Choice Under Uncertainty: Violations of Optimality in Decision Making

by

Kathleen Rodenburg

A Thesis
presented to
The University of Guelph

In partial fulfilment of requirements
for the degree of
Doctor of Philosophy
in
Economics

Guelph, Ontario, Canada

ABSTRACT

CHOICE UNDER UNCERTAINTY: VIOLATIONS OF OPTIMALITY IN DECISION MAKING

Kathleen S. Rodenburg
University of Guelph, 2013

Advisors:
Professor Michael Hoy
Professor Bram Cadsby

This thesis is an investigation of how subjects behave in an individual binary choice decision task with the option to purchase or observe for free additional information before reaching a decision. In part 1 of this thesis, an investigative study is conducted with the intent to sharpen the view to literature concerning corresponding psychology and economics experiments designed to test decision tasks that involve purchasing and observing information from an imperfect message prior to taking a terminal action choice. This investigative study identifies areas of research that warrant further investigation as well as provides enhancements for execution in the subsequent experiment conducted in Part 2 & 3 of this thesis. In Part 2 & 3, I conduct an experiment to test how subjects behave in an individual binary choice decision task with the option to purchase or observe for free additional information before reaching a final decision. I find that subjects' behaviour over time converges toward optimal decisions prior to observing an imperfect information signal. However, when subjects observe an imperfect information signal prior to their terminal choice there is greater deviation from optimal behaviour. I find in addition to behaviour that is reflective of a risk-neutral BEU maximizer, status quo bias, over-weighing the informational value of the message received and past statistically independent outcomes influencing future choices. The subjects' willingness to pay (WTP) to use the additional information gathered from an imperfect message service when making a final decision was on average less than the risk neutral BEU willingness to pay benchmark. Moreover, as the informative value of the message increased, causing the BEU valuation to increase, subjects under-estimated the value of the message signal to a greater degree. Although risk attitudes may have influenced the subjects' WTP decisions, it does not account for the increased

conservative WTP behaviour when information became more valuable. Additionally, the findings from this study suggest that individuals adopt different decision rules depending on both personal attributes (i.e. skillset, gender, experience) and on the context and environment in which the decision task is conducted.

DEDICATION

I wish to dedicate this doctoral thesis to my Father, Dr. Donald E. Grierson, Distinguished Professor Emeritus, Civil Engineering, University of Waterloo (March, 1939- August, 2011). He believed in me and inspired me always to persevere regardless of the obstacle. I am forever grateful.

ACKNOWLEDGEMENTS

I am very grateful to my advisors for their constant commitment and support to the completion of this doctoral thesis. In particular, I wish to thank Michael Hoy for directing me toward the field of experimental economics and providing great feedback on my research intent and direction. I am equally grateful to Bram Cadsby for his great attention to detail, mastery in behavioural economics, direction and research funding. His insight and understanding of the econometric methods used to analyze the experimental data was invaluable. Both Mike and Bram's response rates to my queries were exceptional. I couldn't have asked for more dedication to my research initiatives. Thanks to both David Prescott and Asha Sadanand for assisting in the design of the experiment and the theory underpinning its design.

I wish to thank James Amegashie, Rene Kirkegaard and Alex Maynard for reading my thesis immediately upon request and for providing valuable feedback that enhanced the research immensely.

Thank-you to the professors who supported me and attended my presentations on-going throughout the years, Evie Adomait, Patrick Martin, Chris McKenna, Francis Tapon, Thanasis Stengos, Douglas Auld, Marion Steele and Stephen Kosempel.

A very special thank you to my office mate, colleague and friend, Fraser Summerfield. He provided invaluable insights and support throughout my doctoral thesis. Thanks for attending my presentations and providing good honest feedback, for teaching me key econometric procedures, for listening to my rantings and for drinking copious cups of coffee with me. I am so very grateful.

I am grateful to my many friends in the program, Claudia Schmidt, Arian Khaleghi, Memhet Pinar, Joniada Milla, Jennifer Teng, Dianna Alessandrini and Marcel Oestreich. . They are an incredibly supportive group of individuals.

Finally, but certainly not least, a very special thank-you to my family and friends. My husband, Kent, has supported me and held me up during this experience on many occasions. Thanks for being so dedicated to our family and allowing me to pursue my dreams!! Thank you to my daughters Erin and Mary. Doing homework together will be a memory I will cherish forever. Thank-you to my mother, Colleen O`Shea, and to all my siblings, Joe, Patrick, Larry, Maura, Tamara and Lisa, for always being there.

Table of Contents

List of Tables

List of Figures

Chapter 1

Investigative Study: Experiments in Decision Behaviour

1.0 Introduction

A number of experiments have been conducted in both the psychology and economics fields to test the learning behavior of subjects in decision making tasks. Most recently, psychologists have modeled theories of cognitive learning, where people use cognitive learning techniques to make informed decisions about how to deal with new or similar situations (Casey et. al, 2005). Experimental economics is for the most part not explicitly concerned with the workings of the brain and the cognitive processes of belief learning (i.e., with the exception of the Neuroeconomics field). That said, Bayesian decision-making is classified by psychologists as being one of many cognitive learning techniques, and is widely used in economic models and laboratory experiments. In particular, some behavioural economists focus on comparing a subject's behavior with so-called Bayesian 'optimal' behavior. Although few economists would claim that people always behave optimally, most of the economic literature on learning has a connection to optimization. There are three major observations found in the experimental literature regarding actual (non-optimal) learning behavior: 1)subjects misperceive the diagnostic impact of the data and therefore fail to use the correct likelihood ratios (Griffen &Tversky, 1992; Beach, 1966; Wallsten, 1968; Pitz, 1968; Tversky&Kahneman, 1971); 2) subjects correctly perceive the diagnostic value of the data but fail to aggregate the likelihood ratios properly (Phillips & Edwards, 1966; Edwards, 1968, Grether, 1980); 3) subjects are unwilling to make very high or very low probability judgments and have difficulty dealing with extreme odds (Antoniou, Harrison, Lau & Reid, 2010; Holt & Smith, 2009; Ducharme, 1970).

Reinforcement-learning experiments that study how learning is reinforced through repeated trials are also prevalent in the literature (Suppes & Atkinson, 1960; Malcolm & Lieberman, 1965; O'Neill, 1987; Rapoport &Boebel, 1992; Ochs, 1995; Erev & Roth, 1998; Camerer & Ho, 1999; Feltovich, 2000; Charnes &Levin, 2005). In contrast to optimal decision theory, where individuals are assumed to have perfect rationality, reinforcement learners are assumed to have incomplete knowledge (bounded rationality) of the environment and act in a stimulus-response

way (Izquierdo & Izquierdo, 2006). Furthermore, unlike Bayesian theory, where the algorithm on how a subjective degree of belief should rationally change to account for new evidence is known and agreed upon across disciplines, the algorithm used to benchmark Reinforcement Learning varies by discipline (economics and psychology literature) and within discipline (e.g., see Bush and Mosteller's (1955) linear stochastic model of RL versus Erev and Roth (1998) RL model). Although differences exist, all RL algorithms are based on the same underlying idea that reinforcement learners use the outcomes from past experiences to assist in future decisions[1]. For example, when experiences from past actions result in successful outcomes, reinforcement learners will chose more often these same actions in the future. However, when experiences from past outcomes result in negative consequences, they will avoid these same actions in the future.

Most of the experiments designed to test Reinforcement learning in economics were conducted in the context of a strategic game (O'Neill, 1987; Rapoport &Boebel, 1992; Ochs, 1995; Erev & Roth, 1998; Camerer & Ho, 1999; Feltovich, 2000; Hopkins, 2002). That is, the subject's success or failure is measured based on their strategic choices relative to those made by their opponent. The general findings from these studies are twofold: 1) The Reinforcement Learning model outperformed the equilibrium predictions of classic decision theory both in terms of its description of the subject behavior as well as its predictive power (Thorndike, 1898; Suppes & Atkinson, 1960; Malcolm & Lieberman, 1965; O'Neill, 1987; Rapoport &Boebel, 1992; Ochs, 1995; Erev & Roth, 1998); and 2) subjects used a combination of both reinforcement and belief-based learning i.e., Bayesian Belief-based model (Camerer & Ho, 1999; Feltovich, 2000; Hopkins, 2002). There are few (Charness & Levin, 2005) laboratory experiments conducted in economics that tests Reinforcement Learning relative to Bayesian belief based learning in a non-strategic individual decision task.

Based on laboratory results, researchers have attempted to construct different algorithms that model the behavior of people when confronted with economic decision choices. A number of these algorithmic models are listed in the attached Chart (Appendix 1), which provides a summary of some experiments conducted in both psychology and economics concerning

[1] Learning via Genetic algorithms (GA) is also classified as reinforcement learning. In this type of learning process 'individual solutions are selected through a fitness-based process, where fitter solutions (as measured by a fitness function) are typically more likely to be selected' (http://mathworld.wolfram.com). This category of RL is not discussed in this thesis.

learning behavior in a variety of decision-making scenarios. The results from these experiments serve to both confirm and deny the theoretical predictions of classic decision theory and game theoretical models. As no broadly applicable model of learning behavior has emerged from these and other studies, this may imply that learning behavior depends in small or large part on the context and environment in which the decision making is conducted. That is, learning processes may differ between strategic and non-strategic games, between games of choice versus matching, between multidimensional choices versus binary choices, between one-shot versus repeated-play games, and between scenarios involving simple or complex criteria. In fact, Appendix I illustrates the point that the diverse observations reported among the listed laboratory studies were found for a diverse set of situational criteria used in the experiments. Slembeck (1999) notes that 'it is well known in the literature that decision making is contingent on the environment or situation.' This statement also suggests that the 'type' of learning process may be contingent on the situation. We know for example from the psychology literature that the characteristics of a given decision task determine whether agents learn consciously or unconsciously, and that the processes for these two ways of learning are different (Brenner, 2005).

If economics learning behavior is situational, as is at least somewhat indicated by experimental economics studies conducted to date, this suggests that a potentially fruitful direction for future research is to apply the findings of the wide body of psychology research concerning human learning to well defined economic decision-making problems. In this regard, some possible avenues to explore are: (i) Do people who are otherwise unaware of Bayesian learning tend to cognitively apply this technique when presented with correspondingly correct probabilities, or do they seemingly choose not to apply it? (ii) Do people apply reinforcement-learning cognitive techniques? (iii) What experiments need to be conducted to identify whether or what kind of situational restrictions should be applied to existing economic models so as to enhance their predictive power in both laboratory and real-world settings?

This review chapter is divided into four parts. Section I describes some decision tasks that involve purchasing and observing information from an imperfect message source to assist in the optimization of a terminal action choice, with the intent to sharpen the review focus to literature

concerning corresponding psychology and economics experiments. Section 2 presents a detailed review of this literature and Section 3 discusses areas of research that warrant further investigation identified by the review. Section 4 recommends enhancements to these past experiments for execution in the future experiments found in chapters 2 & 3 of the thesis.

1.1 The Decision Task

While some decisions made by agents in business and industry involve games of strategy, others do not. Indeed, for business executives a non-strategic[2] environment is a pervasive feature of many decision tasks. Furthermore, the option to purchase additional information from a knowingly imperfect message source prior to making a binary terminal decision is not uncommon. For example, a marketing director is continually faced with go/no-go decisions with an option to gather more information at a cost prior to taking the terminal decision. Should we launch this new product and if we do should we gather more information? Should we proceed with the proposed advertising campaign and if we do proceed should we run more consumer impact testing first? [3] Should we go ahead and hire this candidate as the product manager or conduct more interviews? While these decisions may involve strategic elements (e.g., How will my competitors react to my various decision choices? What will be the consequences to my business be as a result?), they also contain non-strategic elements (e.g., Will the distribution channel we have chosen for our new product launch reach our targeted consumer base? Will our advertising campaign resonate with our consumers and achieve the desired recall rates that provide a good return on our advertising investment? Is the potential employee capable of performing at the level that he communicated during the interview process or are there still better candidates yet to be interviewed?). The probability of a yes or no answer to any of these questions is not well known or agreed upon. Furthermore, the cost of purchasing additional information from an uncertain message source is unquantifiable in many cases and often cost

[2] A non-strategic decision environment is defined in this case as an environment where the actions of competitors or other players are not relevant to the decision task.

[3] In advertising managers are concerned with two aspects of an Ad campaign, the 'recall' rate and the 'recognition' rate. Consumers are shown a commercial (listen to a radio clip, observe an Ad board) and then are asked if they can recall what the product is that the advertiser was trying to sell (a recall rate is calculated based on the consumer responses). Consumers are also asked if they recognize the commercial. While a high recognition rate is good, as it indicates that the Ad has a 'catchy' aspect to it, it is worthless if the consumer cannot not recall what the Ad was selling.

prohibitive; at best it can be described as the opportunity cost associated with taking your attention and resources away from one profit enhancing activity in favour of another. As well, how an agent learns from the imperfect information source will affect the final outcome. Therefore, these go/no-go decision choices are made under varying degrees of uncertainty and risk. As such, part of the skill of an expert managing director who is focused on maximizing profits requires learning optimally from the gathered data, knowing when to stop obtaining additional data and when to make a final decision.

In the laboratory experiment in chapters 2 & 3 of this thesis, I wish to enhance our understanding of cognitive learning as a basic process in economic decision making. Specifically, I wish to understand how subjects behave when presented with a binary-choice decision, like those described above, with the option to purchase additional information before reaching a terminal choice. In addition to testing whether subjects possess the cognitive sophistication to apply Bayesian Expected Utility theory when presented with new information, I will also be testing whether they use elements of reinforcement learning when attempting to optimize their action choices.

1.2 Detailed Literature Review

The following review is focused on research studies that conducted experiments designed to test the learning behavior of subjects in a non-strategic binary decision task, with the option to obtain more information either for free and/or at a cost before making a terminal decision. The task required the subjects to use their intuitive probability judgment to decide which of two hypotheses presented to them is correct in order to: 1) from the Economics perspective, make decisions that maximize their expected monetary payoff from the experiment and; 2) from the Psychology perspective, make beneficial decisions to the best of their ability. These studies are highlighted with an asterisk in Appendix 1.

The question as to how people judge the probabilities or likelihoods of uncertain events has been a major focus in behavioral decision research for a number of years. In particular, many experiments have focused on comparing intuitive probability judgments to the rules of statistics. As is observed in Appendix 1, the fact that intuitive judgments often deviate from the laws of

probability is widely accepted (Von Witherfeldt &Edwards 1986). How subjects perform in simple decision tasks relative to the optimal responses predicted by theoretical models is the over-arching theme.

The decision tasks studied by the experiments can be divided into two groups: 1) a matching decision task or 2) a terminal-choice decision task. The matching decision task compares the accuracy of a direct observation of the subject's subjective judgments about probabilistic events with the probabilities derived using Bayesian theory, with a view to understanding whether subjects update their initial beliefs of a certain event when presented with new information using Bayesian-like thinking and, if not, how do their estimates vary from the Bayesian predictions? The terminal choice decision task combines Bayesian theory with expected utility theory; where, the use of Bayes' law is inferred through the terminal choices selected by the subjects. The terminal choice is either motivated by the desire to guess the correct state (in the case of the psychology literature) or to maximize monetary incentives (in the case of the economics literature). For both groups, subjects are given the prior probability of both hypothesis and the probability of observing each event under each hypothesis; therefore, they possess enough knowledge to judge the posterior probabilities of both states.

Typically, the basic set-up for these binary decision task experiments involves two possible 'states', state 1 and state 2. Each state is comprised of different proportions of colored marbles, balls, chips or dies or a like simulation. For each round of the experiment one of the states is selected at random according to a pre-announced probability, which determines the prior probability. Each subject freely observes or has the option to purchase one or a series of draws with replacement (or not), from the selected state. The subject uses these draws: (i) in the case of matching, to assess the probability that either state 1 or state 2 is being used for the draws; (ii) in the case of a terminal decision choice, to select an action which maximizes the expected payoff.

The pay-to-observe task was first studied by Swets & Green(1961). In this experiment the subjects are asked to decide whether a tonal signal had been added to white noise. They were allowed an additional tonal signal for a fixed cost. Payoffs for correct decisions and penalties for incorrect decisions were both varied between treatments. They found that as payoffs and penalties were increased, subjects choose to take additional samples and the proportion of errors

decreased. These experimental results seem to indicate risk aversion by subjects as incentives increased for correct and incorrect responses.

Peterson & Ducharme(1967) asked subjects in a binary choice task to state their posterior probabilities prior to making their terminal choice. Two experiments presented subjects with two different sequences of data. The first sequence of data favored hypothesis 1, while the second sequence of data favored hypothesis 2. Thirty-two subjects participated and were paid a flat fee for taking part. They found a primacy effect; that is, the information received at the beginning of the experiment was more important than information received at the end. As such, subjects took longer to change to hypothesis 2 than what would have occurred if they applied Bayes law properly. The primacy effect was left unexplained; the authors assumed the subjects were applying a form of Bayes law, and that they were taking longer to switch strategies because they gave more weight to earlier information than later information. Grether(1980) in later study described in this chapter, points out that their use of the word 'probability' was open to various interpretations by participants which may have skewed the relevance of their results

Pitz (1968) found that when subjects could purchase, in a binominal sampling task, as many as 10 to 20 data samples, they bought more information than is optimal. When up to only 5 samples were available, the results were more equivalent to optimal policy. The effect of a change to the cost of each additional piece of information (a data sample) was found to be far less noticeable than would be predicted by the optimizing model.

Edwards (1968) uses Bayes' law to update prior probabilities and uses this to compare against subjects' behavior. In his experiments observations are generated by a data generating process known to be in one of two possible states. Hypothesis 1 says that the process is in state 1 and hypothesis 2 says the process is in state 2. He compares the subjects' behavior to a Bayesian decision maker who will decide between the two hypotheses based on the costs of subsequent information, expected payoffs and the probabilities of the two hypotheses. He observes that it takes subjects from 2 to 5 observations to produce a diagnostic prediction equivalent to the Bayesian prediction after one observation. Specifically, the subjects purchased too much information. Subjects displayed a lack of confidence in the power of the observation.

Wallsten (1968) presents an experiment that tests the adequacy of subjective expected utility theory in a Bayesian decision task. Subjects are offered three choices after each observation; select state 1, select state 2 or, buy another observation at a fixed cost. In his experiments he compares the subjective stopping probability of his 14 subjects with the optimal stopping probability (critical value) derived using Bayes theorem. He uses points as an experimental currency to create incentives, with each point representing a preannounced monetary value. He finds that when the average diagnostic impact of an observation was high[4], subjects bought close to the optimal amount of information. However, when the average diagnostic impact was small, subjects tended to buy too few observations. He has difficulty determining whether the subjective expected utility model is adequate or inadequate as he had difficulty separating Subjective Expected Utility theory from the Bayes law updating task.

Fried and Peterson (1969) conducted an experiment involving a binary decision task with the option to purchase additional information prior to taking a terminal choice. They found that people do a nearly optimal job of purchasing information when they must decide prior to observation how much information they wish to buy (called fixed stopping), and that they purchase too little information when they are allowed to decide after each observation whether to purchase another piece of information or not (called optional stopping). To ensure that subjects worked hard, participants had the opportunity to earn cash based on good performance with the potential to lose their own money if they played poorly.

Pitz, Reinhold and Geller (1969) compared three information seeking models with the performance of subjects in a simple binary decision task. Neither the optional strategy nor the fixed stopping strategy was adequate to describe the subjects' behavior. A general finding was that as subjects observed more samples, the critical odds (posterior probabilities associated with the sample draw) required for a terminating decision decreased. The effect was that subjects tended to settle for risky decisions late in the sample sequence.

In a study conducted by Hershman & Levine (1970), 20 subjects were presented with a hypothetical problem in military reconnaissance. Subjects needed to determine from limited

[4] That is the proportion of black to white balls is such that a draw from the urn is very informative.

reconnaissance information which one of two military mixes were used by the adversary. They are told that the two mixes were equally likely and were told the proportion of the two ingredients used (designated by colour) for each of the mixes. Subjects would sample the mix, but with no replacement, and could either make a terminal choice of what mix the sample came from or purchase another sample. Points were used as payment: 100 points if correctly identified and 0 points if not. They were instructed to earn as many points as possible. They were told that purchasing an additional sample would increase their accuracy, and the quality of the 2nd sample information was not tied to the purchase price. Thus an observation that cost 1 point was equally as informative as the message that cost 35 points. As participants arrived they were divided into two groups at random. The first group was the 'informed' group and the second was the 'un-informed' group. Groups were not aware of their designated title. Both the informed and un-informed group were told the initial probability of each event occurring, and then allowed to sample a datum for free. Given the proportional breakdown of the mix for each event, the informed group was provided the Bayesian posterior probability calculation for either event given the message received. The un-informed group was not given the updated probabilities and was required to make their decision without this additional information. Both groups were presented an option to buy additional information at a pre-set price of 1 point. In later rounds subjects were informed that the prevailing price for additional information was 15 points and later it was raised to 35 points. Subjects worked at their own pace and were told that a prize of $10 will be awarded to the participant with the highest number of points for each group.[5] Results from this study show a wide range of purchasing behavior. Of special interest however, was that there were no inter-group differences (between the informed and un-informed) with respect to the number of points spent and the number of correct decisions. In general, purchasing behavior departed significantly from the optimum strategy at each of the three purchase prices. There was significant over-purchasing at each price and only one instance of a failure to gather enough information. Purchases declined as the purchase price increased.

Research conducted by Tversky & Kahneman (1971, 1972, 1973) focused on the matching of subjective probabilities with those estimated using Bayes law. When given a list of descriptors

[5] To put this cash payout into perspective, in California in 1970, you could purchase a dinner for a family of 4 at MacDonald's for $10.00 and still receive $5.00 change.

regarding an individual, subjects were asked the probability that this individual worked in a specific career. The participants generally did a poor job of judging the probabilities when given obvious descriptors (Tversky & Kahneman, 1973) as well as when given fairly obscure descriptors (Tversky & Kahneman, 1971; Kahneman & Tversky, 1972). They found that when revising beliefs, subjects ignored prior or basic information. They suggest that individuals used a particular set of heuristics when making decisions under uncertainty, such as availability[6], anchoring[7], and representativeness[8]. They show that when these heuristics are employed there are predictable and consistent biases in individual judgments concerning the likelihood of certain events.

Grether (1980) sets out to test the accuracy of the heuristic representativeness put forth by Tversky & Kahneman (1972) in an economic setting. He notes that the key descriptive word 'probability' estimates used in both Tversky & Kahneman (1971) and Peterson & Ducharme(1967) are open to various interpretations by participants which may have skewed the relevance of their results. He suggests that in the Tversky & Kahneman study they use stimuli with natural labels that may induce non-monetary utilities. He bases his observation on Hammerton (1973), who found that changing the verbal stimuli to their equivalent mathematical format could lead subjects to different answers. In Grether's study, groups were divided into two groups, with one group receiving a flat rate for participation, while for the other group only one of the subjects' decisions was randomly chosen and if this decision was correct they received $15; otherwise they only received $5. Three bingo cages were used, Cage X, Cage A and Cage B. Cage X contained balls numbered one through six. Cage A contained 4 balls marked with the letter 'N' and 2 balls marked with the letter 'G' and Cage B contained 3 balls marked 'N' and 3 balls marked 'G'. Subjects could not distinguish between the two cages as the balls were all the same colour. A draw from cage X determined which cage would be used for the experiment; i.e., if 1, 2 or 3 were drawn then Cage A would be used; otherwise Cage B is used. Subjects did not see the outcome of the cage X draw. The sample size of 6 was drawn with replacement and 'G'

[6] A person who follows the availability heuristic relies on the elements of the decision that can be easily assessed in memory first.
[7] A person who follows the anchoring heuristic relies too heavily on the first piece of information offered when making decisions
[8] "A person who follows the representativeness heuristic evaluates the probability of an uncertain event by the degree in which it is i) similar in essential properties to its parent population ii) reflects the salient features of the process generated"(Kahneman & Tversky, 1972, p. 431).

or ‘N’ was announced to the group. After 6 draws are completed, subjects are asked to decide the cage they think the balls came from. Grether finds that people give too much weight to new evidence and too little weight to priors. However, priors are not ignored. He finds that as people gain more experience through repetition of the task and if at the same time are financially motivated they will play closer to the Bayesian optimal strategy. Importantly, he finds that incentives did not seem to matter on their own. He also suggests that subjects over-weight the likelihood ratio relative to prior odds; Grether (1992) later disputes this finding in light of a new set of results. He observes that subjects appear to adopt different strategies in response to observing different data. Some experiments included financial incentives for accuracy and some did not. Of those lacking financial incentives a larger proportion of subjects gave absurd responses. He suggests that data from experiments in which there are ‘no financial incentives should be treated as being possibly contaminated’ (Grether, 1992, page54). Finally, the results from his study support the study of Tversky & Kahneman (1973) that people tend to use representativeness heuristic.

Griffen & Tversky(1992) find in their study of 35 subjects that over-confidence in subjective probabilities occurs when the sample proportion is high(a datum from the sample has a greater probability of being similar to essential properties of its parent population i.e., a red ball is sampled from an urn containing a high proportion of red balls) and the sample size is low (can only observe a small sample of the parent population i.e., only one ball drawn from the urn), and that under-confidence occurs when the reverse is true. Subjects were also asked to estimate their probability that the ‘bias’ favoured the over or under-confidence hypothesis. In their study subjects were offered a prize of $20 if they were the person whose judgments were closest to the correct values.

Charness & Levin (2005) test Bayesian updating with expected utility theory (BEU) and Reinforcement learning (RL) in an individual choice task. They refer to the RL heuristic as a ‘win-stay, lose-shift’ strategy. The objective of the experiment is to observe whether subjects continue to use Bayes rule when it is aligned or not aligned with the RL heuristic (‘win-stay, lose-shift’) and whether ‘the propensity’ to use Bayes rule in either case is affected by the introduction of immediate reinforcement after the first draw. Specifically, they wish to test how

behavior differs when subjects are paid for a successful first draw in addition to receiving payment for their performance for the entire round?; where it is intuited that the payment received or not received after the first draw, would be accompanied with feelings of success or failure. They construct an experiment where the BEU strategy follows the 'win-stay, lose-shift' heuristic and where it follows a 'win-shift, lose-stay' heuristic. Therefore for part of the experiment the Bayesian updating rule is such that if a subject is successful in the first stage of the experiment, he will continue to play the same strategy in the second stage and if the subject is unsuccessful in the first stage he will shift his strategy in the second stage. That is, the BEU decision rule and the reinforcement rule are aligned. For another part of the experiment, the parameters change such that, it is more advantageous for the Bayesian player to shift his strategy in the second stage of the experiment if successful in the first stage and keep the same strategy if unsuccessful (win-shift, lose-stay' strategy). To the best of their knowledge, and in my own investigative study, this is the first experiment to study specifically what happens when BEU and RL heuristics work against one another. They were interested in two kinds of choices that a participant makes: 1) 'starting choices ' which refer to the choice of urn from which to draw the first ball[9], and 2) 'switching choices' which refer to the choice of urn from which to draw a second ball, after having experienced the first draw. Therefore, in a two-urn and two-colour ball experiment where the proportion of coloured balls varies between urns, subject who do not know the state of the world are required to make two draws with replacement. Thus each person has the choice of whether or not to stay with the same urn for the second draw. As such, the subjects' second draw is statistically dependent with the first draw. The observations were: 1) both the 'win-stay, lose-shift ' and BEU heuristic are being used by subjects and when there is a conflict between the two heuristics there is a greater divergence from BEU behavior. Specifically, when the heuristics are aligned, subjects behave as BEU maximizers but when they are not aligned 50% of all switching decisions violate the Bayes updating rule, and 2) these switching errors are reduced when subjects are not paid for their initial choice, suggesting that this immediate reinforcement is a key factor in explaining why subjects deviate from Bayesian updating. They also found a gender effect that suggests women are more likely to deviate from BEU behavior.

[9] Note that in this experimental design, subjects decide which urn to draw the initial ball from without being able to see the contents of either urn.

Gabaix, Laibon, Moloche & Weinberg (2006) experimentally evaluated another learning behaviour benchmark called the directed cognition model (DC). This experiment is of interest to this study as it suggests that individuals may solve problems by looking forward as opposed to the Reinforcement learning models which suggest that individuals solve problems by looking back at the outcomes from their past experiences. The model is based on the intuition that cognitive resources of an individual should be allocated like other scarce resources. Therefore, an individual does not insist on identifying with complete certainty the best action, but takes an educated decision when the cognitive costs of analyzing the problem outweigh the expected benefits; saving cognitive resources for future action choices. They begin with a simple choice problem for which it is possible to compute optimal choices. In this problem, subjects choose among 3 uncorrelated investment projects. Agents can sequentially investigate the projects by acquiring additional information about the project to reveal their state (i.e., to reveal the payoff if the project is pursued first, second or third) prior to selecting a project. Information acquisition costs a fixed rate regardless of the project investigated. Agents can stop acquiring information at any time and choose one project. Each of the investment decision included a low probability project, a high probability project and a sure thing. Gabaix et al were interested in understanding how the subject proceeded to make a terminal choice. Specifically, in what order do they chose to investigate each project and how often will they investigate the project before making their final decision. A total of 129 subjects played ten randomly ordered investment games (3 investment project choices per round). Subject behavior is benchmarked relative to the DC model[10]. They found that the laboratory behavior matched the predictions of the directed cognition algorithm; successfully explaining the sequence of information acquisition made by the subject.

In simple probability matching experiments where subjects must predict the probability that a coloured ball was drawn either from urn 1 or urn 2, Holt & Smith (2009) show that there are systematic biases. Specifically, they find that subjects overestimate posterior probabilities when

[10] The DC algorithm consists of a three step process: 1) algorithm evaluates the risk-neutral expected pays of the various investments including the cost of acquiring additional information; 2) algorithm selects the investment with the greatest expected payoff; and 3) algorithm repeatedly cycles through the first two steps changing the order in which the project is investigated and stops when the costs of information acquisition outweighs the expected payoffs.

initial beliefs are less than 1/3, and tend to underestimate posterior probabilities when initial beliefs are more than 2/3. In their study they introduce an incentive structure that elicits truthful revelation of subjective probabilities, and the use of the term 'chances out of 100' in lieu of the term 'probability'.

Antoniou, Harrison, Lau & Reed (2010) in their study find a similar result as Holt & Smith (2009). They find consistent evidence that subjects over-estimate when posterior probabilities are less than 0.5 and underestimate when posterior probabilities are more than 0.5. Of interest is their demonstration of the need for a two-choice task to properly determine a subjects' subjective probability; i.e., one task aimed at identifying the subjects risk attitudes toward environments with known probabilities and another task that elicits subjective beliefs for events. By understanding the risk attitudes of the subjects they are in a better position to compare behavior in the laboratory with the predictions of Bayesian learning, which includes both the application of Bayes law and subjective utility theory. They find in their study that the order in which the risk-attitude test is used has an effect on the outcome. Specifically, the risk-attitude test should be conducted first as this had no effect on the task of eliciting beliefs; however, when done in reverse the eliciting-belief task did have an effect on the risk-attitude test. With this experimental design, they find that subjected Utility theory is a sufficient model for eliciting subjective beliefs.

1.3 Areas in the Literature that Warrant Further Investigation

This review has identified the following six areas that warrant further investigation.

1) While there is pervasive evidence of reinforcement learning behavior in the strategic game context (Erev & Roth, 1998; Camerer & Ho, 1999; Feltovich, 2000), little has been done to investigate reinforcement learning relative to Bayesian learning in the context of simple non-strategic decision problems in an economic laboratory context (as highlighted by Charness & Levin(2005)).

2) There is inconclusive evidence as to why individuals do not accurately apply Bayes law beyond the simple explanation that people ' lack the cognitive sophistication' to do the math. However, Hershman & Levine (1970) found that subjects provided with the

computational results of the Bayes law calculation did not correspondingly optimize their behavior. Hershman & Levine did not find any group differences in rates of correct responses between subjects who were given the calculated Bayes updating information versus those that were just given sufficient parameters (initial probabilities and likelihood information) to make the calculations themselves. There are several hypotheses as to why this may be the case; e.g., the uninformed group made accurate subjective probability assessments on their own thus matching the informed groups decisions or, alternatively, none of the participants believed in the laws of probability and commonly chose to adapt a different heuristic when making decisions. Charness and Levin (2005) identify a need to understand the 'cross over' threshold between simple updating and more complicated updating. Specifically, they suggest that subjects can calculate Bayes law when the math is simple, but have difficulty calculating this when the math becomes more complicated. As such, it is at this point, when the calculations became too complicated, that subjects chose to apply a different heuristic other than Bayes law when making decisions.

3) Although the terminal choice experiment is more representative of the type of non-strategic decision task highlighted in section 2, agents in general do not decide on the basis of probabilistic outcomes of certain events but rather decide on specific action choices. Thus, separating subjective probabilities from other key aspects of decision-making behavior has been problematic for many researchers (starting with Wallesten, 1968). Most studies until recently relied on the fact that subjects are risk neutral (Holt, 1986). It has been shown that the risk-neutrality assumption is inconsistent with laboratory evidence, even with low payoffs (Holt & Laury, 1993; Harrison et al., 2010).

4) In most studies, the willingness to pay for additional information is measured in terms a fixed cost per observation. As the experimenter sets the purchase price, the precise value that a subject placed on the service is not known i.e., only a lower bound can be established (Fried & Peterson (1965), Green & Swets (1966), Edwards (1968), Wallensten (1968), Pritz (1968), Hershman & Levine(1970)). While we can identify the risk neutral Bayesian optimal benchmark for the value of new information, we cannot use this option-to-buy mechanism to assess exactly how much subjects value the new

information gathered from the message service. Furthermore, under this payment scheme we are also unable to observe the degree to which the amount paid for the information influences the subject's decisions.

5) Most of the studies specific to the decision task highlighted in section 2 were conducted using methodologies employed by psychologists. Colin Camerer (1995) notes in the Handbook of Experimental Economics (page 589) that methods employed by psychologists are different than those employed by experimental economists. Specifically, subjects participating in these studies were either not paid for performance, or if paid, were paid in small amounts. A key to ensuring robust data is to design payoffs which ensure the truthful elicitation of subjective probabilities and provide incentives to exert effort for the requested task (see Davies and Holt, 1993, chapter 8 for examples; Offerman and Sonnemans, 2004; Offerman et al., 2007;Holt & Smith, 2009). It is questionable whether the payment structure for many of these studies was faithful to the theory they set out to test. Additionally, these studies use stimuli with natural labels that may have 'induced non-monetary utilities' (Tversky & Kahneman,1972). Furthermore, in many cases the experiments did not involve repetitive tasks under stationary replication. Another key concern is that in some cases subjects may have been exposed to deceptive treatments. Even if subjects are not exposed to deception, Grether (1980) notes that, without transparency, subjects may believe that they are being exposed to deceptive treatments. As such, due to these differences in methodology, if we do find a similar experiment in psychology there would be benefit in replicating the experiment using experimental economic methodology. In fact, Camerer (1995) notes that the 'replication of the psychologists' findings using methods of experimental economics is therefore popular to test the robustness of results.'

6) Many of the studies and consequent findings specific to our decision tasks used extremely small sample sizes of approximately 20 subjects. Moreover, the econometric procedures developed recently were not available; e.g., the ability to estimate individual asymptotic probabilities using all the data (logit models, etc.) at the time these studies were conducted.

1.4 Conclusion and Follow-up

This review focused on the set of research studies where experiments were designed to test the learning behavior of subjects in a non-strategic binary decision task, with the option to obtain more information either for free and/or at a cost before making a terminal decision. As previously highlighted (see numbers 5 & 6 above), replicating the results from these experiments to test the robustness of the findings using recent economic methodologies and larger sample sizes is of value. Furthermore, enhancements to the experimental design will enable the testing of additional theories.

In addition to testing whether subjects' behavior is reflective of Bayesian Expected Utility (BEU) it is of interest to this author to test whether subjects treat past independent events as if they are interdependent with future events when making future decision choices (Reinforcement Learning). As reinforcement learning is representative of non-strategic learning, in that updating new information is not contingent on an opponent's action choices, it seems intuitive that the reinforcement learning algorithms put forth by researchers such as Erev & Roth (1998), Camerer & Ho (1999) to describe strategic behavior in laboratory settings may have application as well to non-strategic[11] decision making environments. The Charness and Levin paper is the only study found that compared both these benchmarks (BEU and RL) in the non-strategic binary decision task. However, in their paper they benchmark subject behavior relative to the RL heuristic across a statistically interdependent decision task. It is of interest to understand how past events influence future decisions when the two decision environments have no statistical dependence. For example, the fact that a marketing manager chose not to conduct further consumer impact testing of an advertising campaign prior to launching it into the marketplace, and the campaign was successful, should not influence her decision to conduct or not conduct consumer impact testing for a different future advertising campaign. It is the contention of this author that indeed individuals treat these two decisions as interdependent events.

Selecting a Reinforcement Learning behaviour benchmark for the experiment is challenging given the many different algorithms used to proxy this behavior. Furthermore, regardless of the

[11] A decision is made with no opponent.

RL heuristic selected many assumptions are required. For example, when selecting a more complicated heuristic, like the one used by Erev and Roth (1998) or Feltovich (2000) an assumption regarding the subject's initial propensity to select one action over another action is required. One option is to assume that the initial propensity to take an action is equal to the initial probability that a decision task is being performed in a particular state. However, the propensity to take an action is not the same as the initial belief that the decision is taking place in one state. Another option is to assign an initial propensity for each action to a subject using a randomized number generator. The rationale for this random assignment is difficult to justify given no prior knowledge of the subjects preferences for either action. Furthermore, an assumption must be made on how much weight, in terms of importance, a subject places on each of her prior outcomes; an assumption must be made regarding the subjects' availability[12] heuristic (i.e., does an outcome from two periods prior impact the current decision to the same degree as an outcome from just one period prior?). Arguably, the RL heuristic employed in the Charness & Levin paper also requires several assumptive parameters, in particular, a rather unrealistic assumption that subjects have a very short memory. That is, current decision choices are predicated on the most recent outcomes. However, the Charness & Levin RL algorithm, has been used as a benchmark in a non-strategic decision environment (versus in a strategic game) in a prior experiment and is simple to model[13]. Hence, using this RL algorithm will allow us to test the robustness of some of their findings. In particular, to further pursue the hypothesis that a 'cross over' threshold exists where the updating task given new information becomes too complicated, causing subjects to abandon Bayes law and rely more on the RL heuristic. They suggest that finding this threshold can be accomplished by varying the distribution of balls within the urns and the payoffs associated with a successful outcome. Additionally, this simple model will provide an initial cursory test of reinforcement learning that can be replaced later with a more complicated RL benchmark to either verify or refute the findings from this study.

In a strategic game the reinforcement learner follows past successes or failures by observing whether they won or loss relative to an opponent. In an independent decision task reinforcement learners make decisions based on their own past performance and therefore must keep track of

[12] An 'availability' heuristic indicates the elements of the decision that can be easily assessed in memory first.
[13] The Charness & Levin (2005) RL heuristic is described in detail in chapter 2

their own winnings and losses. A tracking sheet maintained by subjects during the experiment would provide a history of events within an individual decision task involving no opponent which would enable reinforcement learning to occur if subjects were predisposed to this behaviour.

Hershman & Levine (1970) reveal a very interesting result that warrants further investigation; namely that subjects provided with the computational results of the Bayes law calculation did not correspondingly optimize their behavior. This contradicts the conventional explanation cited by many studies that people 'lack the cognitive sophistication' to do the math and therefore cannot accurately apply Bayes law. It is a result worth testing using a larger sample size (in the H& L experiment, n=20) as the finding could provide some valuable insight into why subjects behavior is not reflective of the Bayesian updating benchmark. An experiment design that separates the subject pool in half across all treatments, assigning subjects to either an informed group (provided with the Bayes law calculation in terms of chances out of 100 that the message received came from one state versus the other) or an uninformed group (subjects are provided with all the parameter values necessary to calculate Bayes law on their own, but are not told the Bayes law calculation) serves to test the theory that individuals do not apply Bayes law because they are unable to do the math.

In these binary choice experiments we observe the terminal action chosen by subjects. From this action we can determine whether the behavior is reflective of BEU decision choices or not. However, we cannot observe the Bayes law component (the updating task) separate from Expected utility theory. An experimental design that first requires subjects make a decision in the absence of new information allows us to observe whether subjects are capable of maximizing expected utility separate from the updating component of the decision rule. Comparing this to a decision task that requires the subject to apply both components of the decision rule (B+EU) in order to follow optimal behavior, allows us to potentially comment on a subjects ability to combine Bayes law in conjunction with expected utility theory.

The willingness-to-pay mechanism to observe additional information is modified from the previous studies in order to provide a truncated method for determining how subjects value additional information. This method allows us to determine whether a subject is willing to

purchase a message and at the same time assess the value placed on this message when making a terminal decision without requiring the subject to identify how many messages at a set price they would be willing to purchase. Furthermore, we are able to access how the subjects' willingness to pay for information influences their decision choices relative to the optimal decision theory. A WTP elicitation method utilizing key aspects of a Vickrey second price auction (Vickrey, 1961) provides a mechanism that is designed to be incentive compatible; thus ensuring that subjects reveal their truthful valuation of the information signal. With this type of mechanism when a subject selects a WTP value that is equal to her independent private value, she maximizes her expected utility by revealing her true valuation of the message received.

Finally, the framing of the experiment is critical in ensuring robust results. First it is essential that subjects know they are participating in an individual decision task and are not in any way playing against an opponent. Second it is essential that subjects do not believe that they are victim to any deceptive experimental practices (Harrison & al., 2010; Grether, 1980). As such, a computerized experiment is abandoned in lieu of paper and pencil. A computerized decision environment may signal to the subject that they are playing a game against an opponent (in this case the computer) and may lead subjects to believe that there is deception in the randomized processes (i.e., the selection of the states, the selection of the colour chip drawn, the decision on which action a subject will receive payment).

Appendix 1

i) Non-Strategic[14]	Experiments	J[15]	Testing	Explanation	Findings
Matching game with option to purchase more information	Edwards & Slovic, 1965	P	Information gathering strategies	Find a unique target cell randomly hidden in a 16 cell matrix. Prize for finding target, and subject was charged for successful looks	About ½ strategies used were optimal and serious divergences from the optimal very rare
*Binary choice game with option to purchase more information prior to terminal choice. Sequential updating	Fried &Peterson, 1965	P	Fixed stopping rule versus optional stopping rule	Fixed- must decide in advance how much info required Optional-can decide as they go	Although optional is optimal, subjects do better under fixed stopping rule
Pay to observe task	Green &Swets, 1966	P	Optimal information gathering	Decide whether a tonal signal had been added to white noise. Additional looks available at a fixed cost.	As payoffs and penalties increased, subject took more samples and error proportion decreased. S's error rates were lower than in comparable fixed rate Tests.
* Binary choice game with free information.	Peterson & DuCharme, 1967	P	Bayesian Learning Primacy effect	Asked to state probabilities prior to taking terminal choice, there are no payoffs for players	Primacy effect- information received at the beginning of the experiment more important then the information received at the end.
*Binary choice game with option to purchase more information prior to terminal choice. Sequential updating	Edwards,1968	P	Bayesian Learning	Statistical rule using 'Bayes law' to update prior beliefs	Takes 2-5 observations to produce diagnostic impact equal to the Bayesian impact of one observation. Exhibit under-confidence in the power of the observation
*Same as above	Wallsten, 1968	P	Application of SEU to Bayesian decision task	Wish to find optimal stopping probability and compare with subjects critical values	When avg. diagnostic impact of Observation was great, bought close to Optimal. When avg. DI was small buy too Few observations. Either SEU adequate or Inadequate model, need to be able to Separate Bayesian updating task.
*Same as above	Pitz, 1968	P	Bayesian learning	Changed problem difficulty from previous experiments by introducing purchase price and the # of samples available for purchase	When subjects can purchase 10-20 data samples they bought more information then optimal. For samples of 0-5 more in line with optimal. Effect of the cost of info was far less noticeable than would be predicted by the optimization model
	Pitz, Reinhold &Geller,1968		Bayesian learning Fixed sample size	Compare 3 information seeking models with subjects performance	As sample size increased, the critical odds Required for a terminating decision de-creased. Effect was that subjects tended to settle for unduly risky decisions late in sample sequence
*Binary choice game with option to purchase more information prior to terminal choice. Sequential updating	Hershman & Levine,1970	P	Bayesian updating and expected utility theory	Subjects are presented with 3 price points for information. For one price point should always pay for info, 2nd price point should only pay sometimes and 3rd price point should never pay for info	✓ Lrg indiv. differences in purchasing behavior ✓ # of purchases declines with increase informational impact of the free sample&the purchase price of 2nd sample ✓ Substantial departures from optimal policy

[14] Non-strategic-Individual play against static nature

[15] Journal(J) article found in either psychology(P) literature or economic(E) literature

	Tversky &Kahneman, 1971		Bayesian Learning		People search too little and Learn too quickly compared to models of optimal sampling. Exhibit over-confidence.
*Matching probability that a statement is describing a certain state	Tversky&Kahneman, 1972	E/P	Simple heuristics for updating ✓ Availability ✓ Representation ✓ Anchoring	Subjects presented with verbal descriptors and must choose the probability that the individual works in a particular career	Subjects use representative heuristic the most-'rule of thumb' i.e. the statement was i) similar in properties to parent population ii) reflects the salient features of the process generated
*Binary choice game with free sampling – subjects asked to select the state	Grether, 1980 Grether, 1992	E	✓ T&K 1972 results in an economic setting ✓ Bayesian learning	Probability estimates are inferred from binary choice. Best guess from group receives highest payoff.	Representative heuristic good for untutored and unmotivated subjects. As people gain more experience and are financially motive they act closer to a Bayesian
Free sampling Matching game	Griffen & Tversky,1992	P	Bayesian Learning Theory	Subjects to match probability estimates	Over-confidence occurs when sample proportion(strength) is high and sample size(weight) is low. Under-confidence when weight is high and strength is low. Find that strength dominates weight.
*Binary choice game with free sampling	Charness & Levin, 2005		BEU with Reinforcement learning	Decision maker makes draw from urn with replacement if draw a black ball receive payment and receive payment if correctly identify state	1) payoffs same, RL and BU same 2) payoffs not same and paid for initial choice, half decisions inconsistent with BU 3) draw provides only information and no payoff, errors occur less.
Tri-choice investment game	Gabaix, Laibon, Moloche & Weinberg(2006)	E	Directed Cognition model(DC)	Investment game, pick sequence of gathering data and executing project	DC model explains sequence of information acquisition. When DC model and GW differ, DC model does a better job of matching laboratory evidence.
*Binary choice game with free sampling using a betting interface- 'multiple price list' for placing bets on which state is true	Antoniou, Harrison, Lau &Reed(2010)	E	Subjective Utility theory with Bayesian Learning(BEU)	1.SEU control for risk attitudes(ra), don't assume Bayes rule-agnostic 2. SEU control for ra, don't assume Bayes law, assume parametric structure on the determination of prob. 3.Weighted combination of 1&2	Subjects over-estimate Bayesian posterior prob. less than ½ and underestimate Bayesian Posterior probabilities greater than ½ Agnostic model(#1) is the best fit
	Selten & Stoecker, 1986 Selten &Buchta, 1999	E	Learning Direction theory	After experience, people contemplate what might be the better decision and then they adjust their behavior in that direction	
*Binary choice game with free sampling	Holt & Smith(2009)	E	An update on Bayesian learning	Simple experiment to test Bayesian updating	For initial beliefs less than 1/3 subjects Over-estimate posterior probabilities for initial beliefs more than 2/3 subjects tend to underestimate posterior probabilities. Developed a model to account for these biases.
Situation ii)Strategic [16]	Experiments	J[17]	Testing	Explanation	Findings
Repeated matrix games with Unique strategy Equilibrium- no opportunities for cooperation	Thordike, 1898 Suppes &Atkinson, 1960 Malcolm &Lieberman,1965 Erev &Roth,1998	P/E	Reinforcement theory	Non-cognitive, non-strategic, no reflection on prior behavior, re-use high paying strategy with increasing probability	One parameter reinforcement learning model from psychology literature out performs the equilibrium prediction model in terms of descriptive and predictive Power

[16] Strategic- play with opponent with changing strategies

[17] Journal(J) article found in either psychology(P) literature or economic(E) literature

	O'Neil,1987 Rapoport &Boebel,1992 Ochs,1995				
Normal Form Games Strategic game	Camerer &Ho, 1999	E	Experience-Weighted Attraction Learning		People use a combination of belief based learning-fictitious play RL. Behavior is on a continuum between the two behaviors therefore the EWA model developed and the exponent determines the amount of RL versus fictitious play by subject.
Multi-stage Asymmetric Information Games	Feltovich,2000 Hopkins, 2002	E	Fictitious & Reinforcement theory	Fictitious play-quasi-bayesian updating mechanism. Players face exogenous, stationary unknown distribution of opponents strategies. Actors choose their best reply to the observed frequency distribution of their opponents	Both fictitious and RL Out-perform N.E. For some criteria RL beats fictitious For some criteria fictitious beats RL. Performance of models depends on the experiment

Chapter 2

Violations in Optimal Decision Theory: An Experiment

2.0 Introduction

The question as to how people judge the probabilities or likelihoods of uncertain events has been a major focus in behavioral decision research for a number of years. The fact that intuitive judgments often deviate from the laws of probability are widely accepted (Harrison & al. 2010, Holt & Smith 2009, Von Witherfeldt & Edwards 1986). However, controversy still exists surrounding both the identification and root cause of systematic deviations from optimal behaviour. The results from several experiments serve to both corroborate and refute the theoretical predictions of classic decision theory. As no broadly applicable model of learning behavior has emerged from these and other studies, this may imply that learning behavior depends in small or large part on the context and environment in which the decision making is conducted. It follows that the proper identification of the situational restrictions that should be applied to existing economic models would enhance their predictive power in both laboratory and real world settings, where a more representative model of decision making processes would serve to enhance economic policy development.

In this laboratory experiment I observe how subjects behave in an individual decision task involving the choices between two different actions (non-strategic 2-action binary decision task) with the option to purchase or observe for free additional information before reaching the final decision (terminal choice). In addition to testing whether subjects' choices follow the predictions of risk neutral Bayesian Expected Utility (BEU) theory, I also test whether they follow the predictions of a Reinforcement Learning (RL) model using a simple RL algorithm employed by Charness and Levin (2005)[18].

Bayesian Expected Utility theory and Reinforcement Learning models are different in how they presume learning progresses and as such, in some cases, the consequent outcomes. Bayesian

[18] See chapter 1 for detail surrounding the selection of this Reinforcement learning algorithm for this study.

learning is widely used in economic models and laboratory experiments where the focus is on comparing a subject's behavior with so-called Bayesian 'optimal' behavior. Agents who learn according to the BEU learning model use the laws of probability to update prior beliefs when observing new information and then use these new updated beliefs to make decisions that maximize their expected utility. In contrast, Reinforcement learning, widely discussed and tested in psychology literature, supposes that agents make decisions based on past positive and negative outcomes. That is, agents' actions that lead to good outcomes in the past are more likely to be repeated in the future, whereas agents' actions that lead to bad outcomes in the past are less likely to be repeated. While decisions based on past successful and unsuccessful outcomes may be relevant and potentially optimal when future and past decision choices are statistically inter-dependent, this may not be the case when future decisions are statistically independent.

Despite the numerous experiments conducted to test behavior in a binary-choice decision task environment, there are in both economic and psychology literature several important issues that warrant further investigation and are subsequently addressed in this study.

While there is pervasive evidence of reinforcement learning behavior in the strategic game context (Erev & Roth, 1998; Camerer & Ho, 1999; Feltovich, 2000), little has been done to investigate reinforcement learning relative to Bayesian learning in a simple non-strategic decision task in an economic laboratory context (as highlighted by Charness & Levin(2005)). It is possible that the incentive mechanisms used in the binary choice decision task may have excluded the prevalence of a subject's use of a reinforcement heuristic as past payoffs were not always realized and if realized not always known after each round of the decision task. It is interesting to note that in the strategic game context with opponents, a subject's winnings are a key benchmark for measuring performance relative to other players. Regardless of whether the subject is paid for each round of play or paid for just one round, the winnings are a critical element used to assess the strategic play of the opponent. This visual cue could be a key trigger to activating the reinforcement heuristic in a subject's behavior in a simple non-strategic binary-decision task context.

Of particular interest to this study is an experiment conducted by Charness & Levin (2005). In their study they test Bayesian updating with expected utility maximization (BEU) and

Reinforcement learning (RL) using a simple RL heuristic design in an individual choice task. The objective of the experiment is to observe whether subjects continue to use Bayes rule when it is aligned or not aligned with the RL heuristic and whether 'the propensity' to use Bayes rule in either case is affected by the introduction of immediate reinforcement after the first decision and prior to the final decision. As an extension to their study they identify a need to understand the 'cross-over threshold' between simple and more complicated updating given the observation that subjects in their study use BEU when decision tasks are easy and RL when decision tasks become more difficult. Furthermore, their study looks specifically at the use of the RL heuristic by subjects in the context of a statistically inter-dependent decision choice. For purposes of this study, understanding whether subjects use the history from previous decision choices which are statistically independent is the main focus.

Furthermore, from the studies reviewed, there is inconclusive and insufficient evidence as to why individuals do not accurately apply Bayes law beyond the simple explanation that people 'lack the cognitive sophistication' to do the math. For the Bayesian Expected Utility Model, separating subjective probabilities from other key aspects of decision-making behavior, i.e. expected Utility theory, has been a difficult challenge for many researchers (starting as early as Wallesten, 1968).

Another disadvantage in past experimental designs is that the willingness to pay for additional information is measured in terms of a fixed specified cost per observation. How much a subject is willing to pay for information is determined by the number of samples purchased. As the experimenter sets the purchase price, the precise value that a subject places on the service is not known (i.e., only a lower bound can be established). While the Bayesian optimal benchmark for the value of new information can be identified, how subjects value new information under the option to buy mechanism presented in these papers cannot be assessed. Furthermore, under this payment scheme how the amount paid for the information influences the accuracy of the decisions made is not observable.

Finally, many of the studies and consequent findings specific to our decision task design[19] occurred prior to 1980 (Fried & Peterson, 1965; Green & Swets, 1966; Peterson & Ducharme, 1967; Edwards, 1968; Wallsten, 1968; Pitz, 1968; Hershman & Levine, 1970). It is questionable whether the payment structure for many of these studies was faithful to the theory they set out to test. For example, in some studies subjects were not paid for performance, were paid very small amounts (Hershman & Levine, 1970), or paid in a manner which may have induced incorrect behaviors (Fried & Peterson,1965[20]; Wallsten, 1968[21]; Grether , 1980, 1992). Additionally, for these earlier studies, the sample size for the experiments were extremely small (i.e. Wallsten, 1968, n=14; Hershman & Levine, 1970, n=20) and several of the econometric procedures and techniques used today were not available.

Therefore, as augmentation to these previous studies, this experiment provides further insight to the following important questions: 1) Do subjects use Bayesian learning to maximize their expected payoffs (BEU) when making decisions? And if so, does a cross-over threshold exist where the task becomes too difficult for subjects to apply BEU decision rules? 2) Do subjects who are provided with the posterior probability calculation deviate less from BEU optimal decision theory than subjects who are left to calculate the posterior probabilities on their own? 3) Do subjects treat independent rounds of the decision task as interdependent events applying a Reinforcement Learning heuristic when making these decisions? 4) Will the paid for observations result in fewer deviations from BEU optimal or RL action choices than the observations which are provided to subjects for free? And, 5) are there systematic deviations from optimal BEU behaviour beyond the partial RL heuristic used in this study?

From this study, I find evidence of both the risk-neutral Bayesian Expected Utility (BEU) and the Reinforcement learning (RL) algorithm adapted from the Charness and Levin study (2005) reflecting subjects' decision choices, with the former being more prevalent. Furthermore, I find

[19] The experiment design in this study is built from an old set of experiments conducted between 1965-1980. Although the Charness and Levin(2005) is most relevant as it compares the BEU decision in conjunction with the RL heuristic as in this study, the design for this experiment is more closely related to the earlier studies (i.e., 2 urns, two colour ball) and therefore, an important aspect of the literature review.

[20] Reward if correct, larger penalty if wrong. Subjects could lose some of their own money

[21] Money illusion- received 1/6¢ per point

that subjects' action choices prior to receiving an imperfect information signal were non-optimal at the beginning of the experiment, even though the optimal action was associated with the first-order stochastically dominant (FOSD) lottery. Eventually, however, these subjects' first action choices converged on the optimal action through repeated rounds of the same decision task. Additionally, when subjects observed a relevant information signal prior to their terminal decision and the consequent lottery associated with the optimal action was no longer first-order stochastically dominant there was a greater deviation from optimal behaviour over all the rounds of this same decision task. There is evidence that suggests the existence of a threshold where subjects' behaviour is less reflective of the BEU model when tasks become more complex (Charness & Levin, 2005). However, in the Charness and Levin's paper they suggest that threshold is at a point where the BEU decision rule is too complicated to calculate. This study contradicts this explanation by simplifying the math component of the BEU decision rule for a sub-set of subjects and finds that there is no improvement toward optimal behaviour. Additionally, Charness and Levin found that when tasks became more complex the subjects' behaviour was more reflective of the RL heuristic. Results from this study, found that the likelihood of BEU decisions increased when the BEU and RL heuristics are aligned (i.e., the BEU decision choice is the same as the RL decision choice for the subject) and the likelihood of BEU decisions decreased when the two heuristics are not aligned (i.e., the BEU decision choice is not the same as the RL decision choice for the subject). Finally, I find two systematic deviations from optimal BEU decision theory in addition to the reinforcement learning model used in this study. First, in a two-action choice binary decision task, where the second action choice is dependent on the information received from an imperfect message, a subset of subjects behaviour reflects an under-weighing of the value of new information when it is contrary to their original choice. Subjects prefer to stay with their first action choice regardless of the message received.[22] Second, there is a smaller sub-set of subjects whose decision choices reflect a consistent over-weighing of the informational value of the message received.

[22] Samuelson and Zeckhauser (1988) refer to this type of phenomenon as a 'status quo bias'.

Section 1 describes the experiment. In Section 2 the results and interpretations are presented. Section 3 provides conclusions and future research opportunities.

2.1 Experimental Design

I conducted 6 different treatments during 12 classroom sessions on the University of Guelph campus, Guelph, Ontario, with 180 students recruited by e-mail from the undergraduate Bachelor of Commerce student population. On average subjects earned $33.60 for a 90 minute session. Each classroom session consisted of approximately 15 students who participated in 24 rounds of an individual task consisting of two (2) binary-choice decisions per round; where the second binary choice decision occurred after observing an imperfect statistically relevant information signal. For a subset of the groups and rounds, subjects had a third decision choice that required them to specify their willingness to pay for this additional information, where the WTP amount specified determined whether a subject's first (before observing the message) or second (after observing the message) action choice was recognized for payment. Therefore, some subjects participated in two types of decision tasks; a decision task with a 'FREE' (FREE) message and a decision task with an 'OPTION TO PURCHASE' (OTP) a message (appendix 2). The remainder of the subjects were placed in the control group and participated in 24 rounds of the FREE message decision task only.

Upon arrival, participants were given a handout explaining the experiment set-up and detailed instructions. The facilitator read the instructions aloud and demonstrated the experiment (appendix 1). The subjects were told that the amount of money that they would earn depends both on their individual choices and on random chance. In addition, they were told that the objective of the experiment is to maximize their earnings. Each subject participated in a practice round for both the FREE message and the OTP message decision task prior to commencing the rounds designated for payment.

For the FREE and OTP message decision task, subjects are shown at the beginning of each round two opaque bags, each containing a combination of red and blue poker chips. The distribution of red to blue chips within the two bags is symmetric with one bag containing a greater proportion

of red chips and one bag containing a greater proportion of blue chips. For example, if bag 1 contains 35 red and 15 blue chips, bag 2 will contain 15 red and 35 blue chips. Subjects are told and shown the precise number and combination of red and blue chips contained within each bag. The step-by-step procedure for the FREE message decision task is outlined in Table 1 and described below.

In step 1, a subject is selected to perform a random draw that determines with equal probability which one of the two bags described above is selected for use during the round. All participants, including the subject performing the random draw, do not learn until the end of the round which bag has been chosen. In step 2, subjects are asked to choose one of two actions (action A or action B), where each action is associated with two different payoff amounts dependent on the bag that was randomly selected in step 1. In step 3, subjects are shown a sample draw of a poker chip (imperfect message) from the selected bag. In step 4 subjects can either maintain the action choice selected in step 2 BEFORE observing the sample draw or change their action choice selection AFTER observing the sample draw.

Table 2 provides the information that is shown and communicated to the subjects prior to taking their first and second action choice decisions for rounds 1-4 when performing the FREE message decision task.

In step 5 a random draw determines with equal chance whether the subjects' first or second action choice is used to calculate earnings. This payment mechanism incentivizes participants to apply effort to both action choices. In step 6, the bag that was used during the round is revealed.

Table 1: Sequential steps for the Free Message Task

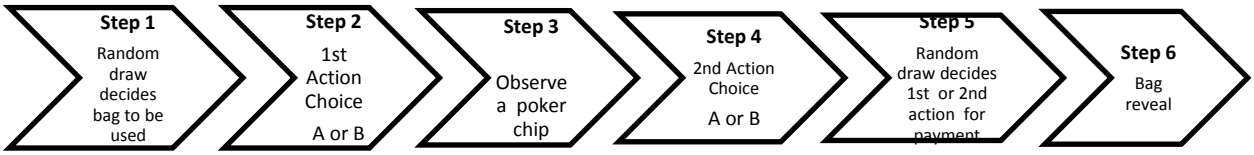

The action that was selected (1st or 2nd) based on the random draw in step 5 determines the size of the payment received by the participant as outlined in table 3. From table 3 for rounds 1-4, if bag 1 is revealed as the bag selected in step 1 of the experiment, the participant will receive $2.00 if they selected action A and $0.50 if they selected action B. However, if bag 2 is revealed as the bag selected in step 1, the participant will receive $0.75 if they selected action A and $1.75 if they selected action B.

Table 2: FREE Message Task Exogenous Parameters

Rounds 1-4

BAG 1		BAG 2	
Red chips	35	Red chips	15
Blue chips	15	Blue chips	35
Total chips	50	Total chips	50

Potential Earnings

Pick Action A:	If the bag chosen by participant was bag 1 you receive	$2.00
	If the bag chosen by participant was bag 2 you receive	$0.75
Pick Action B:	If the bag chosen by participant was bag 1 you receive	$0.50
	If the bag chosen by participant was bag 2 you receive	$1.75

The step-by-step procedure for the OTP message decision task is outlined in Table 3 and is described below.

The first four procedural steps for the OTP message task are identical to the FREE message task. However, after observing the sample draw (imperfect message) in step 3 and selecting an action conditional on this draw in step 4, subjects in step 5 must indicate how much they would be willing to pay in order for their second versus their first action choice (the decision made prior to observing the sample draw) to be used for determining their payment. Once the willingness to pay (WTP) price has been specified, the experimenter in step 6 asks a participant to draw a random price from a box which contains 51 tokens each specifying a unique price point ranging from $0.00 to $0.50. The subjects are unaware of the range of prices contained within the box.

The random price drawn determines the actual price required for using the second versus the first action choice to calculate earnings. In step 7, if the subject's specified WTP is less than the randomly determined price, the initial action choice (first choice) will be used to calculate her earnings and there will be no price deducted from the payoff associated with this decision. However, if her specified WTP is greater than or equal to the randomly determined price, then the revised decision (second choice) will be used to calculate the earnings and the random price drawn will be deducted from the total earnings for the round for the subject. This WTP elicitation method is designed to be incentive compatible; thus ensuring that subjects reveal their truthful valuation of the information signal.

Table 3: Sequential steps for the OTP Message Task

Table 4 provides the information shown and communicated to the subjects prior to making their first and second action choice decisions for rounds 1-4 when performing the OTP message decision task.[23]

[23] Note: During the 'Option to Purchase' (OTP) decision task, subjects must indicate how much they would be willing to pay in order for their second action choice to be used for payment. This is analogous to the following scenario: I book a trip to Florida, I observe that a hurricane is potentially pending; I book a trip to California but am informed that I must pay more to change my reservation. Based on my confidence in the weather forecast, how much would I be willing to pay to make this change? The rationale for this design is threefold: 1) it is easier to execute. All participants in a session observe the chip versus a design where only the subset of permitted participants can observe; 2) The Bayesian calculation for the optimal WTP amount is simplified. The added probability of receiving one of two messages is removed from the calculation providing subjects with an easier optimal WTP calculation; 3) More observations of subjects' second action choices are collected.

Table 4: OTP Message Task Exogenous Parameters

Rounds 1-4

BAG 1		BAG 2	
Red chips	35	Red chips	15
Blue chips	15	Blue chips	35
Total chips	50	Total chips	50

First Action Choice Potential Earnings:

Pick Action A:	If the bag chosen by participant was bag 1 you receive	$2.00
	If the bag chosen by participant was bag 2 you receive	$0.75
Pick Action B:	If the bag chosen by participant was bag 1 you receive	$0.50
	If the bag chosen by participant was bag 2 you receive	$1.75

Second Action Choice Potential Earnings

Pick Action A:	If the bag chosen by participant was bag 1 you receive	$2.00-$P
	If the bag chosen by participant was bag 2 you receive	$0.75-$P
Pick Action B:	If the bag chosen by participant was bag 1 you receive	$0.50-$P
	If the bag chosen by participant was bag 2 you receive	$1.75-$P

In step 8, the bag that was used during the round is revealed.

The action that was selected (1st or 2nd) based on the WTP amount in step 5 relative to the random price drawn in step 6 will determine whether participants will receive payment for their first versus their second action choice. From table 5 for rounds 1-4, if the subject's WTP is less than the random price drawn in step 6, they will receive payment on their first action choice. And if bag 1 is revealed as the bag selected during the round (for rounds 1-4), the participant will receive $2.00 if they selected action A and $0.50 if they selected action B; and, if bag 2 is revealed, the participant will receive $0.75 if they selected action A and $1.75 if they selected action B. However, if the subject's WTP amount is greater than or equal to the random price drawn in step 6, they will receive payment on their second choice. Let's assume the subject's WTP amount is $0.15 and the random price drawn was $0.10. Given that the WTP is greater than the random price drawn, the subject's second choice will be used to determine payment. If bag 1 is revealed as the bag selected during the round (for rounds 1-4), the participant will receive $2.00 minus the random price drawn ($2.00-$0.10), $1.90, if action A was selected and $0.40 ($0.50-$0.10) if action B was selected; and if bag 2 is revealed, the participant will receive $0.65 ($0.75-$0.10) if action A was selected and $1.65 ($1.75-$0.10) if action B was selected.

For both types of decision tasks (FREE and OTP) all subjects are informed each round of their earnings. Subjects are asked to record their first and second action choices, the results of each of the random draws, whether they received payment for their first or second action choice and their actual earnings for each round on the provided tracking sheet. The objective of the tracking sheet is to keep an account of each subject's history of events from past rounds to allow for the potential manifestation of reinforcement learning behaviour (appendix 3).[24]

The exogenous parameters, the distribution of red to blue chips contained within each bag and the payoffs associated with each action choice, change every four rounds and remain constant for 4 consecutive rounds. Table 5 provides the exogenous parameter values for the 24 rounds and the decision rule required to follow the risk- neutral BEU behaviour.

Given the exogenous parameters for this experiment, the risk neutral (RN) optimal action taken prior to receiving an imperfect message is associated with a lottery that first-order stochastically dominates the alternative action's lottery for all rounds. Therefore, any expected utility maximizer with monotonic preferences should select the optimal first action regardless of risk preferences. The rationale for this design is to assist subjects in an easy optimal first choice, allowing for a cleaner assessment of subject behaviour when selecting a second action conditional on an imperfect information signal.

Similarly, the 2nd RN optimal action conditional on the red chip message is also associated with the lottery that first order stochastically dominates the alternative action's lottery for all rounds. Again in this case, any expected utility maximizer with monotonic preferences should select the optimal action regardless of risk preferences. On the other hand, there is no first or second order stochastic dominate lottery associated with either of the action choices conditional on a blue chip message. Although in this case it is now possible for risk preferences to influence choice, the optimal second choice for the risk neutral BEU maximizer continues to be the same optimal

[24]In a strategic game subjects follow past successes or failures by observing whether they won or loss relative to an opponent. This tracking sheet provides a similar history of events within an individual decision task with no opponent.

choice over a wide range of constant relative and absolute risk aversion utility curves.[25] Therefore, given this experimental design, when the message received is a blue chip versus a red chip, the consequent action choice is more suggestive of a subject's ability to follow the BEU decision rules. The willingness to pay (WTP) action to use the information to activate the subject's second action choice for payment results in changes to the lottery parameters, and as such risk aversion may influence the optimal WTP benchmark. This is discussed in more detail in chapter 3.

There is one final note on the choice of the risk neutrality assumption when establishing the BEU benchmark for comparison with subject behaviour. Arrow (1971) demonstrates in his Essays on the Theory of Risk Bearing that expected utility maximizers are (almost everywhere) arbitrarily close to risk neutral behaviour when stakes are arbitrarily small. This is later verified by the Rabin Calibration (Rabin, 2000) which shows that the risk neutral prediction holds not only for small stakes but also for large and economically important stakes.[26]

Following the exogenous parameters identified in Table 5, suppose the subject behaves as a risk-neutral Bayesian Expected Utility maximizer. Each set of four rounds forces a new optimal decision. There are two possible states, represented by S_j, j ϵ {1,2}, where S_1 indicates bag 1 and S_2 indicates bag 2. A risk neutral BEU participant takes an initial action given the unconditional (prior) probability of either state with the objective of maximizing her expected earnings. Let the unconditional probability (initial belief) of playing in state j be, $prob\ (S_j)$, where, $\sum_j prob\ (S_j)$ = 1. Let C($\boldsymbol{a}$,S_j) be the payoff if action $\boldsymbol{a}$ is chosen conditional on the state (S_j), where $\boldsymbol{a}$ ϵ {A,B}. Without any additional information about the probability of the state being the bag with predominately red chips or the bag with the predominately blue chips, the initial decision to choose action A or B is based on the prior probabilities of being in either state, $prob\ (S_j)$, and the state contingent payoffs associated with each action, C($\boldsymbol{a}$,S_j). Specifically, the risk-neutral BEU will choose action A versus action B when:

[25] Under CRRA assumptions, ($u(c) = \frac{C^{1-\delta}}{1-\delta}$) when 0<$\delta < 2$, optimal second choice equals the RN BEU choice for rounds 1-8 & 13-20. Under CARA assumptions($u(c) = -e^{-\lambda C}$) when 0<λ<5, optimal second choice equals RN BEU choice for rounds 1-8 & 13-20. These parameter values, δ and λ are further relaxed during rounds 9-12 & 21-24.

[26] Of course, there are others who argue these results using experimental data. However, these findings in addition to the exogenous parameter choices for this experiment provide good rational for the Risk neutral assumption when establishing this benchmark.

$$EP_{action\ A} = prob\ (S_1)C(A,S_1) + \ prob\ (S_2)C(A,S_2)\ \geq EP_{action\ B}$$
$$= prob\ (S_1)C(B,S_1) + \ prob\ (S_2)C(B,S_2)$$

Therefore, given the parameter values for rounds 1-4 & rounds 13-16 presented in Table 5, the initial BEU action choice in the absence of a message will be action A, as the expected payoff from action A is greater than that of action B.[27] The first action choice by the risk neutral BEU maximizer only requires the application of expected utility theory portion of the decision rule.

Next, the risk neutral BEU maximizer is provided with one of two possible messages in the form of a colour chip drawn from the randomly selected bag. Let the two possible messages be M_k, k ϵ {1,2}, where $M_1$is message 1(indicating a red chip message) and M_2 is message 2 (indicating a blue chip message). The participant is then required to propose a second action choice conditional on the message received. To do this the BEU maximizer will first, update her prior probabilities of being in either state to a new set of probabilities (posterior) using Bayes theorem. Second, she will combine these updated probabilities to determine the expected payoff from taking either action and then choose the action with the highest expected payoffs.

Bayes theorem states that the posterior probability that a risk-neutral BEU maximizer should attach to the state after receiving a message, $prob\ (S_j|M_k)$, is:

$Prob\ (S_j|M_k) \equiv \frac{(prob\ S_j)(prob(M_k|S_j)}{prob(M_k|S_j)(prob\ S_j)+prob(M_k|S_{\neq j})(prob\ S_{\neq j})}$; j=1,2; j≠1,2; k=1,2; (Eqn. 1)

Where the $prob(M_k|S_j)$represents the likelihood of the message (M_k) conditional on state, S_j .

Note that regardless of the message received, one of two states must persist. Therefore,

$prob\ (S_j|M_k) + prob\ (S_{\neq j}|M_k) = 1$ (Eqn. 2)

Using Bayes theorem from Eqn. 1, the probability that the bag selected is bag 1(S_1) given that a red chip (M_1) was drawn is:

$$Prob(S_1|M_1) = \frac{prob(M_1|S_1)prob(S_1)}{prob(M_1|S_1)prob(S_1) + prob(M_1|S_2)prob(S_2)}$$

In short-form notation let,

$prob(S_j) \equiv \ \pi_j$; $prob(M_k|S_j) \equiv q_{k.j}$; $prob(S_j|M_k) \equiv \ \pi_{j.k}$.

[27] From Eqn. 1, .5($2.00) + .5(.75) =$1.37 >.5(1.75)+.5(.50)=$1.125.

Hence, the conditional probabilities of S_j given message M_1(red message) using short-form notation are:

$\pi_{1.1} = \frac{q_{1.1}\pi_1}{q_{1.1}\pi_1 + q_{2.1}\pi_2}$; and from Eqn. 2 $\pi_{2.1} = 1 - \pi_{1.1}$; (Eqns. 3 &4)

And, the conditional probabilities of S_j given message M_2(blue chip) are:

$\pi_{1.2} = \frac{q_{2.1}\pi_2}{q_{2.1}\pi_1 + q_{2.2}\pi_2}$; and from Eqn. 2 $\pi_{2.2} = 1 - \pi_{1.2}$; (Eqns. 5 & 6)

Therefore, from Eqns. 3 & 4, the expected payoff of choosing action A when message 1 (red chip) is received is:

$$EP_{action\ A}|_{M_1} = \pi_{1.1}\ C(A, S_1) + \pi_{2.1}\ C(A, S_2) \quad \text{(Eqn. 7)}$$

The expected payoff of choosing action B when message 1 (red chip) is received is:

$$EP_{action\ B}|_{M_1} = \pi_{1.1}\ C(B, S_1) + \pi_{2.1}\ C(B, S_2) \quad \text{(Eqn. 8)}$$

Given the red chip message (M_1), the risk neutral BEU maximizer will choose action A if the expected payoff is greater than choosing action B given the posterior probabilities conditional on the red chip message.

From Eqns. 7 & 8, the risk-neutral BEU will choose action A if:

$$EP_{action\ A}|_{M_1} = \pi_{1.1}\ C(A, S_1) + \pi_{2.1}\ C(A, S_2) \geq EP_{action\ B}|_{M_1} = \pi_{1.1}\ C(B, S_1) + \pi_{2.1}\ C(B, S_2). \text{ (Eqn. 9)}$$

Table 5: Exogenous Parameters by Round Set

SET	1	2	3
Rounds	1-4 13-16	5-8 17-20	9-12 21-24
State Contingent Payoffs			
Action A **Bag 1 revealed, C(A,S_1)** **Bag 2 revealed, C(A,S_2)**	 $2.00 $0.75	 $1.75 $0.50	 $1.00 $0.50
Action B **Bag 1 revealed, C(B,S_1)** **Bag 2 revealed, C(B,S_2)**	 $0.50 $1.75	 $0.75 $2.00	 $0.75 $2.00
Initial Beliefs **Bag 1/Bag 2 (π_1/ π_2)**	 .5/.5	 .5/.5	 .5/.5
BEU decision rule prior to a message signal(chip draw)	**Action A**	**Action B**	**Action B**
State Characteristics **Total chips bag 1** **# red chips($q_{1.1}$)** **# blue chips ($q_{2.1}$)** **Total chips bag 2** **# red chips($q_{1.2}$)** **# blue chips ($q_{2.2}$)**	 50 35(.70) 15(.30) 50 15(.30) 35(.70)	 50 12(.24) 38(.76) 50 38(.76) 12(.24)	 50 20(.40) 30(.60) 50 30(.60) 20(.40)
Bayes Law Posterior Probabilities **$\pi_{1.1}$** **$\pi_{2.1}$** **$\pi_{1.2}$** **$\pi_{2.2}$**	**.70** **.30** **.30** **.70**	**.24** **.76** **.76** **.24**	**.40** **.60** **.60** **.40**
BEU decision rule After a message signal is received (chip draw)	**If Red : Action A** **If Blue: Action B**	**If Red : Action B** **If Blue: Action A**	**If Red : Action B** **If Blue: Action B**

Given the parameter values in Table 5 for rounds 1-4 & rounds 13-16, and given a red chip draw, one should choose action A, given that, $EP_{action\ A}|_{M_1} \geq EP_{action\ B}|_{M_1}$:

From Eqn. 7,

$$EP_{action\ A}|_{M_1} = \frac{.70(.5)}{.70(.5) + .30(.5)}(\$2.00) + \frac{.30(.5)}{.70(.5) + .30(.5)}(\$0.75) = \$1.625$$

From Eqn. 8,

$$EP_{action\ B}|_{M_1} = \frac{.70(.5)}{.70(.5) + .30(.5)}(\$0.50) + \frac{.30(.5)}{.70(.5) + .30(.5)}(\$1.75) = \$0.875$$

Similarly, if a blue chip is drawn, choose action B, given that, $EP_{action\ B}|_{M_2} > EP_{action\ A}|_{M_2}$:

$$EP_{action\ B}|_{M_2} = \frac{.70(.5)}{.70(.5) + .30(.5)}(\$1.75) + \frac{.30(.5)}{.70(.5) + .30(.5)}(\$0.50) = \$1.375$$

$$EP_{action\ A}\,|_{M_2} = \frac{.70(.5)}{.70(.5) + .30(.5)}\,(\$0.75) + \frac{.30(.5)}{.70(.5) + .30(.5)}\,(\$2.00) = \$1.125$$

For rounds 5-8 & rounds 17-20, following the math from above (see appendix 4), it is always optimal for the BEU decision maker to select action B as the first choice and action B as the second choice if a red chip is drawn and action A as the second choice if a blue chip is drawn. For rounds 9-12 & rounds 21-24, it is optimal to select action B as the first choice and also select action B as the second choice regardless of the colour chip drawn. The second action choice after observing the imperfect message requires the application of Bayes law in conjunction with expected utility theory in order to follow the BEU decision rule.

Prior to the announcement of the distribution of red to blue chips contained within each bag, and given the expected payoffs associated with each action choice, the critical values for the posterior probabilities, $\pi^c_{j.k}$, i.e., the switching rule where the BEU decision switches to the alternative action choice (i.e. A to B or B to A), can be calculated.

For example, a BEU decision maker will switch her choice (from action A to B) conditional on observing a blue chip (M_2) if:

$$EP_{action\ B}|_{M_2} = \pi_{1.2}\, C(B,S_1) + \pi_{2.2}\, C(B,S_2)$$

$$\geq EP_{action\ A}|_{M_2} = \pi_{1.2}\, C(A,S_1) + \pi_{2.2}\, C(A,S_2) \quad \text{(Eqn. 10)}$$

The critical values of the posterior probabilities, $\pi^c_{1.2}$ & $\pi^c_{2.2}$, where a BEU decision maker will switch to the alternative action choice conditional on observing message 2 (blue chip) can be calculated by changing the weak inequality sign to an equality sign in Eqn. 10 and solving for $\pi_{1.2}$ & $\pi_{2.2}$.

$$EP_{action\ B}|_{M_2} = \pi_{1.2}\, C(B,S_1) + \pi_{2.2}\, C(B,S_2)$$

$$= EP_{action\ A}|_{M_2} = \pi_{1.2}\, C(A,S_1) + \pi_{2.2}\, C(A,S_2)$$

Noting that, $\pi_{1.2} = 1 - \pi_{2.2}$, and simplifying gives:

$$(1 - \pi_{2.2})[C(B,S_1) - C(A,S_1)] + \pi_{2.2}\,[C(B,S_2) - C(A,S_2)] = 0 \quad \text{(Eqn. 11)}$$

$$\pi_{2.2}^{c} = \frac{C(B,S_1)-C(A,S_1)}{[C(B,S_1)-C(A,S_1)]+[C(A,S_2)-C(B,S_2)]} \qquad \text{(Eqn. 12)}$$

Given, $(B,S_1) < C(A,S_1)$, $C(A,S_2) < C(B,S_2)$ then $\pi_{2.2}^{c} \in (0,1)$ and $\pi_{1.2}^{c} \in (0,1)$

Next, denote the equation for the difference in expected payoffs between action B and A (left-hand-side of Eqn. 11) by θ, and evaluate the partial derivative of θ with respect to $\pi_{2.2}$,

$$\frac{\partial\theta}{\partial\pi_{2.2}} = -[C(B,S_1) - C(A,S_1)] + [C(B,S_2) - C(A,S_2)] > 0;$$

Given, $C(B,S_1) - C(A,S_1) < 0 \; and \; C(B,S_2) - C(A,S_2) > 0$.

Since $\theta = 0$ at $\pi_{2.2} = \pi_{2.2}^{c} \in (0,1)$ and θ is monotonically increasing in $\pi_{2.2}$, it follows that $\theta > 0$ if $\pi_{2.2} > \pi_{2.2}^{c}$ and $\theta < 0$ if $\pi_{2.2} < \pi_{2.2}^{c}$.

Hence as $\pi_{2.2}$ increases, the expected payoff from taken action B increases. Conversely, as $\pi_{2.2}$ decreases the expected payoff from taking action B also decreases and when $\pi_{2.2} < \pi_{2.2}^{c}$ the expected payoff from selecting action B has decreased to a point where the greater expected payoff is now associated with selecting action A ($EP_{action\ A}|_{M_2} > EP_{action\ B}|_{M_2}$). Therefore, for any $\widehat{\pi_{2.2}} > \pi_{2.2}^{c}$ a BEU decision maker will switch to action B, otherwise she will remain with the initial action A.

As the state contingent payoffs for either action ($A\ or B$) are not symmetrical for any round, the critical values for the posterior probabilities $\pi_{j.k}^{c}$, where the BEU decision switches to the alternative action choice ($A\ to\ B\ or\ B\ to\ A$) will differ depending on the message received. As such the range of posterior probabilities where a BEU decision maker will switch to the alternative action conditional on the message will also be different.

Table 6 provides the ranges of posterior probabilities ($\pi_{j.k}$) where a BEU subject will switch to the alternative action choice conditional on the message received for each set of rounds that is governed by the same exogenous parameters.

Table 6: Posterior Probabilities: Critical Values and Ranges of Estimated Posterior Probabilities by Round A BEU Participant Switches to the Alternative Action

Posterior Probabilities	Rounds 1-4/13-16			Rounds 5-8/17-20			Rounds 9-12/21-24		
$\pi_{j.k}$	Bayes	Critical Value (CV)	Range	Bayes	Critical Value (CV)	Range	Bayes	Critical Value (CV)	Range
$\pi_{bag1.red}$	.70	.40	[0.4, 1.0]	.24	.60	[0.0, 0.6]	.60	.857	[0.0, 0.86]
$\pi_{bag2.red}$	.30	.60	[0.0, 0.6]	.76	.40	[0.4, 1.0]	.40	.143	[0.14, 1.0]
$\pi_{bag1.blue}$	.30	.40	[0.0, 0.4]	.76	.60	[0.6, 1.0]	.40	.857	[0.86, 1.0]
$\pi_{bag2.blue}$	.70	.60	[0.6, 1.0]	.24	.40	[0.0, 0.4]	.60	.143	[0.0, 0.14]

Suppose that the posterior probability using Bayes law indicates that the optimal choice is to change to the alternative action, Table 6 illustrates that there exists a range of posterior probability estimates ($\widehat{\pi_{j.k}}$) where a subject may make a Bayesian updating error ($\widehat{\pi_{j.k}} \neq \pi^*_{j.k}$) yet still arrive at the correct BEU action choice. Additionally, it illustrates that in some cases this range is narrower than in others, allowing for less flexibility in posterior probability estimate ($\widehat{\pi_{j.k}}$) errors that still result in BEU action choices.

For example in rounds 1-4 and 13-16 from Table 6, a subject whose estimated posterior probability ($\widehat{\pi_{j.k}}$) falls within the range listed, assuming the subject is risk neutral and can perform the expected utility portion of the BEU problem, will also select the BEU action choice. However, the available range is reduced from a span of 0.6 points conditional on a red message to a span of 0.4 points conditional on the blue message (see also Figure 1).

Figure 1: Posterior Probability Estimate ($\widehat{\pi_{j.k}}$) Range for Each Action Choice Conditional on the Message Received Rounds 1-4 & 13-16

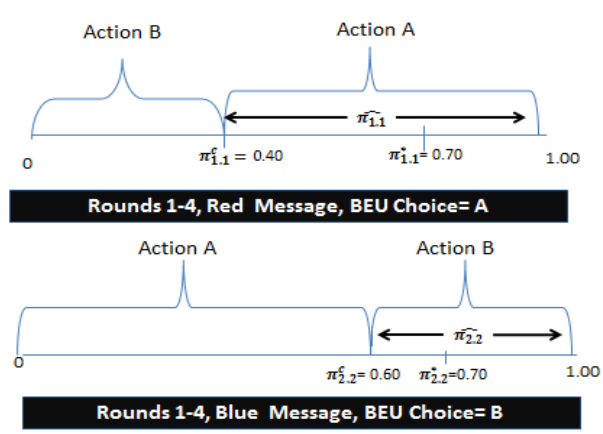

Note from figure 1, that the range for $\widehat{\pi_{j.k}}$ is narrower and the critical value boundary, $\pi_{j.k}^{c}$, where a subject should switch to the alternative action, is closer to the true $\pi_{j.k}^{*}$ when a blue message is received. Therefore, there is a greater likelihood of observing subject behaviour which reflects an under-weighing of the informational value of the blue message (i.e., in rounds 1-4, a risk-neutral subject capable of maximizing expected payoffs who under-weighs the informational value of the blue or red message and estimates a $\widehat{\pi_{1.1}} = \widehat{\pi_{2.2}} = .59$ would have behaviour reflective of BEU when a red message was received and would have non-BEU behaviour when the blue message was received). Furthermore, in an example where the critical value boundary $\pi_{j.k}^{c}$ for switching to the alternative action is greater than the true posterior probability ($\pi_{j.k}^{c} > \pi_{j.k}^{*}$) than we can potentially observe behaviour that is reflective of over-weighing the informational value of the message received. Specifically in this experiment, there is a greater likelihood of observing non-BEU decision choices conditional on a blue chip message that is reflective of under-weighing the informational value of the message for rounds 1-4(13-16) and 5-8(17-20) and over-weighing the informational value of the message for rounds 9-12 (21-24).

The Reinforcement Learner (RL) decision rule used in this study is based on the simple WIN-STAY, LOSE-SHIFT heuristic also used by Charness and Levin (2005). If a subject is successful in the first round of the experiment, she will STAY with this same action choice in the second round (WIN-STAY) and if the subject is unsuccessful in the first round, she will shift to the alternative action choice in the second round (LOSE-SHIFT); where, both RL actions are predicated on the subject experiencing the same past history dictated by both the fixed and random exogenous parameters set by the experiment.

Although the RL heuristic is simple in theory, difficulty arises when attempting to define past successful or unsuccessful action choices. The challenge arises when we attempt to use the financial gain or loss as the win or loss indicator. The necessary mechanisms required to ensure subjects exert best effort at each decision point and provide truthful valuation of a message service results in several possible reasons why a subject won or lost monetarily. Specifically, the reasons why a subject lost monetarily during the OTP decision task are: 1) the willingness to pay (WTP) price is less than the random price(RP) drawn forcing payment to be received on the

subjects first action choice, the first choice incorrectly identified the higher payoff state (the payoff for this action given the state is less than the payoff for the alternative action given the state if payment was extended) and the second choice correctly identified the higher payoff state (the payoff for this action given the state is greater than the payoff for the alternative action given the state if payment was extended); 2) the WTP is less than the RP drawn forcing payment to be received on the subjects first action choice, and the first and second action choice incorrectly identified the higher payoff state; Or 3) the WTP is greater than the RP drawn forcing payment to be received on the second action choice and the second choice incorrectly identified the higher payoff state. Similarly, the reasons why subjects lose financially during the FREE decision task are attributed to: 1) pay draw was first, the first choice was incorrect and the second choice was correct; 2) pay draw was first, the first and second choice were incorrect ; 3) pay draw was second and the second choice was incorrect.

As a result of the above interpretations, I restrict the observation of the RL behavior to the second action choice. I assume that the subject will apply the WIN-STAY heuristic for a current round second action choice when the prior round second action choice correctly identified the state associated with the higher payoffs (WIN-guessed the right bag) and will apply the LOSE-SHIFT heuristic for a current round second action choice when the prior round second action choice incorrectly identified the state associated with the higher payoffs (LOSE-guessed the wrong bag), regardless of the actual amount of payment received. Additionally, the WIN-STAY or LOSE-SHIFT heuristic can only be applied in a current round if the exogenous parameter values experienced by the subject are the same as what was experienced in a prior round. Specifically, if the configuration of red to blue chips within each bag, the consequent payoffs conditional on the action choice taken and the message received in a prior round are the same as the current round.

For example, for the first round a RL participant is presented with a set of fixed exogenous parameters (bag configuration and consequent payoffs) that will remain constant for four consecutive rounds. Given this set of parameter values, regardless of the colour chip observed in round 1, there is no RL behaviour for the subject to follow. However, going forward (for 3 more rounds) subjects accumulate history from prior rounds. Let's assume the following outcomes for

round 1: a red chip was observed, the subject took a second action that was associated with the higher payoff state once the true state was revealed (regardless of whether payment was extended or not) i.e. a WIN outcome. If the RL participant observes a red chip in the second round, she will use the past history gathered from round 1 to determine her second action choice for round 2. Hence, she will STAY with the same decision choice from round 1 in round 2 given that the choice in round 1 resulted in a WIN (WIN-STAY). However, if the RL participant instead observes a blue chip in the second round there is no blue chip history and therefore, no RL heuristic to apply. Now let's assume the following outcomes for round 1: a red chip was observed, the subject took a second action which was associated with the lower payoff state once the true state was revealed (regardless of whether payment was extended or not) i.e. a LOSE outcome. In this case, when a subject observes a red chip in the second round, the subject will SHIFT her second choice decision to the alternative action from the one chosen in round 1 (LOSE-SHIFT).[28]

Table 7 outlines the details of each of the six different treatment groups. Treatments 1, 2, 3 and 4 participated in 12 rounds of a decision task containing a FREE message (FREE) and 12 rounds of a decision task containing the OPTION TO PURCHASE (OTP) message. Treatments 1 & 2 received the FREE message task first and the OTP message task second. These decision tasks are reversed for treatments 3 & 4. Treatments 5 & 6 represent the control groups, participating in 24 rounds of the FREE message task only. Treatments 1, 3 & 5 are designated as un-informed. This group is provided with all the information required to apply Bayes law; however, they are not given the posterior probability calculations. Treatments 2, 4 & 6 are designated as informed and are provided with the posterior probability calculations described in terms of chances out of 100 that the colour chip drawn is either from bag 1 or bag 2.[29]

Table 7-Treatment Group Specification

[28] Given the above interpretation, for the first round of each set (rounds 1, 5, 9, 13, 17 & 21) the histories from the prior sets are assumed irrelevant given the new set of parameters defined by the distribution of blue to red chips contained in each bag, and the change in payoffs associated with each action. A LOSE-SHIFT or WIN-STAY action choice only exists in the second round of each set (rounds 2, 6, 10,14,18 & 22) if the random chip observed is the same colour as in round 1, 5, 9, 13, 17 & 21, respectively (similar histories). Conversely, if the colour draw has never been observed in the prior round, there is no RL heuristic to be applied in rounds 2, 6, 10, 14, 18 & 22. For rounds 3 and 4 (15 and 16); 7 and 8 (19 and 20); 11 and 12 (23 and 24) a RL participant must observe the histories from rounds 1,2 and 1,2,3 (13, 14 and 13, 14, 15); 5,6 and 5,6,7 (17, 18 and 17, 18, 19); 9,10 and 9,10,11 (21, 22 and 21, 22, 23), respectively, in order to determine the second action RL decision choices. Note that when no same history exists, the RL observation is recorded as a blank observation in the data set.

[29] I.e. there are 70 chances out of 100 that the chip drawn came from bag 1 and therefore 30 chances out of 100 that the chip came from bag 2. This description of the posterior probabilities avoids any confusion associated with the term 'probability'.

Treatment	1	2	3	4	5	6
No. Subjects (180)	30	30	30	31	29	31
Order: Round 1-12 : Round 12-24	Free OTP	Free OTP	OTP Free	OTP Free	Free Free	Free Free
Bayes Law	Uninformed[30]	Informed	Uninformed	Informed	Uninformed	Informed

In total I collected 4320 observations of subjects' first and second action choices and 1452 observations of the subjects' willingness to pay decisions.

2.2 Results

The data are analyzed using two different measurement criteria. First**,** subject behaviour is benchmarked relative to the action choices of a risk-neutral Bayesian Expected Utility maximizer and that of a Reinforcement Learner (RL) using the simple WIN-STAY, LOSE-SHIFT heuristic described in section II. In the data set, 'inconsistency rates' describe deviations from these two behavior types. Hence, for each subject in the experiment a Bayesian Expected Utility (BEU) first and second choice inconsistency rate (BDR1 & BDR2)[31] and a Reinforcement Learner (RL) second choice inconsistency rate (RDR2)[32] are calculated.

Second, for each round the subject's sequence of first and second action choices conditional on the imperfect message received are tracked and the proportion of subjects who follow each sequence is calculated. A subject can follow one of eight potential two action choice decision sequences for each round. The sequence that follows the BEU optimal action choice varies depending on the exogenous parameters of the experiment (Table 8).

[30] Uninformed: subjects given enough information to calculate Bayes law on their own
Informed: subjects provided with the Bayes Law calculation(posterior probabilities)
[31] BEU 1st and 2nd choice deviation rate
[32] RL 2nd choice deviation rate

Table 8: Eight Potential Two-Action Choice Sequences

Sequence of Action Choices	1st Choice Action	Message Received	2nd Choice Action	2nd Action BEU or Non BEU by Round Set					
				Rds 1-4	Rds 5-8	Rds 9-12	Rds 13-16	Rds 17-20	Rds 21-24
S1	BEU	RED	STAY with 1st	BEU	BEU	BEU	BEU	BEU	BEU
S2	BEU	RED	Shift to Alternative	Non BEU	Non BEU	Non BEU	Non BEU	Non BEU	Non BEU
S3	BEU	BLUE	STAY with 1st	Non BEU	Non BEU	BEU	Non BEU	Non BEU	BEU
S4	BEU	BLUE	Shift to Alternative	BEU	BEU	Non BEU	BEU	BEU	Non BEU
S5	Non BEU	RED	STAY with 1st	Non BEU	Non BEU	Non BEU	Non BEU	Non BEU	Non BEU
S6	Non BEU	RED	Shift to Alternative	BEU	BEU	BEU	BEU	BEU	BEU
S7	Non BEU	BLUE	STAY with 1st	BEU	BEU	Non BEU	BEU	BEU	Non BEU
S8	Non BEU	BLUE	Shift to Alternative	Non BEU	Non BEU	BEU	Non BEU	Non BEU	BEU

For example from Table 8, when a red message is received sequence 1 (S1) is BEU optimal for all rounds, whereas, when a blue message is received sequence 3 (S3) is BEU optimal for rounds 9-12 and 21-24, and sequence 4 (S4) is BEU optimal for rounds 1-8 and 13-20. Additionally, the table highlights for each sequence whether a subject stays with their initial action or shifts to the alternative as their second choice.

To understand the causes of the observed BEU and RL inconsistency rates, I run logit regressions (both random and fixed effects) with the 1st and 2nd choice BEU inconsistency and 2nd choice RL inconsistency as the dependant variable to determine the odds ratio[33], log odds[34], and the marginal effects[35] of the independent variables on these three outcomes The dependant variable in equation (1) represents a 1st choice inconsistency from the risk neutral BEU decision by round and subject. The dependant variables in equation (2) & (3) represent a 2nd choice inconsistency from the risk neutral BEU decision and the 2nd choice inconsistency from the RL decision, respectively, by round and subject. In all three equations the dependent variable is a dichotomous outcome variable, where 0 represents consistency and 1 represents inconsistency relative to the designated behavior benchmarks. There are three types of variables used to explain the data. First, there is a group of explanatory variables that change over the rounds but are the same for all individuals in a given round. Second, there is a set of explanatory variables that vary both over the rounds and between subject and session. Finally, there are explanatory variables that vary between individuals but do not vary over the rounds.[36] The three equations are presented and described in Table 9.

[33] Odds Ratio = (Proportion of sucesses: positive dependent variable(1)/Proportion of failures: nonpositive dependent variable (0))
[34] The logarithm of the odds ratio
[35] Change in the probability of observing the dependent variable, if the independent variable changes by one unit
[36] The Breusch and Pagan Lagrangian multiplier test for all three equations established that individual effects are present in the data. The Hausman test cannot reject the null hypothesis that the coefficients for the fixed and random effects model are the same; implying that the random effects coefficients are not correlated with the individual error terms. As an additional test, I ran a GLS regression fixed and random

Table 9: Logit Regression Equations & Descriptive Summary of Variables

	Equation 1	Equation 2	Equation 2
Dependent Variable	**1st Action Choice** **BEU Inconsistency** (Prior to observing a message) **1** Choice inconsistent **0** Choice consistent	**2nd Action Choice** **BEU Inconsistency** (After observing a message)	**2nd Action Choice** **RL Inconsistency** (After observing a message)
1. Explanatory Variables: vary over rounds, are the same for all individuals	**Experience** **1** second 12 rounds **0** first 12 rounds	**Experience**	**Experience**
	OTP **1** OTP message task **0** Free message task	**OTP**	**OTP**
	Informed **1** subjects given Posterior **Prob.** **0** subjects not given Posterior Probabilities	**Informed**	**Informed**
	Ex-ante Difference in expected POs for choosing Action A or B[37] -Continuous, changes every 4 rds.	**N/A**	**N/A**
	N/A	**Informative Power of chip draw**[38] -Continuous, changes every 4 rds.	**Informative Power of chip draw** .
2. Explanatory Variables: vary over rounds & between subjects	**N/A**	**Difference in Expected Payoffs between Action A and B conditional on the message received**[39] -continuous, changes every 4 rds. & conditional on chip draw	**Difference in Expected Payoffs between Action A and B conditional on the message received**
	N/A	**Shift Required from 1st choice to be BEU optimal** 1 Shift 0 Stay	**Shift Required from 1st choice to be BEU optimal**
	BEU action not consistent with the higher payoff state in prior round **1** Inconsistent **0** Consistent	**BEU action not consistent with the higher payoff state in prior round**	**BEU action not consistent with the higher payoff state in the prior round**
	Paid Second **1** payment on 2nd action **0** payment on 1st action	**Paid Second**	**Paid Second**
3. Explanatory Variables: same over all rounds but vary by individual	**Female** **1** Female **0** Male	**Female**	**Female**
	English Second **1** English 2nd language **0** English 1st language	**English Second**	**English Second**
	Post Survey **1** Classified as RL **0** Classified as theorist	**Post Survey**	**Post Survey**
	Risk Aversion Continuous-Eckel-Grossman test 1 highest RA to 10 least RA	**Risk Aversion**	**Risk Aversion**
	Econ Math **1** Math/econ/optimization **0** non-math student	**Econ Math**	**Econ Math**

effects model and performed the Hausman test and obtained the same result. Comparisons of the same coefficients from all models show the differences to be minimal; the signs and the statistical significance on the coefficients remain the same. Given the additional degrees of freedom, I report on the results from the random effects model and provide the fixed effects results in the appendix.

[37] $\text{Diff EPO} = \frac{[\pi_1 C(A,S_1) + \pi_2 C(A,S_2)] - [\pi_1 C(B,S_1) + \pi_2 C(B,S_2)]}{[\pi_1 C(A,S_1) + \pi_2 C(A,S_2)] + [\pi_1 C(B,S_1) + \pi_2 C(B,S_2)}$

[38] [(# of red chips -#blue chips)/(#blue chips+#red chips)]

[39] $\text{EP}_{\text{action A}} > \text{EP}_{\text{action B}} \rightarrow \quad [\,\pi_{1.1}\, C(A,S_1) + \pi_{2.1}\, C(A,S_2)] - [\pi_{1.1}\, C(B,S_1) + \pi_{2.1}\, C(B,S_2)]$
$=$ difference between good and bad state conditional on the colour chip draw

2.2.1 Summary Results

Table 10 presents the mean values of first and second choice inconsistency rates relative to the BEU benchmark and the 2nd choice inconsistency rate relative to the RL benchmark for all six treatment groups (appendix 7 shows summary statistics for all variables).

Table 10: 1st and 2nd BEU and RL Inconsistency Rates by Treatment Group

Inconsisteny Rate		Treatment 1 FREE/OTP Uninformed	Treatment 2 FREE/OTP Informed	Treatment 3 OTP/FREE Uninformed	Treatment 4 OTP/FREE Informed	Treatment 5 FREE/FREE Uninformed	Treatment 6 FREE/FREE Informed	ALL
1st Choice BEU Inconsistency Rate	**All Rounds**	**12.5%**	**12.7%**	**14.2%**	**13.3%**	**13.4%**	**16.7%**	**13.8%**
	Rounds 1-12	14.7%	16.4%	17.5%	18.3%	18.4%	20.8%	(0.345)
	Rounds 13-24	10.3%	8.9%	10.8%	8.3%	8.3%	12.5%	
2nd Choice BEU Inconsistency Rate	**All Rounds**	**16.7%**	**14.9%**	**15.6%**	**15.4%**	**18.0%**	**16.0%**	**16.0%**
	Rounds 1-12	14.2%	14.4%	19.4%	17.4%	19.5%	20.8%	(0.367)
	Rounds 13-24	19.2%	15.3%	11.7%	13.4%	16.4%	11.1%	
2nd Choice RL Inconsistency Rate	**All Rounds**	**36.1%**	**46.9%**	**37.7%**	**36.6%**	**41.1%**	**44.8%**	**40.6%**
	Rounds 1-12	34.9%	49.7%	36.1%	33.9%	39.1%	42.8%	(0.491)
	Rounds 13-24	37.2%	44.0%	39.2%	39.2%	43.0%	46.7%	

Informed: subjects provided with the posterior probability calculations conditional on the chip draw
Uninformed: Subjects provided with sufficient information to calculate the posterior probabilities on their own

One sample t-tests comparing subject behaviour to 1st and 2nd choice BEU and 2nd choice RL benchmarks confirms that, in aggregate, subjects do not have an action choice (first or second) inconsistency rate relative to the BEU or the RL heuristic that is equal to zero (99% confidence interval). However, subjects' behaviour is less divergent from the risk neutral BEU benchmark.

The first action choice taken by subjects occurs prior to an imperfect message, where the lottery associated with the optimal action is first-order stochastically dominant (FOSD) to the alternative action lottery. Therefore, from the BEU inconsistency rate, it is not possible to conclude that subjects can or cannot maximize their expected utility, it may be that a subject's first action choice that is consistent with the BEU benchmark simply indicates the subject's ability to recognize the FOSD lottery. The second action choice BEU inconsistency rate gives insight into the subjects' ability to combine both Bayes law and Expected Utility Theory in particular when a blue chip message is observed. The second action choice caused greater diversion overall from the BEU benchmark as evidenced by the increase in subjects' BEU inconsistency rates from

13.8% for the first action choice to 16% for the second action choice[40]. Figure 2 shows that 58 subjects (32.2%) followed precisely the BEU decision rule for their first action choice for all 24 rounds. However, this number drops to 14 subjects (7.8%) who followed precisely the BEU decision rule for their second action choice. Relaxing the BEU decision rule assumptions to include subjects who followed the BEU model prediction within a 95% of the time (23 of 24 rounds are BEU consistent) increases the proportion of subjects following the action choices of a BEU decision maker to 81 subjects (45%) for the first action choice and 36 subjects (20%)for the second action choice.

Figure 2: Proportion of Subjects by BEU 1st and 2nd choice Inconsistency rate-All Treatments

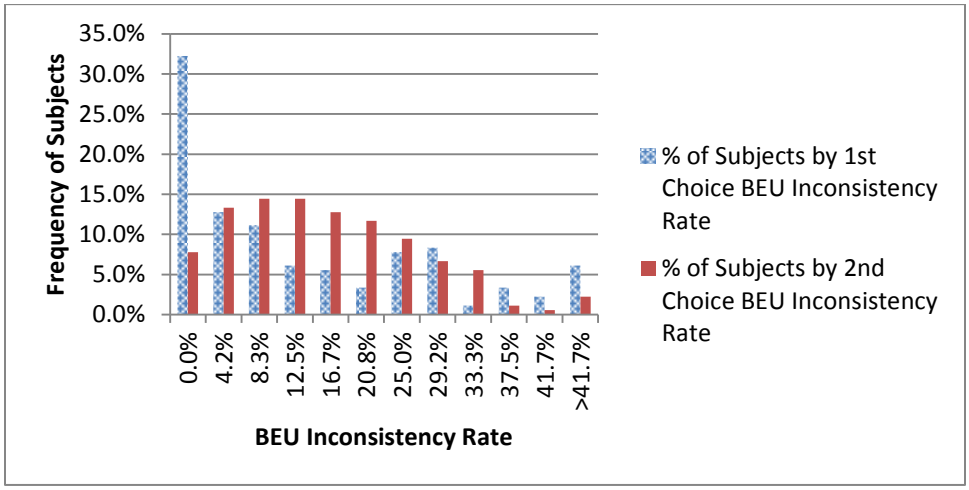

On average, subjects' RL 2nd action choice inconsistency rate is 41% (Table 10). Figure 3 highlights that there are no subjects (0%) who followed precisely the simple RL heuristic for all 24 rounds. Unlike the BEU model, relaxing the RL decision rule model to include subjects who followed the RL model within a 95% confidence interval does not increase the proportion of subjects who followed the RL model predictions.

[40] This difference between 1st and 2nd BEU inconsistencies is statistically significant at the 1% level

Figure 3: Proportion of Subjects by RL 2nd choice Inconsistency rate-All Treatments

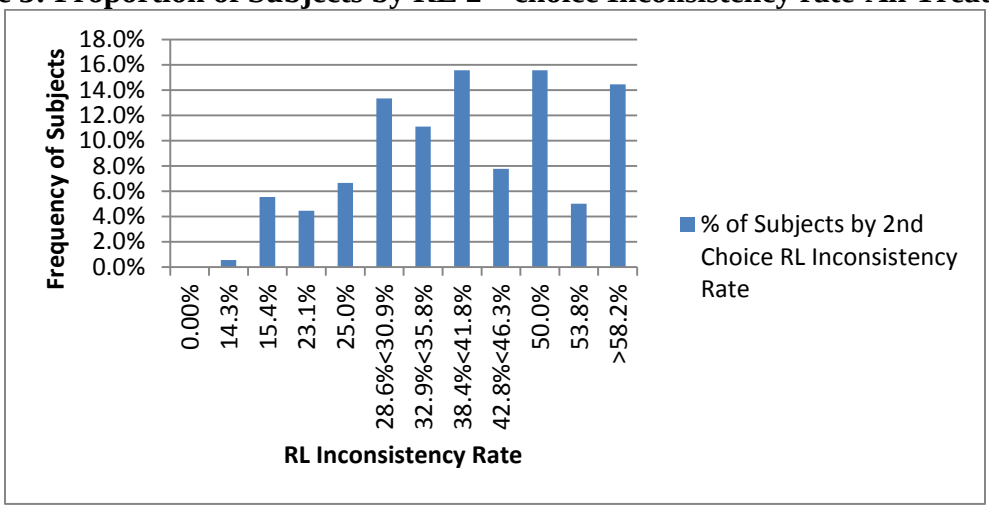

The average financial payoff for subjects over the 24 rounds was $33.50, 4.2% less than the average BEU decision maker ($34.98) and 4.1% more than the average earned if all subjects were Reinforcement Learners ($32.19)[41].

Figures 4 & 5 identify the proportion of subjects who follow each of the 2 action sequence choices (described in table 8) for each round set conditional on observing a red (Figure 4) or blue (Figure 5) chip message. Recall, each round requires a subject to take an action choice prior to a colour chip draw and an action choice after observing a chip draw. These two choices per round conditional on the colour chip draw represent one sequence of decisions.

[41] These payoffs are determined based on subject decision and random choice: 1) the decision choices of the subjects (i.e., first, second and WTP choices) and, 2) random draws (during the Free task a draw decides whether a subject will receive payment on their first or second action choice and during the OTP task a random price draw determines the price a subject will pay for the information observed). It is for this reason that payoffs are not used as an outcome variable in this analysis.

Figure 4: Proportion of Subjects who follow each of the 2-Action Sequence decision choices conditional on a Red Chip message by the sets of rounds sharing the same exogenous parameters. (x=% of subjects, y= round sets)

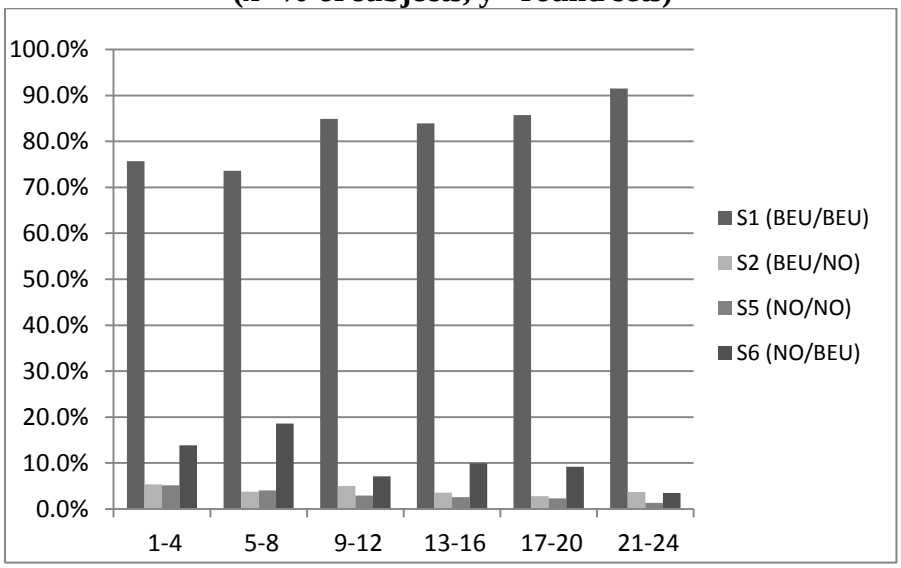

BEU/BEU- 1st Action follows the prediction of the BEU model/2nd Action follows the BEU model etc…

The red message accounted for 55.3% of all the imperfect messages observed by the subjects. Eighty-two percent (82%) of subjects who observed a red chip message chose, on average over the 24 rounds, the two-action sequence that equaled the BEU optimal sequence (see sequence 1, Figure 4).

Figure 5: Proportion of Subjects who follow each of the 2-Action sequence decision choices conditional on a Blue Chip Message by the sets of rounds sharing the same exogenous parameters. (x=% of subjects, y= round sets)

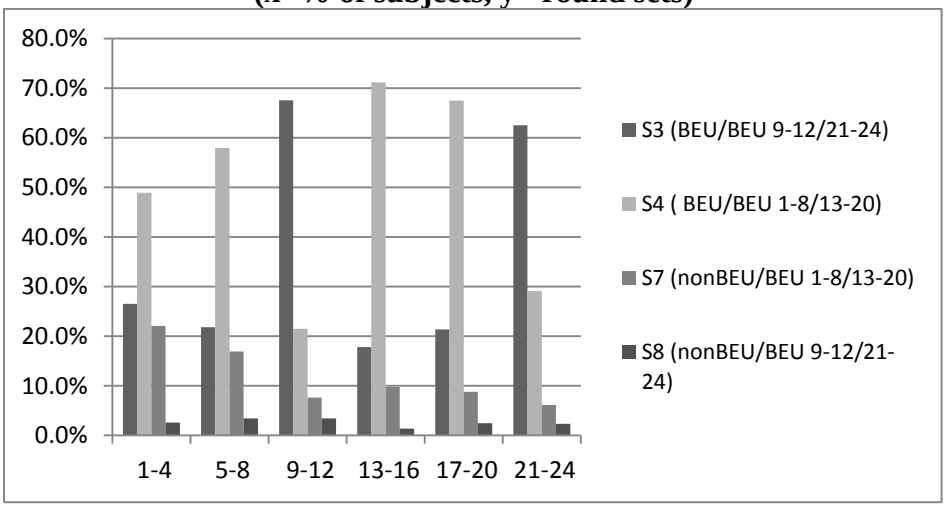

BEU/BEU 9-12/21-24- 1st Action follows the prediction of the BEU model/2nd Action follows the prediction of the BEU model for rounds 9-12 & 21-24 and does not follow the BEU model otherwise etc…

The blue chip message accounted for 44.7% of all the imperfect messages observed by the subjects. Sixty-three percent (63%) of subjects who observed a blue chip message chose, on average over 24 rounds, the two-action sequence that equaled the BEU optimal sequence (see sequence 3, rounds 9-12 and 21-24 & sequence 4, rounds 1-8 & 13-20, Figure 5).

On average the subjects' behaviour is reflective of the BEU sequence of decisions more often when the red message was received (7.1% BEU 2nd choice inconsistency rate). If subjects' behaviour is a consequence of BEU decision rules when observing the red chip, then they assigned probabilities of possible outcomes based on this message and calculated expected payoffs to value these outcomes. They then selected the action associated with the highest payoff. It follows that they should apply this same decision rule when they observe a blue chip. However, in contrast, although some subjects follow the BEU model predictions, there is a higher 2nd choice BEU inconsistency rate across all rounds of 27% when a blue chip is observed (see figure 6).

This difference in BEU inconsistency rates between the red versus the blue chip message may be the result of one or more of the following. First, it is possible that subjects did not follow BEU decision rules, but rather were capable of identifying the FOSD lottery associated with the optimal action given the red chip message. Therefore, as there was no FOSD lottery associated with the optimal action when a blue chip message was received, the decision environment became more complicated leading to more BEU inconsistencies. Second, it is possible that subjects in both decision environments (i.e., red and blue chip message) inappropriately estimated the informational value of the message received (i.e., the posterior probabilities associated with the message) when maximizing expected utility. In this case, the exogenous parameters of the experiment are such that when a red chip message is received their exists a broader range of posterior probability estimates where a subject may make a Bayesian updating error yet still arrive at the correct BEU action choice. On the other hand, the narrower range associated with the blue chip message for making a Bayesian updating error where a subject would still arrive at the correct BEU decision results in a greater likelihood of a BEU inconsistency. Finally, the second action choice conditional on the red chip message requires a

BEU maximizer to stay with the initial optimal action chosen prior to receiving the message for all rounds. However, the second action choice conditional on the blue chip message requires the BEU maximizer to shift to the alternative action from the initial optimal action chosen prior to receiving the message for two-thirds of the rounds (16 of 24 rounds) and to stay with the original optimal action for one-third of the rounds. When the BEU decision rules changes from a shift to a stay with the initial pre-message action choice, the subject's behaviour is more suggestive of how they estimated the informational value of the message received particularly in the absence of a FOSD lottery choice.

When a blue message is observed, in addition to observing behaviour which is reflective of BEU decision rules, I observe during different round intervals behaviour that is reflective of both over-weighing and under-weighing the informational value of the message received.

Figure 6: BEU 2nd Choice Inconsistency rates conditional on chip colour observed by round (x= rounds and y= BEU inconsistency rate)

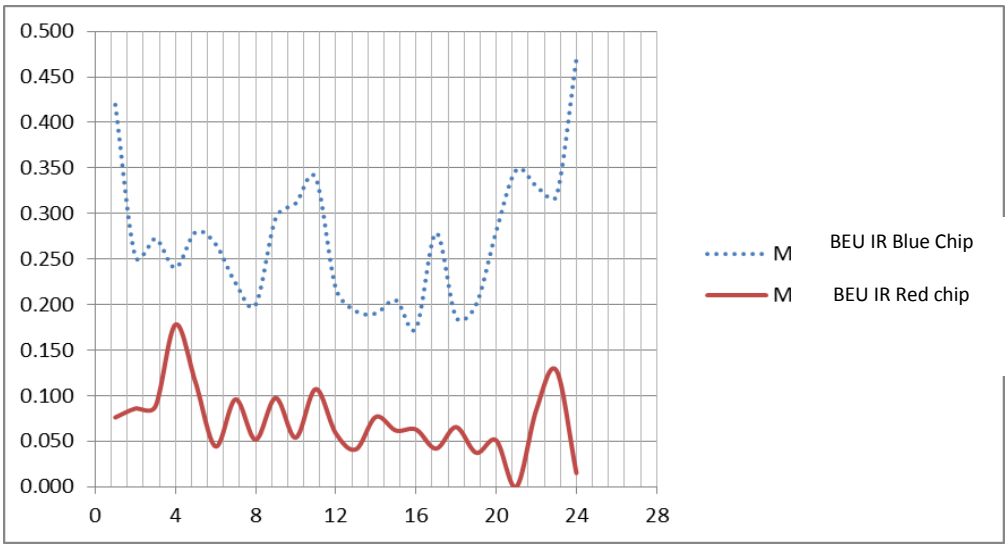

2.2.2 Main Results

Table 11 provides the results from the 3 logit regressions with BEU 1st choice inconsistency (eqn. 1), BEU 2nd choice inconsistency (eqn. 2) and RL 2nd choice inconsistency (eqn. 3) as the

dependent variable. I report the random effects results; odds ratios and marginal effects[42] ; in Table 11 and provide the fixed effects results in the appendix.[43]

Table 11: Logit Regression Results from equation 1, 2 and 3 found in Table 9

	(1) BEU 1st Choice Inconsistency			(2) BEU 2nd Choice Inconsistency			(3) RL 2nd Choice Inconsistency		
Variables (see table 9)	Marginal Effects	Odds Ratio	xtlogit, Re	Marginal Effects	Odds Ratio	xtlogit, Re	Marginal Effects	Odds Ratio	xtlogit, Re
Experience	-0.059***	0.451***	-0.797***	-0.028**	0.778**	-0.250**	0.032	1.147	0.137
	0.010	0.046	0.103	0.014	0.100	0.129	0.025	0.124	0.108
OTP	0.015	1.226	0.204	0.062***	1.77***	0.554***	0.005	0.884	0.024
	0.009	0.154	0.126	0.022	0.35	0.199	0.038	0.152	0.167
Experience*Paid	-	-	-	0.002	0.981	0.018	-0.034	1.156	-0.148
				0.029	0.255	0.260	0.044	0.225	0.194
Informed	0.014	1.226	0.194	0.016	1.15	0.143	0.062***	1.311***	0.271***
	0.020	0.322	0.265	0.018	0.188	0.163	0.025	0.142	0.108
Paid*Informed	-	-	-	-0.065***	0.560***	-0.581***	-0.046	0.818	-0.201
				0.023	0.118	0.210	0.044	0.157	0.193
Ex-ante Difference in Expected Payoffs	-0.342***	0.009***	-4.66***		-	-		-	-
Action A Vs. B	0.057	0.006	0.616						
Informative Power of the Chip Draw	-	-	-	-0.476***	0.014***	-4.287***		0.099***	-2.31***
				0.051	0.006	0.436		0.036	0.367
Difference in Expected Payoffs Action A	-	-	-	-1.027***	0.000***	-9.25***	-0.713***	0.045***	-3.11***
Vs. B Conditional on message received				0.090	0.000	0.769	0.158	0.031	0.699
Shift from 1st action Required	-	-	-	0.128***	3.160***	1.150***	0.034	1.159	0.148
to follow BEU for 2nd action				0.013	0.366	0.116	0.026	0.130	0.112
BEU action not consistent with the	-0.005	0.932	-0.070	0.022***	1.230***	0.205***	-0.151***	0.519***	-0.656***
higher payoff state in prior round	0.008	0.101	0.109	0.013	0.120	0.098	0.023	0.053	0.103
Subject Paid on second action	0.024***	1.377***	0.320***	-0.009	0.922	-0.081	-0.035	0.860	-0.151
in Prior Round	0.008	0.146	0.106	0.011	0.091	0.099	0.022	0.082	0.094
Female	0.027	1.439	0.364	0.021	1.200	0.184	0.022	1.102	0.098
	0.021	0.401	0.279	0.017	0.184	0.153	0.022	0.106	0.096
English 2nd	0.066***	2.458***	0.900***	0.021	1.210	0.189	-0.001	0.997	-0.003
	0.025	0.787	0.320	0.120	0.215	0.177	0.025	0.113	0.113
Post survey: RL	-0.005	0.934	-0.069	0.011	1.110	0.103	0.063***	0.759***	-0.276***
	0.022	0.276	0.295	0.018	0.179	0.161	0.023	0.076	0.099
Risk Aversion	-0.003	0.959	-0.042	-0.002	0.979	-0.021	0.001	0.994	-0.006
	0.004	0.046	0.048	0.003	0.026	0.026	0.004	0.017	0.017
Age	-0.013***	0.845**	-0.168**	-0.001	0.989	-0.011	-0.005	0.979	-0.021
	0.006	0.068	0.081	0.005	0.042	0.043	0.006	0.026	0.026
Econ Math Student	-0.004	0.951	-0.04	-0.027	0.783	-0.245	0.032	1.148	0.138
	0.023	0.295	0.310	0.019	0.135	0.172	0.024	0.120	0.105
Obs.		4320			4320			2265	
log likelihood/R-squared		-1451.85			-1630.61			-1472.77	

***p-value ≤.01 <**p-value ≤.05 < *p-value < .10*

[42] Marginal effects are calculated using Average marginal effects, that is, a marginal effect is computed for each case and the effects are then averaged.

[43] Appendix 8 & 9 shows the Fixed Effects model for the 1st and 2nd choice BEU Inconsistency rate and Appendix 12 shows the Fixed effects model for the 2nd choice RL inconsistency rate.

Result 1: ***Subjects' decisions over time converge toward optimal choices prior to observing an imperfect message.***

A subject's first action choice is made in advance of observing an imperfect message, where there is an equal probability that the decision is being conducted using either bag 1 or bag 2. For all rounds the lottery associated with the BEU first action choice first-order stochastically dominates the lottery associated with the alternative action. Table 11, column 1, highlights the results from the Logit regression with the first choice BEU inconsistency as the dependent variable and provides some insight into the key aspects of subject decision behaviour when choosing a first action.

The predicted probability of a subject committing a first action choice BEU inconsistency is 5.9 percentage points (ppts) greater during the first 12 versus the last 12 rounds of the experiment. Through task repetition of the first decision choice, subjects' behaviour over-time converges on the optimal action. Additionally, as the difference between the FOSD lottery associated with the optimal action and the alternative action's lottery is exaggerated, the odds of a BEU inconsistency are less likely (significant at 1%)[44]. Figure 7 illustrates the subjects' first choice inconsistency rate relative to the BEU benchmark over the 24 rounds. In early rounds, subjects violate both BEU decision rules as well as first-order stochastic dominant choices and only converge on optimal decisions with practise and when the difference between the FOSD lottery and the alternative lottery are exaggerated. This first result does not necessarily imply irrational behaviour on behalf of the subject as the informational knowledge gained from the consequences of either action choice could also be considered a rational learning process.

[44] The difference in ex-ante expected payoffs is the same and remains static during rounds 1-8 and 13-20 and increases by the same amount twice during the experiment; during rounds 9-12 and 21-24.

Figure 7: BEU 1st Choice Inconsistency Rate by Round-All Subjects
(x=1st Choice BEU Inconsistency Rate, y= round)

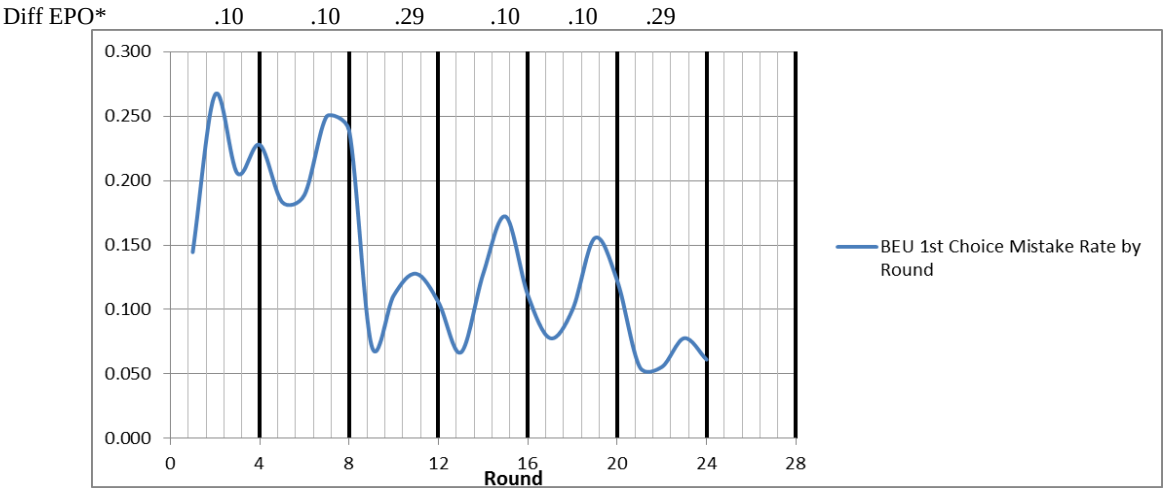

$$* \; Diff\ EPO = \frac{[\pi_1 C(A,S_1) + \pi_2 C(A,S_2)] - [\pi_1 C(B,S_1) + \pi_2 C(B,S_2)]}{[\pi_1 C(A,S_1) + \pi_2 C(A,S_2)] + [\pi_1 C(B,S_1) + \pi_2 C(B,S_2)]}$$

Result 2*: Subjects have a higher BEU inconsistency rate when the BEU decision rule requires the combination of both Bayes Law and Expected Utility theory to arrive at the optimal response. Additionally, this higher inconsistency rate is accentuated when subjects are performing the OTP versus the FREE message decision task.*

Subjects make a second decision, selecting either action A or B, after observing the colour of the chip drawn from the selected bag. To follow the 2nd choice BEU model predictions, a subject applies Bayes law in conjunction with Expected Utility theory. Table 11, column 2, highlights results from the Logit regression with the 2nd choice BEU inconsistency as the dependent variable and provides some insight into the key aspects of subject decision behaviour when choosing a second action. The action selected by subjects after observing an imperfect message results in a higher BEU inconsistency rate than the action selected prior to observing the message. Subjects' 2nd choice BEU inconsistency rate is 16%, 2.2% greater than the first choice BEU inconsistency rate of 13.8%.[45]

When performing the FREE and OTP message task, subjects take two actions; an action before and an action after observing an imperfect message. During the OTP message task subjects'

[45] A two sample t-test identifies this difference to be statistically significant at the 1% level.

make an additional decision and specify their willingness to pay in order to have their second versus their first action choice be used to calculate their earnings. Subjects' 2nd choice BEU inconsistency rate is 17.8% when subjects are performing the OTP message decision task, 2.6% greater than the BEU inconsistency rate when subjects perform the FREE message task.[46] Table 11, column 2, confirms this result and highlights that the predicted probability for a 2nd choice BEU inconsistency increases by 6.2 ppts when subjects are performing the OTP versus the FREE message task. Potential explanations for the higher 2nd choice BEU inconsistency rate associated with OTP message task follow.

It is of interest to investigate whether this result is due to a lack of commitment to the accuracy of the second choice on behalf of the subject when they are required to pay in order to have this choice used to calculate earnings. In this case, the lack of commitment in decision quality of second action choices could result in; 1) a WTP amount of zero or a relatively small WTP bid even though the subjects' second choice is different from their first choice, and 2) a BEU 2nd choice inconsistency rate which is higher for these subjects than for those who specified a more substantial WTP amount. Overall, the BEU 2nd choice inconsistency rate is 24.2% when the second action is different from the first action choice during the OTP message task. This inconsistency rate drops to 20.6% when the WTP is greater than $0.05 and increases substantially to 41.7% when the WTP amount is less than $0.05. This suggests that subjects with a higher willingness to pay amount are more BEU accurate. Although this result supports 'a lack of effort toward decision accuracy in the face of an additional cost' hypothesis, subjects who specified a WTP less than $0.05 only accounted for 16% of the total observations where the first choice did not equal second choice actions.

Samuelson & Zeckhauser (1988) provide another potential explanation for the differences in the BEU inconsistency rate between the FREE and OTP message tasks. They demonstrate through a study of a series of decision making experiments that individuals have a tendency to maintain their previous decision choice even when new incomplete information is acquired that indicates that this decision choice is no longer optimal. They highlight that this 'status quo' bias in some cases leads individuals and firms to partake in fewer information searches than what is required

[46] See footnote 19

to arrive at an optimal decision because they put greater importance on their original decision choice relative to the informational value provided by a new contradictory message. Table 11, column 2, show the predicted probability of a 2nd choice BEU inconsistency increases 12.8 ppts when subjects are required to change their initial action to the alternative action conditional on the chip draw in order to follow the BEU model predictions. To test whether subjects have a bias for the status quo and to determine whether this bias is more prevalent when performing the OTP versus the FREE message task, a comparison is made between the BEU inconsistency rates for both these tasks when the subjects' second action choice is the same as their first action choice (subjects maintaining the status quo). The 2nd choice BEU inconsistency rate when first and second action choices are aligned is 11.7% when subjects perform the FREE message task and 15% when subjects perform the OTP message task. A two sample t-test confirms that this difference is statistically significant at the 1% level. The logic, given this evidence, is that subjects maintain their first decision choice due to a status quo bias which is accentuated when there is an added cost (WTP decision requirement).

Another plausible explanation may be loss aversion. Tversky & Kahneman (1984) suggest that losses are two times more psychologically powerful than gains. Therefore, it may be that the subjects preferred to avoid the losses associated with the cost of using the information over the acquired gains from the informational knowledge the message provided. This is discussed in more detail in chapter 3.

Finally, is it possible that the change imposed by the additional step required to complete the OTP versus the FREE task created confusion or additional complexity to the decision environment? As a facilitator during the experiments I observed that the WTP decision task and the random price draw required for truthful elicitation of the WTP amount caused confusion and resulted in many requests by participants to repeat the instructions. Furthermore, during the OTP task, subjects relied on the posterior probability calculation (informed) as well as past outcomes to assist them with their decision choice.

The predicted probability of a 2nd choice BEU inconsistency is 6.5 ppts less when a subject is informed (provided the probability of being in either state conditional on the chip draw) versus uninformed when performing the OTP message task and this same coefficient has no

significance when subjects are performing the FREE message task.[47] The finding that informed subjects perform no better relative to the BEU benchmark than subjects who were uninformed when performing the FREE message task appears contrary to the hypothesis that subjects do not follow the BEU decision rule because they lack the math skills or cognitive sophistication to perform the Bayes law component of the decision rule. A more plausible explanation may be that the procedural steps to complete the FREE task were easy to master allowing more clarity around odds estimations and less reliance on other cues when making decisions

According to the data, informed subjects performing the OTP message task have the same 2nd choice BEU inconsistency rate as subjects performing the FREE message task[48]. That is, the 2nd choice BEU inconsistency rate difference between the FREE and OTP message task is eliminated when subjects are provided with the Bayes law posterior probability calculation (informed). Additionally, subjects have a higher 2nd choice BEU inconsistency rate during the OTP rounds when the BEU 2nd choice in the prior round was inconsistent with the higher pay-off state for that round.

Given the application of different decision rules (informed and past history) and recognition that OTP task represented a more complex decision environment, these results suggest that learning behavior depends to some extent on the context and environment in which the decision making is conducted.

Result 3: There is evidence that suggests that the RL and BEU heuristics are complementary behaviours and when both are present can either enhance or diminish optimal decisions.

The predicted probability that a subject's behaviour will be reflective of RL model used in this study increases by 6.2 ppts when subjects are uninformed (column 3, Table 11) versus informed. Furthermore, the predicted probability that subjects will apply the RL heuristic in the current round is 15.1 ppts greater when the BEU 2nd action choice in the prior round was inconsistent versus consistent with the higher pay-off state for that round (column 3, Table 11).

[47] To confirm this result, I add the interaction term, OTP*informed and show that the odds of a BEU inconsistency during the OTP message rounds are 1.77 (1/.571) times more likely when subjects are uninformed.

[48] To confirm this result, I test paid-paid*informed=0 and I cannot reject the null.

There is evidence that suggests a subject's 2nd choice RL heuristic can improve and/or diminish behaviour that is reflective of optimal decision theory. The 2nd choice RL heuristic equals the 2nd choice BEU heuristic 39.8% of the time.[49] The BEU inconsistency rate is 13.6% when the RL and BEU heuristics are aligned, 18.9% when the heuristics clash and 16.5% when no RL heuristic exists (see Table 12). The differences between these BEU inconsistency rates are statistically significant at the 1% level (pairwise tests). [50] These results indicate that when a past BEU decision is rewarded (i.e. a WIN) a subject has a greater propensity to apply the BEU decision rule in the future resulting in fewer BEU inconsistencies (the RL and BEU heuristic are aligned). Additionally, if the BEU decision is not rewarded (i.e. LOSE), potentially creating a future decision environment where the subject's RL and BEU heuristic clash, optimal decision behaviour is compromised. However, it does not follow that when the heuristics clash, subjects' behaviour is more reflective of Reinforcement learning (see Table 12). For example, subject behaviour is more reflective of RL when they are uninformed (statistically significant at the 1% level), but is not more reflective of BEU actions when informed (unless performing the OTP task). Similarly, although subjects have a higher 2nd choice BEU inconsistency rate when performing the OTP task, this same coefficient is not accompanied with a lower statistically significant 2nd choice RL inconsistency rate.

Charness and Levin (2005), suggest in their study the existence of a cross-over threshold where subjects' behaviour no longer reflect BEU decision theory when the task increases in complexity and becomes more reflective of this simple reinforcement learning model. However, the implication of the results from this study is that subjects do not substitute BEU with RL decision rules when tasks increase in complexity, but rather complement the BEU decision with the RL

[49] As a result the difference in expected payoffs between the action A and B conditional on the chip draw and the informative power of the chip draw (defined by the distribution of red to blue chips contained within each bag) coefficients found in table 12, column 3, have the same negative sign and statistical significance (at the 1% level) as the BEU Logit regression found in table 11, column 2. Removing the aligned BEU and RL heuristics from the RL inconsistency rate regression found in table 12 (see appendix 12) and observing only the RL heuristics where the BEU and RL are different results in a positive and statistical significant coefficient on the difference in payoffs between the good action and bad action state and the Informative Power of the chip draw coefficient losing its statistical significance.

[50] Subjects are 1.3 times less likely to be BEU inconsistent when their RL heuristic is the same as the BEU optimal choice versus when these two heuristics clash. Similarly, subjects are 1.25 times less likely to be BEU inconsistent when no RL heuristic exists (i.e. they have no prior history) versus when the two heuristics clash.

heuristic when they are aligned in an effort to achieve decision optimality in more complicated environments.

Table 12: BEU & RL Inconsistency Rates conditional on BEU and RL (Not) Alignment

	Aligned	**Not Aligned**	**No RL**
No. of Obs.	1361	905	2055
BEU 2nd Choice Inconsistency Rate	13.6%	18.1%	16.5%
RL 2nd Choice Inconsistency Rate	13.6%	81.9%	NA

Two additional findings of interest pertaining to Reinforcement Learning model utilized in this study: 1) The predicted probability of a first action choice BEU inconsistency is 2.4 percentage points greater when the subjects' second action choice versus their first actions choice receives payment in the prior round (see Table 11, column 1).[51]This suggests that subjects place more weight on optimal decision making at an action choice decision point (1st or 2nd) when this same action choice (1st or 2nd) was rewarded in the past, providing some additional evidence of reinforcement learning; And, 2)subjects who were categorized as reflective/reinforcement learners versus logical/theoretical learners, based on the condensed version of the Honey & Mumford (1986) personality type questionnaire (see appendix 5) conducted at the end of the experiment, were less likely to deviate from the RL benchmark.

Result 4: There are two systematic decision behaviour patterns that deviate from optimal BEU decision theory that are not fully explained by the Reinforcement Learning model:

1. **Behaviour which is reflective of an over-weighing of the informational value of the message received (a.k.a. Over-weigh) ;**
2. **Behaviour which is reflective of an under-weighing of the informational value of the message received (a.k.a. Status Quo[52])**

The design of the experiment, specifically the asymmetric state contingent payoffs associated with either action (A or B), creates intervals of rounds when a blue message is received where behaviour either reflective of over-weighing or under-weighing of the informational value of the

[51] Recall, a first action choice will receive payment as a result of a random draw (FREE message task) or an unsuccessful specified WTP amount (OTP message task).

[52] Subjects display a bias towards staying with original action choices regardless of the message received.

message received can be observed separate from BEU decision rule. In contrast, when a red message is received, subjects who display either of these two behaviour patterns look identical in behaviour to subjects who follow the BEU model.

Figure 8 highlights the proportion of subjects who follow each of the four potential sequence of action choices (1st and 2nd) conditional on observing a blue chip message. The proportion of subjects who followed sequence 3 during rounds 9-12 and 21-24 and sequence 4 during rounds 1-8 and 13-20 provide the main evidence for the two identified alternative behaviour types; Over-weigh (S3) and Status Quo bias (S4), respectively. [53] Note that the subjects who followed the non-optimal decision, sequence 3, during rounds 9-12 are not the same subjects who followed the non-optimal decision, sequence 4, during rounds 1-8 and 13-20.

Figure 8: Proportion of subjects by 2-Action sequence decision choice conditional on a Blue Chip message by sets of rounds sharing the same exogenous parameter values. (x=round set, y=% of subjects)

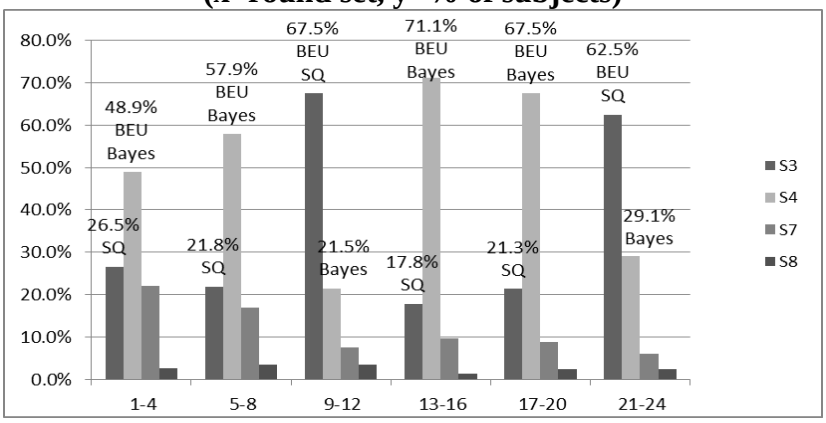

It is assumed that a subject who exhibits either the Over-weigh or the Status Quo decision pattern during the rounds where a blue message is received also applies these same decision rules when a red message is received; even though the behaviour pattern during these rounds is not observable.[54]

To determine group or individual characteristics that contribute to a subject's behaviour which is reflective of either BEU, an over-weighting of the informational value of the message received

[53] While sequence 3 is BEU optimal for rounds 9-12 and 21-24, it is not BEU optimal for rounds 1-8 and 13-20. Similarly, while sequence 4 is BEU optimal for rounds 1-8 and 13-20, it is not BEU optimal for rounds 9-12 and 21-24.

[54] Recall, that subjects who overweigh or play according to status quo rules will look identical to a BEU maximize when a red message is received.

or a status quo decision rule, I run 3 OLS regressions with the proportion of decisions by subjects that reflect each behaviour pattern as the dependent variable and determine the marginal effects of the independent variables on these three outcomes. For all three regressions, the dependant variable is a continuous outcome variable, where the outcome represents the percentage of decisions where subjects followed this behaviour pattern (BEU, Status quo or Over-weight). The explanatory variables include a subset of the variables detailed in Table 9, specifically, the variables that are the same over all rounds but vary by individual (see no. 3 Table 9). Two additional independent variables are added to capture which treatment group that the subject belonged. All three regressions can be found in table 13.

Subject behaviour is more likely to be reflective of the BEU sequence of decisions when they are performing the FREE message task in advance of the OTP task (4.7ppts increase), when they are male (4.1ppts increase) and English is their language of origin (4.6ppts increase).

Subject behaviour is most likely to be reflective of the Status quo decision rule when they are not math or economics students (3.7ppts increase) and classified as a Reinforcement Learner based on the post experiment survey (4.2ppts increase).

Although a proportion of subjects have behaviour reflective of over-weighing the informational value of the message received, there are no characteristics that are statistical significant contributing to this behaviour type.

Table 13: OLS Regression Results[55]

Variables	BEU OLS	Status Quo OLS	Overweight OLS
Free followed by OTP	**0.047*** **	**0.005**	**0.016**
	0.021	0.020	0.029
OTP followed by Free	**0.012**	**-0.024**	**0.033**
	0.021	0.020	0.030
Informed	**-0.006**	**-0.023**	**0.032**
	0.017	0.016	0.024
Female	**-0.042*** **	**-0.012**	**0.031**
	0.019	0.017	0.026
English 2nd	**-0.046*** **	**0.002**	**-0.015**
	0.022	0.020	0.031
Reinforcement Learner Survey	**0.010**	**0.042*** **	**-0.004**
	0.019	0.017	0.027
Risk Aversion	**0.002**	**0.002**	**-0.007**
	0.003	0.003	0.005
Age	**-0.006**	**0.005**	**-0.007**
	0.005	0.005	0.007
Econ Math Student	**0.003**	**-0.037** **	**-0.005**
	0.020	0.019	0.028
Obs.	180	180	180
Adjusted R-squared	0.0702	0.041	0.009

2.3 Conclusions

Subjects performing a relatively simple binary-decision task are adept at selecting optimal choices over time prior to observing additional statistically relevant information. Although this may provide evidence that suggests that subjects are capable of maximizing expected utility, it is also possible based on the lottery choices associated with each action, that subjects choose optimally simply by properly ranking the action associated with the FOSD lottery. When subjects observe a relevant information signal in the absence of a FOSD lottery and are required to combine Bayes law with expected utility theory in order to follow the BEU model predictions, there is greater deviation from this optimal behaviour. Furthermore, when the decision environment changes requiring subjects to perform a decision task which requires an additional

[55] Appendix 12 shows the complete Logit regressions for all independent variables by behaviour type.

step, optimal decision behaviour is further compromised. Specifically, when the decision environment changes from the FREE to the OTP message task subjects rely on past outcomes of success or failure and when available the information provided by the Bayes Law calculation (informed) to assist them with their decision choices.

Although the results are not sufficient evidence to confirm or refute the existence of a cross-over threshold where subjects no longer apply BEU decision rules due to task complexity (Charness and Levin, 2005), it does lend further support to the notion that learning behavior depends in small or large part on the context and environment in which the decision making is conducted. Grether (1989) suggests that in environments of uncertainty individuals use different decision rules in different decision situations. Furthermore, psychologists have also identified this finding and refer to these different decision rules used in different environments as the 'contingent judgement'[56] hypothesis (Payne, Bettman, and Johnson 1992). As observed in this study, when there is no FOSD lottery associated with the optimal action, potentially creating a more difficult decision environment, subjects' behaviour reflects the use of a different set of decision rules.

In addition to Rational decision theory (BEU) and the simple 'Reinforcement Learning' model used in this study, two alternative behaviour patterns emerged: 1) a sub-group of subjects put greater informational value on the message received; And, 2)a sub-group of subjects apply a status quo decision rule; under-weighing the value of new information when it is contrary to their original choice, preferring to stay with their first action choice regardless of the message received.

Samuelson and Zeckhauser (1988) identify regret avoidance and a taste for consistency (Charness and Levin, 2005) as a possible reason for this status quo behaviour. Individuals may have a desire to justify previous commitments, wish to avoid feelings of regret and have a need to feel in control. Kahneman and Tversky (1982), identified that individuals 'feel strong regret from bad outcomes that are the consequence of new actions taken than for similar bad consequences resulting from inaction'. It is reasonable to conclude that a subject's lack of math

[56] Subjects will substitute different decision rules contingent on the decision environment.

skills may create a more risky and uncertain decision environment leading to greater incidences of status quo behaviour.

Risk aversion is ruled out as a possible determinant of the subjects' first and second action choice non-optimal behaviour. When selecting an action prior to observing an imperfect message or selecting an action conditional on receiving a red chip message, the optimal action is always associated with the first-order stochastically dominant lottery. Hence, any subject with monotone utility preferences should select the optimal action regardless of risk preferences. Additionally, although the optimal second action conditional on receiving a blue chip message does not entail a first order stochastically dominant lottery, the differences in the expected payoffs from either action allows for a broad range of CRRA and CARA utility curves (see footnote 9). An attempt was made to determine the risk attitudes of the subjects at the end of the experiment through the administration of the Eckel-Grossman risk task (Eckel-Grossman, 2002; see appendix 6). Regardless of the subjects risk preferences (as defined by the Eckel-Grossman risk test), results from the random and between effects model found in Table 11 & 13, show no statistical significance on the risk attitude coefficient; suggesting that risk preferences did not influence subject behaviour.

The findings from this study suggest that individuals adopt different decision rules depending on both personal attributes (i.e. skillset, gender, experience) and on the context and environment in which the decision task is conducted. Of most interest is understanding whether these deviations from BEU persist in other decision making environments more representative of a real world market setting. As such, further research should be conducted to determine which individual characteristics have a higher propensity for a certain behaviour type and then determine how and when changes to the decision environment influence the choices of the individuals with these characteristics. These changes could also assist in identifying the threshold where subjects no longer apply BEU optimal decision rules but apply different decision making criteria. It may also help identify whether this threshold varies depending on the individual type.

Given the above, the ability to identify individual type, the degree of risk and uncertainty within the decision environment and how individuals behave in these changing environments would be essential in determining the proper mechanism necessary to ensure optimal choices.

Appendix 1

Instructions

Thank you for participating today!

By participating in this experiment, you will have the opportunity to earn money. The actual amount of money you will earn depends both on your choices and on random chance. During this session, we will ask you to make a series of decisions. Please make sure that you completely understand the instructions for each part of the experiment before making any decisions in that part of the experiment. If you have any questions at any point or need clarification, please raise your hand and the experimenter will come to you and answer your question.

In this session there are two rounds that are practice rounds for you to get familiar with two different decision tasks and a series of experimental rounds which will be used to calculate your earnings in a manner to be described in the workbook. You are not allowed to use a calculator, but may write down anything you may need to make your decision on the yellow tracking sheet provided by us.

You will have a ***Workbook*** that will contain the instructions for each round of the game. You will also use the workbook to record all your decision choices for each round of the game. To ensure confidentiality, your workbook is identified only by a participant number, which is never connected to your actual name.

You will also be asked at the end of the session to complete a short questionnaire. Please respond to this questionnaire truthfully and as accurately as possible. The questions provide the experimenter with important data that is of enormous help in organizing and interpreting your decisions. Your decisions and answers to the questionnaires are confidential and will not be revealed to anyone other than the experimenter. The data will only be identified by the participant code assigned to you and will not at any point be connected to your name in any way.

Please make sure that you completely understand the instructions for the experiment. It is important not to make any noises that might disturb others around you. If you have any questions, raise your hand and we will answer your questions individually.

Appendix 2

Practice Round -The Decision task with a 'FREE' Message Signal

As this is a practice round the potential earnings highlighted on the next page will not actually be paid. The intention of this round is to allow you to become familiar with the decision task with a "FREE" message signal.

In all free message signal rounds, there are two bags, each containing a combination of red and blue poker chips. There are 50 poker chips in each bag. However the number of chips that are red and the number of chips that are blue differs between these two bags. You will be informed of the number of red and the number of blue chips contained within each bag. In step 1, a random draw will decide which bag will be selected for use during the round. There is an equal chance that we will be playing the round using bag 1 or bag 2. However, you will not know until the end of the round which bag has been randomly selected for play during the round. In Step 2, you are asked to choose one of two actions and told the financial consequences of taking each of these two actions. In step 3, you are shown a sample draw of a poker chip from the bag. This chip is replaced back into the bag once it has been observed. You can either maintain the decision choice made in step 2 BEFORE observing the sample draw or change your decision choice AFTER observing the sample draw. In step 4, the bag that was used during this round of play will be revealed. A random draw will be made to determine whether your step-2 action choice or your step-4 action choice will be used to calculate your earnings. Therefore, there is an equal chance of receiving earnings calculated based on the action choice you made BEFORE observing the sample draw in step 2 or of receiving earnings based on the action choice you made AFTER observing the sample draw in step 4. You will be informed of your earnings for the round. Please record these earnings on the tracking sheet provided. Since this is a practice round, the earnings for this round will not actually be paid.

For this practice round, the bags contain the following number of red and blue poker chips:

Bag 1	Bag 2
Red chips 35	Red chips 15
Blue chips 15	Blue chips 35
Total chips 50	Total chips 50

Step 1:

A random draw will determine the bag to be used for this round. The procedure is as follows. The experimenter will show you the contents of both bags to verify the number and colour of the poker chips contained within each bag. The number 1 will be pinned to the inside of bag 1 and the number 2 will be pinned to the inside of bag 2. From the exterior of the bag it will be impossible for you to tell which bag is designated 1 or 2. Both of these bags will be placed in a large cardboard box. For each round, a participant will be selected to come forward and reach into the box and select a bag. You will be unable to identify which bag has been selected. There is an equal chance that the round is being played using either bag 1 or bag 2.

Appendix 2 –Continued: Free message signal Task

Step 2:

<u>BEFORE</u> observing the sample draw, please circle below whether you wish to take action A or action B based on the following potential earnings:

Pick **Action A**:	If the bag chosen by the participant was bag 1 you receive	$2.00
	If the bag chosen by the participant was bag 2 you receive	$0.75
Pick **Action B**:	If the bag chosen by the participant was bag 1 you receive	$0.50
	If the bag chosen by the participant was bag 2 you receive	$1.75

Action Choice <u>BEFORE</u> observing a Sample Draw

Circle either Action A or Action B

Action A Action B

Rip off this sheet and place it on the corner of your desk for the research assistant to collect.

Step 3:

The experimenter will ask one participant to draw one poker chip from the bag, show you the colour of the chip and replace it back into the bag.

After observing the sample draw, please circle below whether you wish to take action A or action B based on the following potential earnings:

Pick **Action A**:	If the bag chosen by participant is bag 1 you receive	$2.00
	If the bag chosen by participant is bag 2 you receive	$0.75
Pick **Action B**:	If the bag chosen by participant is bag 1 you receive	$0.50
	If the bag chosen by participant is bag 2 you receive	$1.75

Action Choice <u>AFTER</u> Observing a Sample Poker chip Draw

Circle either Action A or Action B

Action A Action B

Rip off this sheet and place it at the corner of your desk for the research assistant to collect.

Step 4:

A random draw will take place to determine whether you will be paid based on your initial choice as indicated in Step 2 or your revised choice as indicated in Step 3. Therefore, there is an equal chance of receiving payment for either the action choice you made BEFORE observing the sample draw or the action choice you made AFTER observing the sample draw. The experimenter will reveal the bag used for this round. You will then be informed of your earnings for the round. Please record these earnings on the tracking sheet provided.

Practice Round - A Decision Task with an "OPTION TO PURCHASE" a Message Signal

As this is a practice round the potential earnings highlighted on the next page will not actually be paid. The intention of this round is to allow you to become familiar with the decision task with an "OPTION TO PURCHASE" a message signal.

In all the 'option to purchase' a message signal rounds, there are two bags, each containing a combination of red and blue poker chips. There are 50 poker chips in each bag. However the number of chips that are red and the number of chips that are blue differs between these two bags. You will be informed of the number of red and the number of blue poker chips contained within each bag. In step 1, a random draw will decide which bag will be selected for use during the round. There is an equal chance that we will be playing the round using bag 1 or bag 2. However, you will not know until the end of the round which bag has been randomly selected for play during the round. In Step 2, you are asked to choose one of two actions and told the financial consequences of taking each of these two actions. In step 3, you are shown a sample draw of a poker chip from the bag. This chip is replaced back into the bag once it has been observed. You can either maintain the decision choice made in step 2 BEFORE observing the sample draw or change your decision choice AFTER observing the sample draw. However, in step 4, in order to determine whether this revised decision will be used to calculate your earnings, you must indicate how much you would be willing to pay in order for it to be so used.. In step 5, the experimenter will ask a participant to draw a random price from a random price box which will determine the actual price of using your revised decision rather than your initial decision to calculate your earnings. If your specified willingness to pay is less than the randomly determined price, your initial decision will be used to calculate your earnings. Therefore, your earnings will be based on the action choice you made BEFORE you observed the sample draw. However, if your specified willingness to pay is greater than or equal to the randomly determined price, your revised decision will be used to calculate your earnings. Therefore, your earnings will be based on the action choice made AFTER you observed the sample draw and your earnings for this round will be reduced to include the randomly determined price of using the new information. In step 6, the bag that was used during this round of play will be revealed and you will be informed of your earnings for the round. Please record these earnings on the tracking sheet provided.

Step 1 to 3 is identical to the Free message signal task(subjects were walked through these steps again…condensed for the appendix)

Step 4:
Please indicate on the next page the amount of money you would be willing to pay so that the revised action choice you made AFTER observing the sample draw rather than the initial action choice you made BEFORE observing the sample draw is used in order to calculate your earnings.

Your earnings for this round will be determined as follows:

Once you have indicated your willingness to pay to use your revised action choice, a random draw will determine the actual price that you must pay for the action choice you made <u>AFTER</u> observing the sample draw to be used to calculate your earnings. The procedure is as follows. The experimenter will ask a participant to choose a price from a box that contains many possible prices, some low prices and some high prices. If the random price drawn is less than or equal to your specified willingness to pay, the action choice you made <u>AFTER</u> observing the sample draw will be used to calculate your earnings. You will pay the random price selected. Therefore, the action choice used to calculate your earnings will be the one made <u>AFTER</u> you observed the sample draw. The cost of using your revised decision will be subtracted from your earnings. , And, you will earn:

Pick **Action A**:	If the bag chosen by participant is bag 1 you receive	$2.00-$P
	If the bag chosen by participant is bag 2 you receive	$0.75-$P
Pick **Action B**:	If the bag chosen by participant is bag 1 you receive	$0.50-$P
	If the bag chosen by participant is bag 2 you receive	$1.75-$P

Appendix 2-Continued: OTP a message signal decision task

Otherwise, if the random price drawn is greater than your specified willingness to pay, the action choice used to calculate your earnings will be the one made <u>BEFORE</u> you observed the sample draw (i.e., the action choice made in step 2). And, you will earn:

Pick **Action A**:	If the bag chosen by participant is bag 1 you receive	$2.00
	If the bag chosen by facilitator is bag 2 you receive	$0.75
Pick **Action B**:	If the bag chosen by participant is bag 1 you receive	$0.50
	If the bag chosen by participant is bag 2 you receive	$1.75

Please indicate the amount that you are willing to pay in order to use the revised action choice made AFTER the sample draw rather than the initial action choice made <u>BEFORE</u> the sample draw to calculate your earnings.

$_____

Rip off this sheet and place it on the corner of your desk for the research assistant to collect.

Step 5:
The experimenter will now ask a participant to draw a random price from the box.
If the randomly chosen price $P is less than or equal to your specified willingness to pay, then you are paid based on the action choice you made <u>AFTER</u> you observed the sample draw.

If the randomly chosen price $P is greater than your specified willingness to pay, then you are paid based on the action choice you made <u>BEFORE</u> you observed the sample draw.

Step 6:
The experimenter will reveal the bag used for this round and you will be informed of your earnings. Please record your earnings for the round on the tracking sheet provided.

Appendix 3

Participant #_____

Participant Tracking Sheet

Round #	(F) First Action Choice A or B	Chip Colour r or b	(S) Second Action Choice A or B	Action choice for Payment F or S	Bag revealed 1 or 2	Earnings $
Practice						
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						

Round #	(F) First Action Choice A or B	Chip Colour r or b	(S) Second Action Choice A or B	Your Willingness to Purchase Price $WTP	Random Price drawn $P	Action choice for Payment If $P >WTP=F If $P≤WTP=S	Earnings $
Practice							
13							
14							
15							
16							
17							
18							
19							
20							
21							
22							
23							
24							

Appendix 4

First BEU Action Choice

Without any additional information about the probability of the state being S_1 or S_2, the initial decision to choose action A or B is based on prior probabilities, π_s. Specifically, the risk-neutral BEU will choose action A versus action B when:

$$\pi_1 C(A,S_1) + \pi_2 C(A,S_2) \geq \pi_1 C(B,S_1) + \pi_2 C(B,S_2)$$

Given the parameter values from table 4 and prior to a message signal

For rounds 1-4 & 13-16 the initial BEU action choice will be action A, as the expected payoff from action A is greater than that of action B.

.5($2.00) + .5(.75) =$1.37 >.5(1.75)+.5(.50)=$1.125

For rounds 5-8 & 17-20 the initial BEU action choice will be B

.5($2.00) + .5(.75) =$1.37 >.5(1.75)+.5(.50)=$1.125

For rounds 9-12 & 21-24 the initial BEU action choice will be B

.5($2.00) + .5(.75) =$1.37 >.5(1.00)+.5(.50)=$0.75

Second BEU Action Choice

Bayes theorem states: $\boldsymbol{\pi_{S.M}} = \frac{j_{S.M}}{q_M}$; Where, $\boldsymbol{\pi_{S.M}}$ is the conditional (posterior) probability of state S given the message M; $\mathbf{j_{S.M}}$ is the joint probability of state S and the message M; and q_M is the unconditional probability of receiving message M. Therefore, given message 1(red chip), the BEU players chooses action A if the expected payoff is greater than choosing action B given the posterior probabilities associated with message 1. The expected payoff when choosing action A when message 1(red chip) is received is:

$$EP_{action\ A} = \pi_{1.1}\, C(A,S_1) + \pi_{2.1}\, C(A,S_2) = \frac{q_{1.1}\pi_1}{q_{1.1}\pi_1 + q_{1.2}\pi_2}\, C(A,S_1) + \frac{q_{1.2}\pi_2}{q_{1.1}\pi_1 + q_{1.2}\pi_2}\, C(A,S_2),$$

And the risk-neutral BEU will choose action A if:

$$EP_{action\ A} = \pi_{1.1}\, C(A,S_1) + \pi_{2.1}\, C(A,S_2) > EP_{action\ B} = \pi_{1.1}\, C(B,S_1) + \pi_{2.1}\, C(B,S_2).$$

Therefore, given the parameter values in Table 4 for **rounds 1-4 & 13-16**, given red chip draw, **the RN BEU picks action A**,

$$EP_{action\ A} = \frac{.70(.5)}{.70(.5) + .30(.5)}(\$2.00) + \frac{.30(.5)}{.70(.5) + .30(.5)}(\$0.75) = \$1.625$$

$$EP_{action\ B} = \frac{.70(.5)}{.70(.5) + .30(.5)}(\$0.50) + \frac{.30(.5)}{.70(.5) + .30(.5)}(\$1.75) = \$0.875$$

For **rounds 1-4&13-16**,

if a blue chip is drawn, the risk-neutral BEU will choose action B, given that, $EP_{action\ B} > EP_{action\ A}$:

$$EP_{action\ B} = \frac{.70(.5)}{.70(.5) + .30(.5)}(\$1.75) + \frac{.30(.5)}{.70(.5) + .30(.5)}(\$0.50) = \$1.375$$

$$EP_{action\ A} = \frac{.70(.5)}{.70(.5) + .30(.5)}(\$0.75) + \frac{.30(.5)}{.70(.5) + .30(.5)}(\$2.00) = \$1.125$$

For **rounds 5-8 & 17-20 and a red chip draw RN BEU picks Action B**

$$EP_{action\ A} = \frac{.76(.5)}{.76(.5) + .24(.5)}(\$0.50) + \frac{.24(.5)}{.76(.5) + .24(.5)}(\$1.75) = \$0.80$$

$$EP_{action\ B} = \frac{.76(.5)}{.76(.5) + .24(.5)}(\$2.00) + \frac{.24(.5)}{.76(.5) + .24(.5)}(\$0.75) = \$1.70$$

For **rounds 5-8 & 17-20 and a blue chip draw RN BEU picks Action A**

$$EP_{action\ A} = \frac{.76(.5)}{.76(.5) + .24(.5)}(\$1.75) + \frac{.24(.5)}{.76(.5) + .24(.5)}(\$0.50) = \$1.45$$

$$EP_{action\ B} = \frac{.76(.5)}{.76(.5) + .24(.5)}(\$0.75) + \frac{.24(.5)}{.76(.5) + .24(.5)}(\$2.00) = \$1.05$$

For **rounds 9-12 & 21-24 and a red chip draws RN BEU picks Action B**

$$EP_{action\ A} = \frac{.60(.5)}{.60(.5) + .40(.5)}(\$0.50) + \frac{.40(.5)}{.60(.5) + .40(.5)}(\$1.00) = \$0.70$$

$$EP_{action\ B} = \frac{.60(.5)}{.60(.5) + .40(.5)}(\$2.00) + \frac{.40(.5)}{.60(.5) + .40(.5)}(\$0.75) = \$1.50$$

For **rounds 9-12 & 21-24 and a blue chip draws RN BEU picks Action B**

$$EP_{action\ A} = \frac{.60(.5)}{.60(.5) + .40(.5)}(\$1.00) + \frac{.40(.5)}{.60(.5) + .40(.5)}(\$0.50) = \$0.80$$

$$EP_{action\ B} = \frac{.60(.5)}{.60(.5) + .40(.5)}(\$0.75) + \frac{.40(.5)}{.60(.5) + .40(.5)}(\$2.00) = \$1.25$$

Appendix 5
Post-experiment Questionnaire

Participant Code_____.

We would appreciate it if you could provide the following information. Your responses will remain confidential, and you may decline to answer any question if you wish.

1. What is you gender? (Please mark one circle with an X) Male O Female O

2. What is your age? _____.

3. Where were you born?
(Specify one response only, according to present boundaries)
Born in Canada (Specify province or territory): _____
Born outside Canada (Specify country): _____

4. What language do you speak most often at home? (Please mark one circle with an X)
O English O French
O Other – Specify: _____, _____, _____

5. What program are you in at the University? _____

6. What year are you currently in? (1st, 2nd etc..)_____

7. At what grade levels did you take English courses at High School? (please mark with an X)
O Grade 9 O Grade 10
O Grade 11 O Grade 12

8. Which Math courses did you take at High School? If you attended High School outside of Ontario, please mark from the list below the courses that are approximately equivalent to those you studied. (please mark with an X)

O Grade 9 O Grade 10
O Grade 11 O Grade 12- Relations & Functions
O Grade 12-Calculus O Grade 12- Data Management

9. Have you taken any Math or Statistics courses in university (list)? _____
_____

10. Have you participated in an Economics or Psychology experiment before? (check one)
__ Yes __ No

For the questions 11-22 (Honey & Mumford, 1986),

Please circle the degree to which you agree or disagree with the following statements

	Agree Strongly						Disagree Strongly
11. I tend to solve problems using a step-by- step approach.	1	2	3	4	5	6	7
12. I take pride in doing a thorough job.	1	2	3	4	5	6	7
13. What matters most is whether something works in practice.	1	2	3	4	5	6	7
14. I like to relate my actions to a general principle.	1	2	3	4	5	6	7
15. I find it difficult to produce ideas on impulse.	1	2	3	4	5	6	7
16. I prefer to have as many sources of information as possible- the more data to think over the better.	1	2	3	4	5	6	7
17. I am keen to reach answers via a logical approach.	1	2	3	4	5	6	7
18. In discussions I enjoy watching the maneuverings of other people.	1	2	3	4	5	6	7
19. I get along on best with logical, analytical people and less well with spontaneous people.	1	2	3	4	5	6	7
20. I accept and stick to laid down procedures so long as I regard them as an efficient way of getting the job done.	1	2	3	4	5	6	7
21. I am keen on exploring the basic assumptions, principles and theories underpinning things and events.	1	2	3	4	5	6	7
22. In discussions with people I often find I am the most dispassionate and objective.	1	2	3	4	5	6	7

23. How would you describe your strategy for making choices in this study?

_____

_____

_____

Please write any comments you may have:

_____

_____

Appendix 6

Eckel-Grossman, Test for Risk Aversion (2002).

Participant Code_____.

24. Each of these ten circles represents a lottery with two possible prizes. The lotteries are played by flipping a coin. You will earn the amount on the left side of a circle if the flip is a HEAD, while you will the amount on the right side of a circle if the flip is a TAIL. **Please choose the lottery you most prefer by placing an X over it. You will play the lottery you choose for cash.**

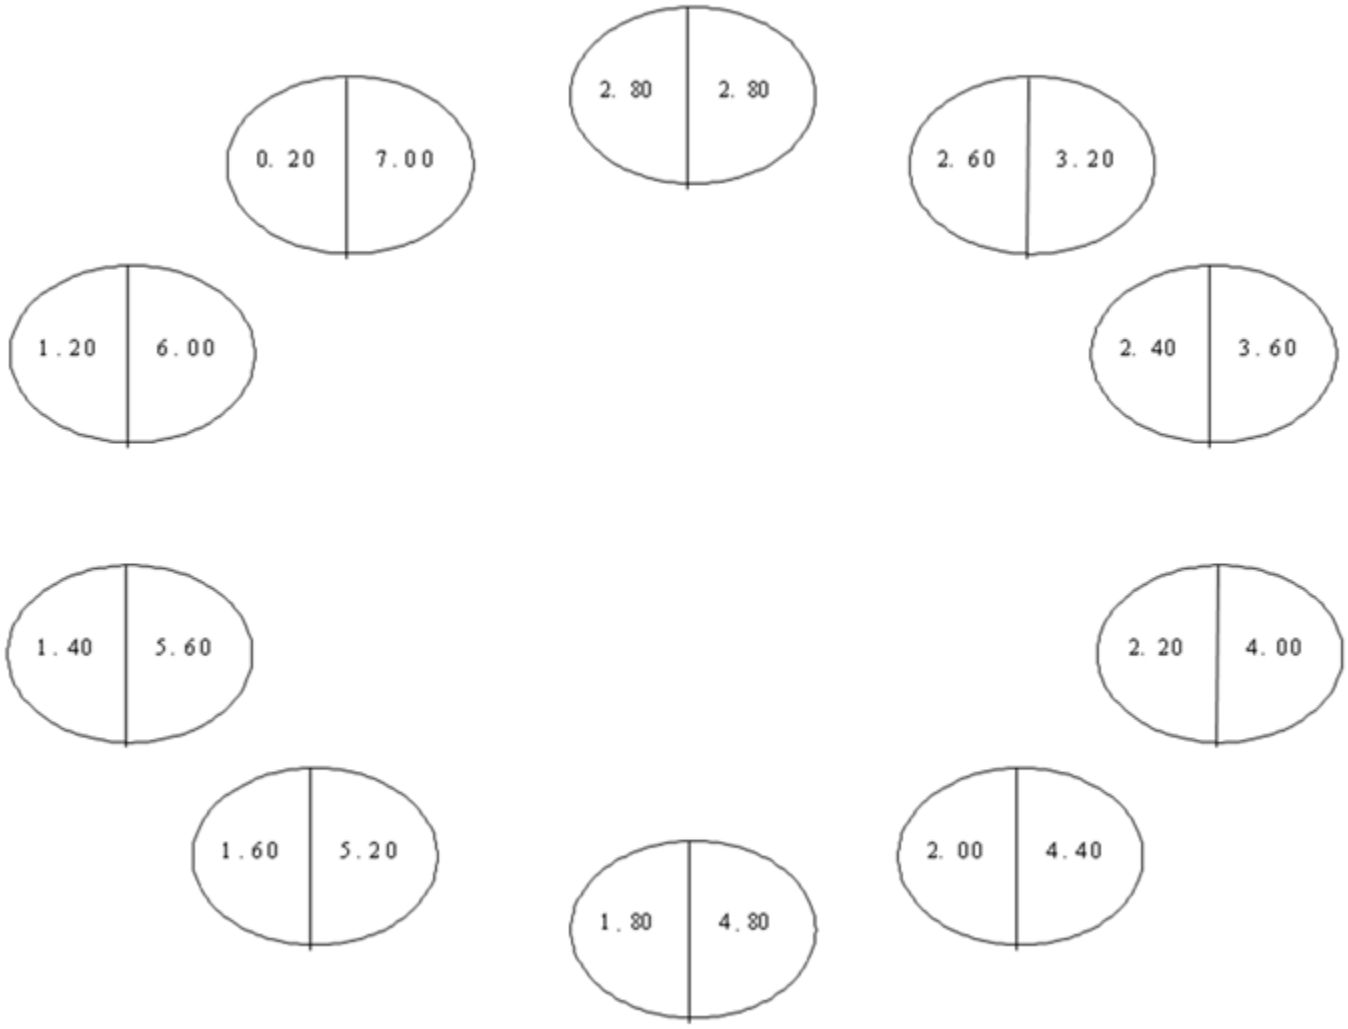

Appendix 7

	BEU First Choice Deviations				BEU Second Choice Deviations				RL Second Choice Deviations			
ORDER	Obs.	**All**	**Un-inform**	**Inform**	**Obs.**	**All**	**Un-inform**	**Inform**	Obs.	**All**	**Un-inform**	**Inform**
Free 1-12	720	15.5%	14.7%	16.4%	720	14.3%	14.2%	14.4%	386	42.2%	34.9%	49.7%
		(0.363)	(0.354)	(0.371)		(0.350)	(.349)	(0.352)		(0.495)	(0.478)	(0.501)
OTP 13-24	720	9.6%	10.3%	8.9%	720	17.2%	19.2%	15.3%	378	40.7%	37.2%	44.0%
		(0.299)	(0.313)	(0.285)		(0.378)	(0.394)	(0.360)		(0.492)	(0.484)	(0.498)
Free/OTP	1440	**12.6%**	**12.5%**	**12.7%**	1440	**15.8%**	**16.7%**	**14.9%**	764	**41.5%**	**36.1%**	**46.9%**
OTP 1-12	732	17.9%	17.5%	18.3%	732	18.4%	19.4%	17.4%	374	35.0%	36.1%	33.9%
		(0.384)	(0.380)	(0.387)		(0.387)	(0.396)	(0.380)		(0.478)	(0.482)	(0.475)
Free 13-24	732	9.6%	10.8%	8.3%	732	12.6%	11.7%	13.4%	374	39.4%	39.2%	39.2%
		(0.294)	(0.311)	(0.277)		(0.332)	(0.321)	(.342)		(0.489)	(0.490)	(0.490)
OTP/Free	1464	**13.8%**	**14.2%**	**13.3%**	1464	**15.5%**	**15.6%**	**15.4%**	748	**37.2%**	**37.7%**	**36.6%**
Control 1-12	708	19.6%	18.4%	20.8%	708	20.2%	19.5%	20.8%	372	40.9%	39.1%	42.8%
		(0.397)	(0.387)	(0.407)		(0.402)	(0.397)	(0.407)		(0.492)	(0.489)	(0.496)
Control 13-24	708	10.4%	8.3%	12.5%	708	13.8%	16.4%	11.1%	389	45.0%	43.0%	46.7%
		(.306)	(0.277)	(0.331)		(0.345)	(0.371)	(0.315)		(0.498)	(0.496)	(0.500)
Control	1416	**15.0%**	**13.4%**	**16.7%**	1416	**17.0%**	**18.0%**	**16.0%**	761	**43.0%**	**41.1%**	**44.8%**
Total	4320	**13.8%**	**13.3%**	**14.2%**	4320	**16.0%**	**16.6%**	**15.4%**	2265	**40.6%**	**38.3%**	**42.8%**
		(0.345)	(.341)	(.349)		(0.367)	(0.374)	(0.361)		(0.491)	(0.428)	(0.495)

Aggregate Summary Statistics By Round Type

	All			Rounds 1-4			Rounds 5-8			Rounds 9-12		
Variable	**Obs.**	**Mean**	**Std. Dev.**	**Obs.**	**Mean**	**Std. Dev.**	**Obs.**	**Mean**	**Std. Dev.**	**Obs.**	**Mean**	**Std. Dev.**
Informed	4320	0.506	0.500	1440	0.506	0.500	1440	0.506	0.500	1440	0.506	0.500
Paid	4320	0.336	0.472	1440	0.336	0.473	1440	0.336	0.473	1440	0.336	0.473
Free 2nd	4320	0.503	0.500	1440	0.503	0.500	1440	0.503	0.500	1440	0.503	0.500
Female	4320	0.467	0.499	1440	0.467	0.499	1440	0.467	0.499	1440	0.467	0.499
Eng. 2nd	4320	0.200	0.400	1440	0.200	0.400	1440	0.200	0.400	1440	0.200	0.400
No HS Math	4320	0.028	0.164	1440	0.028	0.164	1440	0.028	0.164	1440	0.028	0.164
No U Math	4320	0.150	0.357	1440	0.150	0.357	1440	0.150	0.357	1440	0.150	0.357
Reinf Survey	4320	0.683	0.465	1440	0.683	0.465	1440	0.683	0.465	1440	0.683	0.465
Risk Aversion	4320	7.383	2.772	1440	7.383	2.772	1440	7.383	2.772	1440	7.383	2.772
age	4320	19.972	1.733	1440	19.972	1.734	1440	19.972	1.734	1440	19.972	1.734
year	4320	2.316	1.213	1440	2.317	1.214	1440	2.317	1.214	1440	2.317	1.214
Blue Draw	4320	0.447	0.497	1440	0.396	0.489	1440	0.441	0.497	1440	0.504	0.500
Degree of Informativeness	4320	0.373	0.132	1440	0.400	0.000	1440	0.520	0.000	1440	0.200	0.000
WTP	1452	0.404	3.887	484	0.546	4.723	484	0.550	4.768	484	0.117	0.485
Bayes WTP	1452	0.103	0.215	484	0.104	0.190	484	0.177	0.269	484	0.027	0.140
Diff WTP & BWTP	1452	0.302	3.878	484	0.443	4.723	484	0.373	4.753	484	0.090	0.470
Yes Bayes WTP	1452	0.352	0.478	484	0.403	0.491	484	0.442	0.497	484	0.211	0.408
WTP Restricted at $2	1452	0.171	0.384	484	0.197	0.423	484	0.219	0.407	484	0.099	0.304
Diff Restricted WTP & BWTP	1452	0.069	0.375	484	0.092	0.430	484	0.042	0.380	484	0.072	0.306
Random Price draw	1452	0.235	0.147	484	0.233	0.131	484	0.245	0.152	484	0.228	0.157
Exp Bayes PO	4320	1.435	0.212	1440	1.458	0.209	1440	1.473	0.263	1440	1.374	0.125
Paid on 2nd Choice	2868	0.446	0.497	956	0.448	0.498	956	0.447	0.497	956	0.442	0.497
Bag 2 Revealed	4320	0.523	0.500	1440	0.502	0.500	1440	0.552	0.518	1440	0.526	0.500
BR not Bayes Prediction	4320	0.363	0.481	1440	0.356	0.479	1440	0.257	0.437	1440	0.474	0.500
PO	4320	1.396	0.621	1440	1.415	0.611	1440	1.445	0.615	1440	1.329	0.632
Bayes PO	4320	1.457	0.608	1440	1.440	0.606	1440	1.525	0.586	1440	1.407	0.624
Diff PO	4320	-0.061	0.479	1440	-0.026	0.511	1440	-0.079	0.537	1440	-0.078	0.373
% PO maximized	4320	0.534	0.499	1440	0.536	0.499	1440	0.556	0.497	1440	0.510	0.500
% PO Bayes Maximized	4320	0.557	0.497	1440	0.541	0.498	1440	0.603	0.489	1440	0.526	0.499
% lose	4320	0.415	0.493	1440	0.410	0.492	1440	0.367	0.482	1440	0.467	0.499

Appendix 8: 1st Choice BEU Inconsistency Regression

Variables	(1) BEUD1 1st Choice Deviation xtlogit, Re	(2) BEUD1 1st Choice Deviation xtlogit, Fe	(3) BEUD1 1st Choice Deviation GLS, Re	(4) BEUD1 1st Choice Deviation GLS, Fe
Experience	-7970***	-.7873***	-.0754***	-.0753***
	(.1039)	(.1022)	(.0095)	(.0095)
Free followed by OTP	-.0101	(omitted)	-.0121	(omitted)
	(.2768)		(.0237)	
OTP	.2038*	.2467**	.0191*	.0202*
	(.1260)	(.1299)	(.0116)	(.0119)
Informed	.1942	(omitted)	.0124	(omitted)
	(.2652)		(.0224)	
Difference in ex-ante Payoffs	-4.665***	-4.585***	-.4090***	-.4085
	(.6159)	(.6096)	(.0527)	(.0528)
Bag Reveal in Prior round	-.0699	-.0760	-.0044	-.0051
not equal to BEU prediction	(.1060)	(.1082)	(.0102)	(.0102)
Subject Paid on second action	.3200***	.3162***	.0334***	.0330***
in Prior Round	(.1060)	(.1060)	(.0103)	(.0104)
Female	.3629	(omitted)	.0319	(omitted)
	(.2806)		(.0240)	
English 2nd	.8994***	(omitted)	.0828***	(omitted)
	(.3200)		(.0282)	
Reinforcement Learner Survey	-.0678	(omitted)	-.0110	(omitted)
	(.2956)		(.0248)	
Risk Aversion	-.0421	(omitted)	-.0027	(omitted)
	(.0479)		(.0042)	
Age	-.1682***	(omitted)	-.0142***	(omitted)
	(.0808)		(.0065)	
Econ Math Student	-.0510	(omitted)	.0015	(omitted)
	(.3103)		(.0263)	
Obs.	4320	2928	4320	4320
log likelihood/R-squared	-1451.84	-1041.51	.1060	.0339

Appendix 9: Second choice BEU inconsistencies (RE and FE) and Marginal Effects

Variables	(5a) BEUD2 2nd Choice Deviation xtlogit, Re	(5b) BEUD2 2nd Choice Deviation Marginal Effects	(6) BEUD2 2nd Choice Deviation xtlogit,Fe	(7) BEUD2 2nd Choice Deviation GLS,Re	(8) BEUD2 2nd Choice Deviation GLS,Fe
Experience	**-0.257*** **	**-.0285*** **	**-0.261*** **	**-0.025*** **	**-0.025*** **
	0.093	0.010	0.093	0.010	0.010
Free followed by OTP	**-0.167**	**-0.018**	(omitted)	**-0.020**	(omitted)
	0.153	0.016		0.017	
OTP	**0.577*** **	**.0341*** **	**0.682*** **	**0.056*** **	**0.067*** **
	0.151	0.013	0.1636	0.017	0.018
Informed	**0.147**	**-0.008**	**(omitted)**	**0.007**	**(omitted)**
	0.162	0.016		0.017	
Paid*Informed	**-0.5739*** **		**-0.655*** **	**-0.055*** **	**-0.063*** **
	0.210		0.227	0.024	0.025
Difference in Payoffs Good vs. Bad	**-3.570*** **	**-0.397*** **	**-3.665*** **	**-0.360*** **	**-0.369*** **
state conditional on chip draw	0.296	0.035	0.294	0.033	0.033
Shift Required to follow BEU	**1.114*** **	**0.133*** **	**1.049*** **	**0.126*** **	**0.117*** **
	0.117	0.015	0.115	0.014	0.014
Degree informative	**-3.025*** **	**-0.336*** **	**-2.896*** **	**-0.235*** **	**-0.224*** **
	0.425	0.048	0.422	0.045	0.045
Bag Reveal in Prior round	**0.218*** **	**0.025*** **	**0.213*** **	**0.025**	**0.025*** **
not equal to BEU prediction	0.098	0.011	0.098	0.011	0.011
Subject Paid on second action	**-0.078**	**-0.008**	**-0.037**	**-0.018**	**-0.016**
in Prior Round	0.099	0.011	0.101	0.011	0.012
Female	**0.167**	**0.019**	**(omitted)**	**0.014**	**(omitted)**
	0.153	0.017		0.017	
English 2nd	**0.190**	**0.022**	**(omitted)**	**0.014**	**(omitted)**
	0.177	**0.021**		0.020	
Reinforcement Learner Survey	**0.112**	**0.012**	**(omitted)**	**0.014**	**(omitted)**
	0.161	**0.017**		0.017	
Risk Aversion	**-0.021**	**-0.002**	**(omitted)**	**-0.003**	**(omitted)**
	0.026	0.003		0.003	
Age	**-0.011**	**-0.001**	**(omitted)**	**-0.001**	**(omitted)**
	0.043	**0.005**		0.005	
Econ Math Student	**-0.258**	**-0.029**	**(omitted)**	**-0.027**	**(omitted)**
	0.172	**0.019**		0.018	
Obs.	4320	4320	3984	4320	4320
log likelihood/R-squared	-1631.05		-1224.486	0.099	0.094

Appendix 10: Logit regression Free Vs. OTP and OTP uninformed Vs. Informed

xtlogit, Re	(9)	(10)	(11)	(12)	(13)
	BEUD2	BEUD2	BEUD2	BEUD2	BEUD2
	2nd Choice Deviation	2nd Choice Deviation	2nd Choice Deviation	2nd Choice Deviation	2nd Choice Deviation
Variable	All	Free	OTP	OTP Uninformed	OTP Informed
Experience	-0.221***	-0.315***	-0.361	-0.264	-0.466
	0.091	0.143	0.235	0.358	0.346
Free followed by OTP	-0.125	-0.277	(omitted)		
	0.149	0.195			
Informed	-0.07	0.134	-0.534***		
	0.142	0.159	0.235		
Difference in Payoffs Good vs. Bad	-4.418***	-4.431***	-4.650***	-4.757***	-4.435***
state conditional on chip draw	0.271	0.338	0.486	0.630	0.782
Degree informative	-0.934***	-0.708*	-1.5226***	0.670	-4.143***
	0.345	0.432	0.602	0.841	0.921
Bag Reveal in Prior round	0.162*	0.052	0.441***	0.594***	0.146
not equal to BEU prediction	0.095	0.121	0.167	0.227	0.257
Subject Paid on second action	-0.119	-0.074	0.034	0.011	0.180
in Prior Round	0.095	0.117	0.194	0.295	0.267
Female	0.185	0.175	0.255	0.173	0.473
	0.151	0.167	0.255	0.387	0.363
English 2nd	0.189	0.246	0.064	0.029	-0.091
	0.175	0.191	0.300	0.392	0.504
Reinforcement Learner Survey	0.16	0.045	0.366	0.311	0.407
	0.159	0.176	0.272	0.384	0.397
Risk Aversion	-0.018	-0.020	-0.014	0.006	-0.018
	0.026	0.028	0.045	0.067	0.063
Age	-0.0137	-0.015	0.001	0.085	-0.041
	0.042	0.047	0.068	0.099	0.097
Econ Math Student	-0.269	-0.308*	-0.197	0.337	-0.657*
	0.17	0.192	0.272	0.401	0.382
Obs.	4320	2868	1452	720	732
log likelihood/R-squared	-1687.076	-1084.140	-593.102	-306.188	-276.043

Appendix 11- 2nd Choice Reinforcement Inconsistency Rate Logit Fixed and Random Effects

Variables	(14a) RLD2 2nd Choice Deviation xtlogit, Re	(14b) RLD2 2nd Choice Deviation Marginal Effects	(15) RLD2 2nd Choice Deviation xtlogit,Fe	(16) RLD2 2nd Choice Deviation GLS,Re	(17) RLD2 2nd Choice Deviation GLS,Fe
Experience	0.185***	0.042***	0.199***	0.041***	0.043***
	0.089	0.020	0.089	0.020	0.020
Free followed by OTP	0.132	0.030	(omitted)	0.029	(omitted)
	0.098	0.023		0.022	
OTP	-0.088	-0.043	-0.010	-0.020	-0.003
	0.141	0.023	0.162	0.032	0.036
Informed	0.272***	0.048***	(omitted)	0.061***	(omitted)
	0.108	0.021		0.025	
Paid*Informed	-0.194		-0.181	-0.045	-0.037
	0.193		0.223	0.044	0.051
Difference in Payoffs Good vs. Bad	-1.207***	-0.277***	-1.222***	-0.276***	-0.288***
state conditional on chip draw	0.271	0.061	0.285	0.063	0.066
Shift Required to follow BEU	0.144	0.033	0.022	0.032	0.002
	0.114	0.026	0.119	0.026	0.028
Degree informative	-1.901***	-0.436	-1.844***	-0.429***	-0.42***
	0.380	0.085	0.383	0.087	0.088
Bag Reveal in Prior round	-0.653***	-0.146***	-0.815***	-0.147***	-0.188***
not equal to BEU prediction	0.103	0.022	0.107	0.023	0.024
Subject Paid on second action	-0.153	-0.035	-0.186	-0.036	-0.038
in Prior Round	0.095	0.022	0.100	0.022	0.023
Female	0.104	0.024	(omitted)	0.025	(omitted)
	0.096	0.022		0.022	
English 2nd	0.000	0.000	(omitted)	0.000	(omitted)
	0.113	0.026		0.026	
Reinforcement Learner Survey	-0.284***	-0.066***	(omitted)	-0.065	(omitted)
	0.100	0.023		0.023	
Risk Aversion	-0.007	-0.002	(omitted)	-0.002	(omitted)
	0.017	0.004		0.004	
Age	-0.023	-0.005	(omitted)	-0.005	(omitted)
	0.026	0.006		0.006	
Econ Math Student	0.141	0.032	(omitted)	0.032	(omitted)
	0.1048	0.024		0.0242	
Obs.	2265	2265	2265	2265	2265
log likelihood/R-squared	-1472.553		-1124.277	0.0486	0.0382

Appendix 12: Logit Regressions by Behaviour Type

Obs	**35**		**28**		**35**		**18**	
% of sample	**19.4%**		**15.6%**		**19.4%**		**10.0%**	
	BEU		**Status Quo**		**Overweight**		**Non BEU**	
Variables	**Logit**	**OR**	**Logit**	**OR**	**Logit**	**OR**	**Logit**	**OR**
Free followed by OTP	**-0.236**	**0.790**	**0.568**	**1.765**	**0.090**	**1.094**	**-0.654**	**0.520**
	0.488	0.393	0.535	0.929	0.510	0.556	0.651	0.347
OTP followed by Free	**-0.238**	**0.788**	**-1.244***	**0.288***	**0.440**	**1.553**	**-0.169**	**0.844**
	0.545	0.394	0.643	0.181	0.477	0.748	0.603	0.509
Informed	**-0.485**	**0.615**	**-1.012****	**0.363****	**-0.128**	**0.880**	**0.014**	**1.014**
	0.412	0.255	0.512	0.181	0.399	0.350	0.541	0.535
Female	**-0.535**	**0.586**	**0.198**	**1.219**	**0.879****	**2.408****	**-0.349**	**0.705**
	0.492	0.267	0.549	0.640	0.447	1.046	0.545	0.398
English 2nd	**-1.138***	**0.320***	**0.819**	**2.268**	**-0.232**	**0.793**	**0.988***	**2.685***
	0.674	0.213	0.519	1.264	0.522	0.399	0.574	1.510
Reinforcement Learner Survey	**-0.398**	**0.672**	**1.524*****	**4.591*****	**-0.295**	**0.745**	**-0.270**	**0.763**
	0.448	0.285	0.733	2.966	0.443	0.332	0.526	0.442
Risk Aversion	**0.021**	**1.021**	**0.195****	**1.215****	**-0.116***	**0.891***	**-0.126**	**0.882**
	0.076	0.080	0.103	0.121	0.067	0.062	0.094	0.080
Age	**0.120**	**1.127**	**0.158**	**1.171**	**0.012**	**1.012**	**-0.192**	**0.825**
	0.122	0.129	0.134	0.159	0.109	0.115	0.177	0.140
Econ Math Student	**0.922****	**2.513****	**-1.745*****	**0.174*****	**0.037**	**1.037**	**0.500**	**1.648**
	0.451	0.218	0.842	0.148	0.481	0.498	0.532	0.9336
Obs.	180		180		180		180	
log likelihood/R-squared	-79.419		-61.184		-83.427		-53.808	

******p-value ≤.01 <**p-value ≤.05 < *p-value < .10***

Robust standard errors

Chapter 3

The Value of Information: An Experiment Continued

3.0 Introduction

This chapter deals with the subjects' decisions concerning whether and how they choose to improve upon their current knowledge base before taking a terminal action. Although knowledge can be acquired through the experience of completing a task repeatedly, this study is most concerned with the knowledge acquired through direct purchase. Specifically in this experiment, I am interested in understanding how much a subject is willing to pay in order to use the additional knowledge gathered from an imperfect message service when making a final decision.

A disadvantage of past experimental designs is that the willingness to pay for additional information is measured in terms of a fixed specified cost per observation. How much a subject is willing to pay for information is determined by the number of samples purchased. As the experimenter sets the purchase price, the precise value that a subject places on the service is not known i.e., only a lower bound can be established (See chapter 1, Fried & Peterson (1965), Green & Swets (1966), Edwards (1968), Wallensten (1968), Pritz (1968), Hershman & Levine(1970)). While we can identify the risk neutral Bayesian benchmark for the value of new information, we cannot use this option-to-buy mechanism to assess exactly how much subjects value the new information gathered from the message service. Furthermore, under this payment scheme we are also unable to observe the degree to which the amount paid for the information influences the subject's decisions.

In the present experimental design, the willingness-to-pay mechanism is modified in order to provide a shortened method for determining how subjects value additional information. This method allows us to determine whether a subject is willing to purchase a message and at the same time assess the value placed on this message when making a terminal decision without requiring the subject to identify how many messages at a set price they would be willing to purchase. Furthermore, we are able to access how the subjects' willingness to pay for information influences their decision choices relative to optimal decision theory.

This chapter explores the following key questions: 1) Is the price a subject is willing to pay (WTP) for the right to use an observation to update his/her prior decision choice equal to the Bayesian willingness to pay amount (BWTP)?; 2) Are subjects more likely to select the BWTP amount when the informational value of the message service increases?; 3) Are subjects' action choices more consistent with optimal decision theory when their willingness to pay to use the information is consistent with the Bayesian amount?; 4) Is there a difference in how the information is used when subjects pay for its use versus using it for free? In this study we are unable to observe a direct comparison between the subject's value of the information when it is free versus when it is costly. However, the consequent actions by the subject conditional on the message received gives insight into which message service (free vs. costly) leads to more optimal behavior; Finally, 5) does providing subjects with the posterior probabilities conditional on the message received result in a greater number of BEU WTP decisions?

Section 1 describes the experimental design with a detailed focus on the 'willingness to pay' mechanism. Section 2 provides the theoretical framework used to benchmark subject behavior. Section 3 presents the theory underpinning the 5 main hypotheses tested in this study. In Section 4 the results and interpretations are presented and in Section 5 conclusions and future research opportunities are discussed.

3.1 Experimental Design

Recall from chapter 2, the experiment tracked the decision choices of 180 students who were divided into 12 classroom sessions consisting of approximately 15 students each. Each subject participated in 24 rounds of an individual task consisting of two (2) binary-choice decisions per round; the first binary choice decision occurred before and the second binary choice decision occurred after observing an imperfect information signal. For a subset of the groups and rounds, subjects had a third decision choice that required them to specify their willingness to pay in order to use this additional information, where the WTP amount specified determined whether a subject's first (before observing the message) or second (after observing the message) action would represent their terminal choice. In all decision tasks, subjects were told that their decision choices should be motivated by the desire to earn as much money as possible (see Chapter 2; pp. 5-10). Of most interest to the current study are the observations of the subjects who performed

the Option to Purchase (OTP) task. During this task, after observing the sample draw (imperfect message) and selecting an action conditional on this draw , subjects indicated how much they would be willing to pay in order for their second action choice to be recognized for the determination of payment as opposed to the decision made prior to observing the sample draw. Once the willingness to pay (WTP) price was specified, the experimenter asked one of the participants to draw a random price from a box that contained 51 tokens each specifying a unique price point, ranging from $0.00 to $0.50, which determined the actual price required to have their second choice replace their first. The subjects were unaware of the range of prices contained within the box. If the subject's specified WTP was less than the randomly determined price, the initial action choice (first choice) was used to calculate her earnings and no price was deducted from the payoff associated with this decision. However, if her specified WTP was greater than or equal to the randomly determined price, then the revised decision (second choice) was used to calculate the earnings and the random price drawn was deducted from the total earnings for the round.

This WTP elicitation method utilizes key aspects of a Vickrey second price auction (Vickrey, 1961). That is, it is designed to be incentive compatible, thus ensuring that subjects reveal their truthful valuation of the information signal. Essentially subjects are participating in a sealed bid auction against an opponent - in this case a random price draw - where the highest bidder wins, but the price paid is the second highest bid. Therefore, if a subject specifies a WTP bid that is greater than or equal to the random price drawn, the subject wins the bid and gets to use the information acquired to assist in determining her final decision, but only pays the random price (the second price). With this type of mechanism when a subject selects her independent private value, she maximizes her expected utility by revealing her true valuation of the message received (see Appendix 1 for proof of dominance of truthful WTP bidding).

The order of the decisions made during the OTP message task is as follows: 1.) subjects take an action prior to observing a colour chip draw, 2) subjects observe a colour chip draw and take a second action and then, 3) specify a WTP amount to have their second action choice be recognized for payment versus the first choice made prior to observing the colour chip draw. This design is analogous to the following scenario: I book a trip to Florida, I observe that a

hurricane is potentially pending; I book a trip to California but am informed that I must pay more to change my reservation. Based on my confidence in the weather forecast, how much would I be willing to pay to make this change? How confident am I in the message received?

The rationale for this design is threefold: 1) it is easier to execute, i.e., all participants in a session observe the chip versus a design where only the subset of permitted participants can observe; 2) The Bayesian WTP calculation is simplified. It is easier for subjects to estimate the BWTP as the added complication of having to estimate the unconditional probabilities of receiving each one of two possible messages is removed from the calculation; 3) More observations of subjects' second action choices are collected. Whether a subject specifies a WTP amount that gives them the right to implement their second decision or not, we still collect information regarding their second action choice conditional on the message received.[57]

Table 1 outlines the details of the 4 treatment groups relevant to this study. Treatments 1, 2, 3 and 4 participated in 12 rounds of a decision task containing a FREE message (FREE) and 12 rounds of a decision task containing the OPTION TO PURCHASE (OTP) message. Treatments 1 & 2 received the FREE message task first and the OTP message task second. These decision tasks are reversed for treatments 3 & 4. Treatments 1 & 3 are designated as un-informed. This group is provided with all the parameters required to compute Bayes law; however, they are not given the posterior probabilities. Treatments 2 & 4 are designated as informed and subjects are provided with the posterior probabilities described in terms of chances out of 100 that the colour chip drawn is either from bag 1 or bag 2.[58]

[57] This is referred to as the Strategy method and is credited to Selten (1967). This method allows us to collect participant data for all possible information sets (whether or not a subject correctly values the info such that their second choice is implemented, we are able to determine what their 2nd action choice would have been). This method is in contrast to the Standard Direct Response method, where we would only know the participant`s 2nd action choice if they specified the correct price to view the information.
[58] I.e., there are 70 chances out of 100 that the chip drawn came from bag 1 and therefore 30 chances out of 100 that the chip came from bag 2. This description of the posterior probabilities avoids any confusion associated with the term 'probability'.

Table 1 -Treatment Group Specification WTP Decisions

Treatment	1	2	3	4
No. Subjects (121)	30	30	30	31
Order : Round 1-12 : Round 12-24	Free OTP	Free OTP	OTP Free	OTP Free
Bayes Law	Uninformed[59]	Informed	Uninformed	Informed

3.2 Valuing Information using Bayesian Expected Utility Theory

In this experiment design there are two possible states, represented by S_j, j ϵ {1,2}, where S_1 indicates bag 1 and S_2 indicates bag 2. A risk neutral BEU participant takes an initial action given the unconditional (prior) probability of either state with the objective of maximizing her expected earnings. Let the unconditional probability (initial belief) of playing in state j be, $prob\ (S_j)$, where, $\sum_j prob\ (S_j) = 1$. Let C($\boldsymbol{a}$,S_j) be the payoff if action $\boldsymbol{a}$ is chosen conditional on the state (S_j), where $\boldsymbol{a}$ ϵ {A,B}. Without any additional information about the probability of the state being the bag with predominately red chips or the bag with the predominately blue chips, the initial decision to choose action A or B is based on the prior probabilities of being in either state, $prob\ (S_j)$, and the state contingent payoffs associated with each action, C($\boldsymbol{a}$,S_j). Specifically, the risk-neutral BEU will choose action A versus action B when:

$$EP_{action\ A} = prob\ (S_1)C(A,S_1) + prob\ (S_2)C(A,S_2) \geq EP_{action\ B} = prob\ (S_1)C(B,S_1) + prob\ (S_2)C(B,S_2)$$

Next, the risk neutral BEU maximizer is provided with one of two possible messages in the form of a colour chip drawn from the randomly selected bag. Let the two possible messages be M_k, k ϵ {1,2}, where $M_1$is message 1(indicating a red chip message) and M_2 is message 2 (indicating a blue chip message). The participant is then required to propose a second action choice conditional on the message received. To do this the BEU maximizer will first, update her prior probabilities of being in either state to a new set of probabilities (posterior) using Bayes theorem. Second, she will combine these updated probabilities to determine the expected payoff from taking either action and then choose the action with the highest expected payoffs.

[59] Uninformed: subjects were given enough information to calculate Bayes posterior probabilities on their own.
Informed: subjects provided with the Bayes Law calculation (posterior probabilities).

Bayes theorem states that the posterior probability that a risk-neutral BEU maximizer should attach to the state after receiving a message, $prob\ (S_j|M_k)$, is:

$Prob\ (S_j|M_k) \equiv \frac{(prob\ S_j)(prob(M_k|S_j)}{prob(M_k|S_j)(prob\ S_j)+prob(M_k|S_{\neq j})(prob\ S_{\neq j})}$; j=1,2; j≠1,2; k=1,2; (Eqn. 1)

Where the $prob(M_k|S_j)$represents the likelihood of the message (M_k) conditional on state, S_j .

Note that regardless of the message received, one of two states must persist. Therefore,

$prob\ (S_j|M_k) + prob\ (S_{\neq j}|M_k) = 1$ (Eqn. 2)

Using Bayes theorem from Eqn. 1, the probability that the bag selected is bag 1(S_1) given that a red chip (M_1) was drawn is:

$$Prob(S_1|M_1) = \frac{prob(M_1|S_1)prob(S_1)}{prob(M_1|S_1)prob(S_1) + prob(M_1|S_2)prob(S_2)}$$

In short-form notation let,

$prob(S_j) \equiv \pi_j$; $prob(M_k|S_j) \equiv q_{k.j}$; $prob(S_j|M_k) \equiv \pi_{j.k}$.

Hence, the conditional probabilities of S_j given message M_1(red message) using short-form notation are:

$\pi_{1.1} = \frac{q_{1.1}\pi_1}{q_{1.1}\pi_1 + q_{2.1}\pi_2}$; and from Eqn. 2 $\pi_{2.1} = 1 - \pi_{1.1}$; (Eqns. 3 &4)

And, the conditional probabilities of S_j given message M_2(blue chip) are:

$\pi_{1.2} = \frac{q_{2.1}\pi_2}{q_{2.1}\pi_1 + q_{2.2}\pi_2}$; and from Eqn. 2 $\pi_{2.2} = 1 - \pi_{1.2}$; (Eqns. 5 & 6)

Therefore, from Eqns. 3 & 4, the expected payoff of choosing action A when message 1 (red chip) is received is:

$EP_{action\ A}|_{M_1} = \pi_{1.1}\ C(A,S_1) + \pi_{2.1}\ C(A,S_2)$ (Eqn. 7)

The expected payoff of choosing action B when message 1 (red chip) is received is:

$EP_{action\ B}|_{M_1} = \pi_{1.1}\ C(B,S_1) + \pi_{2.1}\ C(B,S_2)$ (Eqn. 8)

Given the red chip message (M_1), the risk neutral BEU maximizer will choose action A if the expected payoff is greater than choosing action B given the posterior probabilities conditional on the red chip message.

From Eqns. 7 & 8, the risk-neutral BEU will choose action A if:

$EP_{action\ A}|_{M_1} = \pi_{1.1}\ C(A,S_1) + \pi_{2.1}\ C(A,S_2) \geq EP_{action\ B}|_{M_1} = \pi_{1.1}\ C(B,S_1) + \pi_{2.1}\ C(B,S_2)$. (Eqn. 9)

The optimal uninformed first decision to take action A or B prior to an information signal (chip draw) is risky. Furthermore, when asked to propose a second terminal decision, the message service reveals the state with uncertainty. Therefore, even after observing a sample chip draw, the second action choice still bears risk. However, the information signal is not worthless. The information must have some value if the message received results in an alternative action choice from the first action choice in order to follow the BEU optimal decision.

Let us assume that the message received results in the optimal action being taken, based on the conditional (posterior) probability, $\pi_{j.k}$. Moreover assume that this optimal action, B, once the message is received is different from the optimal action, A, taken when there is no message. Let us assume the message received is a blue chip (M_2), then, the expected value of the message service (EV_M) for this example can be calculated as follows:

$$EV_{M_2} = EU\big(B;\pi_{j.k}\big) - EU(A;\pi_{j.k})$$

$$EV_{M_2} = [\pi_{1.2}c(B,S_1) + \pi_{2.2}c(B,S_2)] - [\pi_{1.2}c(A,S_1) + \pi_{2.2}c(A,S_2)] \qquad \text{(Eqn. 10)}$$

Below is a numerical example using exogenous parameters[60] from the experiment found in Appendix 2 to calculate the expected value of the information (the worth of the message service). For rounds 1-4 & 13-16 (Set 1), the BEU maximizer chooses action A as their first choice and action B as their second choice when a blue message is received. From Eqn. 4, 5 & 6 the expected value of the information is:

$$EV_{M_2} = \left[\frac{.3(.5)}{.3(.5)+.7(.5)}\$0.50 + \frac{.7(.5)}{.3(.5)+.7(.5)}\$1.75\right] - \left[\frac{.3(.5)}{.3(.5)+.7(.5)}\$2.00 + \frac{.7(.5)}{.3(.5)+.7(.5)}\$0.75\right]$$

$$EV_{M_2} = [(.3)\$0.50 + (.7)\$1.75] - [(.3)\$2.00 + (.7)\$0.75]$$

$$EV_{M_2} = \$0.25$$

[60] The exogenous parameters refer to the distribution of red to blue chips contained within each bag and the payoffs associated with each action choice.

Table 2 calculates the risk- neutral BEU expected value of information for the various scenarios created by the exogenous parameters given by the experimental design (see mathematical Appendix 3 for EV_M for all sets of rounds with identical parameter values)

Table 2: Risk Neutral BEU Value of Information by round

Rounds	1-4 13-16				5-8 17-20				9-12 21-24			
First Action	A*		B		A		B*		A		B*	
Message	RED	BLUE	RED	BLUE	RED	BLUE	RED	BLUE	RED	BLUE	RED	BLUE
BEU optimal Action	A	B	A	B	B	A	B	A	B	B	B	B
BEU Expected Value of Information	$0.00	$0.25	$0.75	$0.00	$0.90	$0.00	$0.00	$0.40	$0.45	$0.80	$0.00	$0.00

*the BEU optimal first action choice

3.3 The 5 Main Hypotheses

The theory and intuition underpinning the 5 main hypotheses tested in this study are described below.

Let us assume the subject's action choice prior to receiving the message is optimal and is action A while their action choice once the message is received is also optimal but is action B. Let us assume the message received is a blue chip (M_2). As this message signals to the BEU maximizer to switch to the alternative action, the expected value of the blue chip message is greater than zero ($EV_{M_2} > 0$). As such, the WTP amount specified by a risk neutral BEU maximizer to use the information will equal the expected value of the message service($WTP = EV_{M_2}$). Therefore from Eqn. 10 the WTP conditional on receiving a blue chip message can be calculated for the risk neutral BEU decision maker.

$$WTP_{M_2} = EV_{M_2} = [\pi^*_{1.2}c(B,S_1) + \pi^*_{2.2}c(B,S_2)] - [\pi^*_{1.2}c(A,S_1) + \pi^*_{2.2}c(A,S_2)] \quad \text{(Eqn. 11)}$$

$$\pi_{1.2} + \pi_{2.2} = 1 \quad \text{(Eqn. 12)}$$

Recall from Eqns. 5 & 6,

$\pi^*_{1.2} = \frac{q_{2.1}\pi_1}{q_{2.1}\pi_1 + q_{2.2}\pi_2}$; represents the Bayes law posterior probability of bag 1, given message 2;

$\pi^*_{2.2} = \frac{q_{2.2}\pi_2}{q_{2.1}\pi_1 + q_{2.2}\pi_2}$; represents the Bayes law posterior probability of bag 2, given message 2.

If $WTP_{M_2} \neq EV_{M_2}$, and assuming that a subject is capable of maximizing expected payoffs, the implied posterior probabilities ($\widehat{\pi_{j.k}}$) assigned by the subject can be calculated.[61]

Given the subject's WTP value, the subject's implied posterior probabilities can be calculated by substituting the true posterior probabilities, $\pi^*_{j.k}$, with the unknown posterior probabilities, $\widehat{\pi_{j.k}}$, in Eqn. 11 in conjunction with Eqn. 12, and simplifying:

$\widehat{\pi_{1.2}} = \frac{WTP+[C(A,S_2)-C(B,S_2)]}{[C(B,S_1)-C(A,S_1)]+[C(A,S_2)-C(B,S_2)]}$; Estimated posterior probability of bag 1 (S_1) , given a blue message (M_2);

$\widehat{\pi_{2.2}} = \frac{[C(B,S_1)-C(A,S_1)]-WTP}{[C(B,S_1)-C(A,S_1)]+[C(A,S_2)-C(B,S_2)]}$;Estimated posterior probability of bag 2 (S_2), given a blue message (M_2).

Given $\mathrm{C(A,S_2) - C(B,S_2) < 0; C(B,S_1) - C(A,S_1) < 0; WTP \geq 0}; \pi_{j.k} \leq 1;$ then $\widehat{\pi_{1.2}} \in (0,1)$ and $\widehat{\pi_{2.2}} \in (0,1)$.

Hence if,

$\widehat{\pi_{1.2}} < \pi^*_{1.2}$, then the subject under-estimates the posterior probability (PP) of bag 1 given message 2, and over-estimates the PP ($\widehat{\pi_{2.2}}$) of bag 2, given message 2 relative to the Bayes law benchmarks,

$\widehat{\pi_{1.2}} = \pi^*_{1.2}$, then the subject appropriately estimates the PP relative to the Bayes law benchmarks,

$\widehat{\pi_{1.2}} > \pi^*_{1.2}$, then the subject over-estimates the PP of bag 1 given message 2, and under-estimates the PP ($\widehat{\pi_{2.2}}$) of bag 2, given message 2 relative to the Bayes law benchmarks.

The Handbook of Experimental Economics (Kagel & Roth, 1995) refers to the underweighting of likelihood information as 'conservatism'.[62] Peterson & Dechane (1967) found that subjects

[61] For WTP =0 a posterior probability can be estimated that represents only an upper bound for the subject. I.e., a subjects posterior probability could be less however, a subject cannot specify a WTP<0.

[62] "Conservatism is a kind of under- confidence that results when people underemphasized the large size of a sample of weak evidence" pp. Hirshleifer & Riley (1992).

took longer to switch to the alternative action after receiving a message (they needed to observe more messages) than if the Bayes law rule had been properly applied. Edwards (1968) observes that it takes subjects 2 to 5 observations to produce a diagnostic prediction equivalent to the Bayesian prediction after one observation (see also, Eger & Dickhaut,1982). McKelvey & Page (1990) conducted an experiment where people observed different parts of a full sample and then reported probability estimates to each other. After hearing the estimates of others, people reported new estimates taking into account the estimates of others and continuing to update estimates over several rounds. They found confirmation of conservatism in the subjects' final estimates relative to the estimates if Bayes law had been properly applied. Harrison et al. (2010) found that subjects over-estimate the value of the message relative to BEU when posterior probabilities are less than 0.5 and under-estimate it when posterior probabilities are greater than 0.5; where the greater than 0.5 result supports a 'conservatism' hypothesis.[63] Hirshleifer and Riley (1992) highlight a useful proposition that is implicit in the Bayesian belief revision process that may help explain some of these findings.

Proposition 1: The higher the prior confidence in initial beliefs the closer the posterior probabilities will be to the prior probabilities for any given message.

This proposition becomes clear if we assume the extreme case where a subject's prior belief is that they are in one state versus the other and that this belief is held with certainty. As such, any message received will have no value, and the subject will undervalue all information signals.

Hypothesis 1 (H1) is based on the findings from these previous studies:

H1: When a message is received, subjects will under-estimate the difference between the posterior probabilities and the prior probabilities relative to the risk-neutral BEU model predictions. Specifically, in this experiment they will under-estimate the probability of bag 2 given message 2 during rounds 1-4 & rounds 13-16 and under-estimate the probability of bag 1 given message 2 ($\pi_{1.2}$) during rounds 5-8 & 17-20.[64]

The informational value of a message can be summarized by the probability of message k given state j ($q_{k.j}$) relative to the probability of the message k given the alternative state ≠j ($q_{k.\neq j}$). The greater the difference in these two probabilities the greater the informational value of the

[63] Initial probabilities for the two states in this experiment were .5/.5.

[64] For example in both these cases, the mean implied estimated posterior probabilities for subjects will lie between the unconditional probabilities (priors) and the conditional Bayesian probabilities (posteriors).

message service. Therefore, given that $q_{k.j} + q_{k.\neq j} = 1$, when $q_{k.j} - q_{k.\neq j} = 1$ the message is completely informative and when $q_{k.j} - q_{k.\neq j} = 0$ the message is completely uninformative.

For example, suppose bag 1 has all red chips and bag 2 has all blue chips. Then a single draw provides conclusive information. A red chip drawn implies with 100% accuracy that bag 1 is in play during the round and a blue chip drawn implies with 100% accuracy that bag 2 is in play during the round. The difference between the probabilities that the message came from bag 1 versus bag 2 is equal to 1 ($q_{k.j} - q_{k.\neq j} = 1$); the chip draw signals the state with certainty. On the other hand, if each bag has precisely half red chips and half blue chips, then a single draw is not informative at all. The difference between the probabilities that the message came from bag 1 versus bag 2 is equal to 0 ($q_{k.j} - q_{k.\neq j} = 0$).

It is predicted in these hypothetical examples, where the difference between the probabilities of the message conditional on the state are at the extreme values (i.e., 0 and 1) and therefore the message provides no information or conclusive information, that subjects will be able to assess the informational value of the message in accordance with BEU theory-or at least better than for intermediate values of $q_{k.j} - q_{k.\neq j}$. Specifically, the subjects reported WTP will be hypothesized to be equal to the risk neutral Bayesian WTP when the message is completely informative versus completely uninformative.

This prediction is based on the following intuition. In the first case where bag 1 contained all blue chips and bag 2 contained all red chips, the first action prior to a chip draw is risky and the action conditional on the chip draw is not risky. The receipt of the information that signals that a subject should switch from their initial action is associated with shifting to a less risky (zero risk) payoff prospect. Furthermore, the optimal action prior to the message and after the message is associated with the first-order stochastically dominant lottery. Therefore, anyone with monotone preferences should select the action associated with the higher expected payoffs (riskless choice) regardless of risk preferences. Given that the message is conclusive and therefore subjects know exactly and for certain the payoff amount associated with either action choice ($0 or $>$0), they should also know the value of the information. In the second case, where each bag contained 50 red chips and 50 blue chips, the action choice prior to and after receiving the message is equally

risky. The message provides no additional information. Therefore, the optimal action before and after the chip draw is the same as in both cases this action's lottery FOSD the alternative action's lottery. Regardless of risk preferences subjects should specify a willingness to pay of zero as they do not need to switch their terminal choice.

However, for intermediate values of $q_{k.j} - q_{k.\neq j}$, it is less obvious whether subjects will be more or less capable of assessing the true value of the information when, for example a random single draw is from a scenario in which the two bags have distributions of red to blue chips of 70/30 versus 30/70 ($q_{k.j} - q_{k.\neq j} = 0.4$, a difference closer to 0) than if the distribution of red to blue chips had been 80/20 versus 20/80 ($q_{k.j} - q_{k.\neq j} = 0.6$, a difference closer to 1). In both cases, the message is informative. However, now the message that signals to a subject to switch to the alternative choice is associated with a lottery that is neither FOSD nor SOSD. Therefore, in addition to potentially complicating the decision environment (i.e., decision cannot be made by simply ranking lottery choices according to first order stochastic dominance), risk preferences may matter. Therefore, across the domain of available posterior probabilities from 0 to 1, but not close to or at O or 1, it is less intuitive how the subject may value the information relative to the risk neutral BWTP benchmark. For example the marginal utility of additional income from selecting the higher payoff state given the message falls faster as risk aversion increases. Thus a more risk averse individual would not be willing to pay as much for the information as someone who is less risk averse.

Therefore, to investigate the open question on how subjects value the information when the message is informative (has value) but is not certain (i.e., the true posterior probabilities conditional on the message are not near or at 0 or 1), I observe subjects WTP behaviour given the message and how this WTP behavior changes when the degree of uncertainty in the message service decreases (i.e. the chip draw is more informative and the difference in the probabilities of the message conditional on the state has increased). Therefore, I set the following as a statistical null hypothesis and test it against the two-sided alternative:

H2: When the informational value of the message received increases, the subjects' implied posterior probability from their WTP declaration will be closer to the risk neutral Bayes posterior probability (i.e., subjects will stipulate the BWTP value more often).

Prior to the announcement of the distribution of red to blue chips contained within each bag, and given the expected payoffs associated with each action choice, the critical values for the posterior probabilities, $\pi^{c}_{j.k}$, i.e., the switching rule where the BEU decision switches to the alternative action choice (i. e. A to B or B to A), can be calculated.

For example, a BEU decision maker will switch her choice (from action A to B) conditional on observing a blue chip (M_2) if:

$$EP_{action\ B}|_{M_2} = \pi_{1.2}\ C(B,S_1) + \ \pi_{2.2}\ C(B,S_2)$$

$$\geq EP_{action\ A}|_{M_2} = \pi_{1.2}\ C(A,S_1) + \ \pi_{2.2}\ C(A,S_2) \quad \text{(Eqn. 13)}$$

The critical values of the posterior probabilities, $\pi^{c}_{1.2}$ & $\pi^{c}_{2.2}$, where a BEU decision maker will switch to the alternative action choice conditional on observing message 2 (blue chip) can be calculated by changing the weak inequality sign to an equality sign in Eqn. 13 and solving for $\pi_{1.2}$ & $\pi_{2.2}$.

$$EP_{action\ B}|_{M_2} = \pi_{1.2}\ C(B,S_1) + \ \pi_{2.2}\ C(B,S_2)$$

$$= EP_{action\ A}|_{M_2} = \pi_{1.2}\ C(A,S_1) + \ \pi_{2.2}\ C(A,S_2)$$

Noting that, $\pi_{1.2} = 1 - \ \pi_{2.2}$, and simplifying gives:

$$(1 - \pi_{2.2}\ )[C(B,S_1) - C(A,S_1)] + \ \pi_{2.2}\ [C(B,S_2) - \ C(A,S_2)] = 0 \quad \text{(Eqn. 14)}$$

$$\pi^{c}_{2.2} = \frac{C(B,S_1) - C(A,S_1)}{[C(B,S_1) - C(A,S_1)] + [C(A,S_2) - C(B,S_2)]} \qquad \text{(Eqn. 15)}$$

Given, $(B,S_1) < C(A,S_1)$, $C(A,S_2) < \ C(B,S_2)$ then $\pi^{c}_{2.2}\ \in (0,1)$ and $\pi^{c}_{1.2}\ \in (0,1)$

Next, denote the equation for the difference in Expected Payoff from taking action B versus action A (LHS of Eqn. 14) by θ, and evaluate the partial derivative of θ with respect to $\pi_{2.2}$.

$$\frac{\partial\theta}{\partial\pi_{2.2}} = -[C(B,S_1) - C(A,S_1)] + [C(B,S_2) - C(A,S_2)] > 0;$$

Given, $C(B,S_1) - C(A,S_1) < 0\ and\ C(B,S_2) - C(A,S_2) > 0$.

Since $\theta = 0$ at $\pi_{2.2} = \pi_{2.2}^{c} \in (0,1)$ and θ is monotonically increasing in $\pi_{2.2}$, it follows that $\theta > 0$ if $\pi_{2.2} > \pi_{2.2}^{c}$ and $\theta < 0$ if $\pi_{2.2} < \pi_{2.2}^{c}$.

Hence as $\pi_{2.2}$ increases, the expected payoff from taking action B increases. Conversely, as $\pi_{2.2}$ decreases the expected payoff from taking action B also decreases and when $\pi_{2.2} < \pi_{2.2}^{c}$ the expected payoff from selecting action B has decreased to a point where the greater expected payoff is now associated with selecting action A ($EP_{action\ A}|_{M_2} > EP_{action\ B}|_{M_2}$). Therefore, for any $\widehat{\pi_{2.2}} > \pi_{2.2}^{c}$ a BEU decision maker will switch to action B, otherwise she will remain with the initial action A.

Table 3 provides the ranges of posterior probabilities ($\pi_{j.k}$) where a BEU subject will switch to the alternative action choice conditional on the message received for each set of rounds that is governed by the same exogenous parameters.

Table 3: Posterior Probabilities: Critical Values and Ranges of Estimated Posterior Probabilities by Round
A BEU Participant Switches to the Alternative Action

Posterior Probabilities	Rounds 1-4/13-16			Rounds 5-8/17-20			Rounds 9-12/21-24		
$\pi_{j.k}$	Bayes	Critical Value (CV)	Range	Bayes	Critical Value (CV)	Range	Bayes	Critical Value (CV)	Range
$\pi_{bag1.red}$	.70	.40	[0.4, 1.0]	.24	.60	[0.0, 0.6]	.60	.857	[0.0, 0.86]
$\pi_{bag2.red}$	.30	.60	[0.0, 0.6]	.76	.40	[0.4, 1.0]	.40	.143	[0.14, 1.0]
$\pi_{bag1.blue}$	.30	.40	[0.0, 0.4]	.76	.60	[0.6, 1.0]	.40	.857	[0.86, 1.0]
$\pi_{bag2.blue}$	.70	.60	[0.6, 1.0]	.24	.40	[0.0, 0.4]	.60	.143	[0.0, 0.14]

Table 3 illustrates that there exists a range of posterior probability estimates ($\widehat{\pi_{j.k}}$) where a subject may make a Bayesian updating error ($\widehat{\pi_{j.k}} \neq \pi_{j.k}^{*}$) yet still arrive at the correct BEU action choice.

Suppose that the true posterior probability estimated using Bayes law conditional on the blue message indicates that the optimal action is to switch from action A to B and that this estimate is greater than the critical posterior probability value ($\pi_{2.2}^{*} > \pi_{2.2}^{c}$) calculated above. Next suppose that a subject estimates a posterior probability that is greater than the critical switching value,

($\pi_{2.2}^{c} < \widehat{\pi_{2.2}}$) but is not equal to the true posterior probability ($\widehat{\pi_{J.k}} \neq \pi_{j.k}^{*}$). This subject, assuming they are capable of estimating the expected utility portion of the BEU decision rule, will then select the optimal BEU action choice even though their estimated posterior probability was not optimal.

When subjects inaccurately identify the BWTP value, a question remains whether this is because subjects have difficulty with just the Bayes law portion of the BEU decision rule, continuing to accurately apply the expected utility portion, or whether they have difficulty with both aspects (Bayes law plus EU) when new information is presented.

Proposition 3: If Subjects are adept at maximizing expected payoffs in accordance with a risk-neutral BEU maximizer then, it is possible that a WTP value that implied a posterior probability that fell within the acceptable range (see table 3) of estimated posterior probabilities (but not the true posterior probability) would still lead subjects to select the optimal terminal action.

Therefore, action choices that deviate from the BEU benchmark will be the same for these subjects as the subjects who specified the optimal WTP amount.

H3: The Second choice BEU inconsistency rate is the same when the WTP equaled BWTP benchmark[65] as when the WTP did not equal BWTP benchmark, but implied a posterior probability that fell within the range of estimated posterior probabilities, that when combined with risk-neutral Expected Utility Theory would still result in an optimal choice.

Hypothesis 4 (H4) originates from the intuition that individuals will exert more effort toward accuracy in their decision choices, if they must pay for information versus observing it for free. This intuition is based on the assumption that the change to the OTP procedural steps requiring subjects to specify a WTP value to have the more informed second action choice be considered for payment (i.e., the subject must now put pay at risk versus no pay at risk in the FREE message task) reframed the task such that there is now more importance placed on the value of the information service. As such, subjects will use the information more optimally than when the

[65] The BWTP benchmark represents a range of acceptable values: -$0.05 < BWTP < $0.05.

information was available for free.[66] Hence, the price an individual is willing to pay during the OTP task to use information that conveys the probable outcomes of alternative actions and the consequent action choices will have a greater correlation with the precision, relevance and reliability of the information than when the message could be used for free (Feldman & March, 1981).

Proposition 4: As WTP choice conveys information about how the information is valued (Stiglitz, 2000), the consequent terminal actions of subjects when they must pay for information would more likely be optimal than when they received the information for free.

Hypothesis 4: When a subject must pay for the information versus observe the information for free, they will deviate less from the BEU optimal decision when taking their final choice.

Hypothesis 5 (H5) originates from chapter 1 of this thesis: There is inconclusive evidence as to why individuals do not accurately apply Bayes law beyond the simple explanation that people 'lack the cognitive sophistication' to do the math (Wallensten, 1968; Fried & Peterson, 1969; Pitz, Reinhold & Geller, 1969; Hershman & Levine, 1970; Hammerston, 1973; Grether, 1980; Charness & Levin, 2005). It follows that providing subjects with the Bayes Law calculation would simplify the math component of the BEU WTP decision rule allowing more subjects to behave optimally (Hershman & Levine, 1970).

Proposition 5: Individuals do not behave optimally because they do not have the math skills to calculate Bayes law.

Hypothesis 5: Providing subjects with the posterior probabilities conditional on the message received will result in more WTP valuations which are BEU accurate.

3.4 Results

The data used for this analysis contain the observations of 121 subjects who participated in 12 rounds of decision tasks where they observed a message for free and 12 rounds that required them to indicate their WTP to use the information observed in order to have their action choice

[66] Tversky &Kahneman(1981) discuss in detail the impact of framing effects on outcomes (the choices that individuals make).

recognized for payment. In total 7260 observations were collected; 2904 observations each of the subjects' first (prior to observing a message) and second (after observing a message) decision choices and most importantly for this chapter, 1452 observations of the subjects' specified WTP prices.

The subjects' willingness to pay decision is benchmarked relative to the willingness to pay decision of a risk-neutral Bayesian Expected Utility maximizer. Hence, for each subject in the experiment, it is determined whether the subject over-, under-. or accurately estimated the value of the information relative to the risk-neutral Bayesian Expected Utility (BEU) model. Additionally, to investigate whether risk preferences influenced the observed WTP behavior found when testing both hypotheses 1 and 2, WTP BEU benchmarks are established assuming both a constant relative and constant absolute risk aversion (CRRA and CARA, respectively) utility function.

The area plot in Figure 1 graphs the pooled WTP decisions of each subject over all rounds and highlights the proportion of rounds that the subject over-valued, under-valued or appropriately valued the information relative to the risk neutral BEU model. Along the horizontal axis, subjects are ranked overall from farthest away to closest to the BEU WTP behavior. On average subjects chose the risk neutral BEU optimal willingness to pay amount 57.5% of the time and either undervalued or overvalued the message received relative to the BEU benchmark 22.6% and 19.9% of the time, respectively.

Figure 1

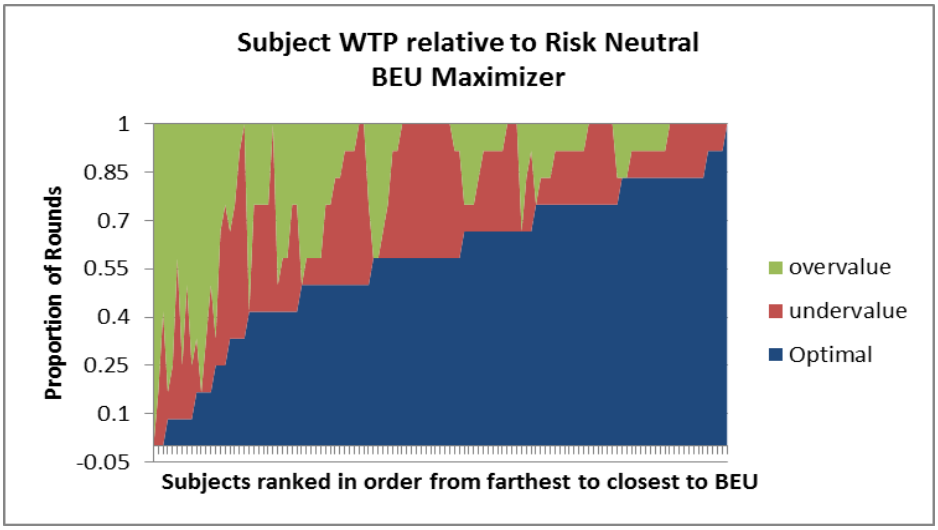

On average, subjects specified a willingness to pay (WTP) amount that was greater than the risk neutral BEU maximizer's WTP amount by 168%[67]; the mean WTP by subjects is $0.40 versus the mean BWTP amount of $0.151 (see Figure 2).

Figure 2

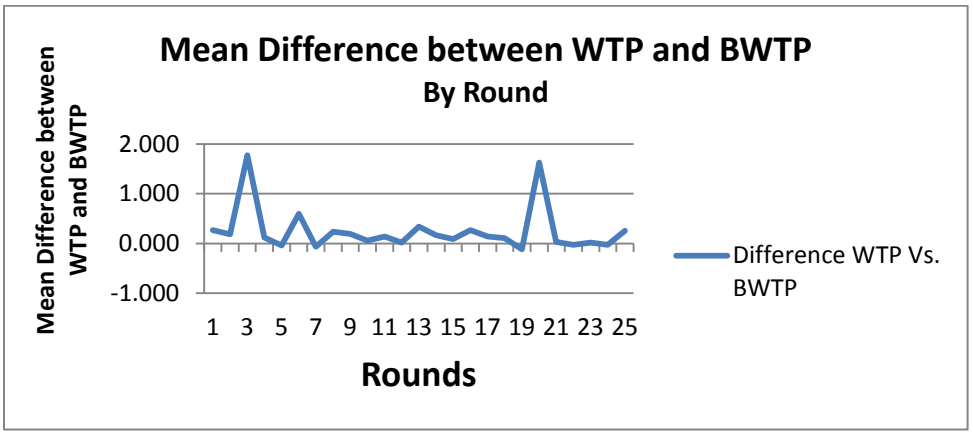

A closer examination of the subjects' WTP observations (specifically, histogram (a) found in appendix 5 as well as the Box graph (a) found in appendix 7) shows 25 outlying WTP data points (1.7% of the total observations: see figure 3). These data points represent implausible WTP observations as they represent values that were greater than the maximum payoff ($2.00) available to a subject for any given round (i.e., $100, $25, & $10) and have a strong and misleading influence on any statistical test based on sample means and variances. The outlying observations are a result of cognitive processes associated with a subject that cannot be identified. As subjects were informed at the beginning of each task that negative payoffs were not possible, the conjecture is that the subjects who specified a WTP amount greater than $2.00 either; i.) knew they could not lose in excess of $2.00 and were willing to pay the maximum amount ($2.00) or, ii.) did not understand the rules of the WTP task. Although, there is no unanimously accepted theoretical framework for the treatment of these outliers (Cousineau & Chartier, 2010; Harmeling et al., 2006), one author notes (Snapp, 2010) that if the outlier has a low probability of repeating in the future, the outlier could be dropped. Using this logic the 25 outlying WTP observations were dropped from the data set and the balance of the observations associated with the 9 subjects who specified these implausible WTP values are maintained. The

[67] A Wilcoxon Rank test verifies that the difference between WTP versus BWTP is statistically significant at the 1% level.

rationale is that the large WTP amounts occurred at the beginning of the OTP task, and plausible WTP amounts were specified later in the rounds, suggesting that subjects eventually learned the rules. The WTP truthful valuation elicitation method added complexity to the subjects decision that was not present when they made either their first or second decision choices.[68] As such, the nine subjects who specified these WTP amounts are kept in the data based on the plausibility of their remaining first, second and WTP decision choices (the probability of these observations repeating themselves was high).

Figure 3

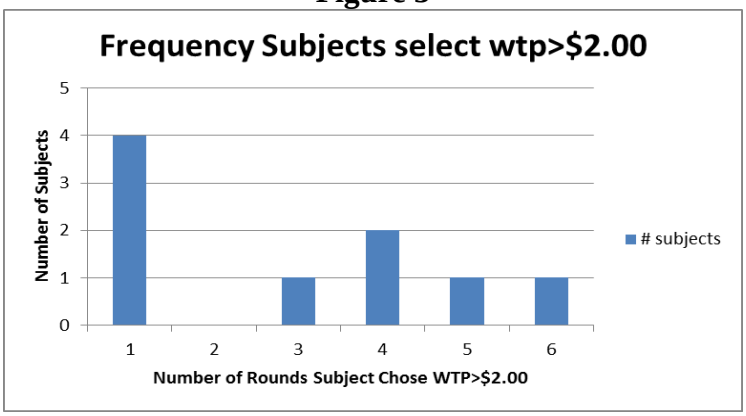

With the outliers dropped, subjects on average undervalued the information relative to the risk neutral BEU maximizer by 6.0% (a mean WTP of $0.139 versus a mean BWTP of $0.148). A Wilcoxon rank test, however, identified this difference as not significant statistically at conventional levels.[69]

Figure 4 compares the mean difference between WTP amounts with the outliers dropped and BWTP by round and compares this mean difference to the mean difference with the data set with the outliers present. Figure 5 compares this mean difference aggregated over all rounds for the two data sets (unrestricted and drop the WTP >$2.00) and illustrates the major influence that the

[68] Another option would have been to drop all observations associated with these nine subjects (first, second action choices, the remaining WTP values in addition to the 25 outlying WTP observations). T-tests were performed on the outcome variables from the data set with all observations included with the exception of the 25 outliers versus all observations dropped associated with the 9 subjects who were responsible for the outliers and there was no significant difference statistically in these outcome variables between the two data sets.

[69] This aggregate result represents only a preliminary analysis of the subjects' WTP behavior, as it does not consider the distribution of WTP values (i.e., some subjects over-valuation may have cancelled out another subjects under-valuation) and serves only to emphasis the impact of the 25 outlying WTP data points. A detailed analysis of subject WTP decisions is conducted in reference to hypothesis 1.

outlying observations (where 9 subjects 'overpaid') had on the comparison between the subjects' WTP values and the BWTP benchmark.

Figure 4

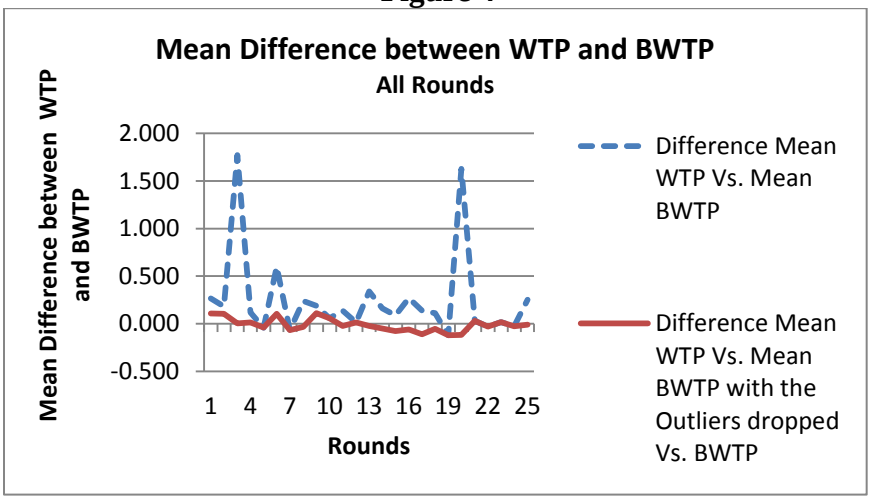

Figure 5

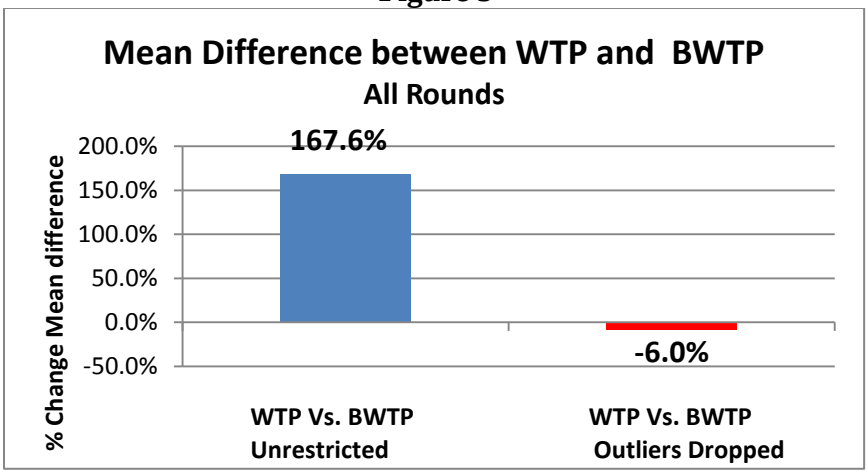

To understand the causes of the observed WTP behaviour, I estimated logit models (a random and a fixed effect model) where the dependent variable was a dichotomous outcome variable, with 0 representing a subject's WTP amount that was not consistent with the BEU amount (BWTP), and 1 representing a subject's WTP amount that was consistent with the BWTP amount, and determined the log odds and marginal effects[70] of the independent variables on this

[70]I report the marginal effect for all 4 equations where ME represents the change in the probability of observing the dependent variable, if the independent variable changes by one unit

outcome (see Table 4, eqn. 1-4). The subject's WTP decision was classified as BWTP consistent if their specified WTP amount fell within the range of + or - $0.05 of the BWTP amount (i.e. $0.05<BWTP<$0.05)[71].

Overall, subjects specified the Bayesian WTP amount in 57.5% of the 1452 WTP decisions observed. However, when the information had no value (BWTP=$0.00) and therefore did not require a subject to shift to the alternative action, the proportion of WTP decisions that were consistent with the BWTP amount rises to 83.6%. On the other hand when the information signaled that a subject should change to the alternative action and therefore, the information had value (BWTP>$0.00), the proportion of WTP decisions that were consistent with the BWTP amount was only 12.4%. Although a WTP decision that equals zero reflects optimal behavior, it can also represent an upper bound for the subject's WTP valuation as it is not possible to select a WTP that under-estimates the informational value of the message received (i.e., WTP cannot be less than zero). To account for this issue an explanatory variable was added to the logit regression found in Table 4 (eqn. 1) controlling for BWTP values that were equal to zero (Table 4, eqn. 2-4).

There are three types of variables used to explain the data for all four regressions. First, there is a group of explanatory variables that change over the rounds but are the same for all individuals in a given round. Second, there is a set of explanatory variables that vary both over the rounds and between subject and session. Finally, there are explanatory variables that vary between individuals but do not vary over the rounds[72]. These variables are presented and described in Appendix 4.

The Breusch and Pagan Lagrangian multiplier test for all four equations established that individual effects are present in the data. The Hausman test rejected the null hypothesis that the coefficients for the fixed and random effects model are the same, implying that the random

[71] Note in order to normalize the range, given different BWTP values dependent on the round parameters the following relative range values were also tested (80%xBWTP)<BWTP<(120%xBWTP). There were no changes to the results regardless of which range is implemented (absolute or relative range values)

[72] These explanatory variables are removed in the fixed effects model as the FE model controls for omitted variables that differ between cases but are constant over-time.

effects coefficients are correlated with the individual error terms. As such in this chapter, I report on the results from the fixed effects model and provide the results from the random effects model in Appendix 10.

As hypothesis 4 and 5 test for the impact of two explanatory variables (experience[73] & informed[74]) that were dropped from the fixed effects model, a between effects OLS model is estimated using a collapsed data set with the dependent variable representing the mean value of the number of times that a subject accurately specified the BWTP amount (condensing 12 observations into 1 for each subject). The explanatory variables used in this regression represent mean values by subject and although vary between individuals do not vary over the rounds. The variables are described in Appendix 4 and the regressions are found in Table 5.

The results from Table 4 and Table 5 are discussed in detail in the subsequent testing of the five main hypotheses proposed in this study.

[73] A sub-set of subjects had experience with the FREE message task in advance of the OTP task and a sub-set did not.
[74] A sub- set of subjects were provided with the Bayes posterior probabilities and a sub-set were not.

Table 4: Logit Regressions (FE)

	(1)	(2)	(3)	(4)
		Marginal Effects		
	Subject WTP= -$0.05<BWTP<$0.05	Subject WTP= -$0.05<BWTP<$0.05 Control BWTP=0	Subject WTP= -$0.05<BWTP<$0.05 Interaction BWTP=0*Set 2	Subject WTP= -$0.05<BWTP<$0.05 Interaction BWTP=0*Set 3
Variable	**logit, FE**	**logit, FE**	**logit, FE**	**logit, FE**
Message Received is Blue	**-0.204*** **	**-0.017**	**-0.013**	**-0.023**
Usually associated with BWTP>$0.00	0.024	0.021	0.021	0.023
Set 2 (rds. 5-8 & 17-20) **Vs. Set 1** (rds. 1-4 & 13-16)	**-0.064*** **	**-0.047*** **	**-0.143*** **	**-0.154*** **
°Informative Power of chip draw increases	0.026	0.019	0.035	0.037
°State contigent payoffs are constant				
Set 3 (rds. 9-12 & 21-24) **Vs. Set 1** (rds. 1-4 & 13-16)	**0.292*** **	**0.073*** **	**0.098*** **	**0.029**
°informative Power of chip draw decreases	0.031	0.024	0.024	0.054
°State contigent payoffs change				
2nd Choice BEU Inconsistency	**-0.430*** **	**-0.247*** **	**-0.243*** **	**-0.254*** **
	0.158	0.024	0.025	0.028
Subject's 2nd action is recognized for payment	**-0.364*** **	**-0.150*** **	**-0.142*** **	**-0.147*** **
WTP≥Random Price drawn	0.030	0.020	0.022	0.024
BEU 2nd action not consistent with the higher payoff state in prior round	**-0.047** **	**-0.020**	**-0.018**	**-0.020**
	0.025	0.017	0.018	0.019
BWTP equals zero (0)		**0.258*** **	**0.201*** **	**0.183*** **
Colour chip draw provided no information		0.012	0.014	0.018
Interaction 1: BWTP equals zero x Set 2			**0.177*** **	**0.199*** **
Colour chip provides no info in Set 2			0.043	0.048
Interaction 2: BWTP equals zero x Set 3				**0.098**
Colour chip provides no info in Set 3				0.066
Obs.	**1416**	**1416**	**1416**	**1416**
Log Likelihood	**-312.89**	**-211.78**	**-201.28**	**-200.02**

*****p-value≤.01<**p-value≤.05<*p-value<.10**

(The smaller numbers beneath the coefficients represent standard errors)

Table 5: OLS Regressions (BE)

	(5)	(6)	(7)
	Proportion of observations Subject WTP= -$0.05<BWTP<$0.05 2nd Choice Error BWTP=0	Proportion of observations Subject WTP= -$0.05<BWTP<$0.05 BWTP=0	Proportion of observations Subject WTP= -$0.05<BWTP<$0.05
Variable	**OLS**	**OLS**	**OLS**
Experience	**-0.067*** **	**-0.072*** **	**-0.084*** **
No Experience with the FREE Message Task	0.028	0.032	0.038
Informed	**0.017**	**0.0225**	**0.039**
Subject was provided with the Bayes law Posterior Probability Calculation	0.022	0.025	0.031
Message Received is Blue	**0.851*** **	**0.814*** **	**0.404*** **
Usually associated with BWTP>$0.00	0.149	0.170	0.195
2nd Choice BEU Inconsistency	**-0.630*** **		
	0.109		
BWTP=0	**0.678*** **	**0.815*** **	
	0.101	0.112	
Subject's 2nd action is recognized for payment	**-1.00*** **	**-1.00*** **	**-1.22*** **
In prior round: WTP≥Random Price drawn	0.109	0.125	0.148
BEU 2nd action not consistent with the higher payoff state in prior round	**-0.298* **	**-0.381** **	**0.346**
	0.170	0.193	0.235
Female	**0.019**	**0.013**	**-0.027**
	0.025	0.028	0.033
English is the Second Language	**-0.013**	**-0.019**	**-0.015**
	0.030	0.034	0.042
Reinforcement Learner	**-0.043* **	**-0.059** **	**-0.062* **
Post-experiment Survey identified subject as a Reinforcement learner versus Theorist	0.025	0.028	0.034
Risk Aversion Score	**0.002**	**0.003**	**0.001**
Holt Laury Risk aversion Lottery choice (1-10)	0.004	0.005	0.005
Age	**0.001**	**-0.001**	**0.005**
	0.006	0.007	0.009
Math Economics Student	**0.048** **	**0.053** **	**0.072*** **
Student has some math optimization training	0.024	0.028	0.034
Obs.	**121**	**121**	**121**
R-sq	**0.7488**	**0.674**	**0.4595**

*****p-value≤.01<**p-value≤.05<*p-value<.10**

(The smaller numbers beneath the coefficients represent standard errors).

H1 predicted that subjects will underestimate the value of the message received relative to the BEU model prediction. Given the aggregate results found in figure 5, we should reject H1; regardless of the WTP data set used for analysis, subjects either over-estimated or optimally estimated the informational value of the message service relative to the BEU benchmark.

However, ending the analysis at the aggregate level would be misleading. A more robust result requires the analysis of the data by sets of rounds exposed to like exogenous parameters (i.e., distribution of colour chips contained within each bag and/or state contingent payoffs were the same across all rounds), as changes to either of these exogenous variables change the BWTP benchmark. Exogenous parameter values are the same for; 1) rounds 1-4 &13-16 (referred to as Set 1), 2) rounds 5-8 & 17-20 (Set 2), and 3) rounds 9-12 & 21-24 (Set 3).

Additionally, when the BWTP value was equal to zero, only WTP decisions that over-estimated or were optimal relative to the benchmark were observed (WTP cannot be less than zero). Including these observations (representing 66% of the total WTP decisions across all rounds) when testing whether the subjects were conservative in their WTP estimates (H1) would bias the results toward optimal and over-estimating WTP behavior.

Furthermore, to follow BEU model predictions, a subject must apply Bayes law in conjunction with Expected Utility theory. Hence, to identify the WTP decision as an over- , an under- or an optimal valuation, it must be established that subjects are capable of computing the Expected Utility portion of BEU decision rule; otherwise, it is possible that the non-optimal WTP value is not due to a miscalculation of Bayes law, but rather a miscalculation of expected payoffs. Establishing that subjects can perform the Expected Utility portion of the BEU problem will be more suggestive that WTP decisions are really indicators of the informational value (in terms of posterior probabilities) that a subject assigned to the message received.

Therefore, removing the observations where the BWTP value was equal to zero (thus eliminating the bias toward over- and optimal valuation),[75] in conjunction with isolating the analysis to the observations where the subjects' sequence of action choices (first & second) for each round followed the BEU benchmark, will allow for a more sensible test of H1.

[75] The following analysis excludes Set 3 as well as all other observations where the BWTP was equal to zero.

During Set 1 (rds. 1-4 & 13-16), BEU subjects selected action A prior to observing an imperfect message and when a blue message was observed changed their decision to action B; where, the optimal BWTP was $0.25 given the posterior probability of .70 of the state being bag 2 given the blue chip message. Subjects who chose BEU actions conditional on the message received specified a WTP amount on average of $0.19[76]; indicating an estimated posterior probability[77] of .68. Figure 6 and 7 illustrate the distribution of the WTP amounts specified by subjects and the corresponding estimated posterior probabilities, respectively. The clear bar in each chart indicates the frequency of observations that equaled (within the ∓$0.05 range) the BEU WTP ($0.25) implying the Bayes law posterior probability conditional on receiving the blue message of.70. The distribution of observations is right skewed (i.e., more observations to the left of the optimal WTP value and consequently the true posterior probability) supporting the hypothesis that subjects will conservatively estimate the probability of the state conditional on the message received (H1). 52% of WTP decisions represented under-estimates of the informational value of the message received relative to the BEU benchmark, while only 26% represented over-estimates.

Figure 6

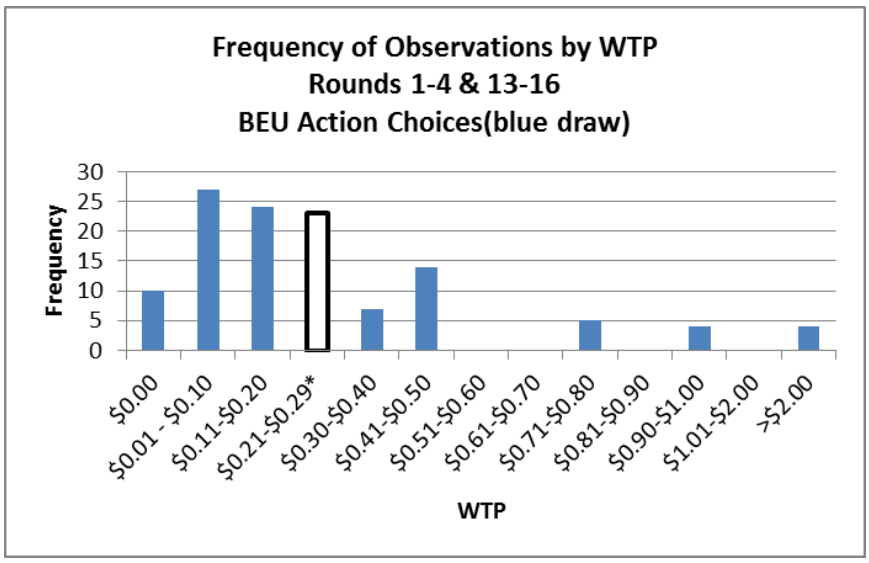

[76] 24% of the Set 1 observations represented subjects who chose the BEU 1st and 2nd choice when a blue message was observed.

[77] Note that for subjects who specified a WTP=0, the implied posterior probability estimate only represents an upper bound (subjects who have an estimated PP that was .50 for example would be required to specify a negative WTP valuable which is implausible given the experimental design).

Figure 7

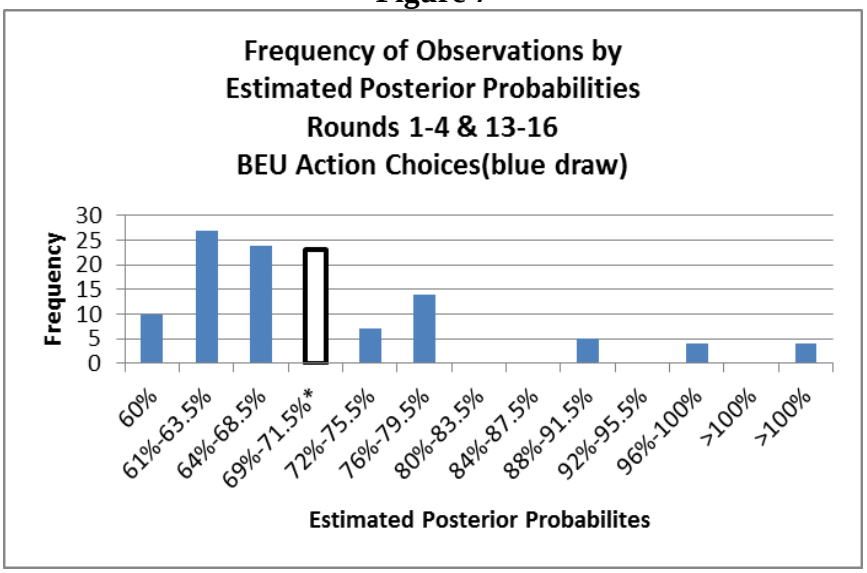

During Set 2 (rds. 5-8 & 17-20), BEU subjects selected action B prior to observing a message and when a blue message was observed changed their decision to action A. The optimal WTP amount for this interval of rounds was $0.40 and the probability of being in state 1, conditional on receiving a blue message was .76. The mean WTP specified by subjects for Set 2 was $0.27; indicating an estimated posterior probability of .705.[78]

Figures 8 and 9 illustrate the distribution of the WTP values and their corresponding estimated posterior probabilities for rounds 5-8 & 17-20 (Set 2). The clear bar in each chart indicates the frequency of subject observations that equaled the BEU WTP ($0.40 in Figure 8) implying the Bayes law estimated posterior probability (.76 in figure 9). The distribution of observations is right skewed (greater proportion of observations to the left of the optimal values) again supporting the hypothesis that subjects will conservatively estimate the WTP values. In this set of rounds, 53.5% (vs. 52% in Set 1) of WTP decisions represented under-estimates of the informational value of the message received relative to the BWTP, while 29.5% (vs. 26% in Set 1) represented over-estimates.[79]

[78] 27.3% of the Set 2 observations represented subjects who chose the BEU 1st and 2nd choice when a blue message was observed

[79] These figures are not consistent with the forthcoming results found in Table 6. The above figures are restricted to include only observations where the BWTP>0 and subjects first and second decision choices were optimal.

Figure 8

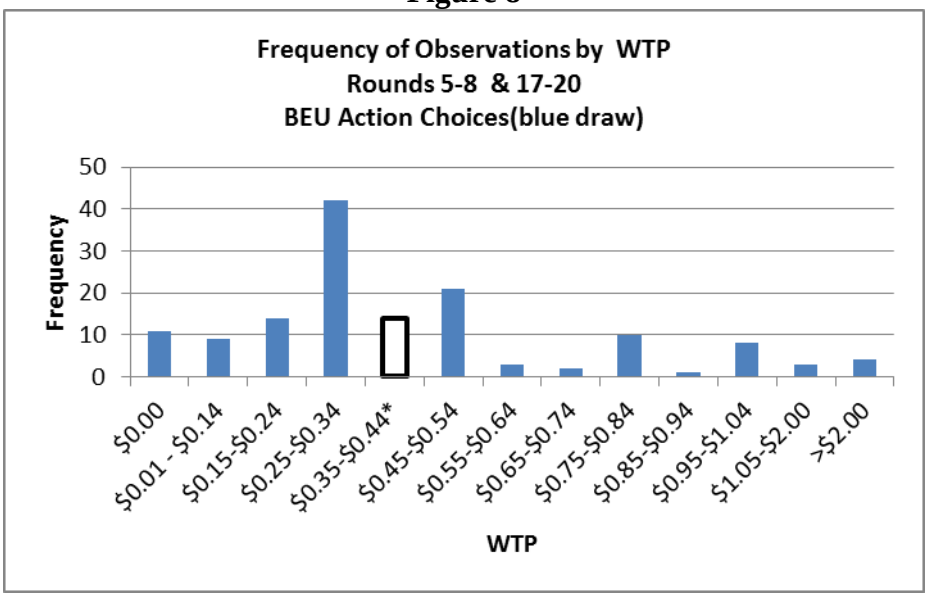

Figure 9

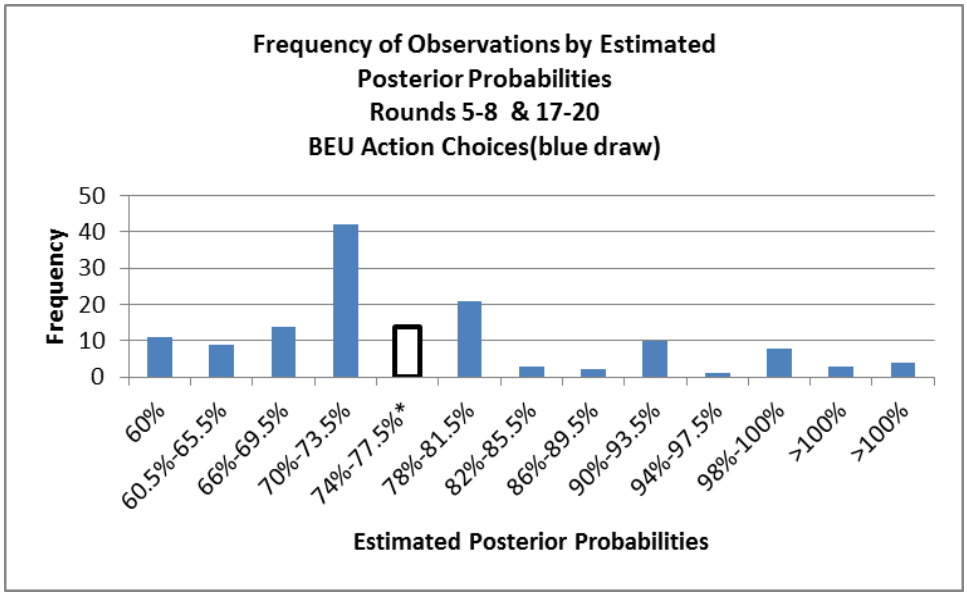

During Set 3 (rds. 9-12 & 21-24), BEU subjects selected action B prior to observing a message and regardless of the message received selected action B as their second choice. The BWTP amount was therefore zero for this set of rounds. 72% of subjects followed the BEU action choices specifying a mean WTP of $0.03. This over-estimation of the WTP value was not statistically significant at conventional levels.

Figure 10, illustrates subject WTP by round relative to the BWTP benchmark when the information signal had no value (i.e., BWTP=0). In these rounds subjects could only over- or optimally value the information service. On the other hand, Figure 11 illustrates subject WTP behavior by round relative to the BWTP benchmark when the informational value of the message received was greater than zero (Figure 11)[80]. This figure also includes the observations of subjects who picked the non-optimal action prior to observing a message and then chose the optimal action once the message was received and further amplifies the subjects' under-valuation of their WTP (implying an under-estimation of the posterior probabilities) relative to the BWTP benchmark.

Figure 10

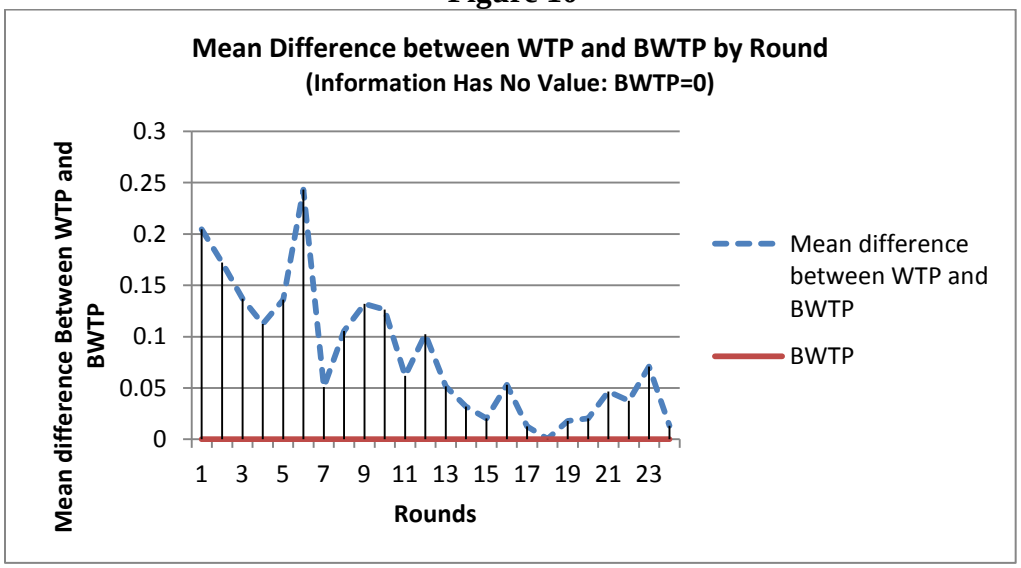

[80] Unlike the previous analysis, this analysis includes subjects that chose the non-optimal first choice but chose the optimal second choice conditional on the message received (the message had an optimal value).

Figure 11

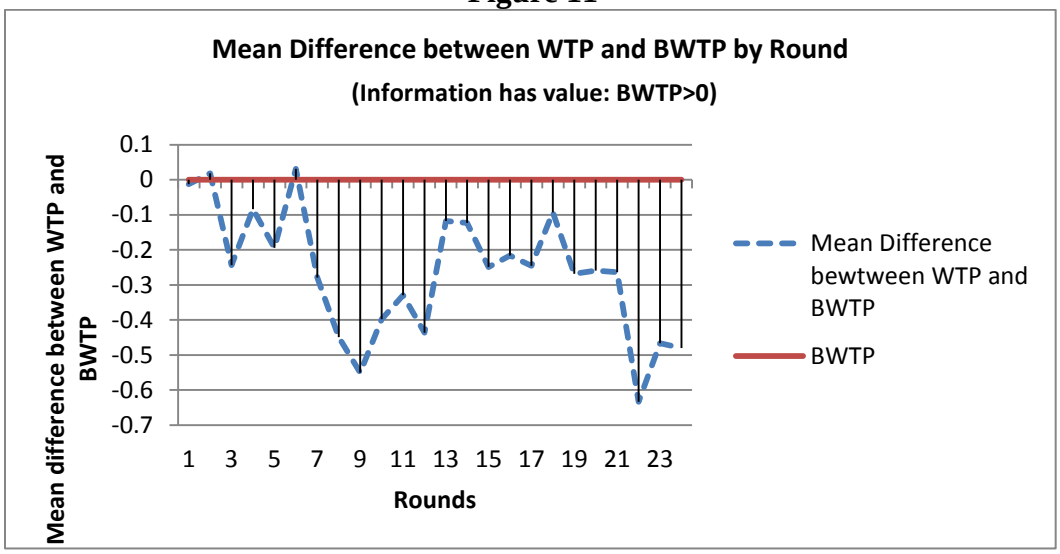

Unlike the initial result based on the aggregate sample, a more detail analysis finds support for hypothesis (H1). Subjects underestimated the value of the message service relative to the risk neutral BEU model prediction.

It is investigated at the aggregate level whether risk preferences may have contributed to the subjects' under-estimation of the WTP values relative to the RN BEU benchmark. The data used to test both hypotheses 1 & 2 restricts the observations to the subjects whose first and second action choices were optimal and to the observations where the information had value. Therefore, only the lottery choices associated with action A and B conditional on the blue chip message are relevant. Hence, in this case neither of the lottery choices are first or second order stochastically dominant regardless of the declared WTP value (i.e., WTP≥0). Therefore, risk aversion may influence WTP behavior. For example, the aggregate declared WTP by subjects was \$0.19 while the risk-neutral BEU value is \$0.25. It is possible for the \$0.19 to be an optimal BEU benchmark if we assume a constant relative risk aversion utility function of the form $u(c) = \frac{c^{1-\delta}}{1-\delta}$ where $\delta = 0.4$. Similarly, under a constant absolute risk aversion utility function of the form $u(c) = -e^{-\lambda c}$ where $\lambda = 1.1$, the optimal WTP is again \$0.19. The under-estimation of the optimal WTP by subjects may in part be attributable to risk aversion; there are risk averse (CRRA and CARA) utility functions that predict an optimal WTP consistent with the subjects' WTP.

H2 examines whether subjects will select the Bayesian optimal willingness to pay price to use information more often when the informational value of the message received increases. Specifically, as the degree of uncertainty decreases as a result of the difference between the proportions of red to blue chips contained within each bag increasing, will the ability for subjects to properly assess the value of the information increase or decrease.

Dividing the observations into sets of rounds conditional on having the same exogenous parameter values (i.e., the same proportion of red to blue chips within each bag and the same state contingent payoffs) allows us to observe how changes to these exogenous parameters influence subject WTP behavior. Table 6 highlights how the frequency of undervaluing, overvaluing and optimal valuation of the message received changed for subjects when the difference in the proportion of red to blue chips within each bag changed and when the state contingent payoffs changed (as in Set 3).

Table 6: Observations divided by sets of rounds sharing the same exogenous parameters

Set of Rounds[81] with Same Parameters	Set 1 Bag 1: R/B[82] 70/30 Bag 2: R/B 30/70 Diff PO pre-message[83]: $0 .25	Set 2 B1: R/B 24/76 B2: R/B 76/24 Diff PO pre-message: $0 .25	Set 3 B1: R/B 40/60 B2: R/B 60/40 Diff PO pre-message[84]: $0 .625	Overall
Optimal[85] WTP	53.3%	45.7%	73.5%	57.5%
Under Value	25.8%	33.7%	8.5%	22.6%
Over Value	20.9%	20.7%	18.0%	19.9%
Total	100%	100%	100%	100%

Moving from Set 1 (rds. 1-4 & 13-16) to Set 2 (rds. 5-8 &17-20) subjects experience one change to the exogenous parameters. That is, the difference between the number of red versus blue chips contained within each bag is accentuated from 40 (70-30) to 52 (76-24), increasing the difference in the probability that a message came from bag 1 versus bag 2 (from $q_{k.j} - q_{k.\neq j} = 0.40$ to 0.60), thus increasing the informative power of the chip draw.

[81] Set 1: Rds. 1-4 & 13-16; Set 2: Rds. 5-8 & 17-20; Set 3: Rds. 9-12 & 21-24

[82] Proportion of red to blue chips contained within the bags

[83] Difference ex-ante expected payoffs Set 1 & 2 = [.50($2.00)+.50($0.75)]-[.50($0.50)+.5($1.75)]

[84] Difference ex-ante expected payoffs Set 3 = [.50($2.00)+.50($0.75)]-[.50($0.50)+.5($1.00)]

[85] This is relative to the Risk-neutral BEU maximizer.

In contradiction to H2, table 6 highlights that when the message becomes more informative in Set 2 versus Set 1, subjects choose the BWTP 7.6 percentage points (ppts.) less often; representing an overall decrease in BWTP choices of 14.3%. The proportion of observations that represented an undervaluing of the message received increased by approximately 7.9ppts; representing an increase of 30.6% for Set 2 versus Set 1 (rds. 1-4 & 13-16). Results from the logit regression found in Table 4 (eqn. 1) verify this result. During Set 2, when the difference in the probability of the message coming from either state increases from 0.40 to 0.52 (increasing the informative power of the chip draw), subjects are 6.4ppts less likely to specify the BWTP benchmark (significant at the 1% level).

The increase in the informative power of the chip draw in Set 2 (rds. 5-8 & 17-20) resulted in an increase to the BEU expected value of the information. Therefore, the optimal WTP decision is more costly than in prior rounds. The finding that subjects in Set 2 are less likely to specify a WTP value that is consistent with the BWTP when the informative power of the chip draw increased does not imply that the subject did not recognize the increase to the value of the message. Evidently, subjects during Set 2 did recognize that the message was more valuable; the mean WTP specified by subjects increased from $0.19 to $0.27 versus the prior set of rounds where the message was less informative.[86] However in Set 2, the mean WTP only represented 67.5% of the BWTP mean ($0.27/$0.40) versus 76% ($0.19/$0.25) in prior rounds. The BEU value increased by 60% ([.40-.25]/.25) but the subjects' value only increased by 42% ([.27-.19]/.19). Furthermore, there were 18% more observations of undervaluing behavior in Set 2 versus Set 1 and the absolute mean difference between the subject WTP values and the BWTP benchmark for the observations representing undervaluing behavior increased to $0.13 from $0.06, respectively (this difference is statistically significant at the 1% level).The finding suggests that as the optimal value of a message increased, subjects' valuations increased at a decreasing rate; the degree of conservatism increased (i.e., people underestimated more the implied probability of the message).

Appendix 9 (a) and (b) show the pooled WTP decisions across the rounds for each subject and highlight the proportion of rounds that each subject over-valued, undervalued or appropriately

[86] Difference is statistically significant at the 1% level.

valued the information relative to the BWTP benchmark for set 1 (graph a) and set 2(graph b). Comparisons of these two graphs illustrate how the increase to the informative power of the chip draw influenced WTP behavior.

Moving to Set 3 (rds. 9-12 & 21-24), the difference in the probability of the message predicting state 1 versus state 2 decreases from 0.40 in Set 1 and 0.52 in Set 2 to 0.20 (60-40), resulting in a less informative chip draw. However, at this same the state contingent payoffs also change. Specifically, the difference in the expected payoffs using the initial probability of being in either state (not allowing for the consideration of the message service) from selecting action A or B is larger than in prior sets, rising from $0.25 in Set 1 and Set 2 to $0.625 in Set 3.[87]As such, for this interval of rounds it is difficult to isolate the cause for change in subject behavior (i.e., what should be attributed to a change in the informative power of the chip draw versus a change in state contingent payoffs). During Set 3, the proportions of optimal WTP valuations rise relative to the second set of rounds to 73.5% from 45.7%. Subjects are 29.9ppts more likely to specify the BWTP amount in Set 3 (Table 4, equation 1, see Set 3 vs. Set 1).

The observation that subject WTP decisions were more reflective of Bayesian WTP behavior when the difference in the probability of the message coming from either state decreased in Set 3, is best explained by the experimental design peculiar to Set 3. In this set of rounds, the difference in state contingent payoffs is exaggerated such that the chip draw for the risk neutral BEU maximizer has no value (BWTP=0). Hence, a conservative individual who selects zero will also have behavior that is reflective of optimal decision theory. In table 4, equation 4, I control for the observations where the BWTP is equal to zero and I add an interaction term to isolate the observations where the BWTP is equal to zero specific to Set 3. The coefficient on Set 3 is now small and insignificant (2.9ppts versus 29.9ppts and significant with no control). There is no difference in BWTP behavior in Set 3 versus Set 1. Controlling for this same variable in Set 2 (adding the interaction term: BWTP=0 x Set 2) further emphasizes the impact that an increase in the informative power of the chip draw had on Set 2 observations relative to Set 1. Subjects

[87]$([.5(\$2.00) + .5(\$0.75)] - [.5(\$0.50) + .5(\$1.75)] = \$0.25\ in\ sets\ 1\ \&\ 2;\ [.5(\$2.00) + .5(\$0.75)] - [.5(\$0.50) + .5(\$1.00)] = \$0.62\ in\ set\ 3).).$

during Set 2 are 15.4ppts less likely than in Set 1 to specify the BWTP amount when the information had value.

As previously discussed, risk preferences may play a role in the subjects' under-estimation of WTP amount relative to the risk-neutral BEU benchmark. However, it does not explain why the subjects' WTP valuations increased at a decreasing rate when the optimal RN WTP increased. For example at the aggregate level in set 2 the RN WTP BEU benchmark is \$0.40 and subjects under-estimated this value and specified a WTP of \$0.27. Assuming that subjects risk preferences don't change over time, and using the same CRRA utility function of the form $u(c) = \frac{c^{1-\delta}}{1-\delta}$ and $\delta = 0.4$ used to predict the subjects' WTP behavior of \$0.19 in set 1, the predicted WTP value for set 2 should be \$0.33 versus \$0.27. Similarly, using the same CARA utility function of the form $u(c) = -e^{-\lambda c}$ and $\lambda = 1.1$ used to predict the subjects' WTP behaviour in set 1, the predicted WTP value for set 2 is again higher at \$0.34 versus \$0.27.

There is no support for Hypothesis 2. Regardless of the risk preferences used in this study to determine the optimal WTP value (RN, CRRA or CARA), as the informative power of the chip draw increased the frequency of BEU WTP decisions decreased and subjects on average underestimated the informational value of the message to a greater degree. However, given the experimental design peculiar to Set 3 (the simultaneous change to the distribution of the colour chips within each bag and state contingent payoffs change) there is insufficient evidence to conclude the reverse statement to be true. That is, as the informative power of the chip draw decreased it is not necessarily the case that subjects are more likely to specify the BWTP value. However given proposition 1 (i.e., the higher the prior confidence in initial beliefs the closer the posterior probabilities will be to the prior probabilities for any given message) and the non-rejection of hypothesis 1 (i.e., subjects underestimated the value of the message received), in conjunction with the fact that the distribution of chips at 60/40 (R/B,B/R) was close to the prior probabilities of 50/50 , it may not be surprising that subjects were more accurate in valuing the information during these rounds than in other rounds. However, this still requires further investigation.

In the future, an experimental design where a treatment exists that holds the state contingent payoffs constant across all rounds while only varying the distribution of colour chips contained within each bag would be beneficial. The behavior of these subjects can be benchmarked against a treatment group where both state contingent payoffs and the distribution of colour chips remain constant. Similarly, a third treatment should be introduced where the distribution of colour chips remains constant and only the state contingent payoffs change. This design would enable observations of subject behavior when only one component of the BEU decision rule changes (Bayes Law or expected utility theory). The design could give insight in the subjects' ability to perform both, one or none of the two components of optimal decision theory. It would also allow us to test further the increase in conservatism in subject WTP behaviour when the message service became more informative.

H3 predicted that subjects' will have the same 2nd choice BEU inconsistency rate whether they specified a willingness to pay that was consistent with the Bayesian WTP amount or whether the WTP implied a posterior probability that fell within the range of posterior probabilities as defined by the critical posterior probability values (see Table 3). From Table 7, hypothesis 3 is rejected. Subjects' mean 2nd choice BEU inconsistency rate is 4.7% when the optimal WTP amount is specified and is 21.6% when subjects specified a WTP amount which was not optimal but implied an estimated posterior probability which fell within the range of posterior probability values (see table 3), that if combined with expected utility theory would still result in the optimal action choices[88]. Additionally from Table 4, eqn. 1, subjects are 43ppts less likely to commit a 2nd choice BEU inconsistency when they have identified a WTP amount which is consistent with the BWTP benchmark.[89] Of course, this change decreases (however, still large and significant) to ≅ 25ppts less likely when we control for the BWTP values that are equal to zero (see eqns. 2-4, Table 4).

[88] This difference is statistically significant at the 1% level

[89] The BEU 2nd choice is made in advance of the subject specifying her WTP price to use the information. A subject who specified the optimal choice presumably used Bayes law in conjunction with EU Theory to arrive the decision choice. As such, this subject should also specify the BWTP amount.

These results suggest that subjects who indicate a WTP value that is not optimal but implies a posterior probability that falls within the acceptable range of posterior probabilities, as defined by the switching rule (see critical value theory, hypothesis 3), are less likely to combine these posterior probabilities with the Expected Utility portion of the BEU decision rule than subjects who indicate a WTP value consistent with the BWTP benchmark.

Table 7: 2nd Choice BEU Inconsistency Rates conditional on WTP

	WTP≠BWTP WTP→$\widehat{\pi_{S.M}}$ → BEU action	WTP=BWTP WTP→$\pi^*_{S.M}$ → BEU action	OTP	FREE
Number of Obs.	546	865	1452	2868
BEU Inconsistency Rate	.216	.047	.178	.147
S.D.	.412	.213	.383	.354

H4 predicted that when subjects must pay for the information rather than observe information for free, they will deviate less from the BEU optimal decision rule when taking their final choice. Table 7 shows the 2nd choice BEU inconsistency rates for subjects who performed the FREE task versus the OTP task. This hypothesis is also rejected. Subjects' second choice BEU inconsistency rate is 14.7% when subjects perform the FREE message task and is 17.8% when performing the OTP message task. A Wilcoxon rank test confirms the difference in 2nd choice BEU inconsistency rates to be statistically significant. Chapter 2, page 23, table 12 highlights that subjects are 3.4ppts more likely to commit a BEU 2nd choice inconsistency when performing the OTP versus the FREE message task.

The potential reasons for this difference are threefold (Chapter 2, page 23): 1) This difference could represent a lack of commitment to the accuracy of the second choice on behalf of the subject when they are required to pay in order to have this choice used to calculate earnings. In this case, the lack of commitment in decision quality of second action choices could result in a WTP amount of zero or a relatively small WTP bid even though the subjects' second choice is different from their first choice. Consequently, these subjects would have a BEU 2nd choice inconsistency rate that is higher than for those who specified a more substantial WTP amount.

Overall, the BEU 2nd choice inconsistency rate is 24.2% when the second action is different from the first action choice during the OTP message task. This inconsistency rate drops to 20.6% when the WTP is greater than $0.05 and increases substantially to 41.7% when the WTP amount is less than $0.05. Although this result supports 'a lack of effort toward decision accuracy in the face of an additional cost' hypothesis, subjects who specified a WTP less than $0.05 only accounted for 16% of the total observations where the first choice did not equal second choice actions; 2) Subjects have a tendency to maintain their previous decision choice even when new information is acquired that indicates that this decision choice is no longer optimal more often when performing the OTP decision task. Samuelson & Zeckhauser (1988) refer to this as a 'status quo bias'. Table 11, column 2, in chapter 2 shows the predicted probability of a 2nd choice BEU inconsistency increased 12.8ppts when subjects were required to change their initial action to the alternative action conditional on the chip draw in order to follow the BEU model predictions. The 2nd choice BEU inconsistency rate when first and second action choices are aligned (subjects are maintaining the status quo) is 11.7% when subjects perform the FREE message task and 15% when subjects perform the OTP message task. A two sample t-test confirms that this difference is statistically significant at the 1% level. The logic, given this evidence, is that subjects maintain their first decision choice due to a status quo bias which is accentuated when there is an added cost (WTP decision requirement); And, 3) it possible that the change imposed by the additional step required to complete the OTP versus the FREE task created confusion or additional complexity to the decision environment. This confusion led to an increase in the likelihood of a 2nd choice BEU inconsistency.

H5 stated that providing subjects with the posterior probabilities conditional on the message received (the Bayes law heuristic), will result in more BEU accurate valuations. From the between effects model found in Table 5, subjects who are provided with the posterior probability of being in either state (informed) did not perform any better relative to the BWTP benchmark than subjects who were left to potentially calculate this statistic on their own (uninformed). There is evidence from this chapter and chapter 2 that suggests that subjects who were provided

with the math component of the Bayes law calculation did no better than other subjects relative to the risk-neutral BEU benchmark.[90]

Hershman & Levine (1970) also found that, even when subjects were provided with the Bayes law calculation they did not correspondingly optimize their behavior. Hershman & Levine did not find any group differences in correct responses between subjects who were given the calculated Bayes updating information versus those that were just given sufficient parameters to make the calculations themselves. There are several conjectures as to why this may be the case. For example, the uninformed group made accurate subjective probability assessments on their own thus matching the informed group's decisions. Alternatively, perhaps none of the participants behave according to models based on probabilities and commonly choose to adapt a different heuristic when making decisions (Chapter 2).

3.5 Conclusions and Further Investigation Requirements

After eliminating a few large outliers, the subjects' willingness to pay to use the additional information gathered form an imperfect message service when making a final decision was on average less than the BEU willingness to pay benchmark. This finding was consistent with several previous studies (Kagel & Roth, 1995; Peterson & Dechane, 1967; Edwards, 1968; Sanders, 1968; Eger & Dichaut, 1982; McKelvey & Page, 1990; Harrison et al., 2010). Hirshleifer & Riley (1992) suggest that this under-estimation of the informational value of the message received is a result of subjects having a higher confidence in their prior beliefs (before observing a message). Additionally, risk preferences may have contributed to the under-estimation WTP behavior, however, this study found that as the informative value of the message increased, causing the BEU valuation to increase, subjects under-estimated the value of the message signal to a greater degree; the degree of conservatism increased. It seems unlikely that the subjects increased their confidence level in prior beliefs as a result of an increase to the informative power of the message. It seems even more unlikely that the subjects risk preferences changed between rounds. Therefore in addition to risk aversion, a potential explanation that may account for increased conservatism in WTP values when the message became more informative

[90] In chapter 2, subjects did not select more often the BEU action choice when they were informed.

may be loss aversion. Tversky & Kahneman (1984) suggest that losses are two time more psychologically powerful than gains. Therefore, subjects may prefer to avoid the losses associated with the cost of using the information over the acquired gains from the informational knowledge the message provided. In set 2 the value of the information service increased, and as such, the BWTP also increased now representing a greater loss to subjects if specified. Recall that subjects did recognize that the information service for these rounds was more valuable, however they did not increase their WTP amounts at the same rate of change as the optimal value.

Additionally this study found that subject WTP decisions that implied the Bayes law posterior probability estimates were correlated with terminal actions that were also reflective of BEU behavior. These subjects appear to be capable of combining Bayes law estimates in conjunction with Expected Utility Theory. However, subjects' WTP decisions that do not imply the Bayes law posterior probability estimates are more likely to be associated with behavior that is less reflective of the BEU benchmark. This non-optimal behavior is also associated with subjects who specified a WTP value that implied a posterior probability that when combined properly with Expected Utility Theory would still result in an optimal choice. This result suggests behavior reflective of an all or nothing decision rule, that is, subjects either apply both or neither component of the BEU decision (Bayes law and Expected Utility Theory) when taking a terminal action.

A possible explanation to why subjects do not value the message optimally is that they lack the math skills necessary to complete the BEU decision task. By providing the subjects with the Bayes law posterior probability estimates conditional on the message received the math component of the BEU decision rule is simplified. However, even with these estimates, the frequency of subjects who made a conservative estimate of the information signal relative to the BEU benchmark did not change, nor did it overall improve optimal decision behavior (Chapter 2 & 3). It appears that subjects in some decision environments do not use Bayes law when making decisions and, as highlighted in this study and previously by Hershman & Levine (1970), subjects may adapt a different heuristic all together.

For example in this study in addition to a 'conservatism' bias when evaluating the message service, there is evidence of subjects using past statistically independent outcomes in future decision choices. Specifically from the regressions found in table 4, when subjects receive payment on their 2nd action choice in the prior round subjects are 36.4ppts less likely to select a WTP value that is optimal in future rounds. A possible explanation again could be loss aversion. Subjects who were paid on their first action choice during the OTP rounds received a mean payoff of $1.41 per round versus $1.21 on their second choice. This is compared to a mean payoff of $1.43 per round on the first choice versus $1.48 on the second choice when these same subjects were participating in the FREE message task. Additionally, Table 5 (eqn. 5) highlights that subjects who were categorized as a reinforcement learner versus a theorist based on the results from the post-survey Honey & Mumford questionnaire (see chapter2) were 6.4ppts less likely to specify the optimal WTP value.

Another question that is worth further investigation is how risk attitudes influence subject behaviour. In this experiment, risk attitudes may have influenced the subjects' WTP decisions. An attempt was made to determine the risk attitudes of the subjects at the end of the experiment through the administration of the Eckel-Grossman risk task (Eckel-Grossman , 2002). The results showed that a majority of the subjects participated in the most risky lottery choices.[91] Harrison et al. (2010) conducted an experiment where they administrated the Holt Laury risk aversion test to a sub-set of subjects prior to the experiment and to the remaining sub-set of subjects after the tasks of the experiment were complete. They found more risky behavior by subjects when they completed the test after the experiment, and more expected behavior (according to theoretic predictions) when they completed the task before the experiment. Although the result suggests that conducting a test for risk-aversion at the beginning of the experiment may be more representative, they warn that this test could inappropriately frame the context of the decision environment for the pending experiment. Regardless of the subjects high risk preferences, results from the between effects model found in Table 5, show no statistical significance for the risk attitudes of subjects (as defined by the Eckel-Grossman risk test) in influencing subject behavior.

[91] Subjects select one of 10 lottery choices, prior to a random draw which determined the outcome. The lottery choices are ranked in order with 1 representing the highest risk aversion and 10 representing the least. 57% of the subjects selected the lottery choice above level 8.

Furthermore, the difference between the reported WTP by subjects versus the BWTP was not significantly statistically different for the subjects who were classified by this test as less risk averse than those subjects classified as more risk averse.

Finally, from the between effects model found in Table 5, subjects are 6.7ppts more likely to specify a WTP amount that was consistent with the BWTP optimal when they had experience with the FREE message task in advance of the OTP message task. Also, students who are enrolled in a math or economics undergraduate program were 4.5ppt more likely to specify a WTP that was consistent with this benchmark. Both findings suggest that subject's WTP behavior improves relative to optimal decision theory with experience (i.e. through repeated trials and with math optimization training).

Further research should be conducted to further separate the key components of the BEU decision rule in an effort to provide insight in the subjects' ability to perform both, one or none of the two components of optimal decision theory. Additionally, from both chapters 2 and 3 there is evidence that suggests that in addition to decision behavior being influenced by the context and environment in which the decision is being made; it is also influenced by subject type. Post-experiment questionnaires designed to identify individual type in more detail would be beneficial.

Appendix 1: Proof of dominance of truthful WTP bidding

The dominant strategy for each subject is to specify their truthful valuation of the message received.

Let $TWTP_i$ be subject i's true value for the message received,
Let WTP_i be subject i's bid to use the information to determine the terminal choice

$$\boldsymbol{The\ Payoff\ for\ subject\ i} \begin{Bmatrix} \boldsymbol{TWTP_i - Random\ Price(RP) drawn\ \ if\ TWTP_i \geq RP} \\ \boldsymbol{0} \qquad\qquad\qquad\qquad \boldsymbol{0therwise} \end{Bmatrix}$$

The strategy for specifying a WTP amount that is greater than a subject's true valuation is dominated by bidding truthfully. Assume that $WTP_i > TWTP_i$

$$\begin{gathered} \boldsymbol{If\ RP <\ TWTP_i}\ then\ subjects\ 2nd\ action\ choice\ is\ recognized\ for\ payment\ . \\ The\ WTP\ amount\ does\ change\ the\ payoff\ so\ the\ payoff \\ is\ the\ same\ whether\ the\ bid\ was\ truthful\ or\ an\ over-bid \\ \boldsymbol{If\ RP > WTP_i}\ then\ subject\ loses\ either\ way, whether\ his\ bid\ were\ truthful \\ or\ it\ represented\ an\ over-bid \\ \boldsymbol{If\ TWTP_i < RP \leq WTP_i}\ \ then\ only\ the\ over-bid\ wins\ the\ auction. \\ The\ payoff\ to\ the\ subject\ would\ be\ negative\ from\ over-bidding\ as \\ RP - TWTP_i > 0\ as\ subject\ pays\ the\ random\ price \\ And\ the\ payoff\ from\ truthful\ bidding\ would\ be\ 0 \\ as\ RP > TWTP_i\ and\ subject\ would\ lose\ bid \end{gathered}$$

Therefore, the strategy of bidding higher than one's truthful valuation is dominated by the strategy of truthful bidding.

The strategy of specifying a WTP that is less than a subject's true valuation is dominated by bidding truthfully. Assume $WTP_i < TWTP_i$

$$\begin{gathered} \boldsymbol{If\ RP > TWTP_i}\ then\ subject\ loses\ the\ 2nd\ action\ choice \\ being\ recognized\ for\ payment\ . \\ The\ truthful\ WTP\ amount\ as\ well\ as\ the\ under-\ bid\ receive\ zero\ payoff \\ So\ the\ payoff\ is\ the\ same\ whether\ the\ subject\ specifies\ the\ true \\ WTP\ or\ under-bids\ this\ amount \\ \boldsymbol{If\ RP < WTP_i}\ then\ subject\ wins\ the\ bid\ whether\ his\ WTP\ amount\ was\ truthful \\ or\ whether\ it\ was\ less\ than\ the\ truthful\ value \\ \boldsymbol{If\ WTP_i < RP \leq TWTP_i}\ \ then\ only\ the\ strategy\ of\ specifying\ the\ truthful\ value \\ wins\ the\ auction.\ The\ payoff\ from\ truthful\ bidding\ would\ be\ zero \\ or\ greater\ than\ zero\ given\ that\ TWTP_i \geq RP \end{gathered}$$

Therefore, the strategy of bidding lower than one's truthful valuation is dominated by the strategy of truthful bidding.

Appendix 2: Exogenous Parameters by Round Set

SET	1	2	3
Rounds	1-4 13-16	5-8 17-20	9-12 21-24
State Contingent Payoffs			
Action A **Bag 1 revealed, $C(A,S_1)$** **Bag 2 revealed, $C(A,S_2)$**	 $2.00 $0.75	 $1.75 $0.50	 $1.00 $0.50
Action B **Bag 1 revealed, $C(B,S_1)$** **Bag 2 revealed, $C(B,S_2)$**	 $0.50 $1.75	 $0.75 $2.00	 $0.75 $2.00
Initial Beliefs **Bag 1/Bag 2 (π_1/ π_2)**	 .5/.5	 .5/.5	 .5/.5
BEU decision rule prior to a message signal(chip draw)	**Action A**	**Action B**	**Action B**
State Characteristics **Total chips bag 1** **# red chips($q_{1.1}$)** **# blue chips ($q_{2.1}$)** **Total chips bag 2** **# red chips($q_{1.2}$)** **# blue chips ($q_{2.2}$)**	 50 35(.70) 15(.30) 50 15(.30) 35(.70)	 50 12(.24) 38(.76) 50 38(.76) 12(.24)	 50 20(.40) 30(.60) 50 30(.60) 20(.40)
Bayes Law Posterior Probabilities **$\pi_{1.1}$** **$\pi_{2.1}$** **$\pi_{1.2}$** **$\pi_{2.2}$**	 **.70** **.30** **.30** **.70**	 **.24** **.76** **.76** **.24**	 **.40** **.60** **.60** **.40**
BEU decision rule After a message signal is received (chip draw)	**If Red : Action A** **If Blue: Action B**	**If Red : Action B** **If Blue: Action A**	**If Red : Action B** **If Blue: Action B**

Appendix 3: The BEU Expected Value of Information (the worth of the message service)

$$EV_m = EU(a_B; \pi_{s.m}) - EU(a_A; \pi_{s.m})$$

For Rounds 1-4 and 13-16 when the first action is A and a blue message is received

$$EV_2 = [\pi_{1.2}c(B, s_1) + \pi_{2.2}c(B, s_2)] - [\pi_{1.2}c(A, s_1) + \pi_{2.2}c(A, s_2)]$$

$$EV_2 = \left[\frac{.3(.5)}{.3(.5) + .7(.5)}\$0.50 + \frac{.7(.5)}{.3(.5) + .7(.5)}\$1.75\right] - \left[\frac{.3(.5)}{.3(.5) + .7(.5)}\$2.00 + \frac{.7(.5)}{.3(.5) + .7(.5)}\$0.75\right]$$

$$EV_2 = [(.3)\$0.50 + (.7)\$1.75] - [(.3)\$2.00 + (.7)\$0.75] = \$0.25$$

For Rounds 1-4 and 13-16 when the first action is A and a Red message is received

$$EV_1 = [\pi_{1.1}c(A, s_1) + \pi_{2.1}c(A, s_2)] - [\pi_{1.1}c(A, s_1) + \pi_{2.1}c(A, s_2)] = \$0.00$$

For Rounds 1-4 and 13-16 when the first action is B and a Blue message is received

$$EV_2 = [\pi_{1.2}c(B, s_1) + \pi_{2.2}c(B, s_2)] - [\pi_{1.2}c(B, s_1) + \pi_{2.2}c(B, s_2)] = \$0.00$$

For Rounds 1-4 and 13-16 when the first action is B and a Red message is received

$$EV_1 = [\pi_{1.1}c(A, s_1) + \pi_{2.1}c(A, s_2)] - [\pi_{1.1}c(B, s_1) + \pi_{2.1}c(B, s_2)]$$

$$EV_1 = [(.7)\$2.00 + (.3)\$0.75] - [(.7)\$0.50 + (.3)\$1.75] = \$0.75$$

For Rounds 5-8 and 17-20 when the first action is A and a blue message is received

$$EV_2 = [\pi_{1.2}c(A, s_1) + \pi_{2.2}c(A, s_2)] - [\pi_{1.2}c(A, s_1) + \pi_{2.2}c(A, s_2)]$$

$$EV_2 = [(.24)\$0.50 + (.76)\$1.75] - [(.24)\$0.50 + (.76)\$1.75] = \$0.00$$

For Rounds 5-8 and 17-20 when the first action is A and a Red message is received

$$EV_1 = [\pi_{1.1}c(B, s_1) + \pi_{2.1}c(B, s_2)] - [\pi_{1.1}c(A, s_1) + \pi_{2.1}c(A, s_2)]$$

$$EV_1 = [(.76)\$2.00 + (.24)\$0.75] - [(.76)\$0.50 + (.24)\$1.75] = \$0.90$$

For Rounds 5-8 and 17-20 when the first action is B and a Blue message is received

$$EV_2 = [\pi_{1.2}c(A, s_1) + \pi_{2.2}c(A, s_2)] - [\pi_{1.2}c(B, s_1) + \pi_{2.2}c(B, s_2)]$$

$$EV_2 = [(.76)\$1.75 + (.24)\$0.50] - [(.76)\$0.75 + (.24)\$2.00] = \$0.40$$

For Rounds 5-8 and 17-20 when the first action is B and a Red message is received

$$EV_1 = [\pi_{1.1}c(B, s_1) + \pi_{2.1}c(B, s_2)] - [\pi_{1.1}c(B, s_1) + \pi_{2.1}c(B, s_2)] = \$0.00$$

For Rounds 9-12 and 21-24 when the first action is A and a blue message is received

$$EV_2 = [\pi_{1.2}c(A, s_1) + \pi_{2.2}c(A, s_2)] - [\pi_{1.2}c(A, s_1) + \pi_{2.2}c(A, s_2)]$$

$$EV_2 = [(.40)\$0.75 + (.60)\$2.00] - [(.40)\$1.00 + (.60)\$0.50] = \$0.80$$

For Rounds 9-12 and 21-24 when the first action is A and a Red message is received

$$EV_1 = [\pi_{1.1}c(B, s_1) + \pi_{2.1}c(B, s_2)] - [\pi_{1.1}c(A, s_1) + \pi_{2.1}c(A, s_2)]$$

$$EV_1 = [(.60)\$0.75 + (.40)\$2.00] - [(.60)\$1.00 + (.40)\$0.50] = \$0.45$$

For Rounds 9-12 and 21-24 when the first action is B and a Blue message is received

$$EV_2 = [\pi_{1.2}c(B, s_1) + \pi_{2.2}c(B, s_2)] - [\pi_{1.2}c(B, s_1) + \pi_{2.2}c(B, s_2)] = \$0.00$$

For Rounds 9-12 and 21-24 when the first action is B and a Red message is received

$$EV_1 = [\pi_{1.1}c(B, s_1) + \pi_{2.1}c(B, s_2)] - [\pi_{1.1}c(B, s_1) + \pi_{2.1}c(B, s_2)] = \$0.00$$

Appendix 4: Logit and OLS Regression Equations & Descriptive Summary of Variables

	Equation 1, 2, 3 & 4 **Logit FE & RE**	**Equation 5, 6 & 7** **OLS BE**
Dependent Variable	**WTP is BEU Optimal** **1** Optimal (-.05<WTP<.05) **0** Non-Optimal WTP	**Count data** **WTP is BEU Optimal** by subject all rounds
Explanatory Variables: vary over rounds, are the same for all individuals	**Experience** **1** second 12 rounds **0** first 12 rounds	**Experience**
	OTP **1** OTP message task **0** Free message task	**N/A**
	Informed **1** subjects given Posterior Prob. **0** subjects not given Posterior Probabilities	**Informed**
	Set 2-rds. 1-4&13-16 Difference in proportion of red to blue chips change from 70-30=40 in set 1 to 76-24=52 in set 2 **A more informative chip draw**	**N/A**
	Set 3-rds. 5-8&17-20 Difference in proportion of red to blue chips change to 60-40=20 **A less informative chip draw** Difference in expected Payoffs using priors increases to $0.62 from $0.25	**N/A**
Explanatory Variables: vary over rounds & between subjects	**Paid Second in Prior Rd.** **1** payment on 2nd action **0** payment on 1st action	**Paid Second in Prior Rd.**
	BEU action not consistent with the higher payoff state in prior round **1** Inconsistent **0** Consistent	**BEU action not consistent with the higher payoff state in prior round** **1** Inconsistent **0** Consistent
	BWTP=0	**BWTP =0**
	BWTP=0 x Set 2	**N/A**
	BWTP=0 x Set 3	**N/A**
Explanatory Variables: same over all rounds but vary by individual	**Female** **1** Female **0** Male	**Female**
	English Second **1** English 2nd language **0** English 1st language	**English Second**
	Post Survey **1** Classified as RL **0** Classified as theorist	**Post Survey**
	Risk Aversion Continuous-Eckel-Grossman test 1 highest RA to 10 least RA	**Risk Aversion**
	Econ Math **1** Math/econ/optimization **0** non-math student	**Econ Math**

Appendix 5: Histogram

WTP –BWTP all observation

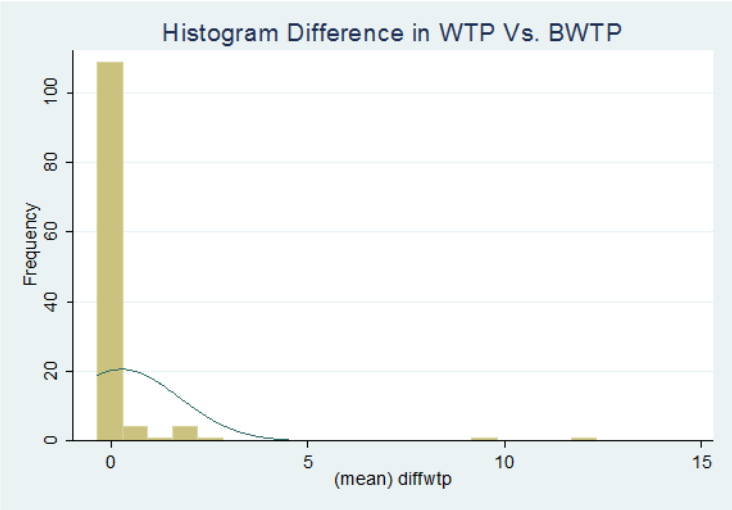

Appendix 6: Kernel Density Estimates

(a)WTP –BWTP all observation;

(b) WTP-BWTP 25 outliers dropped

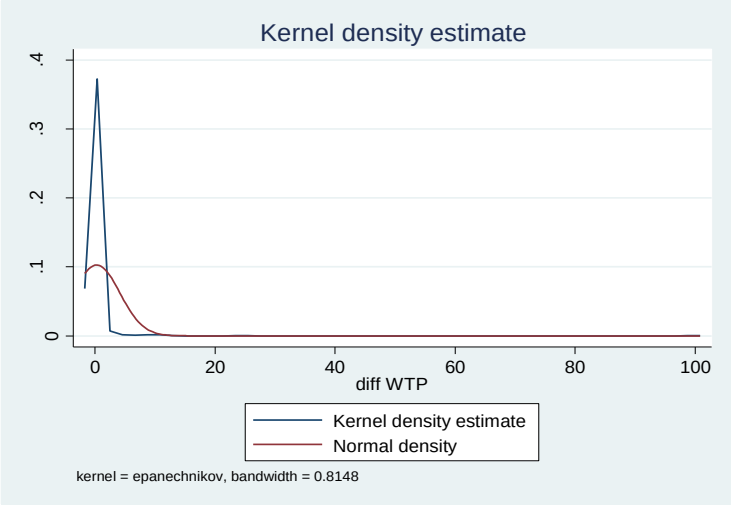

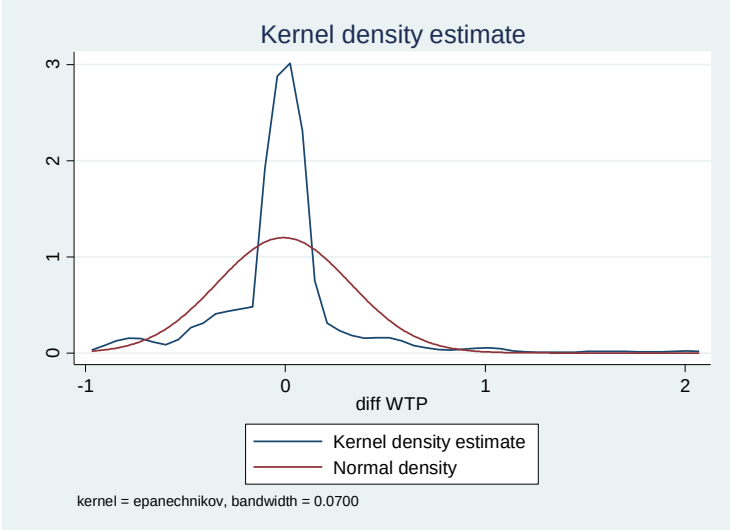

Appendix 7: Box Graph

(a)WTP –BWTP all observation;

(b) WTP-BWTP 25 outliers dropped;

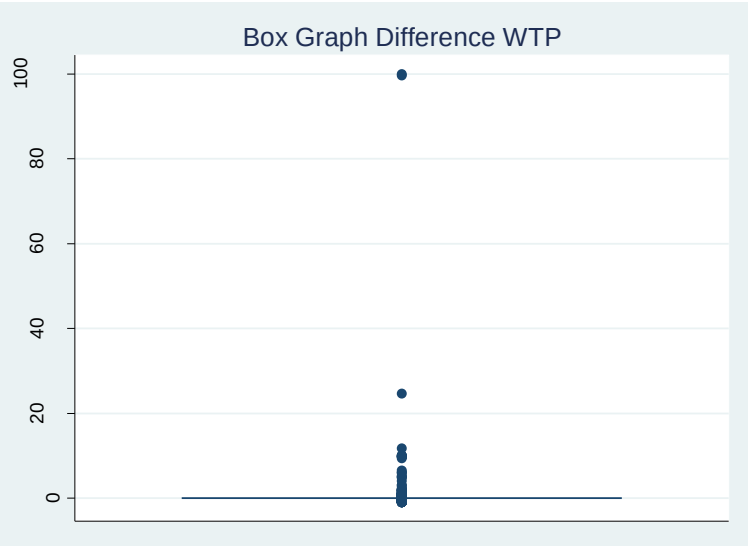

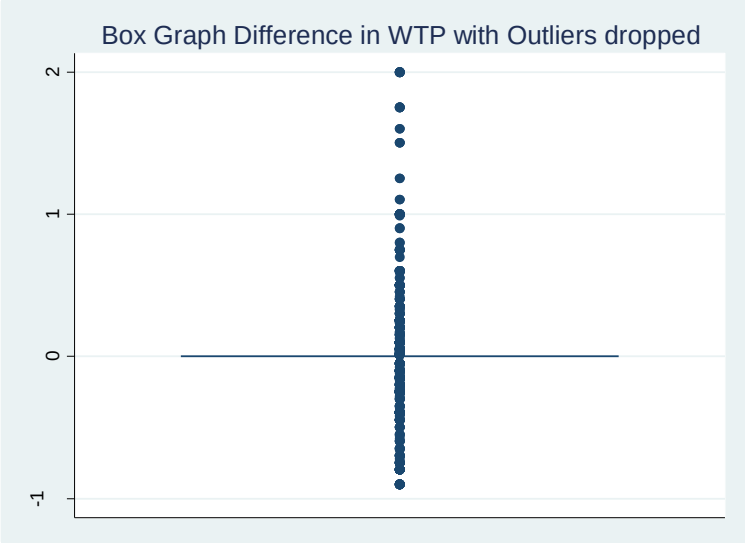

Appendix 8: Normal Probability graph (p-norm)
(a)WTP –BWTP all observation; (b) Restricted WTP-BWTP;

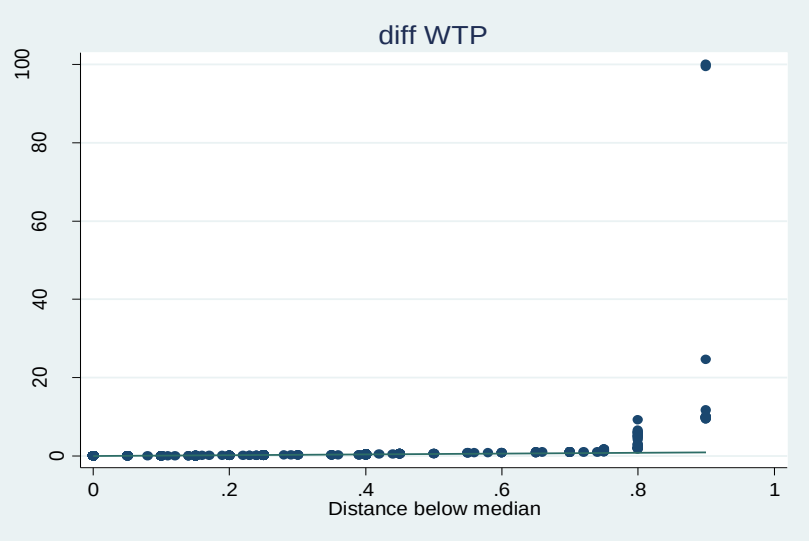

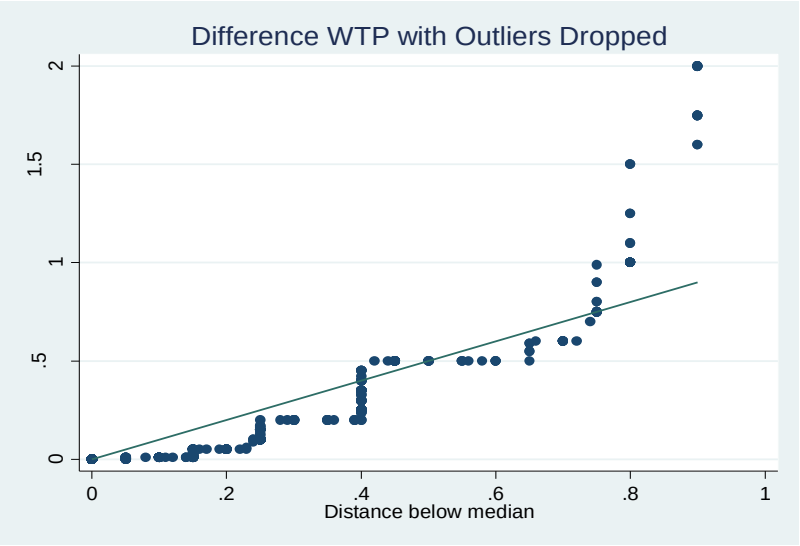

Appendix 9: Area Plots of the Proportion of subject WTP decisions which under, over-estimated or were optimal relative to BWTP

(a)

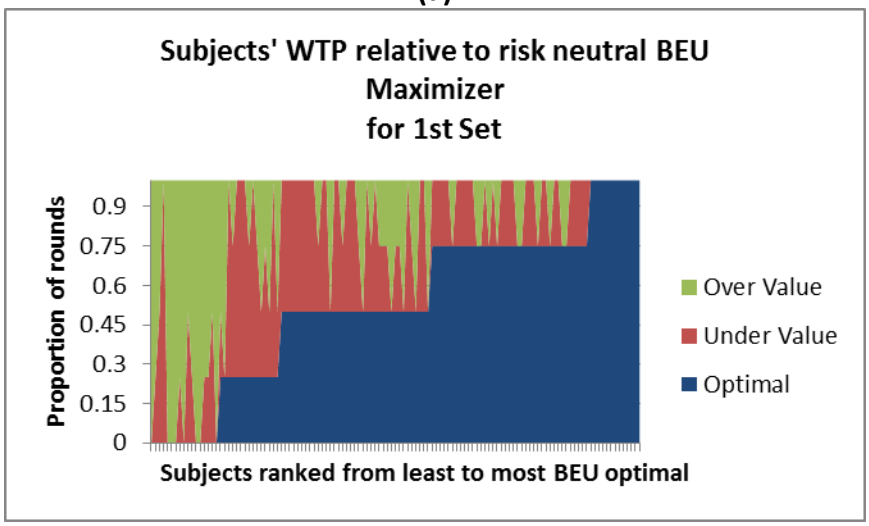

(b)

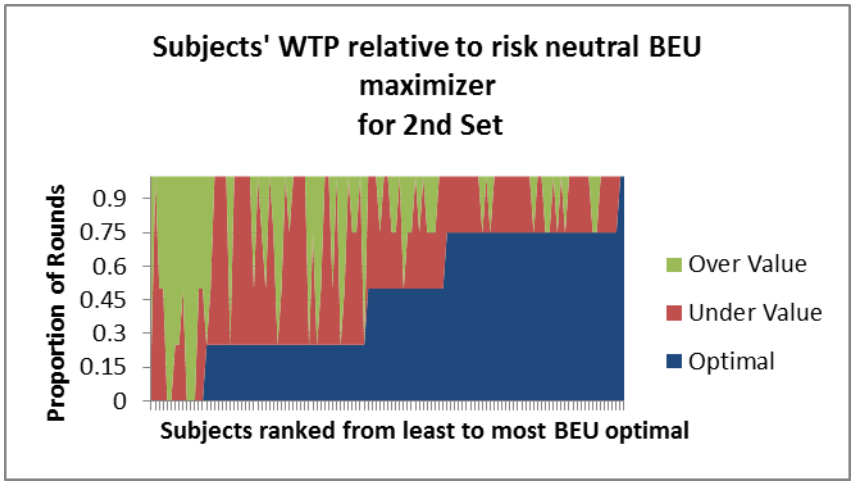

Appendix 10: RE Logit Regression

Variable	(5) Marginal Effects Subject WTP= -$0.05<BWTP<$0.05 logit, RE
Experience	**-0.026**
No Experience with the FREE Message Task	0.023
Informed	**0.032**
Subject was provided with the Bayes law Posterior Probability Calculation	0.024
Message Received is Blue	**-0.151*** **
Usually associated with BWTP>$0.00	0.017
Set 2 (rds. 5-8 & 17-20) Vs. Set 1 (rds. 1-4 & 13-16)	**-0.044*** **
°Informative Power of chip draw increases	0.024
°State contigent payoffs are constant	
Set 3 (rds. 9-12 & 21-24) Vs. Set 1 (rds. 1-4 & 13-16)	**0.216*** **
°informative Power of chip draw decreases	0.024
°State contigent payoffs change	
2nd Choice BEU Inconsistency	**-0.355*** **
	0.020
Subject's 2nd action is recognized for payment	**-0.316*** **
WTP≥Random Price drawn	0.017
BEU 2nd action not consistent with the higher payoff state in prior round	**-0.027**
	0.019
Female	**-0.011**
	0.025
English is the Second Language	**-0.071*** **
	0.027
Reinforcement Learner	**-0.016**
Post-experiment Survey identified subject as a Reinforcement learner versus Theorist	0.027
Risk Aversion Score	**0.000**
Holt Laury Risk aversion Lottery choice (1-10)	0.004
Age	**0.006**
	0.007
Math Economics Student	**0.056*** **
Student has some math optimization training	0.027
Obs.	**1452**
Log Likelihood	**-523.778**
R-sq	

*****p-value≤.01<**p-value≤.05<*p-value<.10**

(The smaller numbers beneath the coefficients represent standard errors).

References

Antoniou, C., Harrison, G., Lau, M. & Read, D. 'Subjective Bayesian Beliefs.' Working Paper, January 2010.

Arrow, K. J., 'Essays on the Theory of Risk Bearing'. Publisher: Markham Pub. Co., Chicago, 1971.

Beach, Barbara, 'Expert Judgment About Uncertainty; Bayesian Decision Making in Realistic Settings', Organizational Behaviour and Human Performance, 1975, 14, page 10-59.

Börgers, T., Sarin, R. ' Learning Through Reinforcement and Replicator Dynamics.' Journal of Economic Theory, 1997, 77, p.p. 1-14.

Brenner, T. 'Agent Learning Representation: Advice on Modeling Economic Learning.' In: The Handbook of Computational Economics, 2005, 2006, 2, p.p. 895-947, Elsevier.

Camerer, C.F. '*Individual Decision Making*' in: Kagel, J.H. & Roth, A.E. (Eds.): The Handbook of Experimental Economics, 1995, 588–703.

Camerer, Colin& Ho, Teck-Hua. 'Experience-Weighted Attraction Learning in Normal Form Games.'Econometrica, 1999, 67(4), pp. 827-74.

Casey, B. J., Tottenham, N., Liston, C., Durston, S., ' Imagining the Developing Brain: What have we learned about Cognitive Development.' Trends in Cognitive Science, 2005, 3, p.p. 104-110.

Charness, Gary, and Dan Levin,. 'When Optimal Choices Feel Wrong: A laboratory Study of Bayesian Updating, Complexity and Affect.' American Economic Review, 2005, 95, pp. 1300-1309.

Choi, J., Laibson, D., Madrian, B., Metrick, A., 'Reinforcement Learning and Savings behavior.' Working Paper, Dec. 2008. Harvard/Yale University.

Cousineau, Denis, and Sylvain Chartier 'Outliers Detection and Treatment: a review.' International Journal of Psychological Research, 2010, 3(1), pp. 58-67.

Davis, D., Holt, C., 'Chapter 8.' In: Experimental Economics, 1993, University Press, Princeton, New Jersey.

Ducharme, W. 'Response Bias Explanation of Conservative Human Inference.' Journal of Experimental Psychology, 1970, 85, p.p. 66-74.

Eckel, C., Grossman, P. ' Forecasting risk attitudes: An experimental study using actual and forecast gamble choices.' Journal of Economic Behaviour and Organization, Volume 68, 2008, 68, p.p. 1-17.

Eger, C., and Dickhaut, J. 'An Examination of the Conservative Information Processing Bias in an Accounting Framework.' Journal of Accounting Research, 1982, 20,p.p. 711-723.

Edwards, W., 'Conservatism in human information processing.' In B. Kleinmuntz (Ed.), Formal representation of human judgment , 1968, pp. 17-52. New York: Wiley.

Edwards, Ward, 'Optimal Strategies for Seeking Information: Models for Statistics, Choice Reaction Times, and Human Information Processing', Journal of Mathematical Psychology, 1965, 2, 312-329.

Edwards, W. & Slovic, P. Seeking information to reduce the risk of decisions. *American Journal of Psychology,* 1965, 78, 188-197.

Erev, Ido & Roth, Alvin E. 'Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria.' Amercian Economic Review, 1998, 88(4), pp. 848-81.

Erev, Ido & Roth, Alvin E. 'Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term.' Games and Economic Behavior, 1995, 8, pp. 164-212.

Feldman, M., and March, J. 'Information in Organizations as Signal and Symbol.' Adminstrative Science Quarterly, 1981, 26, p.p. 171-186.

Feltovich, N. *'Reinforcement-Based vs. Belief-Based Learning Models in Experimental Asymmetric-Information Games',* Econometrica, 2000, Vol. 68, No. 3, pp. 605-641.

Fried, L. S. & Peterson, C. R. 'Information seeking: Optional versus fixed stopping.' *Journal of Experimental Psychology,* 1969, 80, 525-529.

Gabaix, X., Laibson, D., Guilliermo, M., & Weinberg, S. 'Costly Information Acquistion: Experimental Analysis of a Bounded Rational Model.' American Economic Review, Sept. 2006, p.p. 1043-68.

Grether, David M. 'Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence.' Journal of Economic Be-havior and Organization, 1992, 17(1), pp. 31-57.

Grether, David M. 'Bayes Rule as a Descriptive Model: The Representativeness Heuristic.' Quarterly Journal of Economics, 1980, 95(3), pp. 537-57.

Griffin, D., Tversky, A. 'The Weighing of Evidence and the Determinants of Confidence.' Cognitive Psycology, 1992, 24, pp. 411-35.

Hammerton, M., 'A case of radical probability estimation.' Journal of Experimental Psychology ,1973, 101, 252–254.

Harmeling, S., Dornhege, G., Tax, D., Meinecke, F., Muller, K.R. 'From Outliers to Prototypes: Ordering data.' Neurocomputing, 2006, 69, p.p. 13-16..

Heath C., and A. Tversky. 'Preference and belief: Ambiguity and competence in choice under uncertainty. Journal of Risk and Uncertainty', 1991, 4:5-28.

Hershman, R. L. & Levine, J. R. 'Deviations from optimum information-purchase strategies in human decision-making.' *Organizational Behavior and Human Performance,* 1970, 5, 313-329.

Hirshleifer, J., Riley, J. 'The Analytics of Uncertainty and information.' Cambridge Surveys of Economic Literature, Cambridge University Press, 1992.

Holt, C., & Smith, A.,M. ‘An Update on Bayesian Updating.’ Journal of Economic Behaviour & Organization, 2009, 69, pp. 125-134.

Holt, C.A., Laury, S.K., ‘Risk aversion and incentive effects.’ American Economic Review, 2002, 92, 1644–1655.

Holt, C.A. ‘Scoring rules for eliciting subjective probability and utility functions.’ 1986, In: Goel, P.K., Zellner, A. (Eds.), Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti. North-Holland, Amsterdam, pp. 279–290.

Hopkins, Ed, *‘Two Competing Models of How People Learn in Games’*, Econometrica, 2002, Vol. 70, No. 6, pp. 2141-2166.

Kahneman, D. & Tversky, A. ‘Subjective Probability: A Judgment of Representativeness.’ Cognitive Psychology, 1972, 3(3), pp. 430-54.

Kahneman, D. & Tversky, A. ‘On the psychology of prediction.’ *Psychological Review*, 1973, 80, 237-251.

Kahneman, D. & Tversky, A. ‘Rational Choice and the Framing of Decisions.’ The Journal of Business, 1986(1984), 59, p.p. S251-S258.

Loomes, G., ‘Some Lessons from Past Experiments and some Challenges for the Future.’ The Economic Journal, 1999, 109, F35-F45.

Malcolm, D., Lieberman, B., ‘ The Behavior Responsive Individuals play a two-person, zerosum game Requiring the use of Mixed Strategies.’ Psychonomic Sci., 1965, 12, p.p. 373-74.

McKelvey, R., Page, T. ‘Public and Private Information: An Experimental Study of Information Pooling.’ Econometrica, 1990, 58, p.p. 1321-1339.

Ochs, J. ‘Simple Games with Unique Mixed Strategy Equilibrium: An Experimental Study.’ Games and Economic Behavior, 1995, 10(1), p.p. 202-17.

Offerman, T., Sonnemans, J. ’What’s causing overreaction? An experimental investigation of recency and the hot-hand effect. Scandinavian Journal of Economics, 2004, 106, 533–553.

Offerman, T., Sonnemans, J., van de Kuilen, G.,Wakker, P.P., ‘A truth-serum for non-Bayesians: correcting proper scoring rules for risk attitudes.’ Working paper, 2007, University of Amsterdam (CREED).

O’Neill, B., ‘Nonmetric Test of the Minmax theory of two-person zerosum games.’Proc. Natl. Acad. Sci., USA, 1987, 84, p.p. 2106-109.

Peterson, C. R. & DuCharme, W. M. ‘A primacy effect in subjective probability revision.’ *Journal of Experimental Psychology*, 1967, 73, 61-65.

Pitz, G. F. Information seeking when available information is limited. *Journal of Experimental*

Psychology, 1968, 76, 25-34.

Pitz, G. F. 'The influence of prior probabilities on information seeking and decision making.' *Organizational Behavior and Human Performance,* 1969, 4, 213-226.

Pitz, G. F., Reinhold, H. & Geller, E. S. 'Strategies of information seeking in deferred decision making.' *Organizational Behavior and Human Performance,* 1969, 4, 1-19.

Phillips, L.D., Edwards, W. 'Conservatism in a simple Probability Inference Task.' Journal of Experimental Psychology, 1966, 72, p.p. 346-354.

Rabin, Mathew, 'Risk Aversion and Expected Utility Theory: A Calibration Theorem.' Econometrica, Vol. 68, No. 5, September 2000, p.p. 1281-1292.

Rapoport, A., Boebel, R. B. 'Mixed Strategies in Strictly Competitive Games: A further test of the minmax hypothesis.' Games and Economic Behavior, 1992, 4, p.p. 261-83.

Rapoport, A., Tversky, A., 'Choice Behavior in an Optimal Stopping Task.' Organizational Behavior and Human Performance, 1970, 5, p.p. 105-20.

Samuelson, William and Richard Zeckhauser, 'Status Quo Bias in Decision Making'. Journal of Risk and Uncertainty, 1:7-59, 1988

Sarin, R., Vahid, F. 'Predicting How People Play Games: A Simple Dynamic Model of Choice.' Games and Economic Behavior, 2001, 34, p.p. 104-22.

Selten, R. 'Die Strategiemethode zur eingeschrankt rationalen Verjstens im Rahmen eines Oligopolexperiments.' In H. Sauerman (ed.), Beitrage zur experimentellen Wirtschaftsforshung, Tubingen: 1967, Mohr, 136-168.

Slembeck, T., 'A Behavior Approach to Learning in Economics: Toward an Economic Theory of Contingent Learning.' European Economic Association Congress, September 1998, Berlin.

Smith, V.L., Walker, J.M., ' Monetary Rewards and Decision Cost in Experimental Economics.' Economic Inquiry, 1993, 31:2, p.p. 245-61.

Snapp, R., Yu, L., Radermacher, M. 'Probablistic Principal Component Analysis using Expectation Maximization (PPCA-EM) for Analyzing 3D volumes and Missing Data.' Microscopy and Microanalysis, 2010, 16, p.p. 836-837.

Stiglitz, J. ' The Contributions of the Economics of Information to Twentieth Century Economics.' Quarterly Journal of Economics, 2000, 115, p.p. 1141-1478.

Suppes, Patrick and Atkinson, Robert C. 'Markov learning models for multiperson interaction.' Palo Alto: Stanford University Press, 1960.

Swets, J. A. and Green, D. M. 'Sequential observations by human observers of signals in noise.' In C. Cherry (Ed.) *Information theory.* London: Butterworths Scientific Publications, 1961 (in press).

Tversky, A., & Edwards, W. 'Information vs. reward in binary choices.' *Journal o] Experimental Psychology,* 1966, 71, 680-683.

Tversky, A. & Kahneman, D. 'Advances in Prospect Theory: Cumulative Representation of Uncertainty.'1992,5(4), pp.297-323.

Tversky, A. & Kahneman, D. 'Availability: A heuristic for judging frequency and probability.' *Cognitive Psychology,* 1973, 5, 207-232.

Tversky, A. & Kahneman, D. 'Belief in the law of small numbers.' *Psychological Bulletin,* 1971, 76, 105-110.

Von Winterfeldt ,D. ,Edwards, W. ' Decision Analysis and Behavioral Research' Cambridge University Press, New York, 1986.

Wallsten, T. ' Failure of Predictions from Subjectively Expected Utility Theory in Bayesian Decision Task.' Organizational Behavior and Human Performance, 1968, 3, p.p. 239-52.