They show that the gap can be turned on and off by the choice of experimental procedure. They note that “this variation in experimental results undermines the claim that the gap is a fundamental feature of human preferences.” Our result is analogous in showing that the bidding behavior reported in the literature is not a fundamental institutional feature of the first-price private-values auction in the laboratory.

The paper is organized as follows. Section 2 motivates the design choices leading to our presentation of the auction environment. Section 3 describes the experimental protocols, and Section 4 reports the results. Section 5 concludes with a discussion.

PRESENTING AUCTION ENVIRONMENTS IN THE LABORATORY

The first-price auction with a single, indivisible object for sale is generally modeled as a Bayesian game, in the tradition of Vickrey (1961). In the independent private values version, each bidder has a private, idiosyncratic reservation value for the object. These values are independently distributed over some interval, and the distribution is common knowledge. The bidders submit bids simultaneously (in a first-price sealed-bid auction) or using a clock mechanism (in the Dutch auction), and the highest bidder purchases the object. If bidders are risk-neutral, the payoff to the winning bidder is the difference between his reservation value and the price he pays. The utility of the outcome for bidders who do not purchase is normalized to zero.

This environment, with the private values drawn from a uniform distribution, has been studied extensively in the laboratory. In this literature, reservation values are presented using the methods and terminology developed in Coppinger et al., (1980), Cox et al., (1982), and Cox et al., (1988). The instructions describe the reservation value as a cash "resale value." The bidder who purchases the fictitious object on auction sells it back to the experimenter for this amount, and earns the difference between the resale value and the price he pays in the auction. Bidders who do not purchase the object receive monetary earnings of zero.

The resale value protocol uses a direct translation of the utility function from the standard auction model, where the reservation value is motivated by the artifice of the subject selling the object back to the experimenter. This translation is straightforward and clear to anyone familiar with the standard auction model, but may not communicate the nature of the experimental task to a non-specialist in the same way. In the context of their decision task, Plott and Zeiler comment that "[d]ecision theorists might find the language used to describe procedures to be very clear because they are trained to give an operational meaning to technical language." Therefore, we consider a different way to make the concept of a reservation value operational to our non-specialist subjects.

In the field, a reservation value may be determined by the existence of opportunities to purchase a close substitute outside the auction market. Consider a consumer who wishes to purchase an iPod. iPods are frequently sold on Internet auction sites such as eBay. iPods are also widely available at electronics stores. Suppose the consumer has already made the decision to purchase an iPod, but is willing to try an online auction to get a better deal than is available locally. If the consumer fails to win the eBay auction, he then purchases locally. The implied reservation value generated by the possibility of store purchase varies across consumers. Posted prices at stores may depend on geographic location. In addition, consumers differ in the cost of traveling to a store, due to physical distance or opportunity cost of personal time. Thus, consumers have idiosyncratic private reservation values.

Regardless of where he purchases, though, the consumer engages in an economic activity that is essentially the same. In either case, he purchases an iPod at a price lower than his maximum willingness to pay, and he earns positive consumer surplus. The only distinction between winning and not winning the eBay auction is the price he actually pays in the end. Thus, there is a parallel structure between the two outcomes. More generally, if a consumer does not purchase an object in an auction, he will instead participate in some other gainful exchange with the unspent money.

The standard scheme for presenting reservation values does not maintain this parallel structure. Instructions for these experiments necessarily distinguish between how earnings are calculated in the case in which the subject wins the auction, versus when the subject does not. When the earnings for not winning are set to zero, there is a textual difference in the presentation of the earnings calculation. Specifically, when a subject wins, earnings are computed according to a formula like "resale value minus purchase price." When a subject does not win, no formula is needed; his earnings are zero.

Thus, earnings are positive if and only if the subject is successful in increasing consumer surplus. This further emphasizes the dichotomous presentation by segregating the outcomes into those with positive earnings versus those with zero earnings. There is one, and only one, way to earn positive earnings in the experiment: win. Discussions we have led following classroom auction experiments suggest that subjects do take note of the dichotomy and use it as an input in their decision-making process. Despite using neutral terminology, such as "market" instead of "auction" and "purchase" instead of "win," students frequently indicate they chose their bids to "try to win" the auction, or to avoid "getting no payoff This undermines the salience of the tradeoff an expected earnings-maximizing bidder would make between the probability of purchasing the object versus the consumer surplus from that purchase.

Section 3 presents a design for presenting private values which maintains the parallel presentation of the outcomes in the auction. All bidders have an identical maximum willingness to pay for the object. Each bidder receives an idiosyncratic outside price. This outside price serves as the reservation value from the theoretical model. The winning bidder purchases the object in the auction at his bid. The other bidders purchase the object elsewhere at their respective outside prices.

This design presents the outcomes in a way which is both textually and conceptually parallel. In each period, every subject purchases an object. Earnings are always computed using the formula "maximum willingness to pay minus the price paid," that is, the consumer surplus. The only difference between the outcomes is how the price paid is determined. All subjects earn a positive amount each period, so the two outcomes are no longer distinguishable based on whether earnings are positive or zero.

DESIGN

We report results on three sessions in each of four cells of a 2×2 design, for a total of 12 experimental sessions. One dimension manipulates the presentation of the reservation value, one treatment using the standard "resale value" method and the other using an "outside price" frame. The second dimension varies the choice of auction implementation between the sealed-bid and Dutch mechanisms.

Each cohort consisted of 9 subjects recruited from the undergraduate student body at Texas A&M University; no subject participated in more than one session. All interaction among the subjects was mediated via computer in the Economic Research Laboratory at Texas A&M. In each of 60 periods, the subjects were randomly matched into three markets, each with three subjects. The number of periods was announced in the instructions. The matching was done anonymously, and no subject ID numbers or other information about which subjects were participating in which markets in which periods was known to the subjects. At the end of each period, subjects only found out the market price; no information about non-winning bids was revealed.

The design extends the protocol from Turocy et al., (2007). In each period, each subject received a reservation value drawn uniformly from the set {\$0.15, \$0.30, …, \$5.85, \$6.00}. These were drawn independently across subjects and across periods. In sessions using the standard "resale value" (RV) method for presenting reservation values, the instructions read

> *Your Earnings for a period will depend on whether you purchase the commodity in your market, and on the Market Price.*
> *If you purchase a unit of the commodity, your earnings for that period will be calculated according to the equation*
> *Your Earnings = Resale Value - Market Price*
> *If you do not purchase a unit of the commodity, then your earnings for that period will be zero.*

In the "outside price" (OP) treatment, this text was replaced with the language

> *You will purchase exactly one unit of the commodity each period. If you purchase the unit of the commodity in the market, your earnings for that period will be calculated as*
> *Your Earnings = \$6.20 - Market Price*
> *If you do not purchase the unit of the commodity in the market, then you will purchase a unit outside the market at your Outside Price. Your Earnings for the period are then computed as*
> *Your Earnings = \$6.20 - Outside Price*

Bids in the sealed-bid, and clock increments in the Dutch, were restricted to the set {\$0.10, \$0.20, …, \$6.10, \$6.20}. In both the RV and OP treatments, it is a symmetric

Bayes-Nash equilibrium for risk-neutral bidders to bid 2/3 of their signal. To allow comparability across sessions, the same sequence of reservation values and matching into markets was used. In the "resale value" sessions, subjects received the sum of their earnings for all 60 periods. To maintain the same level of expected earnings given the same bidding behavior, for the "outside price" sessions, subjects were paid their earnings from 7 of the 60 periods, which were drawn at random at the end of the session.

EXPERIMENTAL RESULTS

We organize our data analysis around three principal results.

1. Market prices are significantly lower in the outside price sessions.
2. Individual bidding behavior in the sealed-bid is significantly less aggressive in the outside price sessions.
3. Market prices are lower in the Dutch than the sealed-bid using both methods of presenting the reservation value.

Table 4.1 Statistics on market performance for all cohorts. Cohorts labeled RV used the resale value frame; those labeled OP used the outside price frame.

		(a) Market/Theory			(b) Market vs Theory			
Type	Cohort	1-10	6-15	51-60	Above	Equal	Below	Within 0.20
Sealed	RV-1	1.295	1.299	1.296	168	3	9	13
	RV-2	1.356	1.350	1.238	179	2	0	4
	RV-3	1.301	1.261	1.257	175	2	3	13
	Mean	1.320	1.302	1.266	173.8	1.8	4.3	10.0
	OP-1	1.113	1.200	1.144	147	8	25	41
	OP-2	1.130	1.197	1.182	142	7	31	46
	OP-3	1.219	1.198	1.100	117	6	47	42
	Mean	1.154	1.198	1.142	135.5	7.0	34.3	43.0
Dutch	RV-1	1.173	1.185	1.194	160	7	13	31
	RV-2	1.120	1.144	1.123	140	7	33	43
	RV-3	1.195	1.094	1.041	136	5	39	49
	Mean	1.205	1.171	1.169	155	5.5	19.5	41
	OP-1	0.867	0.961	1.089	93	19	68	70.0
	OP-2	1.025	1.029	1.087	117	11	52	59
	OP-3	1.069	1.037	1.019	102	19	59	63
	Mean	0.987	1.009	1.065	104.0	16.3	59.7	64.0

Result 1. Market prices are significantly lower in both sealed-bid and Dutch under OP.

We evaluate market performance relative to the risk-neutral Bayes-Nash equilibrium (RNNE) benchmark. First, for each market in each period, we compute the ratio of the realized market price to the RNNE prediction. If the realized price equals the theory price, this ratio is 1. Figure 4.1 plots ten-period moving averages of the market-to-theory ratio. The RV cohorts are plotted as solid lines, and the OP cohorts as dashed lines. In the sealed-bid, the treatment effect is visually significant. The time series for all cohorts using OP lie everywhere below those using RV. In the Dutch, the effect is most evident early in the sessions. In contrast to the other three treatments, the time trend in the Dutch under OP is upwards.

Table 4.1 reports market-level results for the twelve cohorts. The six cohorts using RV are those reported as cohorts *a* in Turocy et al., (2007). Part (a) of Table 4.1 contains averages of the market-to-theory ratio for each cohort across three ten-period intervals: periods 1-10, periods 6-15, and periods 51-60. The average across cohorts for the market-to-theory ratio is higher for each interval under RV than under OP. We formally test this with the null hypothesis that the average ratio across cohorts is equal, versus the two-sided alternative. For the sealed-bid cohorts, we reject this null hypothesis for each interval (p-values .020 for periods 1-10, .055 for periods 6-15, and .017 for periods 51-60). For the Dutch, we can reject the null hypothesis for periods 1-10 (p-value .090) and 6-15 (p-value .021), but we cannot reject the null of equality late in the sessions (p-value .355 for periods 51-60).

We conjecture that the upward price trend in the Dutch sessions arises because experimentation is relatively inexpensive under OP. There is only a 7 in 60 chance a period will be selected to count towards earnings. Even if it is selected, the bidder still receives positive earnings for that period, irrespective of whether he wins the auction. A savvy first-period strategy is to submit the minimum bid of $0.10 independent of the reservation value the bidder receives. This strategy provides the most information about others' behavior.

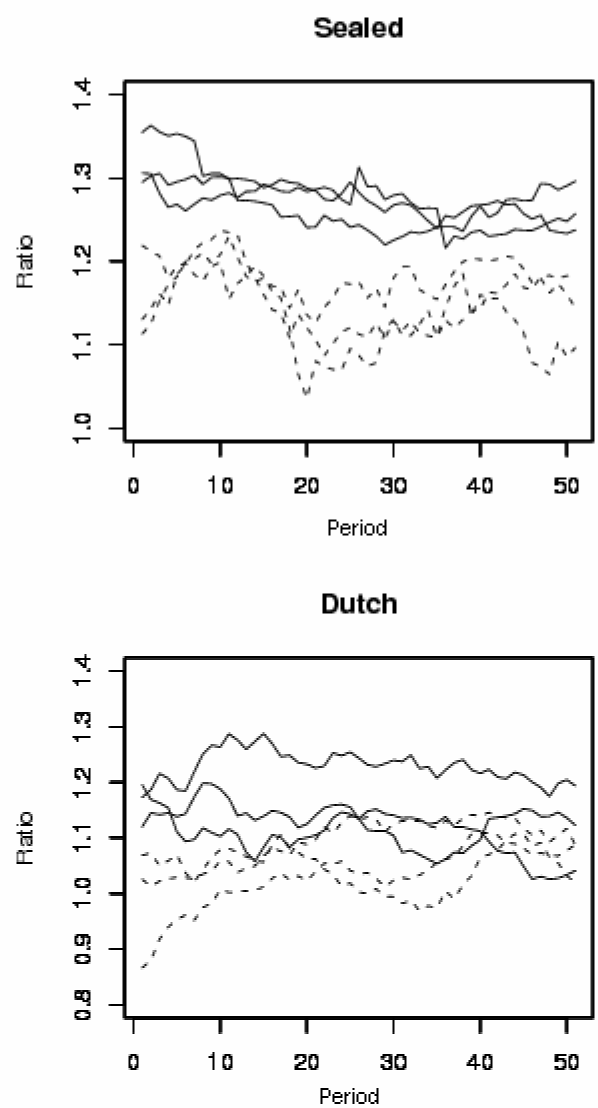

Figure 4.1 Ten period moving averages of the ratio of market price to the risk neutral prediction, by cohort. Solid lines represent cohorts using the RV frame, dotted lines the OP frame.

Because only the winning bid is reported, winning the auction is informationally costly. If the bidder does not win the auction, he learns the maximum of the other two bids submitted in his market; if he does win the auction, he learns only that the maximum of the other two bids was less than his bid. The usefulness of this experimentation might be more transparent in the Dutch implementation, as the clock-based presentation might suggest the idea of playing a waiting game, or game of chicken, in the first few periods.

Another way to compare the observed market prices to theory is to ask how often the market price was greater than or less than the RNNE prediction. For each cohort, counts of these events are presented in group (b) in Table 4.1. In both sealed-bid and Dutch, the RNNE prediction comes closer to being a "median" prediction under OP. To operationalize this, we test the null hypothesis that the average across cohorts of the proportion of markets where price exceeds RNNE is the same between RV and OP, against the two-sided alternative. This null hypothesis is rejected for both the sealed-bid (p-value 0.043) and Dutch (p-value 0.016).

Finally, one can ask how often the RNNE prediction gives a reasonable estimate of the realized market prices. The final column in Table 4.1 counts the number of markets in each cohort in which the market price falls within 20 cents of the RNNE prediction. This is within two bid increments in the sealed-bid, or two clock ticks in the Dutch. The mean number of markets in a cohort satisfying this criterion is higher under OP (p-values 0.0023 for the sealed-bid and 0.029 for the Dutch, using the two-sided alternative).

Result 2. Across subjects, individual bidding behavior in the sealed-bid is less aggressive under OP.

We estimate the linear model

$$\text{bid}_{i\,s\,t} = \alpha_{i\,s} + \beta_{i\,s}\times\text{value}_{i\,s\,t} + \varepsilon_{i\,s\,t} \tag{4.1}$$

for each bidder i in each cohort s, where t denotes the period number. The left panel of Figure 4.2 shows the distribution of estimated bid function slopes β under OP and RV. Under RV, the estimated slopes range from 0.67 to 0.98, with a median of 0.82. This median is consistent with findings in, for example, Cox et al., (1988), who report a "typical" slope of 0.835, and Katok and Kwasnica (2003), who report 0.807. Under OP, the slopes range from 0.47 to 0.98, with a median of 0.71. Only three slopes under RV are less than the median slope under OP, and only one slope under OP exceeds 0.85.

We also propose a more stringent test of the theory, by computing the percentage of periods in which a bidder bids within twenty cents, or two bid increments, of the RNNE prediction. We restrict our count to periods in which a bidder had a reservation value of at least $3.00. This recognizes that for low realizations of the reservation value most or all individually rational bids would fall within the twenty-cent range, as well as restricting attention to scenarios in which a bidder might perceive himself to have a reasonable chance of winning the auction.

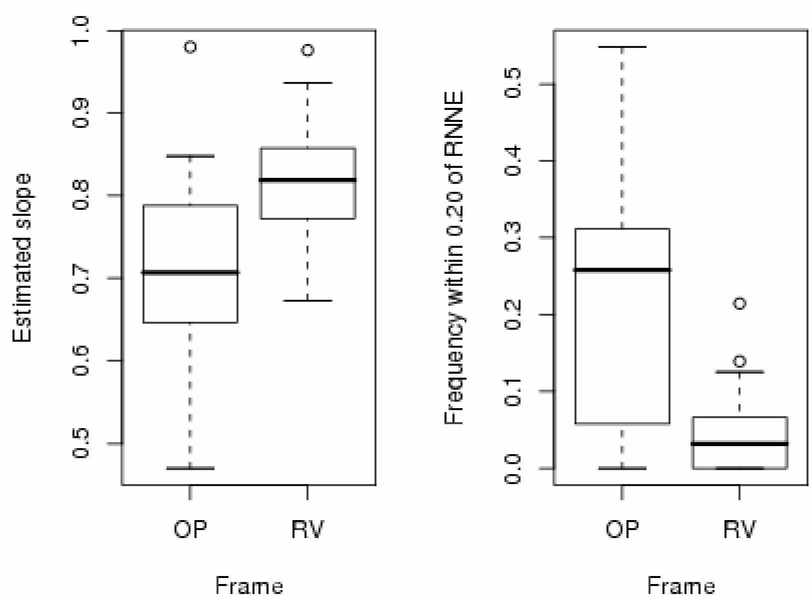

Figure 4.2 Distribution of the estimated bid function slopes for bidders in the sealed bid implementation. Left panel shows the slope. Right panel shows the distribution of the proportion of bids within 0.20 of the risk neutral prediction.

The right panel of Figure 4.2 plots the distribution of these frequencies. Under RV, the median percentage of bids falling within the twenty-cent interval is 3%. Eleven bidders registered *zero* bids within that interval; four bidders exceeded 10%, with a maximum of 21%. Under OP, four bidders had no bids in the interval. The median percentage in the interval was 26%, with a maximum of 55%.

Result 3. Market prices in the Dutch are below those in the sealed-bid in both frames.

The regularity that the Dutch auction gives lower market prices than the sealed-bid remains true under OP. The null hypothesis that the average of the market-to-theory price ratios is equal in sealed-bid and Dutch, tested against the two-sided alternative, can be rejected for all three intervals in Table 4.1. For the sessions under RV, p-values for the tests are .0026 for periods 1-10, .00081 for 6-15, and .023 for 51-60. For the sessions under OP, the p-values are .094 for periods 1-10, .016 for periods 6-15, and .080 for periods 51-60.

DISCUSSION

We show that bidding behavior in private-values, first-price laboratory auctions is sensitive to the presentation of the outcomes. With a frame in which the payoff computation for outcomes is presented in parallel, subjects bid less aggressively, and more often submit bids which are close to the prediction of the risk-neutral Bayes-Nash equilibrium. This is evidence in favor of the proposition that subjects perceive the environments differently, even though a theorist would consider the payoff transformation irrelevant for a risk-neutral bidder.

In a study of field auctions for antique collectible United States coins, Harrison et al., (2007) write

> *We hypothesize that there is a danger that the imposition of an exogenous laboratory control might make it harder, in some settings, to make reliable inferences about field behavior. The reason is that the experimenter might not understand something about the factor being controlled, and might impose it in a way that is inconsistent with the way it arises naturally in the field… (p. 433)*

The reservation value in a private values action may be such a factor. While specialists are comfortable with the idea that every agent has a reservation value, in practice, such a reservation value is latent, insofar as an agent does not stop and formally assess a particular number unless called upon to do so. In imposing controls to establish the private values setting, the concept of the reservation value must be presented in a way that communicates the concept in terms familiar to the subject.

The ultimate goal of absolutely establishing the internal and external validity of any design for an auction experiment is, by definition, unreachable. Our results represent progress in understanding how to map theory and field to the lab by proposing an alternative means of establishing the private value setting. Two translations of the private values environment elicit very different behavioral patterns. Which method – if either – adequately establishes the control that maintains the validity of those mappings is a topic for future consideration and research.

CHAPTER V
CONCLUSIONS

In our exploration of Case-Based Decision Theory, we find that subjects do show some propensity to use case based reasoning in their decision making process. We consider it encouraging that, when able, CBDT organized as large a proportion of subjects' decisions as did the simpler Max-Heuristic. However, the results of this rest upon our assumption that feature-based similarity is the main similarity consideration that subjects are employing. Further research is needed in the area of what similarity functions are appropriate for use in these types of tasks. Also, the monopoly decision making environment used here may not be the one most conducive to case based reasoning. A future area of research may be to explore other environments and decision making tasks.

After careful construction of a graphical interface in which sealed-bid and Dutch implementations of the first-price private-values auction can be presented as similarly as possible, we find that the non-isomorphism persists. Allowing the subjects to engage in 60 independent auctions does not alleviate this result; there appears to be very little, if any, learning over time in the auctions. However, we do find that an implementation which uses a clock mechanism but does not reveal when others have bid elicits market prices between those of the sealed-bid and Dutch implementations. We conclude from this that there is something about the clock mechanism that helps to focus subjects on the price-probability trade-off central to the auction.

In a second attempt to better focus subjects on the price-probability trade-off in an auction environment, we find that presenting the reservation value as an outside price results in market prices that are significantly lower than those in the standard frame. We also find that subjects' bids are now much better organized by risk-neutral Bayes-Nash equilibrium predictions. Continued efforts are needed to pinpoint exactly which behaviors are inherent to the auction mechanism and which are anomalies and artifacts of the experimental design.

REFERENCES

Cason, T. N. & Friedman, D. (1997). Price formation in single call markets. *Econometrica*, *65*, 311–345.

Cooper, A. & Reimann, R. (2003). *About Face 2.0: The Essentials of Interactions Design*. Indianapolis, IN: Wiley.

Coppinger, V. M., Smith, V. L. & Titus, J. A. (1980). Incentives and behavior in english, dutch, and sealed-bid auctions. *Economic Inquiry*, *18*, 1–22.

Cox, J. C., Roberson, B., & Smith, V. L. (1982). Theory and behavior of single object auctions. In V.L. Smith (ed.), *Research in Experimental Economics*, Volume 2. Greenwich, CT: JAI Press.

Cox, J. C., Smith, V.L., & Walker, J. M. (1988). Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty*, *1*, 61–99.

Dorsey, R. & Razzolini, L. (2003). Explaining overbidding in first price auctions using controlled lotteries. *Experimental Economics*, *6*, 123–140.

Fischbacher, U. (1999). z-Tree - zurich toolbox for readymade economic experiments: Experimenter's manual. working paper, University of Zurich.

Gayer, G., Gilboa, I., & Lieberman, O. (2004). Rule-based and case-based reasoning in housing prices. working paper, Tel-Aviv University.

Gigerenzer, G. & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650-669.

Gilboa, I. & Schmeidler, D. (1995). Case-based decision theory. *Quarterly Journal of Economics*, *110*, 605-639.

Gilboa, I. & Schmeidler, D. (1997). Act similarity in case based decision theory. *Economic Theory*, *9*, 47-61.

Gode, D. & Sunder, S. (1993) Allocative efficiency of markets with zero intelligence (ZI) traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, *101*, 119–137.

Goldstone, R. L. & Son, J. (2005). Similarity. In K. Holyoak & R. Morrison (eds.), *Cambridge Handbook of Thinking and Reasoning*. Cambridge, UK: Cambridge University Press.

Guerdjikova A. (2002). Case-based decision theory - An application for the financial markets. working paper, University of Heidelberg

Guerdjikova, A. (2003). On the definition and existence of an equilibrium in a financial market with case-based decision makers. working paper, University of Heidelberg.

Harrison, G. W., List, J. A., & Towe, C. (2007). Naturally occurring preferences and exogenous laboratory experiments: A case study of risk aversion. *Econometrica*, *75*, 433–58.

Holyoak, K. J. & Gordon, P. C. (1983). Social reference points. *Journal of Personality and Social Psychology*, *44*, 881-887.

Isaac, R. M. & James, D. (2000). Just who are you calling risk averse? *Journal of Risk and Uncertainty*, *20*, 177–187.

Jahnke, H., Chwolka, A. & Simons, D. (2005). Coordinating service-sensitive demand and capacity by adaptive decision making: An application of case-based decision theory. *Decision Sciences*, *36*, 1-32.

Kagel, J. H. (1995) Auctions: A survey of experimental research. In J.H. Kagel & A.E. Roth (eds), *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.

Katok, E. & Kwasnica, T. (2003). Time is money: The effect of clock speed on seller's revenue in dutch auctions. working paper, Pennsylvania State University.

Ortony, A., Vondruska, R., Foss, M. & Jones, L. (1985). Salience, similes, and the asymmetry of similarity. *Journal of Memory and Language*, *24*, 569-594.

Plott, C. R. & Zeiler, K. (2005). The willingness to pay/willingness to accept gap, the 'endowment effect,' subject misconceptions and experimental procedures for eliciting valuations. *American Economic Review*, *95*, 530–45.

Podgorny, P. & Garner, W. R. (1979). Reaction time as a measure of inter- and intra-object visual similarity: Letters of the alphabet. *Perception & Psychophysics*, *26*, 37-52.

Polya, G. (1954). *Mathematics and Plausible Reasoning, Vol. 1: Induction and Analogy in Mathematics.* Princeton, NJ: Princeton University Press.

Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, *70*, 1341–78.

Roth, A. E. & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, *8*, 164-212.

Rubinstein, A. (1988). Similarity in decision making under risk (Is there a utility resolution to the allais paradox?). *Journal of Economic Theory*, *46*, 145-153.

Sarin, R. & Vahid, F. (1999). Payoff assessments without probabilities: A simple dynamic model of choice. *Games and Economic Behavior*, *28*, 294-309.

Sarin, R. & Vahid, F. (2004). Strategy similarity and coordination. *Economic Journal*, *114*, 506-527.

Schram, A. (2005). Artificiality: The tension between internal and external validity in economic experiments. *Journal of Economic Methodology*, *12*, 225–37.

Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, *1*, 43-62.

Selten, R. & Buchta, J. (1998) Experimental sealed bid first price auctions with directly observed bid functions. In D. Budescu, I. Erev, & R. Zwick (eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Mahwah, NJ: Lawrence Erlbaum Associates

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99-118.

Smith, V. L. (1982). Microeconomic systems as an experimental science. *American Economic Review*, *72*, 923–955.

Smith, V. L. (2002). Method in experiment: rhetoric and reality. *Experimental Economics*, *5*, 91-110.

Turocy, T. L., Watson, E., & Battalio, R. C. (2007). Framing the first-price auction. *Experimental Economics*, *10*, 37–52.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327-352.

Tversky, A. & Gati, I. (1982). Similarity, separability, and the triangle inequality. *Psychological Review*, *89*, 123-154.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*, 8–37.

WORKS CONSULTED

Cox, J. C., Smith, V. L., & Walker, J. M. (1983). A test that discriminates between two models of the Dutch-first auction non-isomorphism. *Journal of Economic Behavior and Organization*, *4*, 205–219.

Glasnapp, D. & Poggio, J. (1985). *Essentials of Statistical Analysis for the Behavioral Sciences* Columbus, OH: Charles E. Merrill Publishing Company.

Siegel, S. & Castellan, Jr., N. J. (1988). *Nonparametric Statistics for the Behavioral Sciences* (2nd Ed.). Boston, MA: McGraw-Hill.

Turocy, T. L. (2001). *Computation and Robustness in Sealed-Bid Auctions*. PhD dissertation, Northwestern University, Evanston IL.

APPENDIX A

INSTRUCTIONS

This is an experiment in the economics of decision making. Texas A&M University has provided funds for this research. If you follow the instructions and make good decisions, you can earn an appreciable amount of money. At the end of today's session, you will be paid your earnings in private and in cash.

It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc. you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

During this session you will be acting as a firm who is selling a good. You will be selling your good to 30 independent markets. You can think of these as 30 geographically separated islands. In each of the 30 markets (islands) you are the *only* seller of the good. This means that nothing any other seller or firm does can affect you or your market. Each period represents a new market and you will have to make a decision about how many units you want to produce for that market. It is costly to produce this good and if you produce units that do not get sold in that market, you will NOT be able to keep those units for use in other markets. At the end of each period you will earn profits on the units of your good that you do sell in that market.

At the beginning of each period you will receive a *Marketing Report* that contains information regarding some *Market Conditions* for the current market. You can think of this as information about the market that has been gathered for you by the Marketing Department of your firm. Gathering this data is costly to your firm, as such your Marketing Department is not able to gather all information in every market. Therefore, the information that your Marketing Department does collect can vary from market to market. However, nothing you or anyone else does can change what information is gathered in any market.

After gathering the data the Marketing Department sends it to you. Unfortunately there is an error that occurs during that transmission. Instead of receiving the actual data all you receive is a list of the Market Conditions that were collected and a table of symbols representing the actual data. Fortunately the error is consistent. This means that identical symbols for a given Market Condition represent the same actual data. For instance, if the Marketing Department gathers data that says median income is $35,000 and a blue triangle gets transmitted, then whenever the Marketing Department reports $35,000 for median income it will be transmitted as a blue triangle in the Marketing Report. However, a blue triangle can also appear for other Market Conditions, where it does not necessarily stand for $35,000. For example, if the Marketing Department gathers data on a high inflation rate and this information gets transmitted as a blue triangle, then all Marketing Reports with a high inflation rate will have a blue triangle in the table for inflation rate.

After you have received the *Current Marketing Report* you will be asked to choose a *Production Value* of *50*, *100*, *150*, or *200* units. Your *Profits* each period depend on your Production Value and may also depend upon some of the Market Conditions. After you have made your production choice, you will be informed of your Profits for that market. You will then proceed to the next market where you will be given the new market's Current Marketing Report. You will then be asked to choose a Production Value for that market. The session will continue in this manner until you have made production choices for 30 markets.

In order to help you with your decisions, for each market the experimenter has included four different scenarios. In each of these scenarios the experimenter was given a Marketing Report similar to the Marketing Reports that you will be given. The experimenter then chose a Production Value in such a manner that each of the four production choices was chosen once. The Profits reported in these scenarios are the profits that would have been earned in that market given the reported Market Conditions and the chosen Production Value. The Profits reported in these scenarios will NOT be included in your *Total Profits*. Your Total Profits consist only of those profits that you earn during the session, i.e. when *you* are making the production decision. Your Total Profits will be calculated by simply adding up the profits you earn in each of your markets.

Figure A.1 – Decision Screen

Figure A.1 gives an example of the decision screen. At the top right of the screen you will see labeled the current market. At the bottom right of the screen you will see your Total Profits, which will include all Profits you have earned so far. On the left side of the screen you will also see a table with the four scenarios for the current market. In the first column from the left you will see labels for the Market Conditions, the Production Values, and Profits. So in this market, your Marketing Department has gathered information on: the Tourist Population, the Wind, the Humidity, the UV Factor, and the Temperature. Looking at the symbols in the table you can see that there is a blue triangle for Tourist Population in Scenario 1 and Scenario 2. This means that the Tourist Population was the same in both of these scenarios. You can also see that there is a blue triangle for Humidity in Scenario 3. While this is the same symbol that was present in Scenarios 1 and 2 for the Tourist Population, it does not necessarily represent the same thing that it did for Tourist Population. Below the scenarios you will see the Production Values and Profits for those scenarios. Again, the Production Values were chosen so that each value was chosen exactly once. The Profits that you see are the Profits that would have been earned had the given Production Value been chosen with the given scenario. On the right side of the screen you will see the Marketing Report for the current market, in the left hand column are the symbols representing the data from the report and in the right hand column are the labels for the different Market Conditions that are reported. On the bottom of the screen you will see the menu of choices for your Production Value.

In order to select a Production Value simply use your mouse to click in the circle to the left of the value you wish to choose. After clicking in one of the circles you MUST click the Confirm button before your choice will be submitted. If you wish to change your choice you may do so at any time BEFORE clicking the Confirm button. You may change your choice of Production Value as many times as you wish. However, once you have clicked the Confirm button you will NOT be able to change your Production Value for the current market. After you have clicked Confirm a results screen will appear and inform you of your Profits for the current market. Once you have finished viewing these results click Continue to move on. After you have clicked Continue, you will proceed to the next market where you will be given the new market's Marketing Reports and asked to make a production choice for that market.

After you have made production choices for 30 markets the session will be over. A screen will appear informing you of your Total Profits and *Total Earnings*. Your Total Earnings are the amount you will be paid in cash. Your Total Earnings are calculated by dividing your Total Profits by 6,000. In other words for every $6,000 in Profits that you made you will earn $1.00 in cash. For instance, if you earn a Total Profit of $96,000 then your Total Earnings will be $16. If you did not choose to receive a hang tag for parking then you will receive a $5.00 show-up fee in addition to your Total Earnings. In that case your *Total Payment* will be calculated by adding the $5.00 show-up fee to your Total Earnings. So in the above example your Total Payment would be $16.00 + $5.00 or $21.00 in cash. However, if you did choose to take a hang tag for parking your Total Payment will be the same as your Total Earnings. Once the session is over and everyone has viewed their Total Earnings you will be called up, one at a time, to be paid privately and in cash. The session will not be finished until everyone has made decisions for all 30 of their markets. After you have finished please wait patiently for all remaining markets to finish. Are there any questions? If you do have a question, please raise your hand and an experimenter will come to you. Do not ask any questions out loud.

VITA

Name: Elizabeth Ann Watson

Address: Texas A&M University, Department of Economics, 3077 Allen Building, 4228 TAMU, College Station, TX, 77843-4228

Email Address: beth818@tamu.edu

Education: B.G.S, major area Economics, University of Louisiana at Lafayette, 2002

Academic Experience: Graduate Assistant, Teaching, Department of Economics, Texas A&M University, Summer 2003, Summer 2004, Fall 2004, Summer 2005, Fall 2005 Spring 2006, and Summer 2007

Graduate Assistant, Non-teaching, Department of Economics, Texas A&M University, Spring 2003, Fall 2003, Spring 2004, Spring 2005, Summer 2006, and Fall 2006

Awards: Regents Fellowship, Texas A&M University, September 2002 – September 2003

M.J. Groves Scholarship, Department of Economics, Texas A&M University, Private Research Enterprise Center, Summer 2004

Outstanding Teaching Assistant, Department of Economics, Texas A&M University, 2006