An Experiment on Case–Based Decision Making*

Brit Grosskopf   Rajiv Sarin   Elizabeth Watson

Texas A&M University   Trinity University

November 7, 2007

Abstract

We experimentally investigate the disposition of decision makers to use case–based reasoning as suggested by Case–Based Decision Theory (Gilboa and Schmeidler, 1995). Our subjects face a monopoly decision problem about which they have very limited information. Information is presented in a manner which makes similarity judgements according to the feature matching model of Tversky (1977) very plausible. In a 2×2 between–subject design, we vary whether information about the current market is given, in addition to a "history," or not. Furthermore, we investigate the impact feedback has on behavior in several independent markets by either giving immediate feedback about obtained profits or not. The results provide some support for the predictions of Case–Based Decision Theory, particularly when no immediate feedback is given. We also find evidence that other simple heuristics are used.

*We thank Gerd Gigerenzer and the audiences at the Max Planck Institute for Human Development, the Department of Economics at Texas A&M and the ESA meetings in Atlanta, Rome and Tucson for helpful suggestions and discussions. Financial support from the Department of Economics at Texas A&M and the NSF is gratefully acknowledged. Contact: `bgrosskopf@econmail.tamu.edu`, `rsarin@econmail.tamu.edu` and `Elizabeth.Watson2@trinity.edu`.

1 Introduction

According to Hume (1748), "From causes which appear similar we expect similar effects." Hume went on to assert that "This is the sum of all our experimental conclusions." A recent formulation of decision making by Gilboa and Schmeidler (1995, 2001) provides a formalization of the ideas of Hume. Gilboa and Schmeidler's Case–Based Decision Theory (CBDT) suggests that similar actions in similar problems give similar payoffs. Indeed, actions are evaluated according to the similarity weighted sum of payoffs they have yielded in similar problems.

CBDT is not proposed as an alternative to or a generalization of the dominant expected utility theory. Rather, it is seen as being relevant in certain situations. In particular, it is thought to be suitable in environments which are sufficiently complex or if the decision maker's information is severely limited so that usage of the expected utility approach is not plausible. In such situations, the decision maker is more likely to rely upon her past experiences to evaluate current choices and reason in the manner suggested by Hume and CBDT.

While CBDT has been the focus of some applied work in economics (e.g. 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 decision maker (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. In CBDT 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 analysis 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). This specific similarity function seems plausible given the description of the decision problem the subjects face in our experiment. 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 (combined 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 very different experimental design.

In order to test the predictions of CBDT we study production choices of a monopolist who is producing for several independent markets about each of which she has some "history." First, we contrast behavior in situations when CBDT can be used to situations when it cannot be used. This is achieved by comparing choices of a monopolist who has some information about the current market conditions with choices when the monopolist does not have such information. In the latter the monopolist has no information on which to base a similarity comparison. Second, we analyze if immediate feedback about the obtained profits makes the subject more (or less) prone to use CBDT. We do this by either revealing profits from a market right after a production choice is made or delaying such feedback until all production choices have been made.

Our results provide some support for CBDT. Predicted choices coincide more often with actual choices when participants are given information about the current situation alongside a market history. The support for CBDT is stronger when payoff information is delayed until the end of the experiment. We interpret this as suggesting that CBDT (coupled with feature based similarity) is more salient when additional information about the performance of the DMs choices is not immediately provided because such information alters the decision making procedure (and/or the similarity relation) used by the DM. However, we also find support for a simple heuristic being used in the limited information environments that we study. In particular, we find that if payoff information is given immediately after production choices for each market are made, 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 and Appendix B provides additional individual data.

2 Preliminaries

2.1 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 newly established 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

$$U(a) = U_{p,M}(a) = \sum_{(q,a,r)\in M} s(p,q)u(r), \tag{1}$$

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]

2.2 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 webpages. 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

[2]CBDT does not make any distinction between an action that resulted in zero utility and one that simply 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 Rubinstein (1988) and Sarin and Vahid (2004) for previous applications in economics.

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 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 and f is an interval scale that represents the salience

[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 (Tversky 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.

or prominence of various features. 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 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

$$\max_a U(a) = U_{p,M}(a) = \sum_{(q,a,r)\in M} \theta f\,(p \cap q)\,r. \tag{2}$$

3 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. Particular attention is paid to the manner in which we induce the decision maker's memory, how the complexity of the task is created and how the similarity between the current and the past problems is generated.

We chose a monopoly framework as an individual decision making task with an economic context. Since CBDT seeks to describe behavior in situations where a DM never faces the same problem twice, we decided to have subjects make production choices for 30 independent markets.[6] In the instructions, and in the questionnaire following the instructions, we made sure that the subjects understood that the markets were independent from one another and that decisions made for one market would not affect any other market. 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 markets they faced during the experiment. We therefore decided to induce a separate "history" for each of the 30 markets and, consequently, a different memory for each market. For each situation, we displayed four "scenarios" or cases. The scenarios that were displayed were different for each market. In each of the four scenarios, a different production choice was made and that choice and the associated payoff that would have been earned were displayed.[7] Thus, the memory of each subject, for each market, consists of 4 cases, their descriptions, actions chosen and resulting payoffs.

We ensured that in each market all production values were a possible

[6] CBDT can, however, be easily modified to study choice in repeated decision problems (see, e.g., Gilboa and Schmeidler, 2001).

[7] 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.

choice for a case–based decision maker by having each memory contain one observation of each possible production value. A case–based decision maker would calculate the similarity between the current market 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 production choices were randomly generated, the final display was not random. We displayed four scenarios and their corresponding production choices such that predictions for CBDT were different from those of other potential decision guiding rules that we discuss in section 4. While it was possible for payoffs to be negative for a given decision, overall accumulated payoffs could never be negative.

In our experimental design we vary whether subjects are given information about the current market in addition to the four scenarios (cases) in the induced memory or not. This was 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 choices 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]

Another aspect of adequately controlling a DMs memory is to prevent

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

learning across situations. We therefore vary whether subjects are informed about their earnings from each market immediately after they made a production choice for a given market or not. If the feedback is delayed, all profits from all markets are displayed at the end of the experiment. The delayed feedback condition eliminates most kinds of adaptive learning that could have potentially occurred across the independent markets. Table I summarizes our 2×2 between–subject design.

TABLE I:
THE 2×2 EXPERIMENTAL DESIGN

		Information	
		Past + Current	Past Only
Feedback	Immediate	39 subjects	30 subjects
	Delayed	33 subjects	31 subjects

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. In order to generate a complex enough environment, in which it is impossible for subjects to figure out potential states of the world or associated outcomes, we let the profit of the monopolist in a given market depend on many different variables that we call market conditions.

We used a total of 12 different market conditions. Subjects, however, did not know this. In each market subjects were only given information on a random subset of 5 conditions.[9,10] In order not to overload our monopoly

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

[10]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.

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. Subjects never saw the last two conditions: # of potential buyers or gas price. Subjects did not even know that these conditions existed. For each market each condition was randomly chosen to have a value of either 1, 2, or 3.

We generated 15 different payoff functions and used each one twice. The payoff functions varied according to which choice maximized the payoff and the "penalty" for making a non–optimal choice.[11] 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. As mentioned before, two conditions (# of potential buyers and the gas price) were never reported to the subjects. One of these was chosen at random to enter the payoff function in each market.

In order to implement a feature–based similarity approach we displayed all values of the market conditions as symbols.[12] 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

[11] To ensure comparability across treatments, the subjects faced the same payoffs function in the same order.

[12] 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."

the same thing, while it could mean different things for different market conditions. All participants faced the same markets in the same order with the same set of cases given to them.

In summary, in all treatments subjects were told that they were a monopoly firm that has to make production choices for 30 different markets. The production choice could be one of four values, 50, 100, 150, or 200 units. In the treatments in which information about the current market 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 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 in the treatment where immediate feedback was provided and then moved on to the next situation (i.e., market). Otherwise the feedback about the profits made in each of the markets was delayed until the end of the experiment.

The experiments were conducted at the Economic Research Laboratory at Texas A&M University between March 2006 and May 2007. 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 maximum of

$24.56 (this includes a $5 showup fee).[13] 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.

4 Hypotheses

Since EUT is not a reasonable alternative decision making procedure in the environment we consider, we look for simple decision making principles that deliver predictions in our environment and that can pose as alternatives 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 allow people to make decisions by simplifying the complex environment people often face (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 production choice being chosen once), the Max–Heuristic (or, MAX) would predict that a DM chooses 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, MAX seems a natural choice. Note also that even MAX is based on similarity arguments. The similarity notion is, however, very different. MAX is entirely based on the assumption that similar production choices yield similar profits

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

independently of how similar the current market is to the scenario in memory. In fact, a DM using MAX simplifies the environment to a stationary situation that does not change.

In order to distinguish whether our subjects used Case–Based reasoning or MAX, we calibrated the predicted choices such that the choices predicted by CBDT were different from those of MAX in all but 7 situations.[14]

We postulate the following two hypotheses. Hypotheses 1 compares the performance of CBDT across the treatments according to whether current information is available or not. To make the kinds of similarity comparisons implied by CBDT, together with feature-based similarity, the agent needs to know some aspects of the current problem. When such information is not available CBDT does not seem like a plausible alternative. In the absence of current information MAX should perform relatively better than CBDT. That is, MAX should explain relatively more choices than CBDT when it is compared in the treatment when current information is not available than when such information is available. If Hypothesis 1 is satisfied, then this comparison can be stated as our Hypothesis 2.

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

H2: MAX is a relatively better predictor of choices than CBDT in treatments when the current information is not available.

[14]An alternative rule of thumb could be to prioritize similarity, i.e., choose the action that has the most features in common with the current report (as long as its payoff is positive). If there are more than one of those actions, choose the one with the highest payoff from this set. If the action with the highest number of features in common with the current report has a negative payoff associated with it, choose the one that has the next highest number of features in common. Again, if there are more of those, choose the one that has the highest payoff associated with it. In our experimental setup predicted choices of such a heuristic coincide with predicted choices of CBDT in all but three markets. We therefore do not separately analyze its predictive power.

We do not have any hypothesis regarding the timing of the feedback. Intuitively, immediate feedback could interfere with our control over the subject's memory as it could lead to subjects "learning" across situation. Past choices that seemed "reasonable" could be reinforced and hence alter the basis of similarity comparisons.[15] We, therefore, check that no such learning was occurring in our analyses of the data. Since choices in line with MAX gave higher payoffs, the inclination to use this heuristic could get reinforced when immediate feedback is provided.

5 Experimental Results

5.1 Relative Performance of CBDT and MAX

We first analyze how well the predictions of CBDT and MAX 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) 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 (individual data can be found in Appendix B, Table VIII). 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 her MSD would equal 2.

For a first test of Hypothesis 1 we compare the median of the individual MSDs of observed behavior from those predicted by CBDT across the different information treatments. We find that the median of subjects' MSDs is

[15] Note that subjects did not even know the range of payoffs. Any judgement of the reasonability of choices through the evaluation of obtained payoffs is highly subjective.

[16] For a first analysis we assume that the aspiration level of a case–based decision maker is zero. This aspiration level can be easily adapted.

significantly lower in the treatments where the current information is available compared to when such information is not available. This indicates better performance of the CBDT predictions when current information is available and provides support for Hypothesis 1. Interestingly, this effect is stronger in the treatment with delayed feedback. Results of Robust Rank Order Tests are given in Table II.

Table II:
Tests for Differences in CBDT MSDs
(w/ Current vs. w/o Current)

Robust Rank Order Test	**z–statistic**
Immediate Feedback	-2.55^{**}
Delayed Feedback	-7.61^{**}

Note: ** indicates significance on the 5% level

Table III:
Tests for Differences in MAX MSDs
(w/ Current vs. w/o Current)

Robust Rank Order Test	**z–statistic**
Immediate Feedback	2.28^{**}
Delayed Feedback	2.05^{**}

Note: ** indicates significance on the 5% level

Table III indicates that the Max–Heuristic predicts actual choices significantly better when no current information is available as compared to when such information is available. Interestingly, this effect is slightly weaker when feedback is delayed.

Table IV reveals support for Hypothesis 2. It shows that when information about the current situation is not given, MAX predicts choices better than CBDT (i.e., individual MSDs of observed choices from predicted choices are lower for MAX). This holds for both feedback conditions. Table IV, however, also demonstrates that CBDT does not predict a larger proportion of choices when compared to MAX when immediate feedback is given, indicating that equally many participants are guided by CBDT as are guided by MAX. However, when feedback is delayed, CBDT predicts choices better than MAX when information about the current situation is given. Hence, the performance of MAX depends upon the availability of immediate feedback.

TABLE IV:
CBDT VS. MAX

	Robust Rank Order Test	**z–statistic**
Immediate	w/ Current	0.03
Feedback	w/o Current	6.04^{**}
Delayed	w/ Current	-4.27^{**}
Feedback	w/o Current	3.34^{**}

Note: ** indicates significance on the 5% level

5.2 Absolute Performance of CBDT and MAX

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 instance, we would like to know if pure random choice model would do better than either CBDT or MAX. Furthermore, we would like to know what proportion of an individual's choices are predicted correctly in each treatment by these models.

If a subject was choosing randomly, one could interpret this as meaning that her choices coincide with theoretical predictions 25% of the time (since there are four choices to choose from). Given that CBDT as well as MAX 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 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: (1) predicted choices coincide with observed choices at least half of the time (in 15 out of 30 situations), and (2) predictions coincide with observed behavior at least two–thirds of the time (in 20 out of 30 situations). Table V shows the percentage of people for whom CBDT predicts at least as well as the given category, i.e. returns a lower MSD. Table VI illustrates the same for MAX.

Table V:
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)
Immediate	w/ Current	79.5% (31/39)	25.6% (10/39)	15.4% (6/39)
Feedback	w/o Current	70.0% (21/30)	3.3% (1/30)	0% (0/30)
Delayed	w/ Current	97.0% (32/33)	39.4% (13/33)	18.2% (6/33)
Feedback	w/o Current	71.0% (22/31)	0% (0/31)	0% (0/31)

[17] Such a calculation would lead to a 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).

Table VI:
Observed Frequencies of individual MSDs
when comparing choices to MAX predictions

		Max–Heuristic predicts correctly more often than:		
		Random (MSD<1.5)	Half (MSD<1)	Two–Thirds (MSD<0.67)
Immediate	w/ Current	87.2% (34/39)	20.5% (8/39)	10.3% (4/39)
Feedback	w/o Current	93.3% (28/30)	33.3% (10/30)	26.7% (8/30)
Delayed	w/ Current	60.6% (20/33)	9.1% (3/33)	6.1% (2/33)
Feedback	w/o Current	80.6% (25/31)	29.0% (9/31)	12.9% (4/31)

By comparing the last column of Table V with the last column of Table VI, 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 (15.4% vs. 0% and 18.2% vs. 0%, Test of Equality of Proportions: immediate feedback, z–value$= 2.25$, one–tailed $p = 0.0122$; delayed feedback, z–value$= 2.49$, one–tailed $p = 0.0064$).[18] In fact, CBDT cannot be applied in the treatments without the current information. Hence, 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 predicts more than two–thirds of the choices correctly for more subjects than CBDT (26.7% vs. 0% and 12.9% vs. 0%, Test of Equality of Proportions: immediate feedback, z–value$= 3.04$, one–tailed $p = 0.0012$; delayed feedback, z–value$= 2.07$, one–tailed $p = 0.0192$), again supporting Hypothesis 2. MAX predicts choices correctly at least two–thirds of the time for more subjects in treatments when

[18] The specific test statistic is $z = (p_1 - p_2)/S_{p_c}$, where p_i is the proportion in subsample 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 subsample i (see Glasnapp and Poggio, 1985).

no current information is provided (26.7%) compared to when such information is available (10.3%) if immediate feedback is given (Test of Equality of Proportions: z–value= 1.78, one–tailed $p = 0.0375$). This difference is not significant when feedback is delayed (12.9% vs. 6.1%, Test of Equality of Proportions: z–value= 0.9386, n.s.).

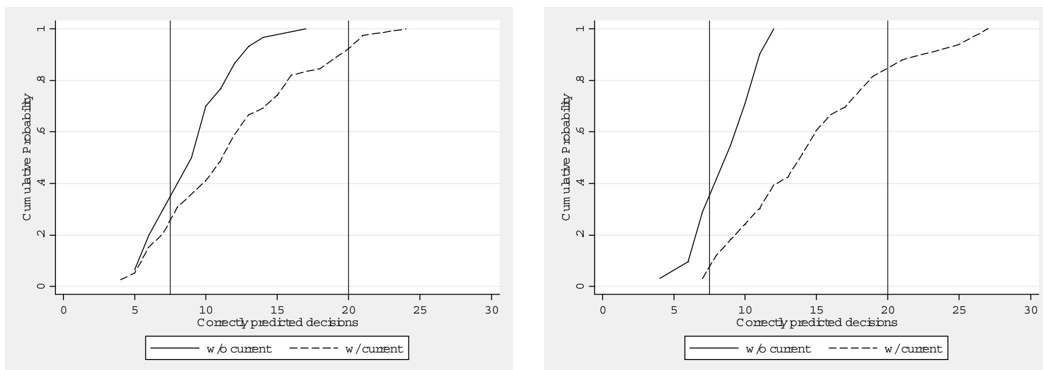

(a) Immediate Feedback (b) Delayed Feedback

Figure 1: Cumulative density of choices in line with CBDT predictions

Figures 1 to 3 graphically show the performance of CBDT predictions. Performance is evaluated by how many of the 30 decisions are predicted correctly, i.e. coincide with actually observed choices. The cumulative density of observed choices coinciding with predicted choices is plotted for the different treatments.[19] To make comparisons easier, two vertical lines, referring to 1/4 and 2/3 of the decisions, are included.[20] The further the lines are to the right, the more often do observed choices coincide with the predictions made by CBDT. The lower the lines are the higher the proportion of subjects for

[19]Each subject is characterized by one number that corresponds to how many of her observed choices coincide with theoretically predicted choices.

[20]The quarter benchmark is chosen because this is the proportion of choices a pure random model would predict correctly. The 2/3 benchmark is supposed to indicate the overwhelming majority of choices.

whom observed choices coincide with the predictions. The figures can be tied back to Tables V and VI. Figure 1 shows that for 20.5% of the subjects CBDT does worse than a random model would when current information is available and feedback is immediately provided. Hence, for 79.5% of the subjects it does better. These percentages are found at the intersection of the dashed line with the vertical line at 1/4. Clearly, observed choices coincide more often with the predicted choices of CBDT when information about the current situation is provided (see Figure 1, panels (a) and (b)). This result is stronger when feedback is delayed. Kolmogorov–Smirnov tests confirm the visual inspection ($p = 0.076$, immediate feedback and $p = 0.000$ delayed feedback).

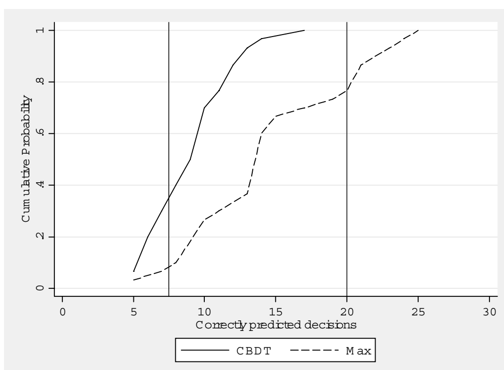

(a) w/o Current

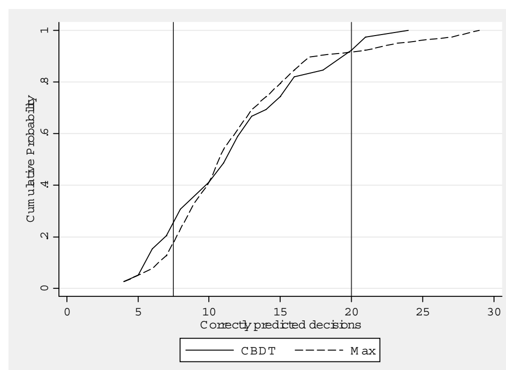

(b) w/ Current

Figure 2: Cumulative density of choices in line with predictions of CBDT or MAX (Immediate Feedback)

Figure 2 compares the performance of CBDT with the performance of MAX when immediate feedback is given. Panel (a) shows that when no information on the current situation is given MAX predicts more choices correctly than CBDT (Kolmogorov–Smirnov test, $p = 0.000$). Performance does not differ when information about the current situation is given (Kolmogorov–Smirnov test, $p = 1.000$).

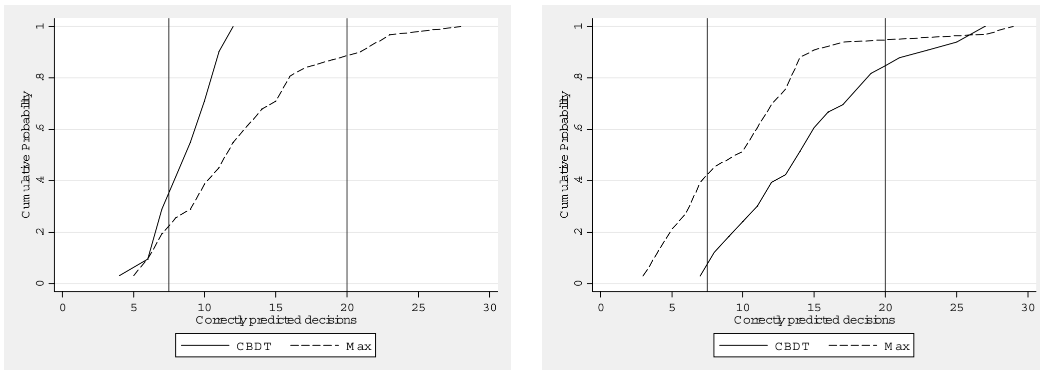

(a) w/o Current (b) w/ Current

Figure 3: Cumulative density of choices in line with predictions of CBDT or MAX (Delayed Feedback)

Figure 3 compares the performance of CBDT with the performance of MAX when feedback is delayed. Again, panel (a) shows that when no information on the current situation is given MAX predicts more choices correctly than CBDT (Kolmogorov–Smirnov test, $p = 0.002$). However, CBDT outperforms MAX when information about the current situation is given (Kolmogorov–Smirnov test, $p = 0.014$).

6 Further Analyses of the Results

6.1 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. In the immediate feedback treatment, in which subjects received feedback about their performance after each decision, it could be conjectured that some kind of reinforcement learning could be taking place.

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 (when calculated with respect to the CBDT predictions) over the second half (see Appendix B for the individual data). We find that 44% (17/39) show a smaller MSD in the second half of the experiment when current information is provided and immediate feedback is given. This percentage is not significantly different from the percentage of those that show an increase in their MSDs, 38.5% (15/39), (Test of Equality of Proportions, z–value= 0.46, n.s.). The same holds when feedback is delayed and information about the current situation is given (18/33 (decrease) vs. 14/33 (increase), Test of Equality of Proportions, z–value= 0.92, n.s.). 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. Interestingly, participants seem to follow MAX more often over time when immediate feedback is given and no information about the current situation is given (21/30 (decrease) vs. 8/30 (increase), Test of Equality of Proportions, z–value= 3.36, one–tailed $p < 0.001$). Immediate feedback is driving this result. The MSDs with respect to MAX do not go down when feedback is delayed (14/31 (decrease) vs. 10/31 (increase), Test of Equality of Proportions, z–value= 1.04, n.s.).

In order to address the second type of learning, we derived predictions of a deterministic payoff assessment learning model (Sarin and Vahid, 1999) and calculated the individual MSDs for each subject. We varied λ, the weight given to the obtained payoff and set initial assessments equal to the payoffs seen in the first four scenarios of market 1.[21,22] While the payoff assessment model does better in both treatments when immediate feedback is given compared to when feedback is delayed, it only predicts choices more than half of the time for one subject out of all treatments. Its general performance is close to random (MSD = 1.5).[23] We therefore feel confident that there is not much learning going on across the situations, i.e. markets.

6.2 Decision Times

While theory is silent on how long it takes to make decisions, it seems intuitive that users of MAX should reach their decisions significantly quicker than users of CBDT. We turn now to investigate this aspect of the individual decision making data.

Figure 4 shows the average time in seconds that it took participants to confirm their choices.[24] As there were no significant differences in decision times among participants that where "classified" as CBDT or MAX,[25] we pool all decision times for all subjects. When no information about the current situation is given, participants decide quicker when they get immediate feedback compared to when feedback is delayed (Robust Rank Order test:

[21] 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 $q_{jk}(t+1) = (1-\lambda)q_{jk}(t) + \lambda x$, while for all other possible choices $l \neq k$, $q_{jl}(t+1) = q_{jl}(t)$.

[22] Given that our paper aims at explaining individual behavior, we did not simulate a probabilistic reinforcement learning model (Roth and Erev, 1995) which would either have led to comparing choices with a probability vector or analyzing population means.

[23] Given that there was not much variance in individual MSDs, we determined overall performance by calculating the average MSD for each treatment. Interestingly, when no current information is given, the model that fits "best" gives a weight of 0.1 to the obtained payoff independently of whether immediate feedback is given or not. Behavior seems more backwards looking. When current information is given, 0.2 returns the lowest average MSD when immediate feedback is given. However, when feedback is delayed the lowest average MSD is obtained with a weight of 0.9.

[24] This time does not include the time it took participants to look at their payoff screen.

[25] A subject is "classified" as CBDT (MAX) user if CBDT (MAX) predicts at least 2/3rds of her choices correctly.

$U = -1.8214$, one–tailed $p = 0.0344$). These quicker participants also decide much quicker than their peers who are given information about the current situation and feedback is immediate (Robust Rank Order test: $U = -6.7416$, one–tailed $p < 0.01$). The difference in decision time is not significant for participants who get either immediate or delayed feedback and see information about the current situation (Robust Rank Order test: $U = -1.4778$, n.s.). Subjects decide marginally quicker when they are given information about the current situation and feedback is delayed compared to when it is immediate (Robust Rank Order test: $U = 1.7070$, one–tailed $p = 0.0436$).

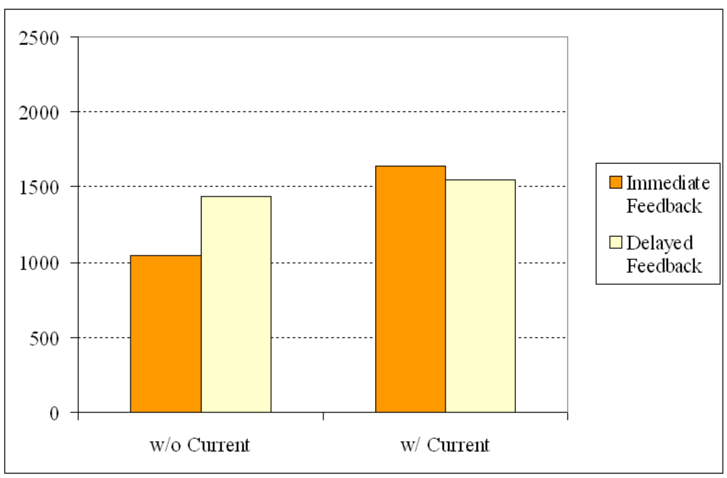

Figure 4: Time in seconds before participants confirmed their choices

7 Conclusion

Very little is known about decision making procedures used in complex, limited information environments. Expected utility theory does not seem plausible in such environments and, furthermore, can not be used to make meaningful predictions. The theory of case based decision making applies in such environments and generates reasonable restrictions on behavior when coupled with a suitable similarity function.[26]

We designed an experiment to test whether subjects use case–based reasoning combined with feature based similarity in an individual decision making environment. Our results provide some support for Case–Based decision theory.

It seems reasonable to suppose that people use a variety of procedures or decision making algorithms in complex environments in which the decision makers information is severely limited. Our results also suggest that people employ a (fast and frugal) heuristic that is much simpler to use in our experimental design. As considerable decision making in economics occurs in such complex, informationally constrained environments, it appears desirable to theoretically and experimentally further study decision making procedures agents use in such environments.

[26] See Matsui (2000) for a formal relation between expected utility theory and CBDT.

References

1. Fischbacher, U. (1999): "z–Tree – Zurich toolbox for readymade economic experiments: Experimenter's manual," mimeo.

2. Gayer, G., I. Gilboa and O. Lieberman (2004): "Rule–Based and Case–Based Reasoning in Housing Prices," mimeo.

3. Gigerenzer, G. and D. G. Goldstein (1996): "Reasoning the Fast and Frugal Way: Models of Bounded Rationality," *Psychological Review*, 103, 650–669.

4. Gilboa, I. and D. Schmeidler (1995): "Case–Based Decision Theory," *Quarterly Journal of Economics*, 110, 605–639.

5. Gilboa, I. and D. Schmeidler (1997): "Act Similarity in Case Based Decision Theory," *Economic Theory,* 9, 47-61.

6. Gilboa, I. and D. Schmeidler (2001): *A Theory of Case-Based Decisions.* Cambridge University Press.

7. Glasnapp, D. and J. Poggio (1985): *Essentials of Statistical Analysis for the Behavioral Sciences.* Merrill, Columbus.

8. Goldstone, R. L. and J. Son (2005): "Similarity," in K. Holyoak and R. Morrison (Eds.) *Cambridge Handbook of Thinking and Reasoning.* Cambridge, England: Cambridge University Press p. 13-36.

9. Guerdjikova A. (2002): "Case–Based Decision Theory – An Application for the Financial Markets," mimeo.

10. Guerdjikova, A. (2003): "On the Definition and Existence of an Equilibrium in a Financial Market with Case–Based Decision Makers," mimeo.

11. Holyoak, K. J. and P. C. Gordon (1983): "Social Reference Points," *Journal of Personality and Social Psychology*, 44, 881–887.

12. Hume, D. (1748): *An Enquiry concerning Human Understanding*, http://eserver.org/18th/hume-enquiry.html.

13. Jahnke, H., A. Chwolka and D. Simons (2005): "Coordinating Service–Sensitive Demand and Capacity by Adaptive Decision Making: An Application of Case–Based Decision Theory," *Decision Sciences*, 36, 1–32.

14. Matsui, A. (2000): "Expected Utility and Case Based Reasoning," *Mathematical Social Sciences*, 39, 1–12.

15. Ortony, A., R. Vondruska, M. Foss, and L. Jones (1985): "Salience, Similes, and the Asymmetry of Similarity," *Journal of Memory and Language*, 24, 569–594.

16. Podgorny, P. and W. R. Garner (1979): "Reaction Time as a Measure of inter- and intraobject visual similarity: Letters of the Alphabet," *Perception & Psychophysics*, 26, 37–52.

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

18. Roth, A. E. and I. Erev (1995): "Learning in Extensive–Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164–212.

19. 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.

20. Sarin, R. and F. Vahid (1999): "Payoff Assessments without Probabilities: A Simple Dynamic Model of Choice," *Games and Economic Behavior*, 28, 294–309.

21. Sarin, R. and F. Vahid (2004): "Strategy Similarity and Coordination," *Economic Journal*, 114, 506–527.

22. Selten, R. (1998): "Axiomatic Characterization of the Quadratic Scoring Rule," *Experimental Economics*, 1, 43–62.

23. Siegel, S. and N. J. Castellan, Jr. (1988): *Nonparametric Statistics for the Behavioral Sciences*, 2nd edition, McGrawHill.

24. Simon, H. A. (1955): "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69, 99–118.

25. Tversky, A. (1977): "Features of Similarity," *Psychological Review*, 84, 327–352.

26. Tversky, A. and I. Gati (1982): "Similarity, Separability, and the Triangle Inequality," *Psychological Review*, 89, 123–154.

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 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.

Figure 1: Screen Shot

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.

Appendix B – Individual Data

TABLE VII: INDIVIDUAL MSDS FOR CBDT AND MAX PREDICTIONS

Immediate Feedback				Delayed Feedback			
w/ Current		w/o Current		w/ Current		w/o Current	
CBDT	Max	CBDT	Max	CBDT	Max	CBDT	Max
0.933	1.333	1.533	0.533	1.067	1.733	1.467	1.067
1.200	1.400	1.400	1.067	0.733	1.267	1.200	1.333
1.600	0.867	1.333	1.067	0.600	1.667	1.467	0.467
0.933	1.467	1.333	0.733	1.200	1.333	1.333	0.933
0.667	1.467	1.267	1.267	0.667	1.667	1.400	0.933
1.667	1.733	1.600	0.600	0.933	1.067	1.333	1.467
1.333	1.333	1.667	1.067	0.333	1.667	1.533	0.133
1.533	0.467	1.533	1.333	0.733	1.067	1.333	1.267
1.533	0.200	1.333	0.600	1.467	0.067	1.533	1.600
1.133	1.067	1.333	0.400	1.000	1.600	1.467	1.533
1.333	1.267	1.400	1.333	1.267	1.200	1.200	1.133
1.267	1.533	1.600	1.067	1.333	1.533	1.267	1.400
1.467	0.867	1.333	1.133	1.067	1.533	1.533	1.600
1.267	1.133	1.133	0.667	1.133	1.000	1.267	1.000
1.000	1.200	1.067	1.067	1.200	1.533	1.400	1.533
1.200	1.000	1.200	1.533	0.467	1.533	1.333	1.333
0.600	1.467	1.333	1.000	1.467	1.200	1.467	1.067
0.667	1.200	1.200	0.867	1.000	1.133	1.600	1.467
1.467	0.600	1.467	1.067	1.333	1.067	1.400	1.200
1.333	0.933	1.200	1.333	1.400	1.133	1.267	1.533
1.333	1.533	1.667	0.333	0.867	1.267	1.267	1.267
1.200	1.000	1.533	0.467	1.533	1.333	1.400	0.600
1.467	1.267	1.267	1.067	1.467	0.200	1.533	1.667
0.400	1.600	1.467	1.200	1.267	1.067	1.267	1.333
1.733	1.333	1.400	1.333	0.200	1.733	1.533	0.933
1.600	0.933	0.867	1.667	1.400	1.467	1.267	1.200
0.667	1.400	1.467	1.000	1.067	1.800	1.600	0.733
1.467	1.400	1.600	1.467	1.000	0.867	1.733	0.467
0.600	1.267	1.133	1.333	0.733	1.600	1.533	0.867
1.000	1.600	1.600	0.600	0.933	1.200	1.200	1.200
0.933	1.267			1.200	1.267	1.333	1.333
1.133	1.200			0.200	1.467		
1.600	0.067			0.733	1.733		
1.067	1.067						
1.267	1.400						
0.800	1.467						
1.200	1.133						
1.600	1.133						
1.133	1.267						

Table VIII: Blocks of 15 rounds of Individual MSDs for CBDT and MAX (w/ Current)

Immediate Feedback				Delayed Feedback			
CBDT		Max		CBDT		Max	
1–15	16–30	1–15	16–30	1–15	16–30	1–15	16–30
0.800	1.067	1.467	1.200	1.333	0.800	1.733	1.733
1.200	1.200	1.467	1.333	0.400	1.067	1.600	0.933
1.600	1.600	1.067	0.667	0.267	0.933	1.467	1.867
1.067	0.800	1.200	1.733	1.067	1.333	1.600	1.067
1.200	0.133	1.467	1.467	0.933	0.400	1.867	1.467
1.867	1.467	1.733	1.733	0.667	1.200	0.933	1.200
1.467	1.200	1.467	1.200	0.400	0.267	1.600	1.733
1.467	1.600	0.800	0.133	1.067	0.400	0.400	1.733
1.467	1.600	0.400	0	1.333	1.600	0.133	0
1.067	1.200	0.933	1.200	1.333	0.667	1.733	1.467
1.600	1.067	1.600	0.933	0.933	1.600	1.600	0.800
1.467	1.067	1.467	1.600	1.200	1.467	1.467	1.600
1.467	1.467	1.200	0.533	0.800	1.333	1.733	1.333
1.333	1.200	1.333	0.933	1.333	0.933	0.800	1.200
0.933	1.067	1.200	1.200	1.067	1.333	1.733	1.333
1.333	1.067	1.067	0.933	0.533	0.400	1.600	1.467
0.800	0.400	1.600	1.333	1.333	1.600	0.933	1.467
0.800	0.533	1.200	1.200	1.333	0.667	1.067	1.200
1.467	1.467	0.800	0.400	1.600	1.067	1.067	1.067
1.330	1.333	0.800	1.067	1.333	1.467	1.333	0.933
1.200	1.467	1.467	1.600	0.933	0.800	1.467	1.067
1.067	1.333	1.200	0.800	1.467	1.600	1.067	1.600
1.600	1.333	1.200	1.333	1.600	1.333	0.133	0.267
0.667	0.133	1.733	1.467	1.067	1.467	1.333	0.800
1.867	1.600	1.333	1.333	0.267	0.133	1.733	1.733
1.467	1.733	1.067	0.800	0.933	1.867	1.333	1.600
0.800	0.533	1.200	1.600	1.333	0.800	1.733	1.867
1.067	1.867	1.467	1.333	1.067	0.933	1.067	0.667
0.533	0.667	1.333	1.200	0.800	0.667	1.867	1.333
0.933	1.067	1.733	1.467	1.200	0.667	0.933	1.467
1.200	0.667	1.067	1.467	1.200	1.200	0.933	1.600
1.200	1.067	1.333	1.067	0.400	0	1.333	1.600
1.600	1.600	0.133	0	0.933	0.533	1.867	1.600
1.067	1.067	1.333	0.800				
1.200	1.333	1.333	1.467				
1.200	0.400	1.467	1.467				
1.067	1.333	1.067	1.200				
1.333	1.867	1.333	0.933				
0.800	1.467	1.333	1.200				

Table IX: Blocks of 15 rounds of Individual MSDs for CBDT and MAX (w/o Current)

Immediate Feedback				Delayed Feedback			
CBDT		Max		CBDT		Max	
1–15	16–30	1–15	16–30	1–15	16–30	1–15	16–30
1.467	1.600	0.800	0.267	1.200	1.733	1.467	0.667
1.333	1.467	1.333	0.800	1.467	0.933	1.200	1.467
1.333	1.333	1.333	0.800	1.600	1.333	0.400	0.533
1.200	1.467	0.933	0.533	1.333	1.333	0.800	1.067
1.333	1.200	1.200	1.333	1.200	1.600	0.933	0.933
1.333	1.867	0.800	0.400	1.200	1.467	1.467	1.467
1.467	1.867	1.333	0.800	1.467	1.600	0.267	0
1.467	1.600	1.600	1.067	1.333	1.333	1.333	1.200
1.067	1.600	0.800	0.400	1.467	1.600	1.600	1.600
1.333	1.333	0.400	0.400	1.467	1.467	1.467	1.600
1.333	1.467	1.200	1.467	1.333	1.067	1.067	1.2
1.600	1.600	1.333	0.800	1.333	1.2	1.333	1.467
1.067	1.600	1.733	0.533	1.467	1.600	1.733	1.467
1.200	1.067	0.800	0.533	1.333	1.200	0.667	1.333
0.800	1.333	1.467	0.667	1.333	1.467	1.600	1.467
1.333	1.067	1.467	1.600	1.333	1.333	1.200	1.467
1.067	1.600	1.333	0.667	1.200	1.733	1.467	0.667
1.067	1.333	0.800	0.933	1.600	1.600	1.467	1.467
1.333	1.600	1.633	0.533	1.467	1.333	1.067	1.333
1.333	1.067	1.067	1.600	1.067	1.467	1.467	1.600
1.600	1.733	0.533	0.133	1.467	1.067	1.200	1.333
1.467	1.600	0.667	0.267	1.467	1.333	0.533	0.667
1.200	1.333	1.200	0.933	1.600	1.467	1.867	1.467
1.467	1.467	1.333	1.067	1.333	1.200	1.600	1.067
1.333	1.467	1.200	1.467	1.733	1.333	1.067	0.800
1.067	0.667	1.467	1.867	1.333	1.200	1.333	1.067
1.467	1.467	1.200	0.800	1.600	1.600	1.067	0.400
1.600	1.600	1.733	1.200	1.733	1.733	0.800	0.133
1.200	1.067	1.467	1.200	1.600	1.467	0.933	0.800
1.600	1.600	0.533	0.667	1.333	1.067	0.933	1.467
				1.200	1.467	1.200	1.067