RISK ASSESSMENT IN MICE AND MEN

by

FUAT BALCI

A Dissertation submitted to the

Graduate School-New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Psychology

written under the direction of

Professor Charles R. Gallistel

and approved by

_____

_____

_____

_____

New Brunswick, New Jersey

October, 2007

ABSTRACT OF THE DISSERTATION

RISK ASSESSMENT IN MICE AND MEN

by FUAT BALCI

Dissertation Director:

Dr. Charles R. Gallistel

Uncertainty is a ubiquitous property of both physical and mental realms. Goal-directed actions that take place under these conditions thus probabilistically predict their consequences. Traditional decision-making research has shown that particularly humans are non-normative decision-makers under uncertainty. On the other hand, considering the cognitive system as an output of evolutionary history, it is not unlikely that it models the uncertainties that partly determine the consequences of its actions. It is also natural to assume that the same system uses these models of uncertainty originating from multiple stochastic processes along with its metric representation of the consequences in planning its actions. Indeed, more recent research has shown closer to optimal performance in decision-making tasks in which the uncertainty was experienced and/or originated from the sensori-motor system. In this research, we investigated this very process in the context of temporal decision-making in both human and mice subjects. We further used this experimental context to answer the essential questions regarding the functional architecture of mind. This questioning specifically targeted the degree of representational

and computational power needed to account for decision-making under uncertainty. In order to answer this question, we conducted computer simulations providing different degrees of representational substitution/power and compared their outputs to the empirical data. We conclude that both human and mice are optimal decision-makers under uncertainty that originates from extrinsic and intrinsic (mental) stochastic processes and observed performance can be better explained by information-processing rather than associative frameworks of mind.

ACKNOWLEDGEMENT AND DEDICATION

I am most grateful to my advisor Charles R. Gallistel for his valuable guidance, continuous support, and comforting patience throughout my studies. His wisdom, passion and commitment to the highest standard of research have inspired my development as a scientist. Without his help and encouragement this dissertation would not be completed. I feel very privileged to have known and worked with him. My gratitude also extends to my committee members, Jacob Feldman, Manish Singh, and Michael Littman for their helpful comments and encouragement. Many thanks to Rochel Gelman and Zenon Pylyshyn for stimulating conversations, which helped me develop my ideas during the time I was at RUCCS. Many thanks to David Freestone, who went above and beyond in helping with this research. David not only provided valuable insight to the project but also conducted experiments, analyzed data, and ran simulations. The nights we worked together in the Gallistel Lab further fed my passion for my research. Jackie Gibson also put in many helpful, long hours and contributed significantly.

I must also acknowledge the unfailing patience and emotional support of my colleagues and friends, Efstathios Papachristos, Carlos Montemayor, Omer Ileri, Kemal and Mustafa Karakayali, Ceyhun Taymez, Sara Cordes, Kristy vanMarle, Harry Haladjian, Kimberly Brenneman and Kevin Reed. They have put up with me during the preparation of this dissertation. I am particularly grateful to Patricia Ramey for her care, support, encouragement, and generous love.

Many thanks to PsychoGenics Inc. scientists, particularly Daniela Brunner,

Emer Leahy, Barbara Caldarone, Taleen Hanania, Sylvie Ramboz, Afshin Ghavami, Bassem El-Khodor, Jessica Malberg, Elliot Ludvig, Steven Oakshot, Oxana Litvine, Liliana Menalled and staff for their constant support and patience. My thanks also extend to RUCCS and Department of Psychology staff, Jo'Ann Meli, Sue Cosentino, Suzanne Eng, and Anne Sokolowski who always rescued me when I was lost within administrative processes. Going back in time I must thank to Cigdem Kagitcibasi, Sami Gulgoz, Zeynep Aycan, and Aylin Kuntay for their encouragement for doing graduate study at Rutgers University. Finally, I would like to thank my parents Halil and Hayriye Balci and my siblings Nevin, Senay, and Nedim for the unconditional love they gave to me. They have put up with and made everything possible and most enjoyable for me.

This dissertation is dedicated to my parents Halil and Hayriye Balci and my undergraduate professor Dr. Zeynep Aycan.

TABLE OF CONTENTS

Abstract of the Dissertation ……,…….. ii

Acknowledgement and Dedication ..……..…….. iv

Table of Contents …….……. vi

List of Tables …….……..viii

List of Figures …….…….. ix

Section 1: Introduction …….…….…….…1

Section 2: Information-processing vs. Associationist Frameworks …….…… 32

2.1. Major Differences between two Approaches …….……...… 33

2.2. Basics of Associative Learning Approach …….……...…… 35

Section 3: Method …….……... 40

3.1. Human Experiments …….…… 40

3.2. Mouse Experiments ……..…….…….……. 47

Section 4: Optimal Decision-Making in Temporal Domain …….…….. 54

Section 5: Data Analysis …….…….…….….... 61

Section 6: Results of Human Experiments …….……. 70

Section 7: Results of Mice Experiments …….…… 90

Section 8: Simulations …….……...……... 105

8.1. Simulation Procedure …….……. 108

8.1.1. Associative Weight Updating Rule …….……. 108

8.1.2. Sampling Rule …….……. 110

8.1.3. Associative Weight Updating Rules with Different Sampling Rules …….…….…… 114

8.1.4. Associative and Representational Manipulations …… 116

8.2. Comparison with Empirical Data ……....122

Section 9: Results of Simulations …… 124

Section 10: Conclusion …… 133

Literature Cited …… 153

Curriculum Vita ……... 158

LIST OF TABLES

Table 3.1. Tabulation of Task Parameters ……..……...… 45

Table 3.2. Task Parameters and Their Expected Effects …….……. 46

Table 3.3. Task Conditions Tested …….…….. 47

Table 3.4. Task Conditions Tested by Subject and Stimulus Type ……..…. 48

Table 4.1. Notations used in the formulation of the optimal decision-making model …….…….……..…… 57

Table 8.1. Simulation Parameters …….…….……. 121

Table 9.1. Absolute temporal distance between simulation and empirical estimations …….…….……...… 130

LIST OF FIGURES

Figure 1.1. Subjective value and probability ……..……... 7

Figure 1.2. Illustration of stimuli used in motor planning tasks …….……..... 18

Figure 1.3. Movement end points of hypothetical optimal decision-makers with different task parameters and motor variability …….……..… 26

Figure 1.4. Certain and stochastic configurations used in Maloney et al. …….. 29

Figure 3.1. Illustration of the human analogue of switch task …….……...44

Figure 3.2. Illustration of switch task as it is used in mice …….……...… 51

Figure 3.3. Cumulative distributions of the switch latencies …….…… 53

Figure 4.1. Illustration of integration of external uncertainty, payoff matrix, and internal uncertainty in decision-maker's mind …….……...… 55

Figure 4.2. Cumulative departure and arrival distributions …….……..… 58

Figure 4.3. The effects of the relative probability of short and long trials versus the effects of different levels of timing variability …….……..… 59

Figure 4.4. The effect of a 9-fold difference in gain magnitudes …….……. 60

Figure 5.1. Illustration of calculation of liberal and conservative proportions of MPEG …...…….…….…… 64

Figure 5.2. Expected gain and likelihood functions …….…….. 67

Figure 5.3. A closer look at the highly likely parameter space …….……. 69

Figure 6.1. $\hat{T}_o$s as a function of $\hat{T}$s …….…….……. 70

Figure 6.2. Temporal distances from $\hat{T}_o$ …….…….….. 72

Figure 6.3. Linear regression lines fitted on the CVs …….…… 73

Figure 6.4. Estimated T_o plotted against $\hat{T}$ …….…….. 74

Figure 6.5. CVs presented for individual subjects ……..……..... 76

Figure 6.6. Absolute temporal distances between $\hat{T}$ and T_o as a function of T_o …… 77

Figure 6.7. Frequency of sessions as a function of likelihood ratios between T_o ….. 78

Figure 6.8. Independence of successive switch latencies …….…….. 80

Figure 6.9. Superposed relative likelihood contours and relative gain contours ...… 83

Figure 6.10. Frequency of sessions under which different proportions of MEG (e.g. x > .99) were observed …….…….……. 85

Figure 6.11. Frequency of sessions in which different F-ratios were observed ……... 87

Figure 6.12. The proportion of sessions in which there was a significant change in absolute temporal distance to the optimal switch latency …… 88

Figure 7.1. Linear regression lines fitted on the CVs observed in individual phases 91

Figure 7.2. T_o as a function of $\hat{T}$ for mice subjects …….……..... 92

Figure 7.3. Average $\hat{T}$s to T_os as a function of different probabilistic conditions … 93

Figure 7.4. Histogram of directional and absolute temporal distances from T_o …… 94

Figure 7.5. Absolute temporal distances from T_o as a function of T_o for mice subjects …….…….…….…. 95

Figure 7.6. Frequency of sessions as a function of likelihood ratios for mice Subjects ...…….…….…….. 96

Figure 7.7. Independence of successive switch latencies for mice subjects …….…. 98

Figure 7.8. Likelihood functions superimposed on the relative expected gain function …….…….……..... 100

Figure 7.9. Observed proportion of MEG in each phase …….…… 101

Figure 7.10. Frequency of phases in which different F-ratios were observed with mice subjects ……. 102

Figure 7.11. The proportion of phases in which there was a significant change in absolute temporal distance …….……...… 103

Figure 8.1. Effects of different values of α and β on the rate of acquisition …….. 117

Figure 8.2. Illustration of comparison of statistical distance ……..…… 123

Figure 9.1 The proportion of simulations the exhibited successful predictions ….. 125

Figure 9.2. Proportion of successful instances as a function of different sampling rules ……..……..…….. 127

SECTION 1
GENERAL INTRODUCTION

One day early in the morning, John arrives at the bus stop where he used to take the bus to school. Although, he runs after the bus that just departed, he misses it and starts waiting for the next one. The next bus may arrive either after 15 minutes or 30 minutes depending on the frequently changing bus schedule. As some time elapses, John expects the bus to arrive any time soon if it was on a 15-minute schedule. He starts to get uneasy since there is an important meeting that he may miss if the bus is instead on a 30-minute schedule. However, as more time elapses he becomes more and more confident that the bus may be on a 30-minute schedule. On the other hand, he also knows that he is not a perfect timer and he maybe off in his estimation of time. Additionally, he knows that the bus may come a few minutes early or late depending on the traffic. However, the bus still has not arrived and John starts thinking of getting a cab instead in order to catch the meeting. As tends to happen in such situations, John decides to get a cab. As the cab takes off he watches the newly arriving bus from the rear window. Being annoyed, he tells to himself "If it weren't for that meeting, I would be on that bus rather than this cab!"

People experience decision-making situations similar to the one presented above on a daily basis. Yet mostly people do not have access to the underlying processes that determine their behavior in such cases. Despite this lack of awareness, the computations and variables utilized in such simple cases also constitute the processes and variables used in many other decision-making tasks such as buying a lottery ticket, playing soccer,

taking one job over another one, etc. These processes constitute the research interest of many cognitive psychologists, economists and ecologists.

One can obtain a general understanding of the processes underlying these decision-making situations by predicting the change in decisions under different cases. Here, the differences between cases are defined as a function of differences in the values of the attributes defining the decision-making problem. In order to employ such an approach to the scenario presented above simply put yourself in John's shoes. Assume that buses actually run on 30 minute schedule and try to predict how differently you would behave (in terms of the waiting time for the next bus) under the following conditions:

1) Buses follow the schedule precisely.

2) John wears a watch.

3) John wears a watch and the buses follow the schedule precisely.

4) John is going to a class that he hates sitting in.

5) It is summer time and buses usually run on a 30 minute schedule.

6) John discovers the bus schedule at the bus stop, which shows that buses run on 15-minute schedule.

One would naturally expect differences in the waiting time for the next bus under these different conditions. For instance, John would wait longer for the next bus if the consequence of being late were not as detrimental as that of missing an important meeting. This would be the case if John was instead being late to a boring class. This commonsensical prediction simply demonstrates the role of magnitude of consequences on the decisions made. Another common aspect of decision-making situations can be

demonstrated by differences in waiting times when John does not wear a watch versus when he does, given the knowledge that buses follow their schedule precisely. In the latter case, at any given time after 15 minutes, John would be less likely to wait for the next bus. This prediction is based on the fact that John's uncertainty about the time elapsed since the last bus would be eliminated. The role of another source of uncertainty can be demonstrated with the difference in the waiting times during summer versus fall time. Since during summer time 30 minute schedules may be more frequent than 15 minute schedules, compared to fall time John would again be less likely to wait for the next bus. These last two cases emphasize the role of uncertainty in defining most decision-making tasks.

Briefly, such an analytic approach even to a simple and mundane decision-making task demonstrates how the underlying mental processes can involve multiple variables and their integration. In this research, we focus on the processing of information regarding different kinds of uncertainty and their integration. The following distinctions will be critical:

1) Whether the source of uncertainty is intrinsic or extrinsic. Intrinsic uncertainty originates from stochastic processes within the animal, such as variability in the motor system, fuzziness in the representation of quantities, etc. Extrinsic uncertainty originates from stochastic processes that occur in the physical world, such as rolling a dice. By wearing a watch, John would eliminate his intrinsic uncertainty about the duration elapsed since the last bus. On the other hand, the different frequency of bus schedules in different seasons would involve extrinsic uncertainty, as would the variability in actual bus arrival times about their scheduled arrival time.

2) The manner in which the information about uncertainty is gathered: experienced versus described uncertainty. John may be told the probability of bus being on a 15 minute schedule by a second person. This would be a described uncertainty. Alternatively, John may have experienced the buses on arriving different schedules himself. This would be an experienced uncertainty. This two-way distinction results in four possible cases/types of uncertainty: *Described Extrinsic Uncertainty*, *Experienced Extrinsic Uncertainty*, *Experienced Intrinsic Uncertainty* and *Described Intrinsic Uncertainty*. Although the first three types are fairly common in daily life, the Described Intrinsic Uncertainty is not common. Thus, our review of literature focuses on the first three cases. We consider these cases in different kinds of decision-making tasks. A key question is how normative the underlying processes are in taking into account and integrating the relevant information.

Historically, most of the research conducted in the area of decision-making in the face of risk has focused on the non-normative phenomena demonstrated with extrinsic uncertainties (e.g. Kahneman & Tversky, 1979). This class of research was followed up by others that investigated the processing of experiential versus symbolic forms of external uncertainty (e.g. Barron & Erev, 2003). This second class of research yielded different results for experiential and symbolic external uncertainty. On the other hand, more recently an independent group of researchers asked similar questions for decision-making tasks which mainly involved intrinsic uncertainty (e.g. Trommershäuser, Maloney, & Landy, 2003a). The results from these three different classes of research constitute a controversy regarding the normativeness of human decision-making.

The classical view of decision-making research produced a large amount of evidence, which showed that under risky conditions, humans' judgments violate the predictions of normative decision-making models (e.g. Kahneman and Tversky, 1979). In these tasks, the subjects were given a verbal description of the choice problem in terms of outcomes and their probabilities. Then they were asked to make a choice across the options described. As an example of non-normative decision-making phenomena in these tasks, humans were found to prefer [Lottery 1 = (1 × $ 5000), a certain $ 5000] over [Lottery 2 = (.8 × $ 10,000; .2 × - $10,000)], even though the expected gain of Lottery 2 is higher than Lottery 1; 6000 > 5000.

The normative solution of these tasks first requires weighting of different outcomes with their probability of occurring. Summation of all weighted outcomes for a given option then determines the expected gain from that option. The better option is chosen by comparing the expected gains and choosing the option with the higher expected gain. For instance, in the example presented previously the expected gain of Lottery 1 is $1 \times 5000 = 5000$ and the expected gain of Lottery 2 is $.8 \times 10{,}000 + (.2 \times -10{,}000) = 6000$. The normative decision maker should therefore favor Lottery 2 over Lottery 1.

Kahneman and Tversky (1979) outlined the non-normative risk attitudes by a fourfold pattern; "*risk seeking for gains*" and "*risk aversion for losses*" of low probability, and "*risk aversion for gains*" and "*risk seeking for losses*" of high probability. Risk aversion refers to being more reluctant to prefer an option with an uncertain payoff over an option with a certain payoff, although the expected payoff of the uncertain option may be equal to or larger than the certain one. Conversely, risk seeking

refers to being more reluctant to prefer an option with a certain payoff to an option with an uncertain payoff, both having the equal expected value. As stated previously, risk preference changes systematically as a function of the sign of the outcome (gain or loss) and its probability. For instance, humans choose a lottery that yields a gain with a low probability over a lottery that yields an equal expected value for certain, namely risk seeking for gains of low probability.

On the contrary, humans choose a lottery that yields a certain loss over a lottery that yields loss with low probability with equal expected value, namely risk aversion for losses of low probability. The reverse pattern is observed when the probabilities of outcomes are high. More specifically, humans choose a lottery with a certain gain over a lottery that yields a gain with a high probability with equal expected value, namely risk aversion for gains of high probability. Conversely, humans choose a lottery that yields a loss with a high probability over a lottery that yields equal expected value for certain, namely risk seeking for losses of high probability.

Tversky and Fox (1995), Kahneman et al. (1979) and Tversky and Kahneman (1992) explained the non-normative decision-making phenomena assuming two key elements, a reference dependent value function (see Figure 1.1A) and a non-linear subjective probability function (see Figure 1.1B). The value function maps outcomes to utilities and its form predicts that humans are risk-averse over gains and risk-seeking over losses. The steeper utility function for losses further predicts the empirical fact that losses have bigger effect than gains. For instance, when subjects are presented with the following two options: a) Gain $ 50 with p = .5 and lose $ 50 with p = .5 b) Nothing with certainty, they tend to prefer the second option. This preference is explained by the

steeper value function for losses. However, in this review, we will focus on the second key element since it is critical for explaining the four-fold non-normative pattern presented above. Tversky et al. (1995), Kahneman et al. (1979), and Tversky et al. (1992) claimed that four-fold risk attitude pattern was due to a non-linear transformation of the probability scale (but see von Neumann & Morgenstern, 1947 for attribution of the same pattern to expected utility function). This function was claimed to overweight small probabilities and under-weight moderate and high probabilities (Figure 1.1B).

For instance, this probability function predicts risk seeking for gains of low probability over an option with a certain gain with a value that is equal to the expected value of the first option. This phenomenon emerges since the uncertain outcome is weighted with a probability value, which is higher than the objective probability. As a result, the subjective expected value of the uncertain option exceeds the value of the certain gain. A similar logic explains other three forms of risk-preference.

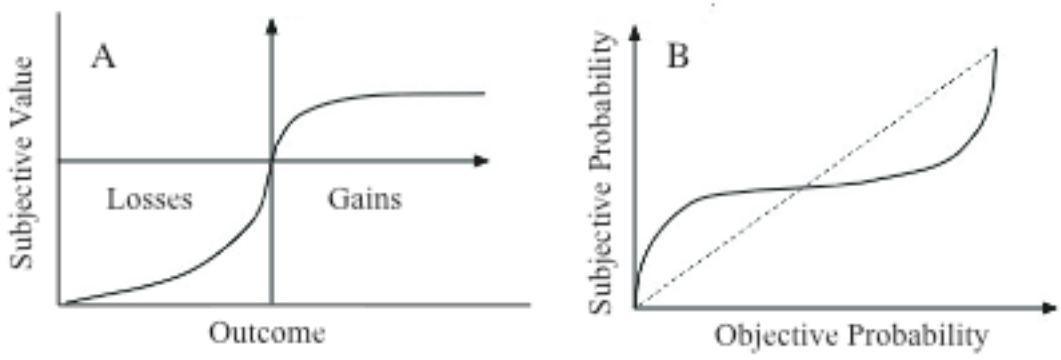

Figure 1.1. Subjective value and probability. Left panel (A) illustrates how subjective value changes as a function of objective losses and gains and right panel (B) illustrates how subjective probability changes as a function of objective probability.

This model was originally developed for situations in which the observer knew the probabilities of different outcomes. Tversky et al. (1995) extended this phenomenon

to cases of uncertainty in which the observer did not know the probabilities of the outcomes. They found that when an increase in probability turned impossibility into possibility or possibility into certainty, it had more impact compared to when the same increase made a possibility higher or lower (with equal degree or more). Other forms of violations of normative models can also be found in literature. Kahneman and Tversky (1973) reported another non-normative phenomenon, neglecting the base-rate information. Assume that a diagnostic test can determine if a given person has a rare disease with 90% accuracy. When asked about the probability that a person has this disease given that they tested positive, most of the subjects give a high probability value that is very close to if not equal to the accuracy rate of the test. On the other hand, the normative answer would be gathered by multiplying the accuracy rate of the test with the occurrence rate of that disease in the population, which would give a much lower probability value than the hit rate of the test (if the disease is observed in every 1000^{th} person, this probability value is .001 x .9 = .0009).

Briefly, Kahneman and Tversky's approach can be interpreted as arguing that our minds may not be built to work by the rules of probability (see Samuels, Stich, and Bishop, 2002). On the other hand, there are two important issues that should be kept in mind while evaluating this view. Firstly, as mentioned above there are different sources of uncertainty and these researches only considered a single source, namely the extrinsic uncertainty. The second issue relates to the form in which the information about the uncertainty was gathered by the subject (described versus experienced uncertainty). In all of these reported tasks the uncertainties were described by probabilities; they were not experienced frequencies. Focusing particularly on this second aspect, Barron & Erev

(2003) and Hertwig, Barron, Weber, and Erev (2004) referred to these classical decision-making tasks as *decisions from descriptions* (see also Weber, Shafir, and Blais, 2004). They further noted most real life situations rely on personal experience (rather than on the explicit sources of information). Under this rationale, several studies by these researchers focused on decisions based on experienced uncertainty instead of described uncertainty.

In Barron et al. (2003), subjects had no prior knowledge regarding the payoff distribution of different options they were given. Instead, they gathered information about the outcomes and their probabilities through repeated choices and experiencing the outcomes of their choices. More specifically, in this experiment subjects were presented with two options (illustrated as buttons) on the computer screen and were asked to choose across these options for 200 trials. The outcome a subject experienced on a trial was drawn at random from the outcome distribution of the option they chose on that trial. At the end of each trial the subjects were presented with both payoff on that trial and cumulative payoffs. The aim of the subjects in this task was to maximize the pay-off points, which were converted into monetary reward at the end of the experiment. This task was different from description-based decision-making tasks since the decision did not depend on descriptions of outcomes and their probabilities but rather on the experience of these attributes by the subjects themselves. In this sense, in this experiment, the subjects learned about the payoff distribution of different options through repeated choices.

In these experiential situations, consistent with Thaler, Tversky, Kahneman, & Schwartz (1997) loss aversion was found to be a robust phenomenon. Loss aversion led their subjects to obtain lower gains than the maximal amount one could gain in Barron et

al. (2003). On the other hand, other findings of this experiment contrasted with the well-established findings of description-based decision-making tasks (ex: Kahneman et al., 1979). For instance, relative to description-based decision-making tasks in the experiential task, a certain outcome was less attractive than the uncertain option with higher expected value. For example, when choice of option A on a computer screen led to an outcome of $ 4 (.8) or $ 0 (.2) -expected value of $ 3.2- and choice of option B led to a certain outcome of $ 3, 63 % of the subjects chose A over B. In other words, the option with the higher expected value but also with a higher risk of gaining nothing was chosen more than half the time over an option that yields a lower outcome. By contrast, the percentage of subjects who were reported to choose alternative A over B in Kahneman et al. (1979) was 20 %. Further, there was a stronger risk aversion tendency in the loss domain than in the gain domain. Note that, this contrasts with the pattern predicted by the value functions; that is risk aversion over the gains and risk seeking over losses. Lastly, in the experiential tasks, the rare outcomes were found to be underweighted. Barron et al. concluded that their findings were consistent with underweighting of low probabilities in decisions from experience (increasing A's attractiveness), which contrasted with overweighting of low probabilities in decisions from descriptions (decreasing A's attractiveness). This particular finding is inconsistent with the non-linear function relating subjective and objective probabilities stated by the Prospect Theory. Consequently, in the experiential tasks subjects were closer to maximizing their gain compared to subjects tested in classical decision-making tasks with described rather than experienced outcomes.

The authors explained these findings by a simple model that assumed a tendency

to rely on recent outcomes (e.g. weighting information as a function of temporal proximity) in decision-making. Thus, their model explained the underweighting of rare events by their tendency to be underrepresented in the small samples of recent outcomes that subjects tend to rely on. On the other hand, in description-based tasks, since all information is presented in a single number containing no temporal sequence prior to decision, one cannot talk about such a recency (or any other time-dependent) effect in those tasks.

Hertwig et al. (2004) conducted a follow-up experiment. In their experiment, one group of subjects were given descriptions of prospects while the other group experienced the outcomes and likelihoods directly. Their findings were consistent with those of Barron et al. (2003); 88 % of subjects in the experience group chose the option with higher expected value while it was chosen by 36 % of subjects in the description group. We emphasize that the preference for such a prospect was reported to be 20 % in Kahneman et al. (1979). Thus, their results also supported underweighting of low probabilities in experiential tasks. Hertwig et al. used the sampling paradigm in their experiments. In their experiment, the subjects were given two different options with different outcome distributions. They were allowed to draw samples from both options. After sampling from both options, the subject was asked to choose one across two options. After they made their choice, they moved on to a different pair of options with new payoff distributions. This is different from making repeated choices across same options with non-changing payoff distributions across trials and getting feedback, as in Baron et al., (2003). Since underweighting of rare events was also reported in their sampling paradigm, Hertwig et al. concluded that this phenomenon was due to direct

experience of outcomes and likelihoods per se rather than repeated decisions.

Weber et al. (2004) also focused on the differences between decisions derived from descriptions versus experience using a more extensive group of subjects and analyses. Work of Weber et al. (2004) contained the meta-analysis of human classical decision-making and non-human animal foraging data derived from work of Shafir (2000). The data considered in the meta-analysis of animal data were gathered from wasps, bees, fish and birds. In each experiment analyzed, the subjects had to choose between an option that provided a constant reward or an option that provided a variable reward, whose expected value was equal to the constant reward. Since expected values were constant within each variable-constant reward pair, the choice of the constant reward was claimed to be a function of variance in the variable reward. Consistent with this reasoning, Barron et al. (2003) reported lower cumulative gains (farther from the maximum possible gain) with higher payoff variability.

Barron's reasoning regarding this fact was that the subjects had a less accurate estimation of the actual expected values with higher payoff variability. On the other hand, Shafir's (2000) analysis of animal data additionally showed that coefficient of variation (CV), which was computed by dividing the standard deviation of outcomes by their expected value, was a better statistical predictor of proportion of risky choices/risk sensitivity than variance (or the standard deviation which are un-standardized measures) and expected value. Basically, his analyses of animals' risky choice data suggested that animals were using the CV as the measure of risk. In particular, Shafir's analysis showed that when the expected value exceeded the required caloric intake, larger CV was associated with greater risk aversion, while when it was below the required caloric intake,

it was associated with risk seeking.

Other theories such as Scalar Utility Theory (see Marsh & Kacelnik, 2002; Kacelnik and Bateson, 1996) similarly predict that risk sensitivity should be a function of CV rather than un-standardized measures such as outcome variance or standard deviation. This theory suggests that each magnitude is represented by a Gaussian probability density function and in these representations standard deviation and mean are proportional. In other words, like Scalar Expectancy Theory (Gibbon, 1977), the Scalar Utility Model assumes that CV remains constant in the representation of different perceived or estimated outcomes, in accord with Weber's Law. Based on this basic assumption (along with further assumptions regarding the representation of risk and choice making, which I won't refer to here) CV is assumed to be a better predictor of choice of risky alternatives than un-standardized measures of variability (Marsh et al., 2002).

In the meta-analysis of human data, Weber et al. (2004) analyzed those choice situations that were similar to animal choice situations analyzed by Shafir (2000). The nature of choice tasks was the same except that in all 226 situations the information about outputs and likelihoods were presented to the subject with one-shot descriptions. The analysis of human data showed that the CV was only a marginally better predictor than other conventional measures such as variance, particularly in the gain domain. In order to investigate the cause of this difference between human and non-human animal data, Weber et al. conducted two experiments. In Experiment 1 human subjects learned outcome value and probability by experience. They were given 2 decks and allowed to sample from them until they thought they had a good idea about which deck was the better one. Once, subjects choose the deck to "gamble" from after a number of draws,

they moved on to a new pair of decks. In a given pair, one of the decks always contained a certain output ($ X), while the other deck's payoffs were variable. The expected values of pairs were equal. The experimenters chose the pairs such that 3 of the variable payoff decks had the same variance but different CVs and 3 of the variable payoff decks had the same CV but different variances.

In Experiment 2, subjects were given full summary information as in classical decision-making tasks rather than experiencing the outputs and their likelihoods. As a result of these two experiments, Weber et al. found that risk sensitivity in humans and non-human animals had common characteristics, such as CV being a better predictor of relative risk of uncertain choice alternatives than (un-standardized) variance or standard deviation. More specifically, they found that risk aversion increased with CV rather than variance or standard deviation. However, interestingly this was found to be true more for Experiment 1 rather than Experiment 2. Based on this difference across tasks and consistent with the interpretation of Hertwig et al. (2004), Weber et al. concluded that it may be the experiential processes per se that might have contributed to a stronger relationship between CV and risk sensitivity. This resembles the similarity between human and non-human animals in the psychophysical properties of representation of discrete quantities, when counting is done non-verbally. Similarly, in the case of verbal counting these psychophysical properties change and similarities disappear (see Cordes, Gelman, Gallistel, & Whalen, 2001). More specifically, in the case of non-verbal counting one observes scalar variability, which is common to human and non-human animals' representation of numerosities (see Mechner, 1958). On the other hand, in the

case of verbal counting one observes binomial variability rather than scalar variability (Cordes et al., 2001).

Based on these findings, we can conclude that when information about external uncertainty is gathered through experience/feedback the violation of normative reasoning about probabilities observed with explicit external uncertainty (uncertainty communicated with one-shot descriptions) diminishes. Therefore, humans are closer to maximizing expected gain, when their decisions are based on experience of different risky alternatives rather than on their descriptions. Further analyses of Weber et al. and Shafir showed that in experiential cases, human decision-making resembles non-human animal decision-making more closely in terms of the quantitative properties underlying decision-making processes.

In the cases presented up to now, presented explicitly or directly experienced, the uncertainty was in the outer world, external to the observer. We did not discuss the representational state or state of any other system in the Observer in any of these cases. A different kind of uncertainty we will consider is the intrinsic uncertainty that is due to the representational state or state of the motor system of the observer. In other words, this second form of uncertainty is due to the stochastic processes that may underlie the representational system in capturing physical attributes and/or the implementational system in acting on the physical world (e.g. motor system). Regarding the representational side of the coin, signatures of uncertainty is a well-studied aspect of preverbal representations of quantities such as time, number, space, area, length, etc. These studies using different psychophysical tasks have shown that scalar variability is a prominent signature of these magnitude representations (e.g. Whalen, Gallistel, &

Gelman, 1999). The scalar nature of uncertainty is assumed to be the intrinsic property of these metric representations. There is no such uncertainty that is intrinsic to the conventional symbolic forms (e.g. Arabic numerals) representing these quantities. For instance, the binary representations of quantities (e.g. 3, 4, 10) are equally discriminable from each other. There is nothing intrinsic in their form that makes 10 more discriminable from 3 compared to 4. They are just different symbols, 0011 vs. 0100 vs. 1010. The analogue form of these quantities however reflects the distances between these quantities in their very form (e.g. ------ vs. -------- vs. ---------- under a linear model) and for a given level of noise 10 is more discriminable from 3 compared to 4. This uncertainty would be expected due to stochastic processes in the accumulation (e.g. timing) or retrieval of this information. Internal uncertainty may alternatively be due to other processes that utilize representations. Under this alternative, the representational system can be assumed to be noise-free while some form of noise can be assumed to be introduced at the level at which these representations are put to use (such as the motor system using spatial representations). Lastly, intrinsic uncertainty may very well be jointly due to representational and execution system of the organism.

For either case in this second form of noise, uncertainty is intrinsic to the Observer and it can be derived from experience through a form of feedback. This would basically be the information gathered regarding the discrepancy between what one represents (or aims to implement) and its observed extension in the real world when the representation is put in use. This form of uncertainty does not require one to explicitly (e.g. verbally) represent the noise. On the contrary, almost in all such cases it may be very hard to communicate it to another agent. However, one can intuitively tell that this

form of uncertainty could be much more accessible to our cognitive system as the input of uncertainty, since the form of representation of this noise itself may be more suitable (or relevant) as instantiations of variables in particular operations of mind. I should again note that the information regarding this form of uncertainty can be derived from experience or arguably it can be emulated in one's mind. On the other hand, it is different from the experiential cases presented previously since in those cases the uncertainty was in the World rather than being intrinsic to the Observer.

Our actions derived from our representations usually have consequences. Those consequences depend on our actual rather than planned actions. Thus in order to be a successful agent in one's environment, the processes underlying decisions should take into account not only the consequences of alternative actions and relevant statistical properties of the environment s/he lives in, but also the uncertainty intrinsic to the system. It is this intrinsic uncertainty that causes the discrepancy between what one represents and its extension in the world or what one aims at and what it achieves.

Recently, some researchers (Maloney, Trommershäuser, & Landy, 2006; Trommershäuser, Maloney, & Landy, 2003a; 2003b) have developed motor tasks that focus on intrinsic uncertainty. These tasks are computationally as complex as the tasks reported in classical decision-making tasks (see Kahneman et al., 1979) regarding the optimal / normative solution. They involved magnitudes and uncertainty and required their combination for the normative/optimal solution, as in Kahneman and Tversky's tasks. However, in these tasks the uncertainty in the task was introduced naturally by the intrinsic noise in the subject's motor system rather than from a source external to the subject.

Subjects were presented with a number of circular reward and penalty regions at different configurations and amounts of overlap between them (see Figure 1.2 for examples). The subjects' task was to hit a reward region trying to avoid the overlapping penalty region within a limited amount of time. They were rewarded for touching the reward region within the critical duration by the amount associated with that region and penalized for touching the penalty region by the amount associated with that region. When they touched the overlapping region, the subjects were rewarded with the sum of amounts associated with the two regions. If the subjects touched the configuration later than the critical period (700 ms), they were penalized by a huge amount (time-out penalty). The subjects tried to maximize their gain. Basically, in these tasks the decision was made about where to aim in the configuration in order to maximize gain. Subjects were tested over many trials in these experiments.

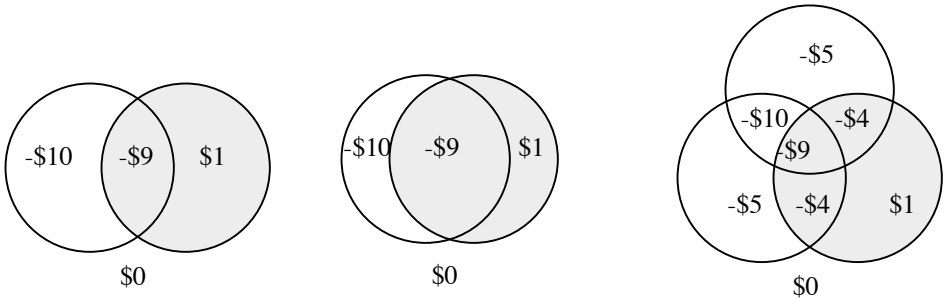

Figure 1.2. Illustration of stimuli used in motor planning tasks (Trommershäuser et al. 2003a, 2003b).

The response time constraint in these tasks ensured a level of variability in motor responses due to the speed-accuracy trade-off associated with pointing. This assumption is based on Fitts' Law, which is a prominent model of human psychomotor behavior (Fitts, 1954). It predicts the rapid aimed movements based on time and distance

parameters. Fitts' Law says that movement time (T) is a logarithmic function of distance (D) when target width (W) is constant or is a logarithmic function of target size when distance is constant.

$$T = a + b\log_2(D/W + 1)$$

where a and b are empirical constants. Taking $\log_2(D/W + 1)$ as index of movement difficulty, the formula can be rewritten as $T = a + bDI$, DI being the index of movement difficulty.

Harris & Wolpert (1998) discussed Fitts' law in terms of variability. They predicted a speed/accuracy tradeoff (Fitts, 1954) by signal-dependent noise that increases as a function of neural control-signals. Since control-signals would be larger in faster movements, the signal-dependent noise that deviates the trajectories from the desired path would increase. They claimed that accuracy could be increased by moving slowly since it would require smaller control-signals. Indeed, in their experiment they found that humans chose trajectories that minimized variability in the final position of arm, which predicts the speed-accuracy trade-off for a given level of accuracy or a given speed of movement. Under this reasoning, imposing a response time constraint to the task (constraining T parameter in the formula above), Maloney et al. (2006) and Trommershäuser et al. (2003a, 2003b) ensured a level of variance in the pointing behavior (intrinsic uncertainty).

There are two important aspects of these motor tasks, external rewards and penalties related to each action and uncertainty in the motor implementation of a strategy chosen by the decision-maker. By defining decision-making under risk as decision making when one has access to the probabilities induced by each possible strategy, these

tasks can be considered as tasks of decision-making under risk. This assumes that the subjects anticipate the uncertainty in their movement (Maloney et al., 2006). Thus, these motor planning tasks resemble those cases used in classical decision-making tasks. As classical decision-making tasks involved normative verbal solutions, the motor tasks described above also involved normative solutions in terms of where to aim in the configuration. The normative solution in these tasks is defined as the mean movement end-point that would maximize the gain. The interesting point about the normative solution to these motor problems is that the dependence of normative solution on the task parameters such as payoff matrix, degree of overlap between regions, geometry of the regions is mediated by the degree of variability in subject's motor system.

In order to predict the optimum point of responding in these tasks, Trommershäuser et al. (2003a, 2003b) developed an optimal movement planner[1]. This model was based on the principles of Statistical Decision Theory (SDT). SDT is a normative model, in the sense that using its principles one can gather strategies that

---

[1] In the optimal movement planner, the expected gain of the motor strategy in these tasks was calculated in the following form:

$$\Gamma(S) = \sum_{i=0}^{N} G_i P(R_i \mid S) + G_{timeout} P(timeout \mid S) + \lambda B(S)$$

where $P(R_i \mid S)$ is the probability of hitting R_i given the strategy (S) before timeout, $\lambda B(S)$ is the biomechanical gain function and $P(timeout \mid S)$ is the probability that a strategy will exceed the time limit. The last two parameters can be considered as being constant across different trials and therefore the formula above can be reduced to:

$$\Gamma(S) = \sum_{i=0}^{N} G_i P(R_i \mid S)$$

where the probability of hitting target region (R_i) given strategy (S) is computed in the following way:

$$P(R_i \mid S) = \int_{R_i,timeout} P(\tau \mid S) d\tau$$

where "R_i, timeout" is the trajectories that pass through R_i at some point in time in between the start of trial and timeout. Basically, according to this model, the consequence of S is the probability density assigned to possible trajectories, $P(\tau \mid S)$.

minimize loss and maximize gain (Maloney, 2002). SDT assumes two parties, Observer and World. The Observer, who should choose between a number of actions, is viewed as an imperfect gatherer of information regarding the true state of the world. S/he acts upon the world and is rewarded or penalized as a result of the true state of the world and his/her actions. The Observer tries not only to estimate the true state of the world (the geometry of the circles-where they are and how they overlap—and the points to be obtained or lost) but also to maximize gain. Briefly, SDT tells which actions to choose from a possible set of actions given the information the observer has about the state world in order to maximize the gain.

Uncertainty that is a determinant of the prediction of SDT can either be viewed as uncertainty in the knowledge of the true state of the world (e.g. How long it took you to read the paragraph up to here?) or in the process that acts on the world (e.g. Touch the exclamation mark within 700 msecs! Where did you actually end up touching? I assume that you have normal or corrected vision) or both. For an in-depth conceptual and formal treatment of statistical decision making theory reader can refer to Maloney (2002) and Mamassian, Landy, & Maloney (2001).

In the motor planning tasks used by these researchers, the payoff matrix was well defined in terms of the location of stationary objects, their geometry and the degree of overlap between them. The uncertainty was considered to be in the functioning of the motor system that executed the action (pointing on the computer screen) given an aim point that was computed by the cognitive system (which was not attributed any uncertainty). In their optimal movement planner, Trommershäuser et al. (2003a, 2003b) estimated the motor variability of individual subjects from their pointing data. Then using

the estimate of this empirical variability, stimulus configurations and the payoff matrix as parameters of their model, they estimated the optimal point of responding (optimal mean movement end point). For the in-depth formal treatment of the model, see Trommershäuser et al. (2003a, 2003b). Thus, the predictive power of this model largely depended on the degree of accuracy in modeling subject's uncertainty. In their model (developed for the motor task), uncertainty was defined by the standard deviation of the Gaussian that is equal in all directions (isotropic).

This model of uncertainty was based on the empirical findings that once the subjects were practiced, the movement end points were distributed around the mean end point as an isotropic bivariate Gaussian distribution. The isotropic nature of the uncertainty about where a point will end reduces the problem to a single dimension. This dimension is the line connecting the centers of the reward and penalty circles in configurations containing two regions. In configurations containing more than one penalty region, this dimension can be generalized to the geometrical symmetry axis (but see Wu, Trommershäuser, Maloney and Landy, 2006 for cases in which the problem does not reduce to the geometrical symmetry axis).

In order to illustrate the behavior of optimal movement planner under these conditions, we simulated the performance of hypothetical optimal subjects under three different payoff matrices (+100 vs. 0, +100 vs. –100, +100 vs. –500), four different levels of motor variability (σ = 4, 8, 14, and 20 pixels) and a single degree of overlap between two regions (see Figure 1.3). In Figure 1.3, the green and red circles represent the reward and penalty regions, respectively. The payoff associated with each of these regions is presented via numerals (white for rewards and black for penalties) within the

corresponding region. The payoff of the intersection area between reward and penalty regions is the sum of the reward and the penalty (signed). The blue dots represent the 4000 movement end points of the hypothetical optimal subject with different levels of motor variability (modeled as isotropic Gaussian distribution).

As one can see, when the penalty associated with the penalty region is 0, no matter how variable the subject's motor variability is, the optimum point of responding is always the center of the reward circle. As one increases the penalty of hitting the penalty region (from top to bottom panels) and/or as one increases the motor variability when the penalty of hitting the penalty region is larger than 0 (from left to right panels for two bottom panels), the optimum point of responding moves to the left, namely farther away from the penalty region.

Trommershäuser et al. (2003a, 2003b) and Maloney et al. (2006) reported that their subjects responded consistently with the predictions of optimal movement-planning model[2]. Namely, they performed optimally or near optimally in these kinds of motor planning tasks by aiming at a point that maximized the expected gain given the configuration. Overall, subjects' responses were found to be within 8% of optimal responding (also see Trommershäuser, Landy, & Maloney, 2006). The standard deviation in the distribution of points about their mean was nearly constant across different conditions.

The high correspondence between the model's prediction and subjects' performance suggests that subjects too should have used stimulus configurations,

---

[2] Optimal movement planner also predicts shifts as a degree of overlap between reward and penalty regions, which was not demonstrated in our simulation. This prediction of the model was also supported by the empirical data.

rewards/penalties and the estimate of their own motor uncertainty while planning their actions. Consistent with these findings, in a simpler task and using simpler measures Fitts and Petersen (1964) also showed that humans take account of their sensori-motor variability when they are planning their movements. Consistent with the overall findings, as suggested above the speed/accuracy tradeoff reported by Fitts et al. (1964) and Harris et al. (1998) also suggested that humans take into account their sensori-motor variability while planning their movements. I should again note that the time constraint imposed on the motor task was based on this very relation between movement time and accuracy.

Further studies that used the same task showed that the motor system could also take appropriate account of task-relevant variability that is introduced externally (e.g. Trommershäuser, Gepshtein, Maloney, Landy and Banks, 2005). In these experiments, subjects estimated the new variability (introduced by the experimenters on the touch screen monitor) in less than 120 trials and adjusted their mean movement end point optimally in order to compensate for the increase in the task relevant variability. They were further found to use this new estimate in novel situations (see also Trommershäuser et al., 2003b). The one condition in which subjects deviated from the optimal point was when the optimal mean end point predicted by the model was outside the target circle. Evidently, subjects view aiming outside the reward circle as violating the instructions or the definition of the task.

The optimal performance of subjects in the decision-making tasks that involved intrinsic uncertainty was also observed when they were asked to choose between two configurations with different expected gains and presented simultaneously (Trommershäuser et al., 2006). In a majority of the trials subjects preferred the

configurations with higher expected gains within a very short amount of time (less than 700 msec). The expected gains from different configurations were again determined taking into account empirical noise gathered from each subjects' pointing data. These results suggested that movement under risk is based on very rapid judgments regarding the expected gain of different configurations and that these judgments are based on the internal ordering of the configurations derived from the estimates of expected gain for each configuration. The rapidity of this process was further demonstrated by the findings of Trommershäuser, Mattis, Landy, & Maloney (2006). Their results showed that humans could optimally integrate the relevant reward information into their movement plan within 200 to 400 msecs prior to the end of actions (also see Smyrnis, Evdokimidis, Constantinidis and Kastrinakis, 2000).

Another interesting finding of this experiment was that while subjects were choosing between simultaneously presented configurations, the mean end point of their selection movement was indeed the point, which maximized the expected gain. In terms of the task involved in Trommershäuser et al. (2006), which is choosing between two simultaneously presented configurations, this experiment is more comparable to classical decision-making tasks[3]. Under this reasoning, consistent with earlier findings this study showed that humans could be optimal decision-makers also in selection tasks when they involve intrinsic probabilities contrasting with their sub-optimal performance in choosing

[3] Each configuration that constitutes an alternative in subject's selection task can be viewed as a lottery for a given strategy (S): " $L(S) = (P(R_1|S), G_1; P(R_2|S), G_2; P(R_3|S), G_3; P(R_4|S), G_4)$ " (Trommershäuser et al. 2006, pg. 983) where $P(R_i|S)$ is the probability of hitting the i^{th} region and G_i is the gain associated with that region. Here, subject's task is not only to plan the optimal strategy, which can also be viewed as choosing from infinite number of lotteries given a configuration but also to choose between different configurations (lotteries).

between risky prospects defined by described extrinsic probabilities.

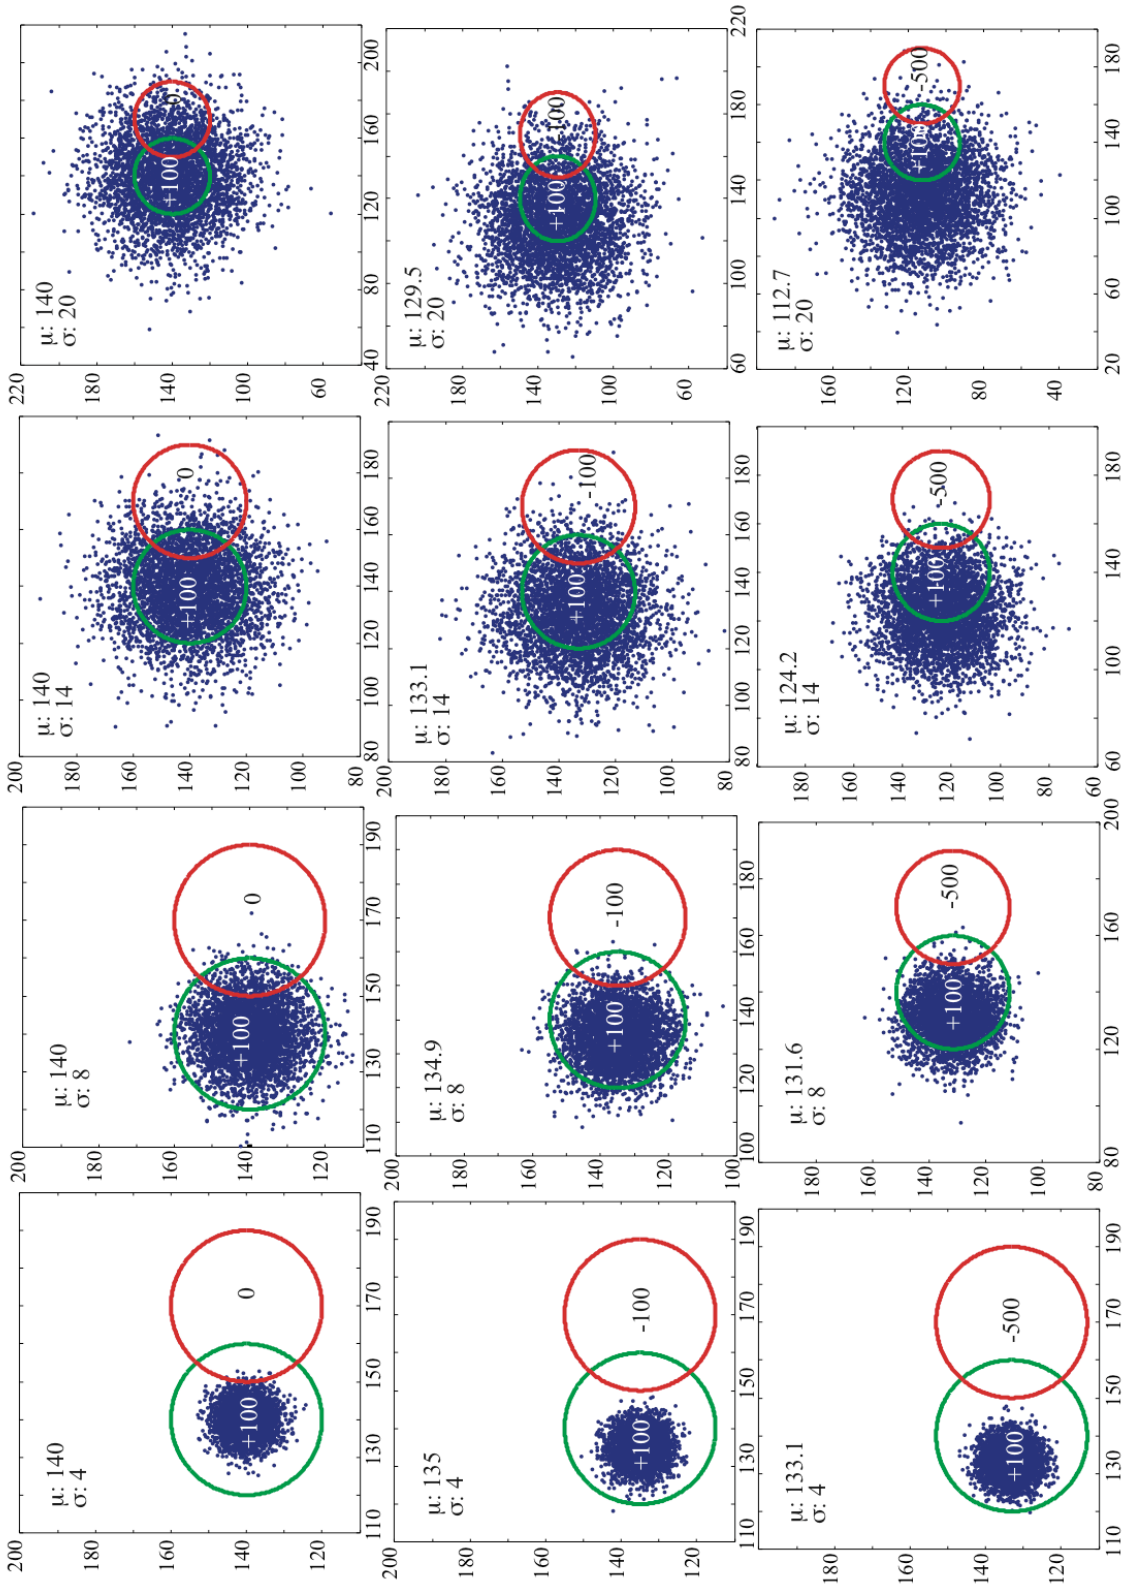

Figure 1.3. Movement end points of hypothetical optimal decision-makers with

different task parameters and motor variability. Green and Red circles are reward and penalty regions, respectively which are associated with payoffs indicated within them. The expected gain from the intersection of two circles is the sum of the reward and penalty (signed). Blue dots represent 4000 movement end points of hypothetical optimal decision-makers with different levels of motor uncertainty.

Briefly, contrary to the non-normative decision-making phenomena demonstrated with classical decision-making tasks, human subjects were found to be normative in the tasks involving intrinsic uncertainty. In order to explain these findings, Trommershäuser and her colleagues basically assumed that in choosing a visuo-motor strategy, the system imposes a probability density on the possible movement trajectories that could occur once the strategy is implemented. Payoff is calculated based on the distribution of anticipated end points and penalty/reward magnitudes. Thus, the expected cost is minimized by taking into account the magnitude of penalties and rewards and the probability of having them, which is defined by the intrinsic uncertainty.

Overall, these findings showed that subjects could represent their intrinsic movement uncertainty, update these estimates when external factors increased task relevant variability, and take this it into account when planning their movement and selecting in between different configurations. Considering that the consequences are determined as a function of one's actual responses rather than the aimed ones, this set of faculties in dealing with the stochastic nature of motor actions (intrinsic uncertainty) is a necessity to maximize the gain, which human subjects were found to entertain. On the other hand, we also saw that there may be some factors such as the tendency to stay in the target region or delay in gathering task-relevant information during movement that may impose non-normative constraints on the decision-making process.

Recently, Hudson, Maloney, and Landy (under review) tested decision-making in the domain of motor timing under symmetrical and asymmetrical payoff matrices. In their experiment, the subjects were asked to touch a spatial target within a target temporal window. If they could touch the spatial target within the temporal window, they were rewarded and if they touched the spatial target within the penalty temporal window or failed to touch the spatial target, they were penalized. In addition to the spatial target, subjects were presented with a timer bar that indicated the rewarded and penalized temporal windows. The experimenter also added a 25 ms of noise to the temporal end points. They found that their subjects could correctly estimate the temporal uncertainty in their motor system and compensate for it along with the experimentally imposed temporal uncertainty. They further anticipate how this uncertainty they estimate interacts with the rewards and penalties associated with different temporal points. Briefly, these results showed that similar to decision-making in the spatial domain, subjects were optimal decision-makers in timing their responses in ms range. Note that the experimenters tested the motor-timing in this experiment, that is temporal uncertainty inherent in the motor system.

At this point, we have discussed two types of uncertainty; described and experienced & extrinsic and intrinsic. The type of uncertainty in Kahneman and Tversky tasks were described extrinsic; they were chosen and communicated to the subjects by the experimenter. On the other hand, the type of uncertainty in the motor planning tasks was intrinsic and experience-based. Considering the studies reported here we can say that humans cannot reason normatively under risky situations when the risk is determined by described information about probability distributions while they are optimal when the

source of uncertainty is intrinsic. However, this discrepancy across conditions is derived from the findings of independent experiments.

A further attempt focused on testing the decision-making with intrinsic and described extrinsic uncertainty in a single task. In this experiment (Maloney, Trommershäuser, & Landy, 2006), in addition to intrinsic motor variability, the experimenters introduced explicit probability cases (see Figure 1.4). The results of this experiment showed that the performance of subjects was near optimal when the uncertainty originated solely from their own motor system (intrinsic probability) and thus under certainty conditions. On the other hand, their performance was sub-optimal when external uncertainty about rewards and penalties was introduced verbally (described extrinsic uncertainty) representing classical pencil and paper decision-making tasks (Maloney, et al., 2006).

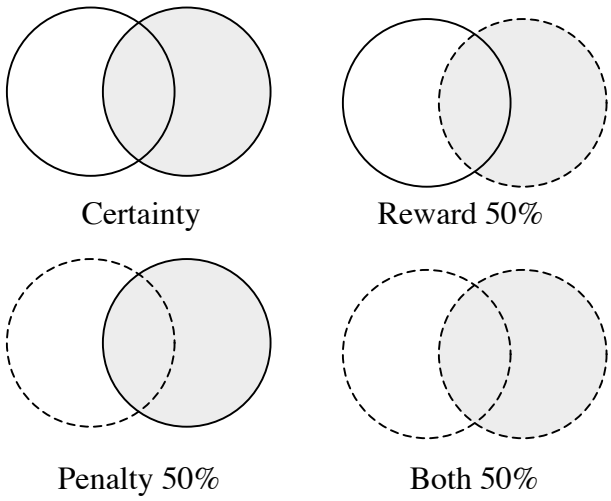

Figure 1.4. Certain and stochastic configurations used in Maloney et al. (2006). Dashed vs. Solid circles represent certainty vs. probabilistic (p=.5) consequences, respectively. Shaded vs. White circles represent reward vs. penalty regions, respectively. Figure is adapted from Maloney et al. (2006, pg. 308).

On the other hand, as stated earlier, Barron et al. (2003), Hertwig et al. (2004) and Weber et al. (2004) found that subjects who experienced outcomes and probabilities were closer to maximizing their gain than when they were presented with the information explicitly. Note that in Maloney et al., although subjects were told about the uncertainty conditions they were also experiencing the stochastic cases by getting feedback for their pointing responses for a number of trials (36). Given the findings reported by Barron et al., Hertwig et al., and Weber et al. and that subjects experienced extrinsic uncertainty in Maloney et al., one might expect that the introduction of extrinsic probabilities should not have disrupted their near optimal performance originally observed in certainty conditions. There are several possibilities that may have caused this discrepancy from optimal responding individually or jointly. Firstly, even though they may be closer to maximizing gain by experiencing the uncertainty, humans might never be optimal decision-makers when uncertainty is extrinsic. Consequently, when external uncertainties are introduced (regardless of their being described or experienced) to tasks where humans are known to be optimal, extrinsic uncertainties may result in sub-optimal performance. A second possibility is that experiencing extrinsic uncertainty may result in accurate representation of those probabilities; however, this information may not refer to a built-in variable in motor planning that defines the task in this experiment. The third possibility is the subjects might exhibit optimal responding after they have an accurate estimate of extrinsic probabilities (later in the block) however, the analyses considered the data as a uniform set. In either case, the results of this particular study consistent with classical decision-making tasks reveal non-optimality of humans in representing/integrating extrinsic uncertainties.

Maloney et al. (2006) further discussed the fact that classical pencil and paper decision making tasks do not resemble most real life situations; Information is presented in numeric or graphical format and the subjects are not given any justification regarding the state of affairs. There is a limited number of alternatives and there is no time constraint imposed on the subjects. The authors concluded that paper and pencil tasks don't resemble most of our daily situations and that "… the human capacity for decision making bears the same relation to the economic tasks of classical decision making as human language competence bears to solving the Sunday crossword puzzle." (pg. 26)

The optimality of decision-making performance in tasks that involved intrinsic/motor uncertainty motivated us to investigate the same question in the temporal domain where the intrinsic uncertainty originates from the representation of temporal intervals. Also, the discrepancy of performance in the processing of intrinsic vs. extrinsic uncertainties in the literature motivated us to investigate performance under these two types of uncertainties separately and jointly. Briefly, in this research, we investigated statistical decision-making that involves intrinsic and extrinsic uncertainty and different payoff matrices in the temporal domain. Different from most of the research done in this area, in this research we used both human and non-human subjects in the same task. Before presenting the details, it would be informative to talk about the implications of this research for different conceptual frameworks in contemporary psychology. In the following section we will talk about two main frameworks that are used in order to explain mental phenomena.

SECTION 2

INFORMATION-PROCESSING vs. ASSOCIATIONIST FRAMEWORKS

As reported in Section 1, risk assessment can be investigated in both humans and animals (e.g. Shafir, 2000). In the case of humans, information processing is a natural explanatory approach to this process under the assumptions of contemporary psychophysics. When this very question is transformed to animal learning however, one comes across a different approach, namely the associative learning. These approaches are the two sides of a major theoretical controversy regarding the functional architecture of mind. Exploiting its translational nature, in this research we used risk assessment as the cognitive context in which we compared and contrasted these two explanatory approaches. In specific, we tested and compared the performance of representational / computational and associative learning models in accounting for the quantitative and qualitative properties of the empirical data (gathered with a simple decision-making task with well-defined task relevant parameters).

A point of considerable interest –and one that motivates our contrastive modeling- is that the associative learning approach tries to minimize or even eliminate the representational assumptions that are inherent in information-processing models. Thus, in using both associative and information-processing models, we focus on the role of representational assumptions and the question whether they can be dispensed with. This question makes particular sense in cognitive science as it relates to the question of representational power needed to explain the cognitive processes. In the following subsection, we present the major differences between these two approaches to cognition.

2.1. Major Differences between two Approaches

The associative learning approach is different from representational / computational approach to mental phenomena in the sense that although different factors determine association formation and its strength, the associations do not encode any of the information that goes into the conditioning experience and the associations cannot be put into combinatorial operations. Briefly, according to these models conditioning process does not operate at the level of information but at the level of independent trials. It is the events rather than information that increment the associative values (see Gallistel, 2002 for further discussion). Under this view, associations cannot be treated as symbols since symbols contain information and can be put into combinatorial operations to form new ones. On the other hand, from an information processing point of view, conditioning is a mental phenomenon that is driven by information. Thus, in computational (information-processing) models of conditioning (Gallistel & Gibbon, 2000; Gibbon, 1977) the agent encodes the information that goes into conditioning experience such as the interval between events, puts them into combinatorial operations and use this information in order to decide for instance when to respond.

As one can note, the assumptions of computational view is fairly different from the associative learning framework. The computational view assumes a read-write memory, which encodes the information about the conditioning experience. Under this view, information encoded in the memory can be further put into combinatorial operations to generate new representations. Thus information-processing view of mental processes have two important components, namely representation and operation (input → representation → processing → output). Representations can be defined as mapping of

information and relations in the world to symbols in the mind. They are encoded, modified and updated in the read-write memory. They are selective in terms of the information they contain such that from a single experience, one can end up many representations each of which refers to different attribute of the experience. These symbolic representations can be put into combinatorial operations, which constitutes the computational part of this framework.

The optimal movement planner developed by Trommershäuser and colleagues represents an information-processing (or representational/ computational) model for a cognitive process. In their model, subjects are assumed to represent the metric properties of the surface area they are responding upon, the payoff matrix associated with different areas of this surface, internal uncertainty that originates from the sensory-motor system and extrinsic uncertainty that is imposed by the experimenters. Further, their model assumes that subjects could arithmetically integrate the information they represent, which results in empirical approximation to the optimal decision-maker (except in the case of integrating extrinsic uncertainty). Thus their optimal movement model satisfies both the representational and computational assumptions of the information-processing framework.

This framework is supported by the lack of convergence to the optimal point of responding over the course of trials, which would emerge if the subject randomly chooses a strategy and through rewards and penalties s/he adjusts her/his response over trials to maximize the gain. When Trommershäuser et al. (2003b) examined the end points over successive trials, they found that end points randomly fluctuated with constant variability around the mean – the optimal point of responding itself. This pattern of responding

shows that subjects solved the optimization problem from the very start of the experiment such that their course to the optimum solution was not charted from the empirical data. Consequently, having such a nature of responding over trials falsifies the reinforcement-learning explanation for subjects' performance. In the following section, we will introduce the basics of associative learning approach.

2.2. Basics of Associative Learning Approach:

Experimental psychology has been dominated by theories of learning that adopt associative processes as the main underlying mechanism of "mental phenomena". This influence started with Pavlov's experiments who found that previously neutral stimulus elicits conditioned response when it is repeatedly paired with unconditioned (intrinsically reinforcing or penalizing) stimuli (1927, 1928). This conditioned response represents the unconditioned response that unconditioned stimulus would normally elicit (such as salivation, freezing, etc.). Earlier than Pavlov, Thorndike in United States found that a response that resulted in pleasant outcome was more likely to occur again in future while a response that resulted in unpleasant outcome was less likely to occur again in future (1911). This latter phenomenon is called the "law of effect". The former experimental phenomenon is referred to as classical or Pavlovian conditioning and the latter one is referred to as operant, instrumental or Thordikean conditioning.

Based on these experimental phenomena, experimental psychology in between 30s and 50s was dominated by animal learning research and theories that adopted an associative learning framework (Hull, Skinner, Spence, Tolman, Rescorla & Wagner). In the second half of the 20th century, the associative framework was to a substantial extent replaced by information-processing framework in the area of human psychology, while it

continued to be a dominant framework in the area of animal learning. Recently the associative learning regained interest from researchers and theoreticians in the area of neuroscience and cognitive science (i.e. connectionist models). This resurrected form of associative learning theories tries to explain complex cognitive phenomena ranging from grammar acquisition to visual perception via associative processes. In this zeitgeist, it is important to consider this framework in modeling simpler cognitive phenomena.

Although classical and operant conditioning involve different events (CS → US and RESPONSE → REINFORCER, respectively), the elementary rules that govern the association formation in both cases are virtually same. Regarding classical conditioning, associative-learning models assume that upon experiencing CS-US contingencies, an association occurs between their "representations" (or "traces"). Therefore, later when the representation of CS is activated also the representation of US is activated. The logic that underlies stimulus-stimulus or stimulus-outcome associations (Pavlovian/classical conditioning) also applies to response-outcome associations, which is as mentioned above called "law of effect" (Thorndike, 1911). This simple principle was used in explaining operant conditioning as well. Under this view, it is assumed that an association is formed between operant response and its consequence(s) and the discriminative stimulus signals the validity of these associations.

According to associative learning models, learning emerges through basic rules that refer to the relation between the occurrences of events or stimuli. The primary law of association formation was considered to be the law of contiguity. This law states that two events are associated when they co-occur in space and time, that is there has to be a temporal pairing between CS-US. This law suggests that there is a critical temporal

window during which association between the "paired" stimuli can occur. This view also directs most neurobiological investigations of learning (Gluck & Thompson, 1987; Hawkins & Kandel, 1984; Fanselow, 1993). For instance, in neurobiology associations are explained by the temporal pairing of pre and post-synaptic activity (Gallistel & Gibbon, 2001).

Secondary laws of associative learning determine how strong the association will be. These secondary laws include the frequency, duration, recency and saliency of the co-occurrences. For instance, the associative strength between the representations of CS and US is claimed to increase as a function of the number of times the CS-US contingency is experienced.

Given the commonality of the rules that govern classical and operant conditioning, we will explain the functioning of the associative learning models via an associative learning rule that is well-accepted in explaining classical conditioning, namely Rescorla-Wagner model (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972) for adjusting the strength of an association. This formula has been called the general linear model (Yeschiam & Busemeyer, 2005) because of the frequency with which it has been employed in the quantitative modeling of the associative learning process.

$$\Delta V = c_{ij}(US - V)$$

$$V_{new} = V_{old} + \Delta V = (1 - \alpha)V_{old} + \alpha US$$

where V is the current strength of the association between CS and US, ΔV is the change in the associative value of the CS during a trial, α represents the learning rate ($0<\alpha<1$), and US is 1 or 0 depending on whether the CS is or is not followed by a US on any given

trial. We used this formula in adjusting the associative weights of CS in our associative models of the process underlying.

Note that according these models the effect of increments (when US = 1) is relatively higher when the weight of CS is relatively low and the effect of decrements (when US = 0) is relatively higher when the weight of CS is relatively high. This quantitative property of the model has been used to explaining the traditional learning curve (logarithmically increasing – see Figure 8.1 for examples), which has been empirically and conceptually challenged by Papachristos & Gallistel (2005).

In order to compare the performance of these opposing frameworks, we ran simulations that executed associative learning algorithms with different levels of substitution for representational power. In our task, the representational power refers to encoding of different information about multiple task relevant attributes. One of the critical task relevant attributes was the temporal interval. For instance, for a system that does not represent temporal intervals, there is no information that would aid the system to constrain its actions as a function of a temporal reference point. On the other hand, if the system could represent temporal intervals, it would know where its actions fall on the time line with respect to other events. One should note that the metric representation of temporal interval not only involves defining/constraining the action space but also gives a signed parametric redirection to the actions, as a function of their consequences. For instance, if one gets penalized for being late, with the representation of time that agent would have the information that would aid him/her to be earlier next time (if s/he wants to avoid the penalty). Another attribute of the task that would require representational power for maximizing the gain is the representation of the *magnitudes* of reward and

penalties. For a system that cannot represent these magnitudes, any magnitude of reward or penalty would have the same effect on the association that involves constituents of that particular experience (except the sign of change in the form of increment or decrement by a multiplicative factor). We will be simulating these associative models and investigate their behavior given different levels of substitution for representational power with respect to empirical data.

SECTION 3

METHOD

3.1. Human Experiment:

Subjects

The subjects were undergraduate student, graduate student, postdoctoral volunteers at Rutgers University. There were 7 subjects, four males and three females. The age of subjects ranged between 21 and 34.

Apparatus

The temporal parameters, presentation of the stimulus, and recording of the responses were controlled with a Macintosh running OS X 10.3.9. The experiments were written in MATLAB 7.04 using the Psychophysics Toolbox extensions (Brainard, 1997).

Procedure

In the human analogue of switch task, subjects were asked to catch a target that occurred either on the left side of the screen after a short duration (i.e. 2 seconds) or on the right side of the screen after a long duration (i.e. 3 seconds). The trial type (short or long) was not signaled to the subjects and thus subjects relied solely on their sense of time in deciding where the target might appear. Subjects started responding at the short-latency location to catch the short referent; when they judged that the short time had passed without the short-referent appearing, then they switched to the long-latency location to catch the long referent, assuming that the target would occur on the right side of the screen. The switch points in the long trials constituted the unit of analysis. The following section describes the details of the procedure.

Training Phase 1**.** The appearance of a cross in the middle of the screen signaled

the onset of a trial, while its disappearance signaled the offset. The duration of a trial varied unpredictably between two values, 2 or 3 seconds. The subjects were initially shown the two durations twice and told which one was the short and which one was the long one. Then they were presented with the short and long durations for 10 trials in a random order (5 Short and 5 Long trials). The subjects were asked to indicate if the signal duration they experienced was the short or the long one by pressing “S” and “L” keys, respectively. The subjects were given feedback immediately after each trial. Wrong responses were also signaled by a brief tone. This phase ensured that the subjects learned/re-experienced the reference durations before starting their test trials.

Practice Phase. The task in this phase was exactly the same as in the testing phase. Any given trial was unpredictably either a short or a long one. If it was a short trial, a square appeared to the left of the cross 2 seconds after the cross appeared; if it was a long trial, a square appeared to the right of the cross 3 seconds after the cross appeared. During each trial, the subject could control a frame by using “V” and “B” keys. When the subject pressed “V” key the frame would appear to the left of the cross, while it would appear to the right of the cross when the subject pressed “B” key (see Figure 3.1).

The subjects were instructed to "catch" the square by locating the frame on the correct side of the screen, namely the side where the target square would appear. The subjects were told to start a trial with a short response (at the left side of the screen) and when they thought the short duration has elapsed without an appearance of the target, to switch to long response (to the right side of the screen). The program prompted an error message with a relevant explanation (e.g. “Do not start with a long response”, “Do not release your finger”, “Do not press B after V”) if subject started with a long response, did

not press any keys and/or switched to short response after long response. If the frame was to the left of the cross when the square appeared after 2 seconds, it was a hit, while if the frame was to the right of the cross when the square appeared after 2 seconds it was a miss (we call this an *early switch*). Similarly, if the frame was to the right of the cross when the square appeared after 3 seconds, it was a hit, while if the frame was to the left of the cross when the square appeared after 3 seconds it was a miss (called *late switch*). In the case of a miss the square appeared in red and in the case of a hit the square appeared in green. The misses were also signaled with a brief tone. 10 trials of training block were presented to the subjects some of them being short trials and the rest being long trials. The order of presentation of short and long trials was randomized. The probability of short and long trials in the training block was equal to the probabilities that were used for that particular test session. By the nature of the task, the subject could observe his/her performance. In each trial subjects (re)experienced the critical short and long durations because the targets would appear either to the left side of the screen after the short duration or to the right side of the screen after the long duration irrespective of subjects response[4]. In this phase, the subjects were expected to form an estimate of their representational noise in the unlikely event that they did not bring one with them to the experiment.

Both Training and Practice blocks were presented at the beginning of each session. This allowed subjects to update their knowledge about the temporal parameters of the task. In the first session of the experiment, the subjects were presented with an

[4] In the pointing tasks summarized in the introduction, the estimate of variability was gathered by the subjects by observing the distance between where they aimed at and where they hit on the screen (e.g. Maloney et al., 2006; Trommershäuser et al., 2003a; 2003b).

animated video that demonstrated different responses in different kinds of trials (short and long) and their consequences (catching the short and long latency target, early switch and late switch).

Testing Phase. The design was the same as in the Practice Phase, except that the subjects were given their trial based and total scores/performance at the end of each trial. There were two kinds of targets and two ways to present trial-based and cumulative performance. The target could be a square and the gains and losses could be presented with numerals. More specifically, when the subject hit or missed the target, s/he was presented a positive or negative score related to the type of hit and miss, respectively. Then this score was added to the total score accumulated from the preceding trials.

The target could alternatively be a set of circles. In this case, the trial based performance was presented with the corresponding number of circles that were color coded with green and red for rewards and penalties, respectively. The cumulative performance was presented via a line that got longer with additional rewards and shorter with additional penalties. If the cumulative performance was below zero the color of the line was red and if it was above zero the color of the line was green.

The goal of the subject was to maximize his/her score. Experiment was run via one-hour long daily sessions. In each session, there was a total of 5 blocks, each consisting of 100 trials. Short and long trials were presented in according to the predetermined probabilities for that particular session. At the end of every 100^{th} trial, the subject was prompted to take a break. Subject was also allowed to take a break at the end of every 10^{th} trial.

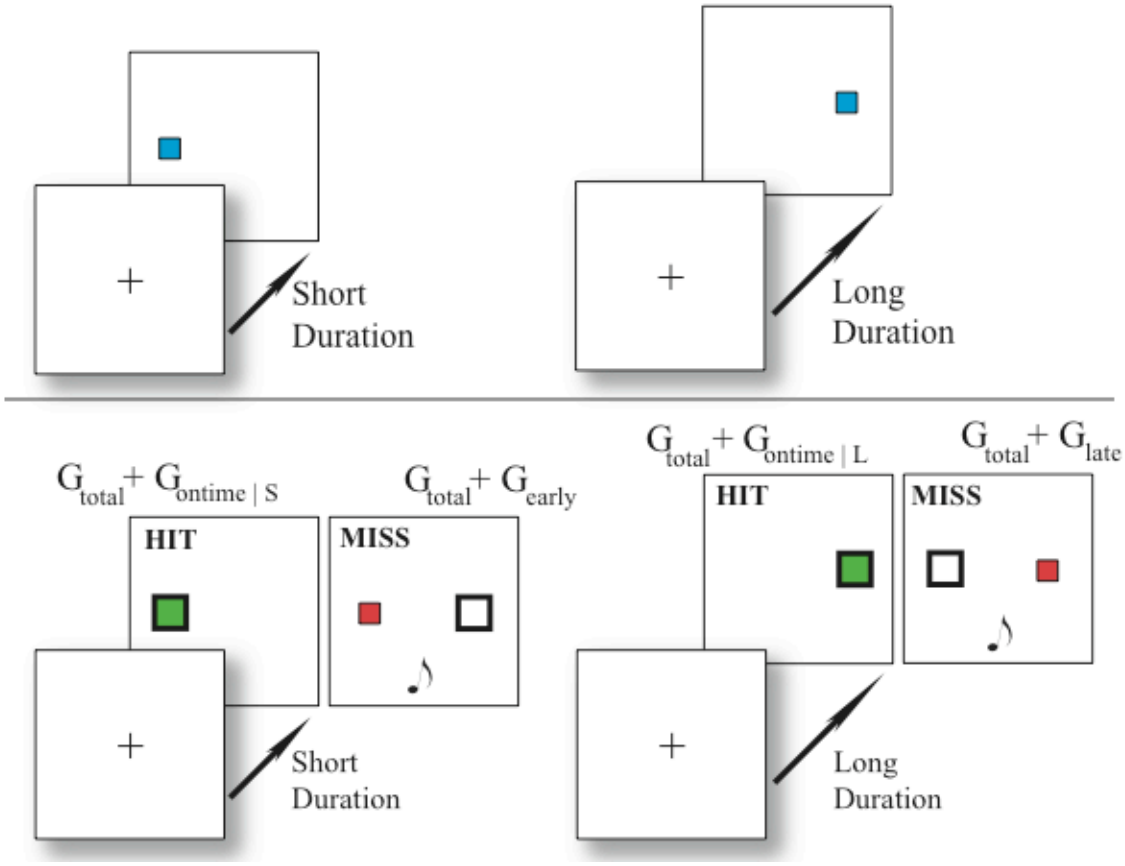

Figure 3.1. Illustration of the human analogue of switch task. The target (i.e. square) appears to the left of the cross after short duration, while it appears to the right of the cross after long duration. When the subject catches the target with the frame, gain associated with being on-time is added to the total gain. Similarly, in the case of switching late and early, the penalty points associated with switching late and early are subtracted from the net gain. Gain is presented with letter "G" in this plot. G_{total} stands for the total score that subject has accumulated since the beginning of the test session.

There are two parameters, which were varied in the expectation that they would have an effect on the location of switches with respect to trial time. The following are two of these parameters that were varied in this particular study.

1) Probability of short and long trials occurring

2) Payoff Matrix: Rewards and penalties associated with hits and misses on the two kinds of trials.

At the beginning of each session, subjects were informed about the payoff matrix but not the probabilities of different kinds of trials. Thus subjects would gather an estimate of probability of different trials during test trials. Table 3.1 illustrates the conditions that arise from the cross-tabulation of these parameters for a short-latency trial.

Note that in the top-left condition both the probability and reward magnitude would be expected to favor the long-latency outcome in planning the switch latency. The bottom-right condition on the other hand, both the probability and reward magnitude would be expected to favor the short-latency outcome in planning the switch latency. Specifically, these conditions would be expected to result in relatively early and late mean switch latencies, respectively.

Table 3.1. Tabulation of Task Parameters

Relative Probability of Short Trial	Relative Magnitude of Short Reward		
	Low Mag.	Equal Mag.	High Mag.
Low Prob.	Low P. & Low M.	Low P. & Equal M.	Low P. & High M.
Equal Prob.	Equal P. & Low M.	Equal P. & Equal M.	Equal P. & High M.
High Prob.	High P. & Low M.	High P. & Equal M.	High P. & High M.

In our human experiment, there was at most a 9-fold difference between the probabilities ($p = 0.1$ vs. $p = 0.9$) and at most a 50-fold difference between the magnitudes (e.g. -50 vs. –1, 50 vs. 5, -9 vs. –1, 1 vs. 9, 5 vs. 5, etc.). The conditions that

we used and their expected effects on the switch points are listed in Table 3.2.

Table 3.2. Task Parameters and Their Expected Effects

Parameters	Expected Gain
p(short) = p(long) & Gain(Short) = Gain(Long):	Neutral
p(short) = p(long) & Gain(Short) < Gain(Long):	Earlier Switch
p(short) = p(long) & Gain(Short) > Gain(Long):	Later Switch
p(short) > p(long) & Gain(Short) = Gain(Long):	Later Switch
p(short) > p(long) & Gain(Short) < Gain(Long):	Neutralized
p(short) > p(long) & Gain(Short) > Gain(Long):	Later Switch
p(short) < p(long) & Gain(Short) = Gain(Long):	Earlier Switch
p(short) < p(long) & Gain(Short) < Gain(Long):	Earlier Switch
p(short) < p(long) & Gain(Short) > Gain(Long):	Neutralized

Each condition was tested over 500 trials in a single session. Further, some conditions were tested twice with both types of targets, namely square vs. set of circles. Thus, these conditions were tested with both numerical and graphical presentation of performance. The order of conditions and target types was varied randomly for each subject. It was only data from the long trials that were analyzed. Therefore, the number of analyzable trials was less than 500 for each condition. The performance of the subjects was converted to monetary value and the subjects were notified about the conversion rate. Table 3.3 summarizes all the experimental conditions and Table 3.4 illustrates the task conditions tested with individual subjects along with the type of stimulus used. Task conditions were sorted in a random order for each subject.

Table 3.3. Task Conditions Tested

Task ID	P(short)	Reward(short)	Reward(Long)	Penalty(Early)	Penalty(Late)
1	0.1	0	0	1	9
2	0.1	1	9	0	0
3	0.1	1	9	1	9
4	0.1	5	5	0	0
5	0.1	9	1	0	0
6	0.1	9	1	1	9
7	0.5	0	0	1	9
8	0.5	1	9	0	0
9	0.5	1	9	5	5
10	0.5	5	5	0	0
11	0.5	5	5	5	5
12	0.5	9	1	0	0
13	0.5	9	1	9	1
14	0.9	1	9	0	0
15	0.9	5	5	0	0
16	0.9	9	1	0	0
17	0.5	5	5	50	5
18	0.1	5	50	5	5
19	0.5	5	5	1	50
20	0.1	5	1	50	5
21	0.5	9	1	50	5
22	0.5	5	50	9	1
23	0.5	5	50	50	5
24	0.5	5	5	5	50
25	0.1	50	5	5	5
26	0.5	50	5	5	5
27	0.1	9	1	50	1
28	0.1	50	5	50	5

3.2. Mouse Experiment

Subjects

Twelve naïve C57BL/6N female mice (Harlan, Indianapolis IN) were used in this experiment. They were eight weeks old upon arrival. The mice were housed individually in polypropylene cages. The cabinets that held the cages were lit on a 12:12 light/dark cycle (lights on at 8:00 p.m.). The experiments were run during the dark cycle. The mice

were maintained at 85% of their free feeding weight by being fed lab chow after each session. During the sessions, they were fed 20 mg Noyes food pellets (PJAI-0020) as a reinforcer. Water was available ad libitum in the home cages and experimental chambers.

Table 3.4. Task Conditions Tested by Subject and Stimulus Type. The numbers represent the corresponding conditions presented in Table 3.3. For each subject we present the task condition tested along with the stimulus type used in the session. The order of task conditions was randomized for each subject (not reflected in Table 3.4). "N" corresponds to numerical, "G" corresponds to graphical stimulus. NG and GN refer to those cases in which the subject was tested with the corresponding task condition using both symbolic and non-symbolic stimulus. The order of G-N pair shows the order in which the stimulus type was tested relative to the other stimulus type in that particular task condition.

Subject 1	1	7	9	10	12	13	15	16	23	26	28					Task ID
	G	GN	GN	GN	GN	GN	G	G	N	N	N					Stimulus Type
Subject 2	1	3	4	5	8	9	10	14	15	16	17	18	19	20	22	Task ID
	GN	GN	GN	GN	GN	GN	GN	GN	GN	GN	N	N	N	N	N	Stimulus Type
Subject 3	2	3	4	5	6	7	8	12	13	14						Task ID
	NG	NG	NG	NG	G	NG	N	NG	N	NG						Stimulus Type
Subject 4	1	9	10	13	15	16										Task ID
	G	NG	G	NG	G	G										Stimulus Type
Subject 5	2	3	4	5	6	7	8	12	17	18	23	24	25			Task ID
	N	N	NG	NG	N	N	N	N	N	N	N	N	N			Stimulus Type
Subject 6	11	19	21	27												Task ID
	N	N	N	N												Stimulus Type
Subject 7	1	3	4	5	8	9	10	14	15	16						Task ID
	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG						Stimulus Type

Apparatus

Six operant chambers (Med Associates, ENV 307-W: 21.6 cm x 17.8 cm x 12.7 cm) inside ventilated, sound attenuated boxes (Med Associates, ENV-018-M: 55.9 cm x 55.9 cm x 35.6 cm) were used in this experiment. Two opposing sidewalls were made of metal, the other two, and the ceiling, of clear Plexiglas. Each chamber had a grid floor with twenty-five evenly spaced metal rungs. Two pellet dispensers (ENV-203-20) were able to deliver food to two cubic hoppers (Med Associates ENV-203-20, 24mm on each side) spaced five inches apart on the side wall. On the opposite wall, another cubic hopper was used to initiate trials. Each hopper was illuminable, and equipped with an infrared beam that detected nose pokes. Along that same wall, a water bottle protruded into the chamber. A white noise generator was able to deliver an 80dB, flat 10-25,000 Hz white noise for programmable durations. The experiment was controlled by a software (Med-PC IV, Med Associates) that also logged and time-stamped the events—the onsets and offsets of interruptions of the IR beams in the station, the onsets and offsets of white noise and the delivery of food pellets. Event times were recorded with a resolution of 20 ms.

Procedure

Familiarization. Three days prior to the experiment, the mice were food deprived to 85% of their free-feeding weight. In addition, on each of these 3 days the mice were placed in the experimental chambers for 30 minutes, and five 20 mg Noyes food pellets (PJAI-0020) were left in their home cages to familiarize them with the pellets used in the experiment.

Training. This procedure is an adaptation of the switch paradigm proposed by

Balci, Papachristos, Gallistel, Brunner, Gibson, Shumyatsky (2007b). The task is summarized in Figure 3.2. The illumination of the control hopper (the hopper opposite the two feeding hoppers) signaled that a trial could be initiated. In order to initiate each trial, the mouse had to poke in this control hopper. This ensured that the mouse was in a fixed location at the start of every trial. Each trial started with both of the feeding hoppers illuminated, and with the onset of the white noise. The hopper lights turned off after a nose poke on either hole at or after the termination of the noise. The white noise stayed on for a preprogrammed duration (termed "short" or "long") with a predetermined probability.

Each of the two feeding hoppers was associated with either the short or long duration. A reward was only delivered if the first nose poke at or after the termination of the noise was in the appropriate hopper. After each trial, there was a 30 second fixed delay plus a variable interval of 60 seconds (drawn from an exponential distribution) before the next trial could be initiated. There were 3 different phases in which the probability of a given trial being short were manipulated. Each phase continued until the performance of the mice stabilized.

Phases One through Three. Our previous results suggested that mice could not discriminate 1:2 duration ratio when they constitute their initial training set (Balci et al., 2007b). Thus, we started out with 1:3 duration ratio, where the short duration was 3 seconds and the long duration was 9 seconds. After training with this pair of temporal intervals, and after several parameter manipulations (i.e. changing the probability of a given trial's occurrence), we decided to reduce the ratio to 1:2; 3 second short trial and a 6 second long trial. This ratio was used throughout the rest of the experiment.

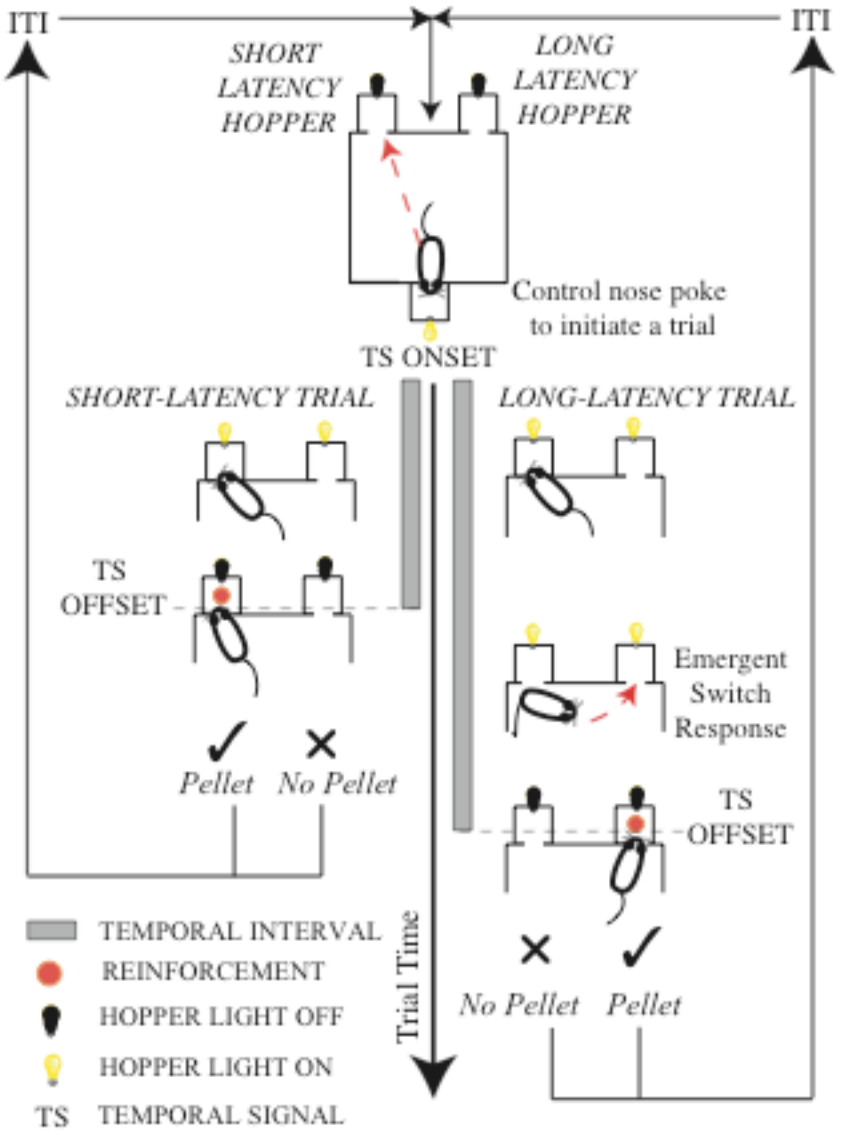

Figure 3.2. Illustration of switch task as it is used in mice adapted from Balci, Papachristos, Gallistel, Brunner, Gibson, & Shumyatsky (2007b). Mice self-initiated their trials by nose poking into the illuminated control hopper (top panel). Nose poke to the control hopper immediately initiated the auditory stimulus that signaled the critical temporal interval. A given trial was either a short-latency (left panel) or a long-latency trial (right panel). In almost all trials, mice started nose poking into the short-latency hopper. In long trials, somewhere in between short and long intervals, mice departed from the short-latency hopper for the long-latency hopper. They received pellets for the correct responses and incorrect responses were not reinforced (bottom panel). Each trial was followed by a variable intertrial interval.

Phase Four. In phase four, the probability of a given trial being short was .5. This

phase was run for 22 daily sessions.

Phase Five. In phase five, the probability of a short trial was .25 for half the mice (Group 1) and .75 for the other half (Group 2). This phase was run for 20 daily sessions.

Phase Six. In phase six, The probability of a short trial for Group 1 was .9, while for Group 2 the probability was .1. This session was run for 22 daily sessions.

Data Analysis

Data from each subject was pooled across sessions for each phase. This produced a rich data set of switch latencies (the point, in seconds, the mouse left the short hole to go to the long hole) for each subject. Switch latencies that occurred after 9 seconds and the switches that were followed by extraordinarily large travel times (maximum travel time that has been observed before 6 seconds plus 3 seconds) were excluded from the data. These data points constituted a small portion of the whole data set (e.g. % 2) and ensured that the data set did not include those switch points where the mouse clearly was not participating in the task.

In several individual mice, the cumulative distribution of switch latencies tended to reveal a mixed distribution. This strongly suggests that two processes were at work in the brain of these particular mice. Earlier switch distribution can be attributed to impulsivity (Balci & Freestone, in preparation) and the later switch distribution can be attributed to timed responses. Figure 3.3 shows two examples of such a distribution (Left Panel). In these cases, a mixed distribution, consisting of a Weibull and a Gaussian distribution, were fit to the data using the maximum likelihood estimates. The psychophysical function was modeled by the Gaussian component of the mixes distribution (Figure 3.3, right panel).

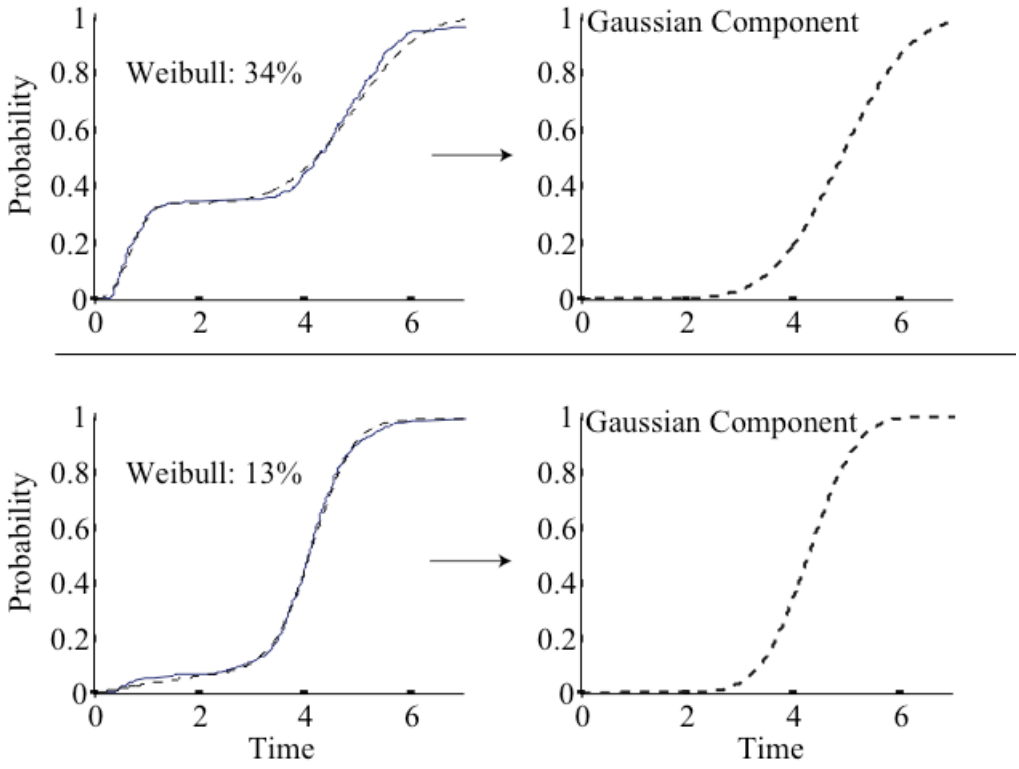

Figure 3.3. Cumulative distributions of the switch latencies. In the Left panels, the cumulative distributions of the switch latencies (solid curves) were fit using a mixed Weibull and Gaussian distribution (dashed curves). The right panel plots only the Gaussian component of the best fitting mixture distribution. The mean and standard deviation of this Gaussian component provided our estimates of the subject's target switch time ($\hat{T}$) and timing variability ($\hat{\sigma}$).

SECTION 4

OPTIMAL DECISION-MAKING IN TEMPORAL DOMAIN

An optimal temporal decision-maker was developed to predict the optimum point of switching during signaled temporal intervals (trial duration) that would maximize one's gain. Although, in real life it is generally difficult to define a common metric for gains/losses for decision-makers, our experimental task allowed us to define these quantities. The calculation of the optimum point of responding involved integration of three sets of parameters (see Figure 4.1 for illustration): 1) Extrinsic Uncertainty (probability of different trials), 2) Payoff Matrix (rewards/penalties associated with switching at different points during different trial types), and 3) Intrinsic Uncertainty (precision in representing temporal intervals in the memory or timing itself). To further clarify the relevance of last parameter, intrinsic uncertainty results in discrepancies between when the subject intends/aims to respond (with respect to trial time) and when s/he actually responds (see Figure 4.2). For example, in our task, the subject might have planed to switch at 2.5 s. [a point that lies between the short (2 s.) and long duration (3 s.)], however, the actual response might occur at points that lie outside this critical range (e.g. 1.9 s.). Note, the shorter this range (i.e. 2 vs. 2.5 rather than 2 vs. 3 s.), the more likely a subject with a given degree of internal uncertainty will fall out of this critical range.

The optimal temporal decision-maker computes the mean switch latency that would maximize the expected gain for different task-relevant conditions and their combinations (extrinsic uncertainty, intrinsic uncertainty and the payoff matrix). The

computation executed by optimal decision-maker might result in intuitive predictions about the optimum point of responding. For instance, if the probability of a short trial is much higher than the long trial (along with the assumptions of symmetrical payoff matrix and uncertain representation of temporal intervals), the subject would maximize its gain by aiming at a point that is farther beyond the end of the short reference duration.

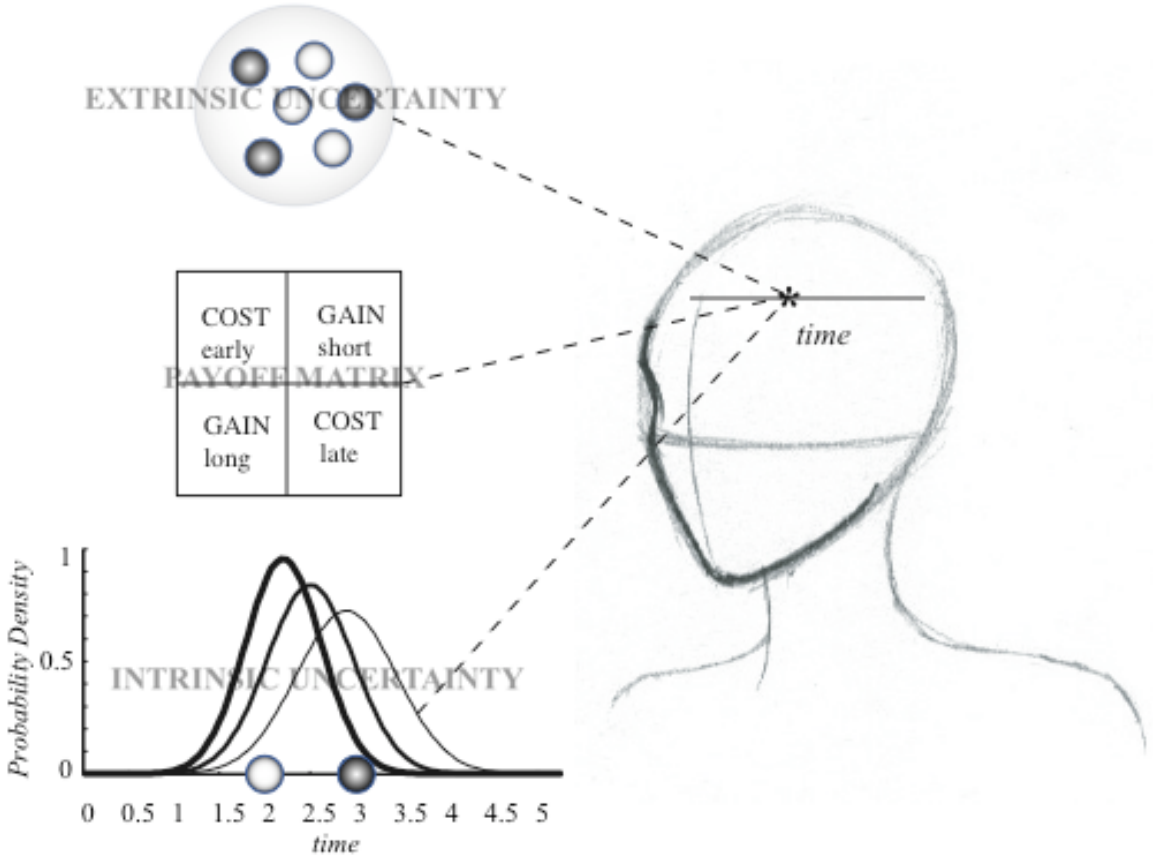

Figure 4.1. Illustration of integration of extrinsic uncertainty, payoff matrix, and intrinsic uncertainty in decision-maker's mind. In this particular illustration, the probability of short trials (light circles) and long trials (dark circles) is equal (top set). The middle cross-tab illustrates the all-possible outcomes, namely payoff matrix. In the bottom panel, short (2 s) and long (3 s) trials are located on the corresponding locations on the time line. The very left distribution that is centered around the short-latency (light circle) illustrates the error distribution around the subject's representation of the short temporal interval (see the bottom panel). The very right distribution that is centered around the long-latency (dark circle)

illustrates the error distribution around the subject's representation of the long temporal interval. The middle distribution represents the error distribution around the subject's representation of a temporal interval that lies in between short and long temporal intervals. As one can note, the error distributions get flatter as they are centered on the longer durations. This illustrates the scalar variability around the representation of temporal intervals. The coefficient of variation (standard deviation/mean) is constant across three distributions. Thus, when these distributions are normalized by their means (relative time scale), they would superpose. More importantly, this figure illustrates the integration of extrinsic uncertainty, payoff matrix and intrinsic uncertainty by the subject in planning the switch latency that would maximize his/her gain.

In another case, the probability of short and long trials might be equal but the penalty of switching late (x>Long) might be much higher than the penalty of switching early. Subject in this case would maximize its gain by aiming at a point that is closer to the short reference duration. These shifts in the optimal aim points might be magnified or canceled out by asymmetry in both probability and payoff matrix in the same or different directions, respectively. For instance, the probability of short latency trial might be 0.9 whereas the reward of catching long referent might be 9 times larger than for catching short latency. In such a case, the effect of the asymmetry in rewards would cancel out the effect of the asymmetry in the probabilities on the optimum point of responding.

Our optimal decision-maker model gives a quantitative prediction about the optimum point of responding by integrating all of these parameters, namely extrinsic uncertainty, payoff matrix, and intrinsic uncertainty. The model computes the optimum point of responding by calculating the expected gain for different temporal aim points. The algorithm used in this model is presented in the following equation and notations used are explained in Table 4.1.

Expected Gain Equation

$$
\begin{aligned}
E(g) = {} & g(d < S)\,p(s)\,\Phi(S,T,wT) && \text{a}\\
& + g(d > S)\,p(s)\left(1 - \Phi(S,T,wT)\right) && \text{b}\\
& + g(a < L)\,p(L)\,\Phi(L,T+s,wT) && \text{c}\\
& + g(a > L)\,p(L)\left(1 - \Phi(L,T+s,wT)\right) && \text{d}
\end{aligned}
$$

Note. a) Expected loss from premature departures b) Expected gain from appropriately delayed, hence aborted departures c) Expected gain from in time arrivals d) Expected loss from late arrivals. Differentiating with respect to T and solving for T with derivative $E'(g)$ set to 0 gives the analytic expression for optimal point. However, we determined the optimal point by numerical/graphical methods.

Table 4.1. Notations used in the formulation of the optimal decision-making model.

d	Departure time
a	Arrival time
$s = a - d$	Time required to switch
S	Short fixed interval
L	Long fixed interval
T	Target switch time (decision criterion)
w	Weber fraction (subject's coefficient of variation)
$p(s)$	Probability of a short trial
$p(l) = 1 - p(s)$	Probability of a long trial
$g(d < S)$	Loss from early departure
$g(d > S)$	Gain from on time departure
$g(a < L)$	Gain from on time arrival
$g(a > L) =$	Loss from late arrival
$p(d < S) = \Phi(S,T,wT)$	Probability of an early departure
$p(d > S) = 1 - \Phi(S,T,wT)$	Probability of an on time departure (not early)
$p(a < L) = \Phi(L,T+s,wT)$	Probability of an on time arrival
$p(a > L) = 1 - \Phi(L,T+s,wT)$	Probability of a late arrival

Note that the expected gain computation takes 10 quantities as its arguments: 4 gains/losses, the durations of the short and long intervals, the subject's target time and variability, the time required to make the switch, and the probability of a short trial.

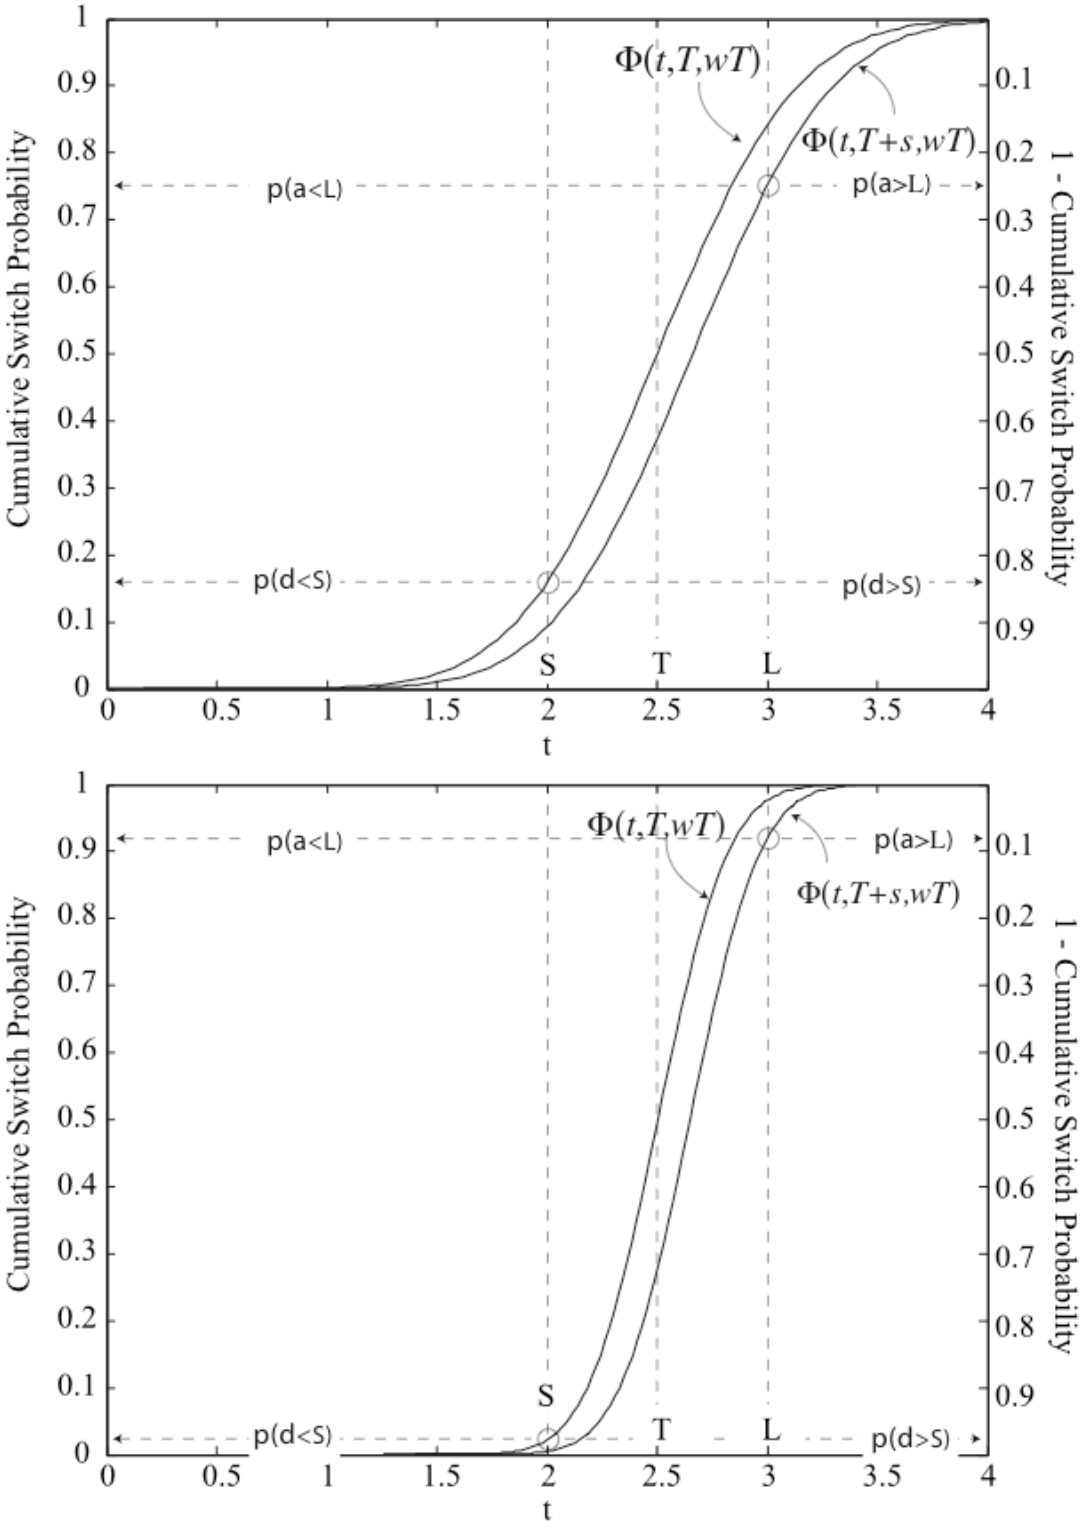

Figure 4.2. Cumulative departure and arrival distributions for w =.2 (top) and w = .1 (bottom), with the relevant probabilities, $p(d < S)$, $p(d > S)$, $p(a < L)$, $p(a > L)$. These probabilities are obtained by evaluating the cumulative normal distribution

functions $\Phi(t,T,wT)$ and $\Phi(t,T+s,wT)$ at $t = S$ and $t = L$, respectively. $\Phi(S,T,wT) = \int_{-\infty}^{S} \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(t-T)^2}{2w^2T^2}} dt$, the cumulative normal distribution.

The expected gain was computed for multiple Ts that were uniformly distributed between 0 and L+1 (L+3 for mice) seconds at small increments (e.g. 0, .005, .01, …, 3.995, 4 for humans). T that resulted in the maximum expected gain was defined as the optimal point of switching. The optimal T thus depends on 9 quantities. Figure 4.3 illustrates the output of this computation for different levels/conditions of external and internal uncertainty (2 of these 9 critical quantities).

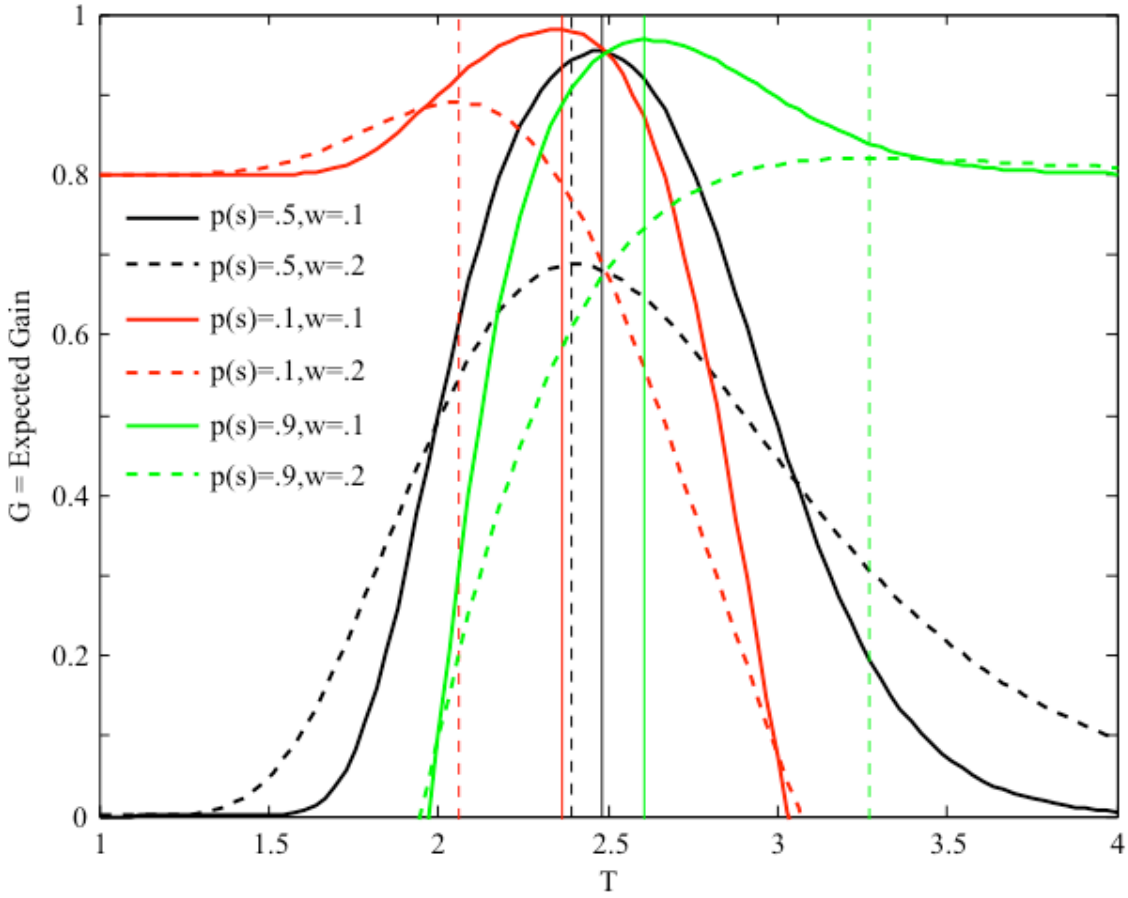

Figure 4.3. The effects of the relative probability of short and long trials versus the effects of different levels of timing variability (noise, uncertainty) on the expected gain curve. Vertical lines show the modes (locations of the optimal target time, T_o). The computations were made with $g(d < S) = -1$, $g(d > S) = 1$,

$g(a < L) = 1$, $g(a > L) = -1$. The locations of the optimal targets do not change if the negative gains are made 0. The solid curves are computed with $w = .1$, that is, +/- 10% temporal precision; the dashed curves are computed with $w = .2$, that is, +/- 20% temporal precision. The different colors are for different relative probabilities of short and long trials: red = p(s) = .1, that is, long trials are 9 times more likely than short trials; black = p(s) = .5, that is, the two kinds of trials are equally likely; green = p(s) = .9, that is, short trials are 9 times more likely than long trials. Note that the effect of a twofold change in timing imprecision on the location of the optimal target (vertical lines) is greater than the effects of an order of magnitude change in the relative likelihoods of the two kinds of trials. Thus, accurate internal noise representation (accurate representation of the variability in one's timing) is critical to the determination of an optimal switch criterion.

Figure 4.4 illustrates the interactive effect of complimentary asymmetries in the payoff matrix and external uncertainties.

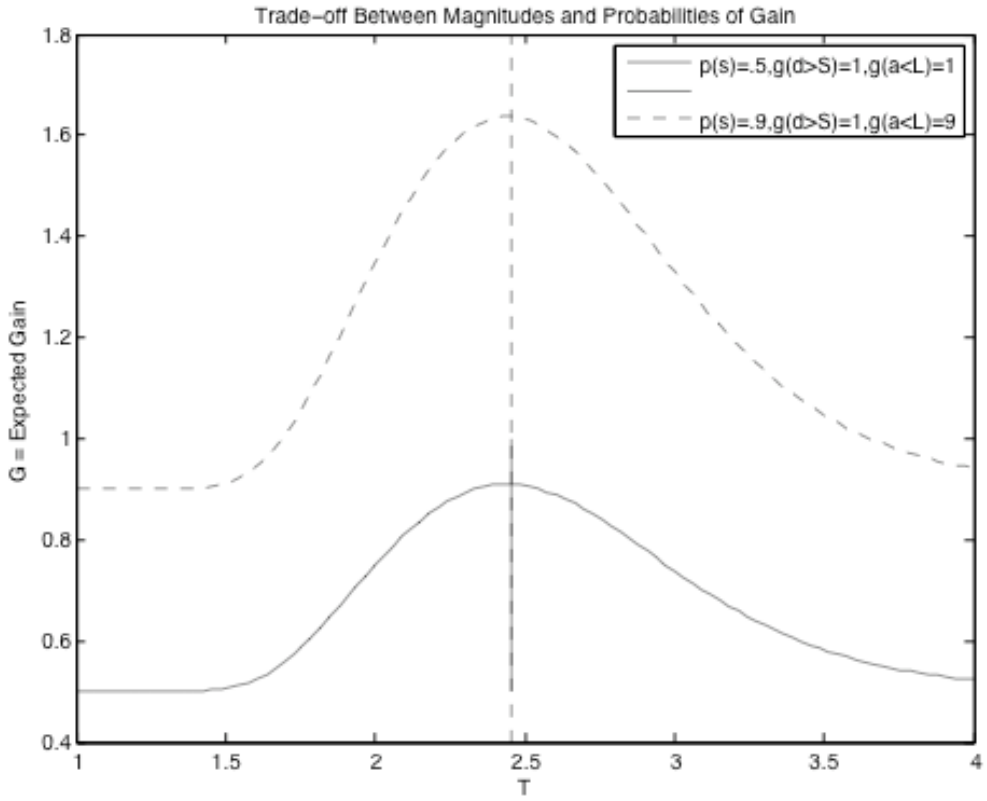

Figure 4.4. The effect of a 9-fold difference in gain magnitudes exactly offsets a 9-fold difference in probability in the opposite direction, but only if both the magnitude ratio and the probability ratio are accurately represented and combined multiplicatively.

SECTION 5

DATA ANALYSIS

With our data analyses, we sought answer the following questions:

1) How close is the subject's average switch latency to the optimal point? In making this assessment we take into account both the uncertainty regarding the location of the optimal point (due to empirical uncertainty about the true value of the subject's variability) and uncertainty regarding the location of the subject's target switch latency (due to uncertainty about the true value of the mean).

2) How optimal the subject is when optimality is defined as the proportion of the maximal possible expected gain (MPEG) for a subject with a given level of timing variability?

3) How do the performance of the subject change as a function of task parameters: such as the external probabilities, the pay-off matrix, and the manner in which gains and losses are communicated?

4) Does the variability of the switch latencies change as a function of task parameters?

The optimal value, T_o, of the target switch time, T, is a function of our estimate, $\hat{\sigma}$, of the standard deviation of the subject's distribution of switch latencies. This estimate may come either directly from our within-subject sample of the switch latencies under a given set of conditions. Or, it may come from our estimate of the subject's Weber fraction, $\hat{w} = \hat{\sigma}/\hat{T}$, using the $\hat{w}$ derived from the data across all the sessions, on the assumption of scalar variability, that is, a constant proportion between σ and T. The target switch time, T, is estimated by $\hat{T}$,which is the mean of our sample of the

subject's switch times under a given set of conditions. Thus, the likelihood function for the parameter estimates in our statistical model of the subject's switch times is a likelihood surface above the 2D plane of parameter estimates, $\langle \hat{\sigma}, \hat{T} \rangle$. The mode of this likelihood surface (the location of the summit of the likelihood hill) is the likelihood maximizing point in the plane of possible parameter values[5]. The function, $T_o = f(\hat{\sigma})$ is a curve in this same plane. Suppose we make contours (level curves) at octave intervals around the summit of the likelihood hill, so that each successively more remote contour represents a reduction by a further factor of 2 in relative likelihood, and we project these contours onto the $\langle \hat{\sigma}, \hat{T} \rangle$ plane. If the $T_o = f(\hat{\sigma})$ curve intersects the first (innermost) contour, then we can say that $\hat{T}$ is statistically indistinguishable from our estimate, $\hat{T}_o$, of the optimal value. If, on the other hand, the $T_o = f(\hat{\sigma})$ curve intersects the fifth contour but not the innermost four, then we could say that our estimate of the subject's target switch time is statistically distinguishable at something like the .05 level from the optimal. If $T_o = f(\hat{\sigma})$ curve misses the 6 innermost contours, then we could say that it is clearly statistically distinguishable from the optimal point. In the cases where $\hat{\sigma}$ derived

[5] In calculating the likelihood function for humans, we took account of the fact that the data were censored by the fact that subjects could not switch after the long reference duration. This means that on trials when subject demonstrated late switches, it is not known where the subject would have switched if the trial had not ended. To account for this, we replaced the data points that corresponded to late switches with their estimated counterparts. In order to implement this correction for censored data set, we fitted Gaussians to the data using Matlab's, Distribution Fitting Toolbox, with the late switches flagged as censored points. We then drew samples from this best-fitting distribution and replaced the late switch points with the points that were equal or larger than the long reference duration in our data set.

from our cross-session $\hat{w}$, we used $\hat{\sigma} = \hat{w}\hat{T}_{ses}$, where $\hat{T}_{ses}$ is the sample mean for that session**.** Since our data was fit by a normal distribution better than a lognormal distribution, in gathering our measures we instead used a parameter space that was composed of $\hat{\sigma}$ and $\hat{T}$ instead of $\hat{w}$ and $\hat{T}$.

The measure of statistical distinguishability may not be a meaningful measure from the subject's point of view, however, because a switch time clearly distinguishable statistically from the optimal might nonetheless lead to only a trivial reduction in the subject's expected gain. Another, in some ways more instructive, analysis is to plot the *expected gain* surface above the $\langle \hat{\sigma}, \hat{T} \rangle$ parameter plane, make percent reduction contours around its summit (99% of maximum, 98% of maximum, 97%, etc), project them down onto the parameter plane, and see what the highest contour is that is intersected by the $T_o = f(\hat{\sigma})$ curve. This constituted our second measure of optimality.

In constructing this measure, we used two different measures of the proportion of MPEG (Maximum Possible Expected Gain). In the first case, the proportions were computed simply as a function of the maximal possible expected gain (see Figure 5.1, bottom panel). Thus, these proportions could take both positive (e.g. 0.95) and negative values (e.g. –0.25). They would take negative values when MPEG (at T_o) had a positive value (e.g. 25) and the expected gain (EG) at $\hat{T}$ had a negative value (e.g. –50 → -50/25 = -2). In most cases, this constituted the liberal definition of the optimality, in the sense that the proportion of the maximum possible expected gain computed in this way was closer to 1. In the second case, these proportions were computed by dividing the interval between a given expected gain and the minimal expected gain by the interval between the maximal and minimal possible expected gains:

$$\frac{EG_{\hat{T},\sigma} - \min(EG_{\sigma}^{2\to3})}{\max(EG_{\sigma}^{2\to3}) - \min(EG_{\hat{\sigma}}^{2\to3})}$$

where EG is the expected gain (see Figure 5.1 top panel).

In this latter case, we considered only those expected gains that were expected to occur for the range of 2-3 s, which were the short and long temporal referents of the task. In most cases, this constituted the more conservative definition of optimality.

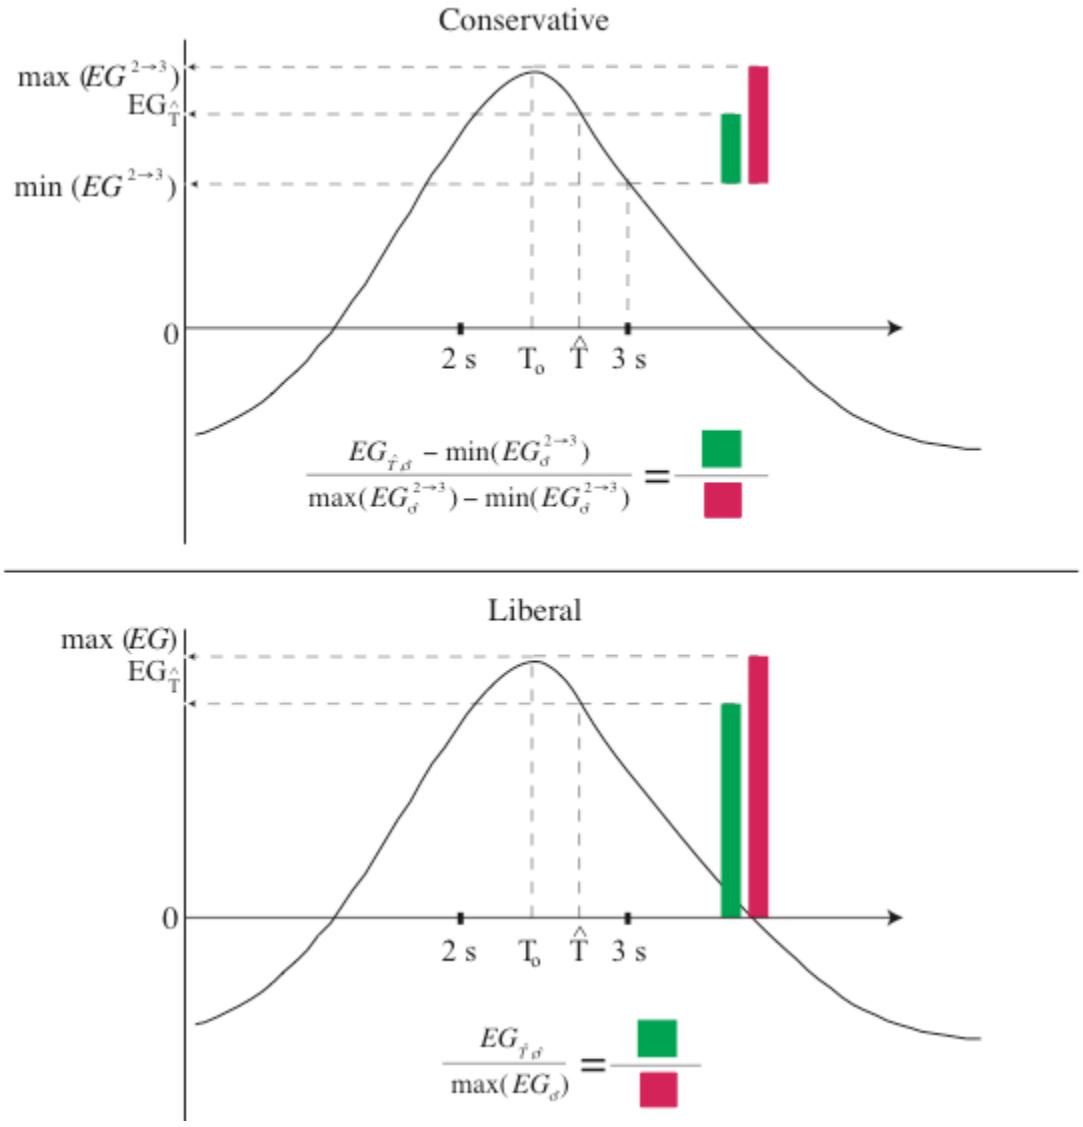

Figure 5.1. Illustration of calculation of liberal (bottom panel) and conservative (top panel) proportions of MPEG. In the case illustrated in this figure, EG within

2 and 3 seconds was always positive. In these situations, calculation illustrated in the top panel refer to a relatively conservative estimate of proportion of MPEG and the calculation illustrated in the bottom panel refer to a relatively liberal estimate of proportion of MPEG. When the EG turns negative at any point within this interval, the calculation illustrated at the top panel gives a relatively liberal estimate of proportion of MPEG and the calculation illustrated at the bottom panel gives a relatively conservative estimate of proportion of MPEG. Given the higher frequency of the former cases in our data set, we adopted the terminology illustrated in this figure.

Figure 5.2 demonstrates the superimposition of likelihood function and the EG function on the parameter plane of $\langle \hat{\sigma}, \hat{T} \rangle$ using actual data gathered from a single session. The contours of the likelihood function are indicated by concentric blue circles. The parameter estimates that maximized the likelihood of data are indicated by the cross within the inner most circle. Each contour delimits the likelihood function at 1 natural log unit decrements. Brown contours superimposed on the same parameter plane are level curves of the expected relative gain function for a subject with perfect timing. Expected relative gain is the gain expected relative to the maximum possible expected gain--for a subject with perfect timing precision.

These contours allow us to judge how well a subject is doing relative to the perfect subject, the subject with no timing variability. In the example shown in the top panel of Figure 5.2, the mode of the likelihood function (the x) falls between the .90 and .95 relative gain function. Thus, we estimate that this subject had an expected relative gain of about 93% of what the perfect subject could expect. However, we assume that the subject's timing variability is not a controllable parameter of their behavior, unlike their target switch latency. Thus, what one really wants to estimate is how close the subject's target switch latency, T , is to the optimal target, T_o, for a subject with our subject's

temporal variability. The heavy brown curve labeled “ridge” in the top panel is the locus of optimal target latencies as a function of assumed subject timing variability.

It specifies what the target switch latency should be for a subject with a given level of timing variability. As the given level of timing variability increases, the maximum gain a subject can expect decreases relative to the gain that a perfect subject could obtain. Thus, this curve traces a relative expected gain ridge. If our estimate, $\hat{T}$, of a subject’s target latency is on this ridge line, then we estimate that the subject’s target latency is at the maximum relative to what a subject with (our estimate, $\hat{\sigma}$ of) that timing variability could obtain. The real measure of how a subject is doing in terms of expected gain is how far down the subject is from this ridge (rather than from the peak on which the perfect subject sits). Thus, in the bottom panel, we plot these relative gain contours (thin brown curves).

In our example, the modal point of the likelihood function (the point whose coordinates are our best estimates of the subject’s variability and target latency) is essentially on the ridge. The blue likelihood contours around this point allow one to judge how far the subject might be from the ridge. The first such relative likelihood contour falls almost entirely between the two .99 relative gain contours, so odds are $1/e^{-1} = 2.7:1$ that the subject’s target latency would produce a gain that was within 99% of the maximum gain a subject with his/her level of variability could obtain. In what follows, when we report estimated relative gain values, they take into account our estimate of the subject’s timing variability. In other words, they specify how high the subject is relative

to the elevation of the nearest point on the relative gain ridge.

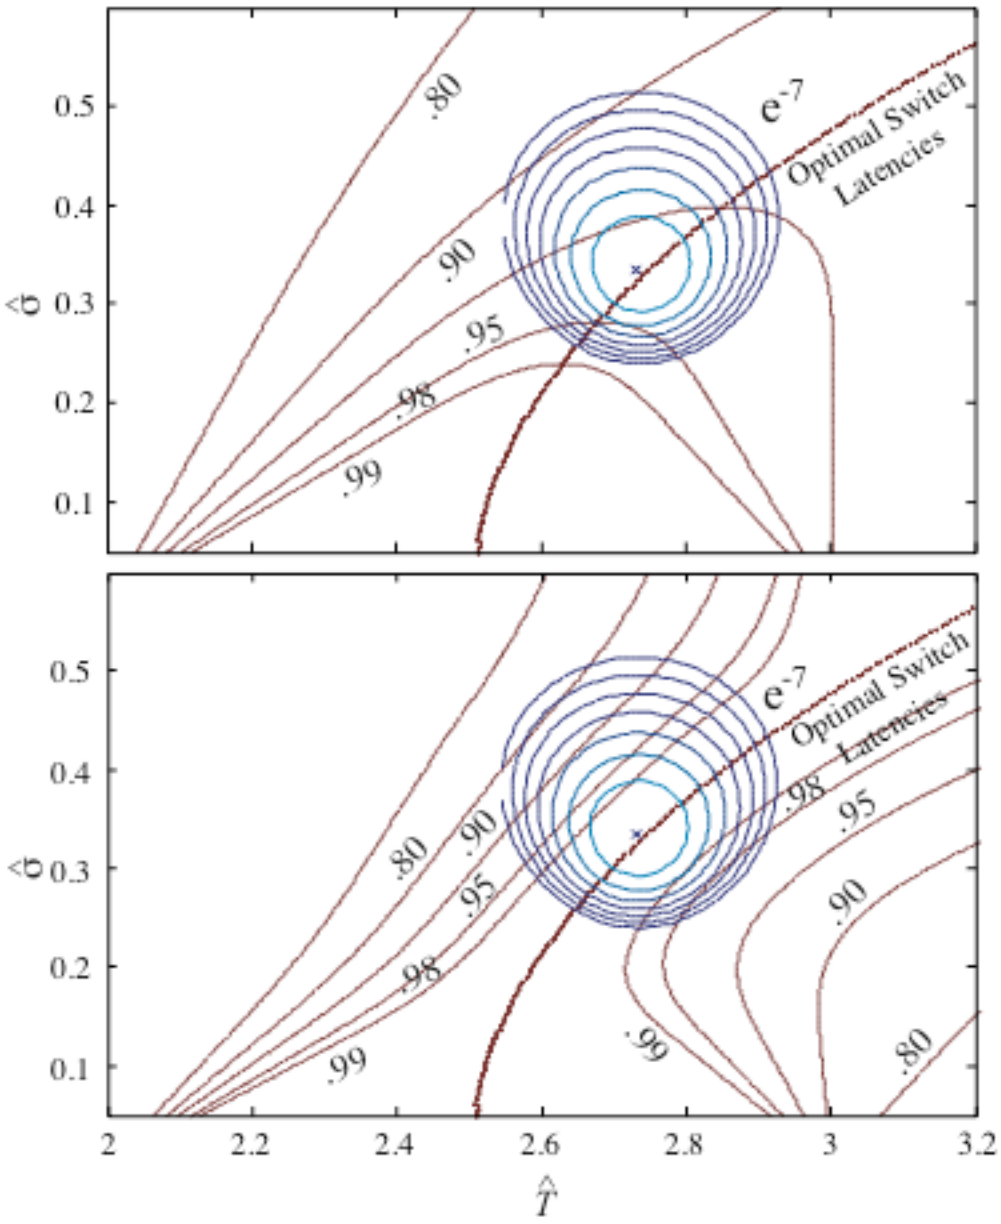

Figure 5.2. Expected gain and likelihood functions. Top. The thick brown line is the locus of the optimal points of responding at different estimates of timing variability ($\hat{\sigma}$). Concentric blue ovals are contours (level curves) of the likelihood function at natural log unit intervals. Thus, the innermost contour delimits the space of parameter values within which the relative likelihood of the data is e^{-1} = .37 of the maximum likelihood; the next contour delimits the space of relative likelihoods > e^{-2}, and so on. The thin brown curves are contours (level curves) of

the expected gain function when there is no noise in the estimation of time. The bottom-most relative-gain contour delimits the space (combinations of variability and decision criterion) within which the subject obtains more than 99% of the maximum possible gain. The remaining contours delimit the 98%, 95%, 90% and 80% regions of the parameter plane. Bottom panel. Same as top panel, except that now the gain contours have been normalized with respect to the maximum possible gain at a given level of variability. Geometrically speaking, the thick brown line plots the locus of the gain ridge, and the contours delimit percents of this ridge. As the subject's variability increases, the height of the gain ridge decreases, because the subject's variability sets a limit on how well the subject can do relative to a subject that can estimate time perfectly. We are interested in the proportion of the maximum gain that a subject with a given assumed level of variability in fact obtains, given our estimate of the subject's target switch latency.

In order to determine how far in temporal units the subject's mean switch latency was from that subject's optimal switch latency, we calculated the distance $\hat{T} - \hat{T}_o$), where $\hat{T}$ is our best (maximum likelihood) estimate of the subjects target switch latency and $\hat{T}_o$ is our estimate of the optimal target given $\hat{\sigma}$, which is our best (maximum likelihood) estimate of the subject's timing variability (the length of the red "Temporal Distance" line in Figure 5.3). This is an estimate of the subject's systematic temporal error.

The subject's statistical distance from optimality is estimated by the relative likelihood of the most likely point on the relative gain ridge. In Figure 5.3, the green line runs from the best estimate of the subject's parameters (the peak of the likelihood function) to the most likely point on the relative gain ridge. The likelihood of that point as a fraction of the maximum likelihood is our measure of a subject's statistical distance from optimality.

Finally, we estimate the subject's *relative gain loss*. This is the expected gain at $\langle \hat{\sigma}, \hat{T} \rangle$ relative to the expected gain at the $\hat{\sigma}$ point on the relative gain ridge (the point where the ridge crosses a horizontal line drawn through $\langle \hat{\sigma}, \hat{T} \rangle$).

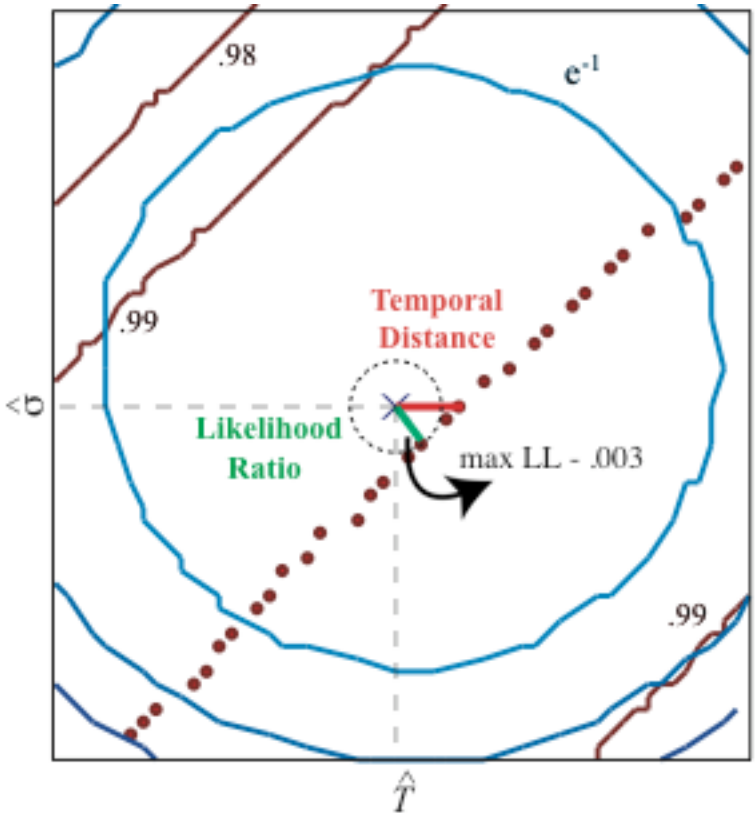

Figure 5.3. A closer look at the highly likely parameter space on the bottom panel of Figure 5.2. On this plot, we indicate our critical measures of optimality: i) the systematic temporal error (the length of the red line); ii) the statistical distance of the subject from the statistically nearest optimal point, which is the point where the green line meets the relative gain ridge. This latter measure is the relative likelihood of the most likely point on the ridge—that is the likelihood of the point at the other end of the green line relative to the likelihood of the data at the point marked by the X; iii) Relative gain loss, is the expected gain at the X end of the red line segment, relative to the gain at the other end, where it intersects the ridge line.

SECTION 6

RESULTS OF HUMAN EXPERIMENTS

General Results:

For each session, we estimated the mean (target) switch latency ($\hat{T}$) and the optimal switch latency ($\hat{T}_o$), given the estimate of $\hat{\sigma}$ for that session. To allow visual comparison of the $\hat{T}$s and $\hat{T}_o$s, Figure 6.1 plots T_os as a function of $\hat{T}$s.

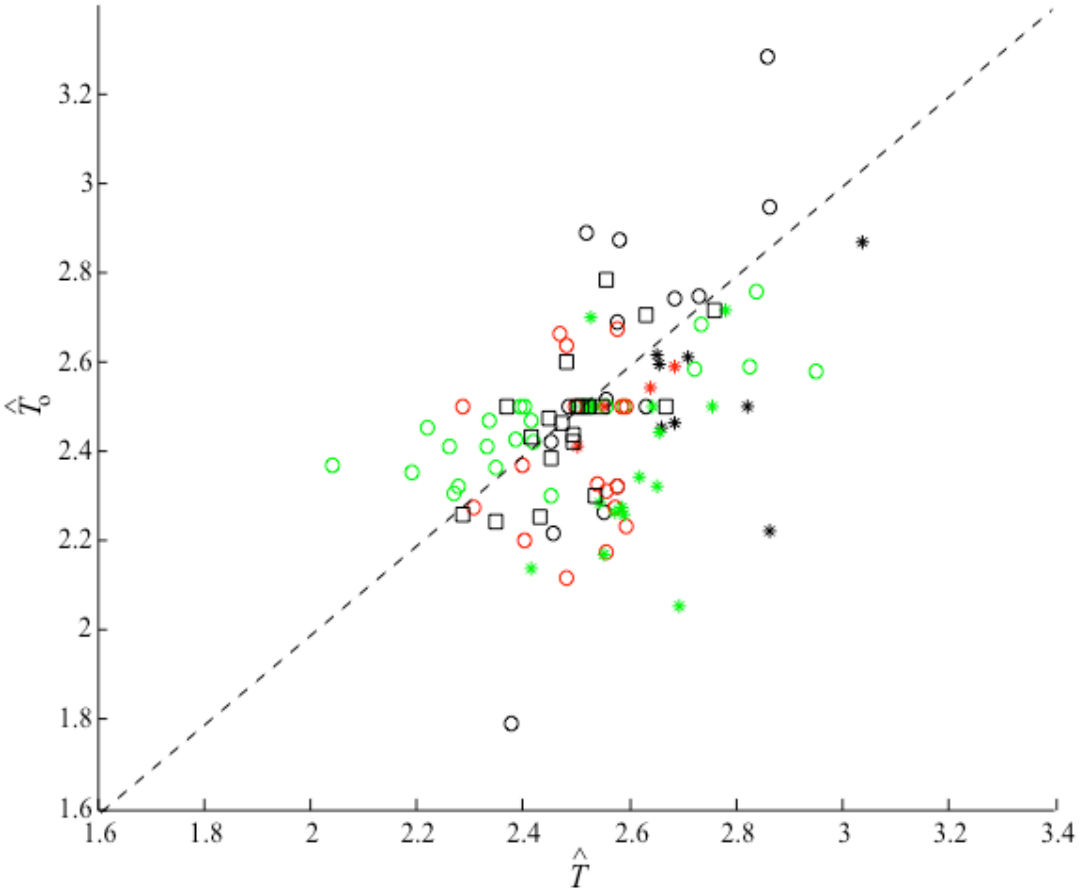

Figure 6.1. $\hat{T}_o$s as a function of $\hat{T}$s. Each symbol represents $\hat{T}_o$s plotted as a function corresponding $\hat{T}$s for a single session. Each color and shape combination refers to a different subject. For the purposes of comparison, we have drawn the line (diagonal dashed line) of slope 1 through the plot of $\hat{T}$. A slope of 1 means that the average change in subjects' target switch point is the same as the average change in the optimal switch point.

The mean switch latencies, which are our estimate of the subject's target switch latency, are generally close to the optimal switch latencies, and, as the optimal switch latency increases, so do the observed mean switch latencies: the slope of the mutual regressions of $\hat{T}$ and T_o calculated for each subject and averaged across them was significantly greater than 0 (mean slope = .72, t(6) = 2.44, p = .05) and not significantly different from 1 (p = .37). In other words, on average subjects' adjustments of their target switch time in response to changes in circumstances approximately matched the changes in the optimal switch point.

The closeness of $\hat{T}$s to $\hat{T}_o$s (temporal-distance between $\hat{T}$ and the corresponding $\hat{T}_o$) can be quantified in two different ways: directional (e.g. +20 ms, -50 ms, etc.) and absolute (e.g. |+20| ms, |-50| ms). When the temporal distances were considered directionally and averaged across sessions and subjects, switch latencies were on average 80 ms later than the optimal switch latency (77 ± s.e. 19 ms). When the temporal distances were considered in terms of their absolute values and averaged in the same way, subjects' switch latencies were on average 160 ms farther away (non-directional) from the optimal switch latency (158 ± s.e. 14 ms). These temporal distances between $\hat{T}$ and T_o correspond to around 3% and 5% of 3 s range of possible values that $\hat{T}$ could take in our experiments. The histograms of both absolute and directional temporal discrepancies are depicted at Figure 6.2. When these distances were calculated for individual subjects and averaged across them, these values were 98 ± s.e. 38 ms for directional and 162 ± s.e. 27 ms for absolute temporal distances.

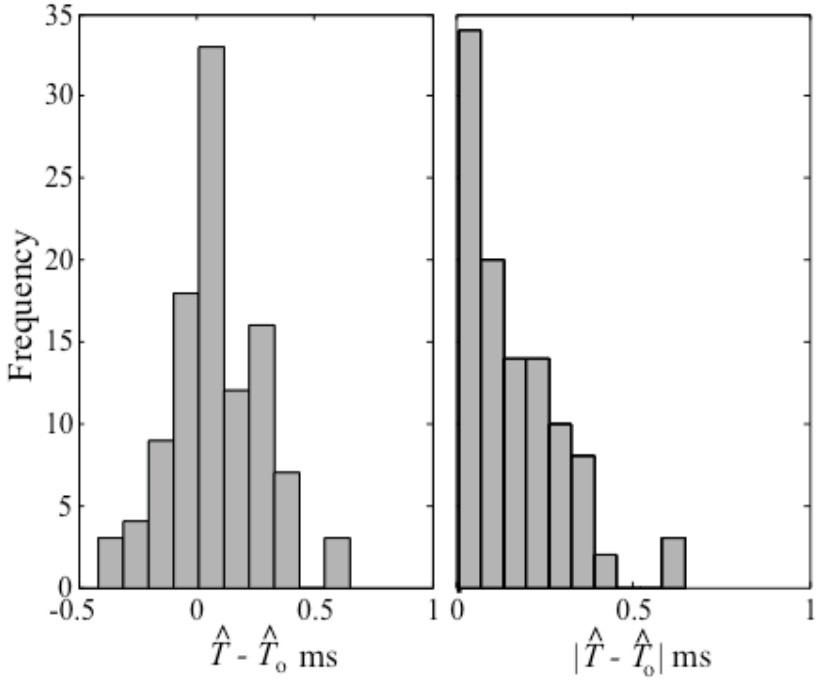

Figure 6.2. Temporal distances from $\hat{T}_o$. Left-panel: Histogram of the signed differences between $\hat{T}$ (subject's estimated target switch latency) and $\hat{T}_o$ (estimate of the optimal switch latency). Right-panel: Histogram of the absolute temporal differences.

Figure 6.3 depicts the CVs plotted as a function of $\hat{T}$s for individual subjects. In the same plots, we also fitted the linear regression lines. In only 2 (Subject 2 & 4) out of 7 subjects did regression analysis reveal significant regression of CVs on $\hat{T}$. The significant slopes were small and in opposite directions, and the cross-subject average slope did not differ significantly from 0. In the light of these analyses and the well-established scalar variability of timed behavior, which predicts constant CVs, we assumed that temporal uncertainty as measured by the Weber fraction did not change as a function task parameters within individual subjects. This motivated us to also use a single estimate of temporal variability in analyzing data gathered from different sessions but with the same subject. This overall temporal uncertainty (hereafter referred to as subject-

specific – as opposed to session specific- temporal uncertainty) was calculated based on the switch latencies collapsed across all sessions of a particular subject for each stimulus type (numerical or graphical). The estimate of subject-specific temporal uncertainty could differ from the session-specific temporal uncertainties. Consequently, for a particular session both $\hat{T}$ and $\hat{T}_o$ could change when we considered subject vs. session-specific temporal uncertainty.

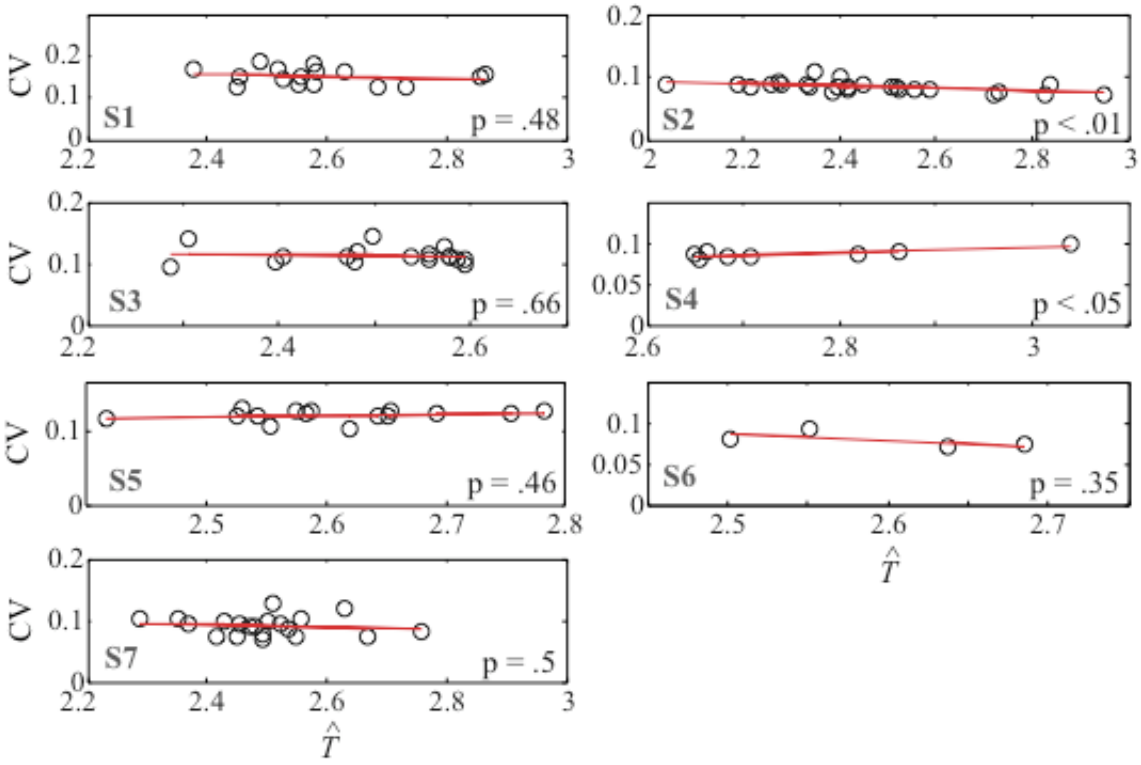

Figure 6.3. Linear regression lines (red lines) fitted on the CVs observed in individual sessions by individual subjects. CVs were plotted as a function of $\hat{T}$s.

In Figure 6.1, we used $\hat{\sigma}$ from each single session (session-specific temporal uncertainty) in estimating $\hat{T}_o$ for that session (i.e., for the conditions prevailing in a given session). The fact that the CV is more or less constant allows a different approach to computing $\hat{T}_o$, in which we used the estimate of the subject-specific CV from all the

sessions within a stimulus type combined. Figure 6.4 is the histogram of the $\hat{T}$-$\hat{T}_o$ differences when $\hat{T}_o$ is estimated from the overall estimate of the subject's CV. Our analyses also revealed that the slope of mutual regressions computed for each subject and averaged across subjects was significantly higher than the slope of 0, $t(6) = 3.12$, $p < .05$. In this particular analyses, we also found that the slopes averaged across subjects was significantly lower than the slope of 1, $t(6) = 3.53$, $p < .05$. (mean slope = .47).

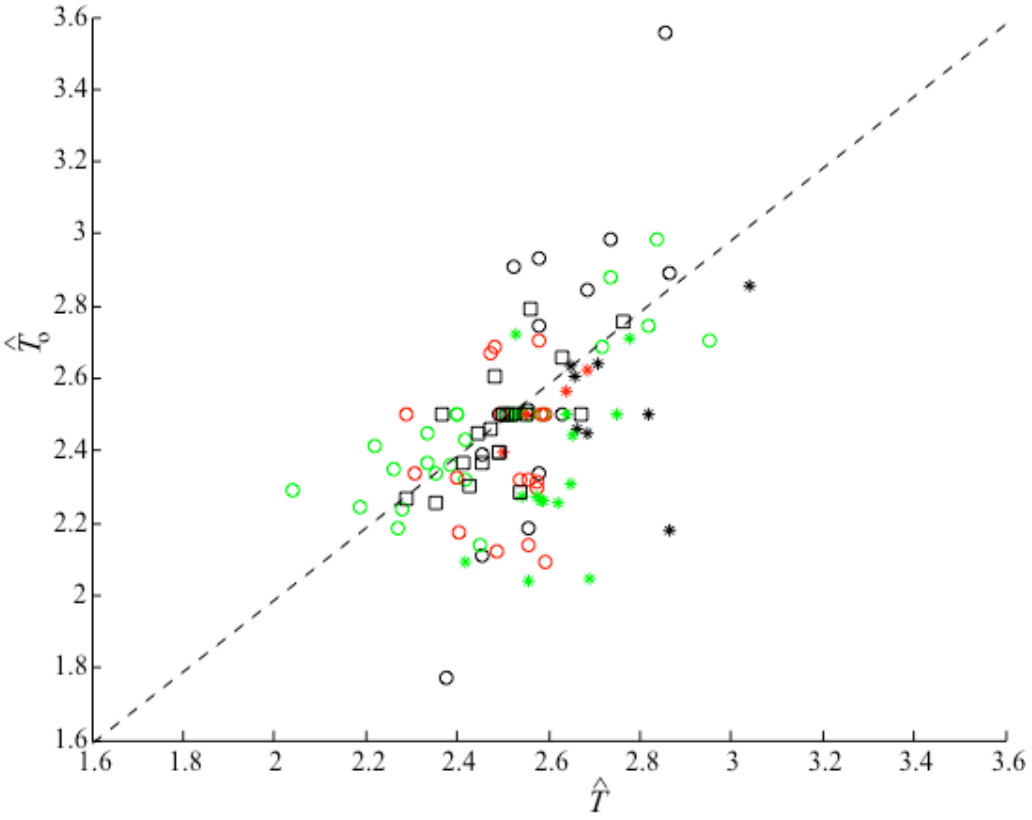

Figure 6.4. Estimated T_o plotted against $\hat{T}$, as in Figure 6.1, but with the estimates of T_o based on the cross-session CV, rather than on the session-specific $\hat{\sigma}$.

The close similarity between the plots in Figures 6.4 and 6.1 implies that it does not much matter whether one uses a cross-session estimate of variability or a session-

specific estimate of a subject's variability in estimating the optimal switch latency for that subject in a given condition. In the case of subject-specific temporal uncertainty, subjects were found to be on average (across sessions) within 80 msec proximity of $\hat{T}_o$s (78 ± s.e. 21 ms as opposed to 77 ± s.e. 19 ms when temporal distance was considered directionally). When temporal distance was considered in terms of its absolute value, subjects were found to be on average (across sessions) within 170 ms proximity of $\hat{T}_o$s (169 ± s.e. 15 ms as opposed to 158 ± s.e. 14 ms). When the temporal distances were calculated for individual subjects and averaged across them, these values were 97 ± s.e. 41 ms for directional and 172 ± s.e. 33 ms for absolute temporal distances. Majority of the results that will be presented in this section will involve measures that were estimated under both session-specific and subject-specific temporal uncertainty.

Although temporal uncertainty did not change across sessions within a subject, it varied considerably between subjects ($F(6,98) = 57.45$, $p < .001$ - see Figure 6.5).

Finally, in Figure 6.6 we present the $\hat{T}$-$\hat{T}_o$ for each session as a function of $\hat{T}_o$. We present these measures both calculated by assuming session (top panel) and subject-specific (bottom panel) temporal variability.

Assuming both session-specific and subject-specific temporal uncertainties, we calculated the regression of absolute temporal distance from the optimal switch latency on $\hat{T}_o$ for each individual subject. In both cases, slopes averaged across subjects were not found to be significantly different from the slope of 0. Thus, absolute temporal distance from $\hat{T}_o$ (optimum switch latency) was constant for a wide range of $\hat{T}_o$ values. A closer look at Figure 6.6 (both top and bottom panel) on the other hand also suggests that as $\hat{T}_o$s

were closer to the referent durations, the temporal distances tended to increase. When we tested quadratic regression of absolute temporal distance from optimal point on $\hat{T}_o$, under both session and subject-specific temporal certainty we found significant regression in 4 out of 7 subjects. Thus, we concluded that for some subjects temporal distance increased as $\hat{T}_o$ was closer to the referent durations. Overall, based on these findings, we concluded that the absolute temporal distance from $\hat{T}_o$ (optimal switch latency) is constant for a large portion of response interval range, but that our subjects, like Trommershäuser et al's subjects (2003b) resist setting a target at or beyond the ostensibly "legitimate" range.

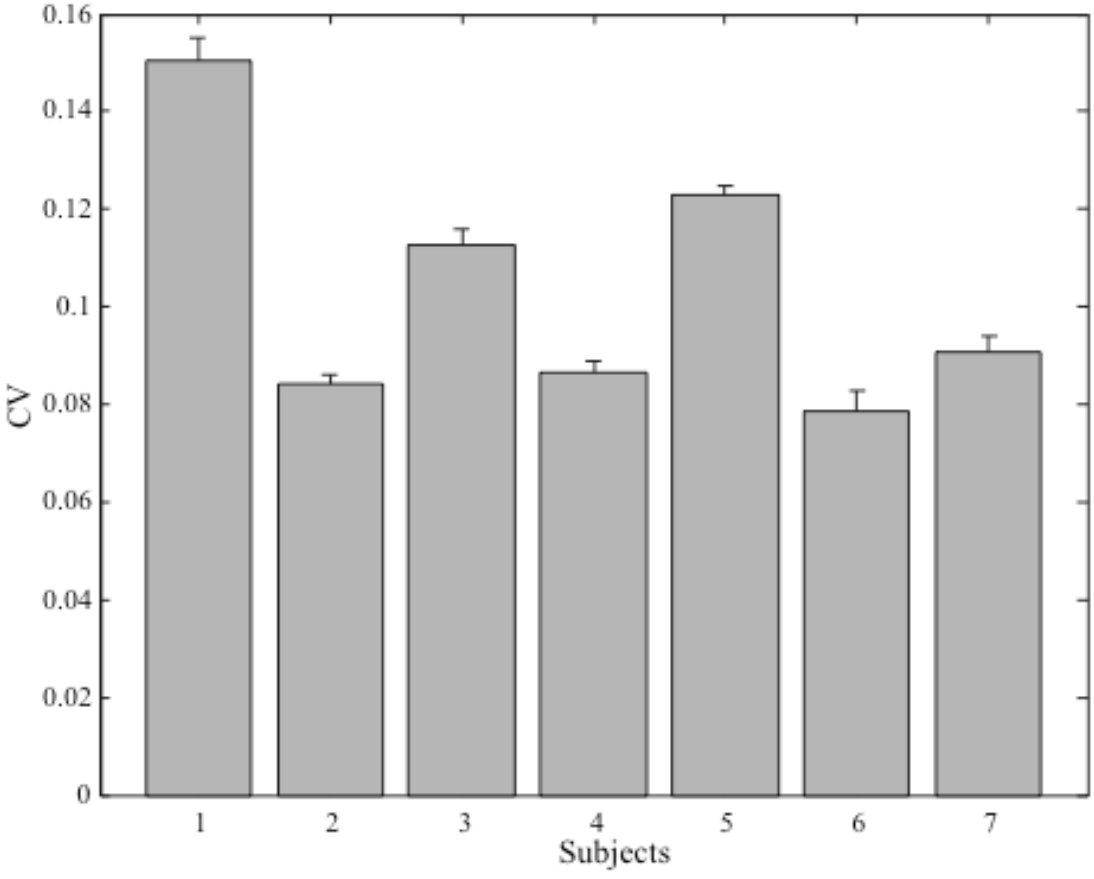

Figure 6.5. CVs presented for individual subjects. Errors bars represent the standard errors of estimations of CVs across all sessions of a given subject.

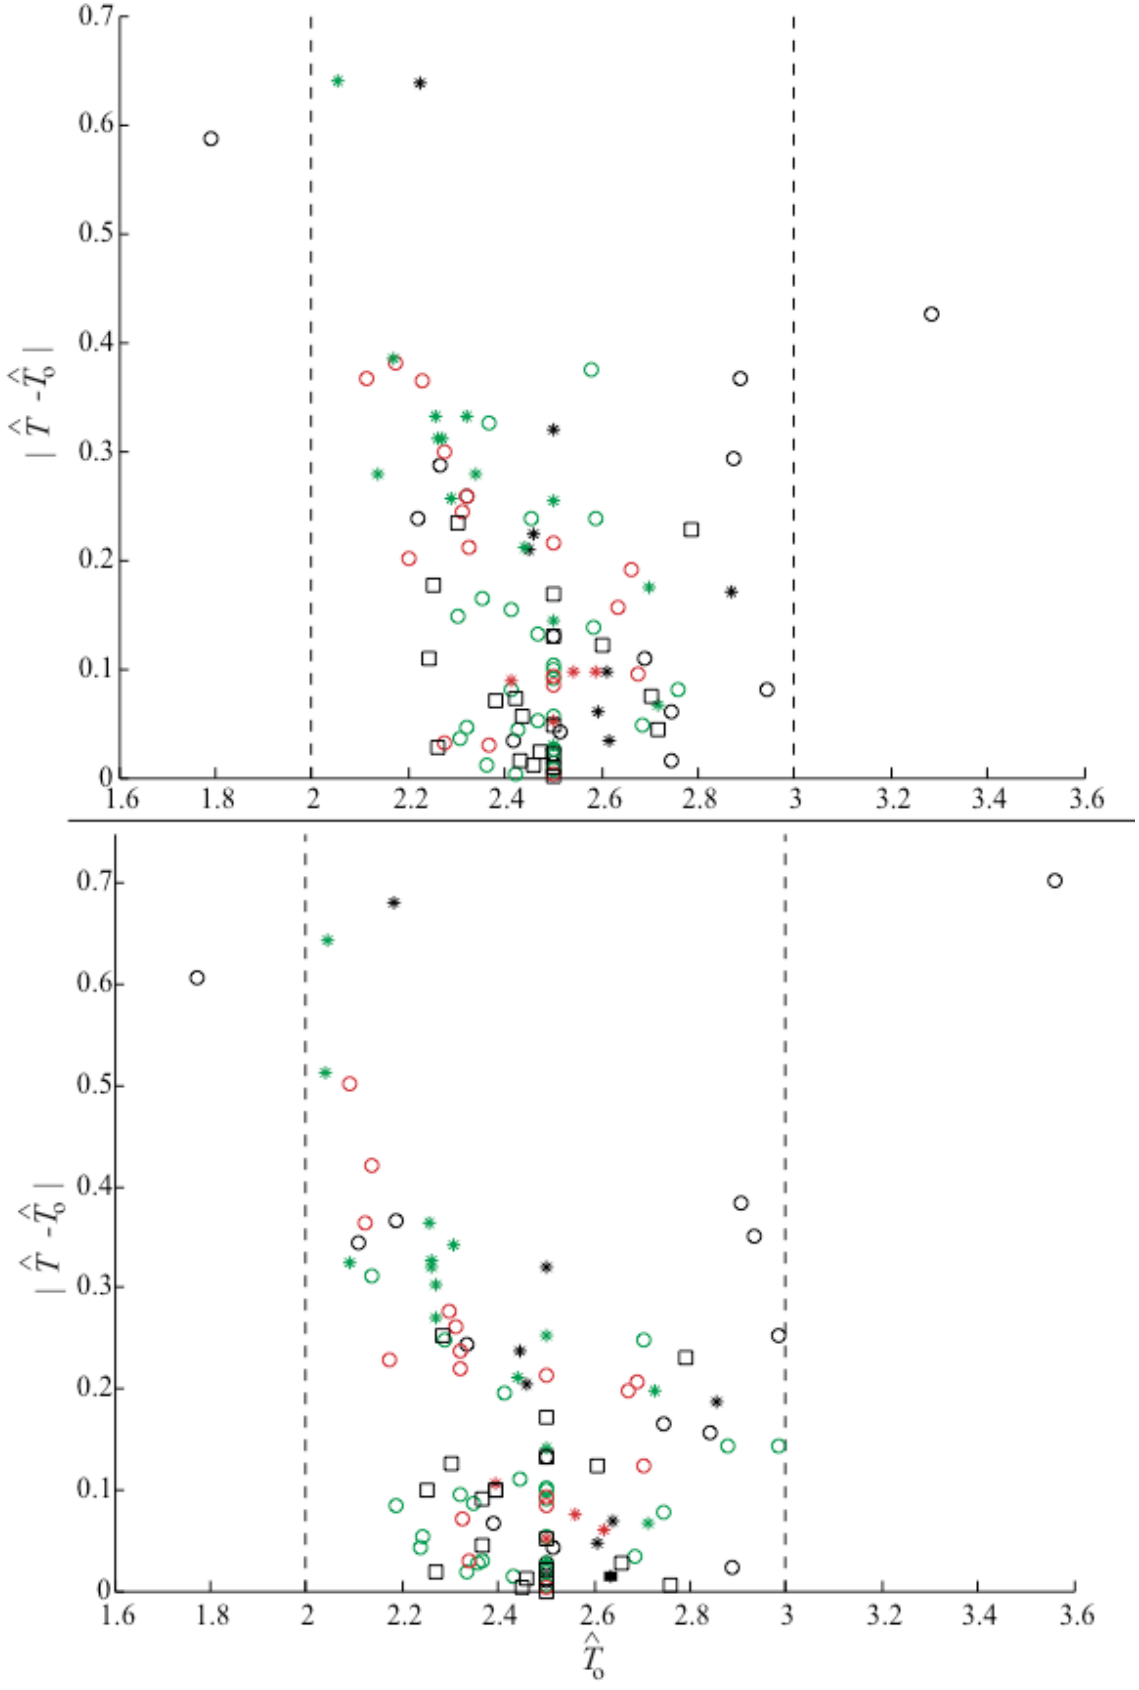
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
$|\hat{T} - \hat{T}_o|$
1.6
1.8
2
2.2
2.4
2.6
2.8
3
3.2
3.4
3.6
$\hat{T}_o$

Figure 6.6. Absolute temporal distances between $\hat{T}$ and $\hat{T}_o$ as a function of $\hat{T}_o$. Measures depicted at the top panel are derived assuming session-specific temporal variability and measures depicted at the bottom panel are derived assuming subject-specific temporal variability. Sessions that were run with different subjects are illustrated by different symbols.

Finally, we have looked at the frequency of the likelihood ratios (see Figure 6.7) of the optimal switch latency given the empirical data. Figure 6.7 suggests bimodality in the frequency of sessions as a function of different likelihood ratios. This can be worded as in a given session, the optimal switch latency was either very likely under the empirical data or it was more than 1000 times less likely. The high frequency of the latter cases is very likely to be caused because of very precise estimates in the model parameters.

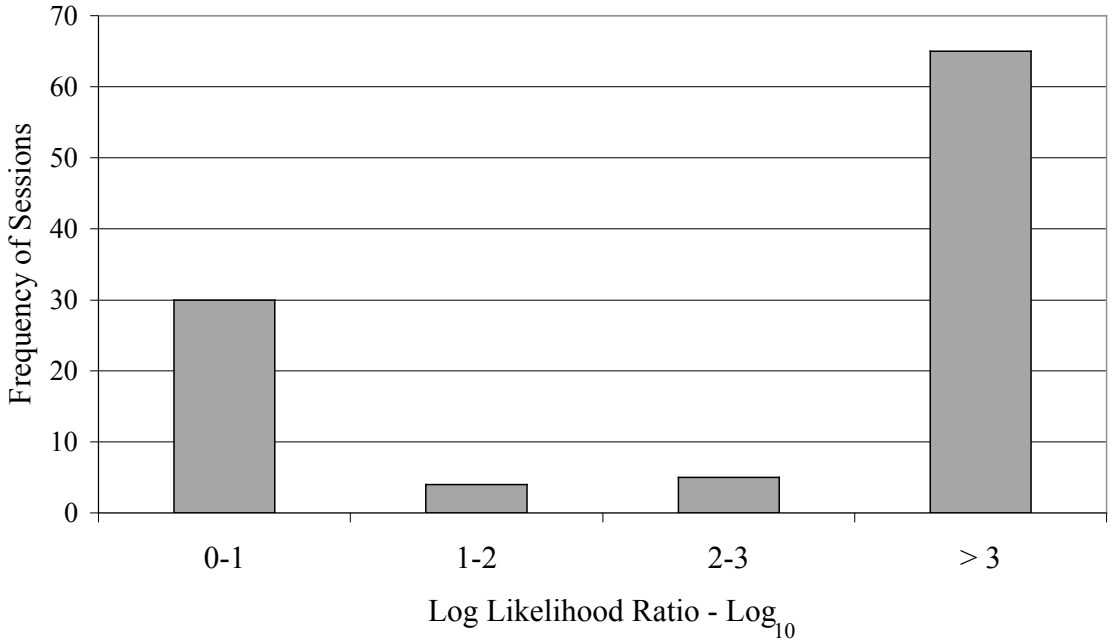

Figure 6.7. Frequency of sessions as a function of likelihood ratios between $\hat{T}_o$ and $\hat{T}$.

Independence of Successive Switch Latencies:

As a check on the independence of successive switch latencies, we conducted autocorrelations on the switch vectors (observed in a given session) at different signed lags. Prior to running the autocorrelations, we detrended the switch vectors. In other words, we removed the possible underlying upward and downward trends in the switch vectors gathered from individual sessions. Autocorrelation was applied on a randomly selected five different subjects' sessions. Figure 6.8 shows the output of our analyses for these 5 subject/sessions. In none of these 5 cases, there was considerable autocorrelation between successive switch latencies and thus one can conclude that the decisions made from trial to trial were not dependent on the decision made prior to them.

Optimality: When and How Optimal?

Up to this point, we have reported the descriptive statistics that revealed general properties of our data set. In this context, we have demonstrated simple correspondences between $\hat{T}$ and $\hat{T}_o$. On the other hand, an important question that we have sought to answer remains to be answered at this point, that is "How good our subjects were in maximizing their gain?" One of the direct ways to approach this question is considering the statistical properties of the empirical switch latencies along with the expected gain functions (given those statistical properties) in the following manner.

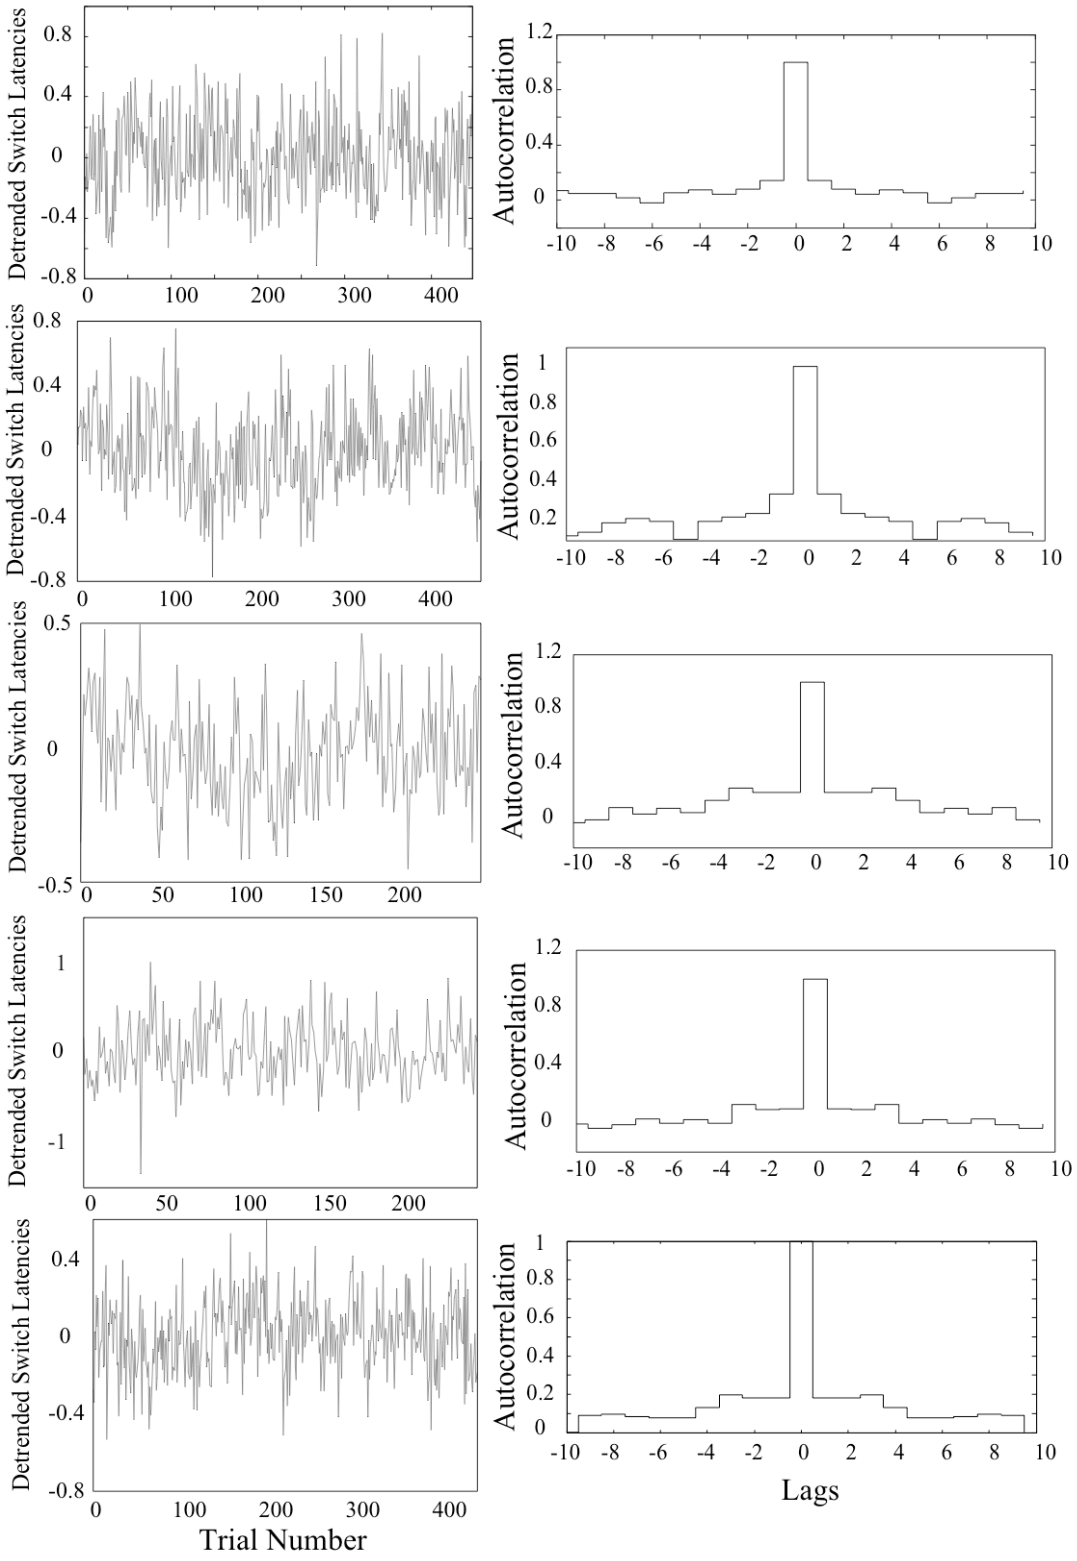

Figure 6.8. Independence of successive switch latencies. Left Panel: Detrended trial-by-trial switch latencies. Right Panel: Output of autocorrelation run on detrended switch vector.

In the parameter plane composed of $\hat{T}$ and $\hat{\sigma}$, we defined an expected gain function given the task parameters. Using optimal decision-making model together with the task parameters (probabilities of short and long trials and the pay-off matrix), we determined expected gains for a finely spaced m by n grid of points in this parameter plane. The expected gain at a point is the average gain that a decision maker with that noise level ($\hat{\sigma}$) in his/her representation of elapsed time and that decision criterion ($\hat{T}$) would obtain. With both the relative likelihood function and the expected gain function mapped onto the common parameter plane, subject's optimality can be measured by the relative likelihood of the closest point on the ridge of the gain function (statistical distance—see Figure 5.3) and by the relative gain of the maximum likelihood estimate of the subject's target (relative gain loss—see Figure 5.1).

The second measure (relative gain loss) indicates the cost that the subject incurs by using a non-optimal target latency. From the subject's perspective, this is the meaningful measure of optimality. A subject whose expected gain is 99.99% of the gain to be expected at the optimal target latency would rationally be indifferent to the possibility that our statistical distance measure would reveal that his target latency was "very significantly" different from the optimal target latency.

Figure 6.9 shows level curves of the likelihood function and of the proportion-of-maximum-gain function for six different sessions with considerably different task parameters. The relative-gain contours are at 99%, 98%, 95%, 90%, and 80% of the MPEG (for a subject with the degree of timing variability given by the $\hat{\sigma}$ coordinate of a point along the contour). The relative-likelihood contours are at the natural log unit levels (e^{-1}, e^{-2}, etc)

Depending on the task parameters, the ridge of the expected gain function exhibits a direction. For instance, if a task parameter involves a relatively high penalty for switching late, the ridge bends to the left as $\hat{\sigma}$ increases, favoring shorter switch latencies, hence less chance of incurring the severe penalty attendant on late switches. Where the concentric relative likelihood contours fall on the expected gain contours depicts our relative certainty about how optimal the subject's target switch latency was.

Expected gain functions depicted in Figure 6.9 (both left and right panels) were normalized by MPEGs of each $\hat{\sigma}$. Consequently, rather than a single optimal point of responding defined for the whole $\hat{T}$ and $\hat{\sigma}$ parameter plane, one gathers an "optimal path" of responding across parameter dimension of σ. This way the optimality of the subjects is evaluated based on $\hat{T}$ only, which as mentioned earlier is the only decision criterion controlled by the subjects.

Cases presented in Figure 6.9 suggest that even though $\hat{T}$ could be unquestionably different from the optimal point of responding, their expected gain was still within the 1% loss from the MPEG. In some other cases, $\hat{T}$ was found to well coincide with $\hat{T}_{\mathrm{o}}$. This pattern characterizes the temporal distance and likelihood ratio quantifications of optimality as more conservative measures compared to the measures of proportion of MPEG. In the following sections, we report the quantification of these two types of comparisons.

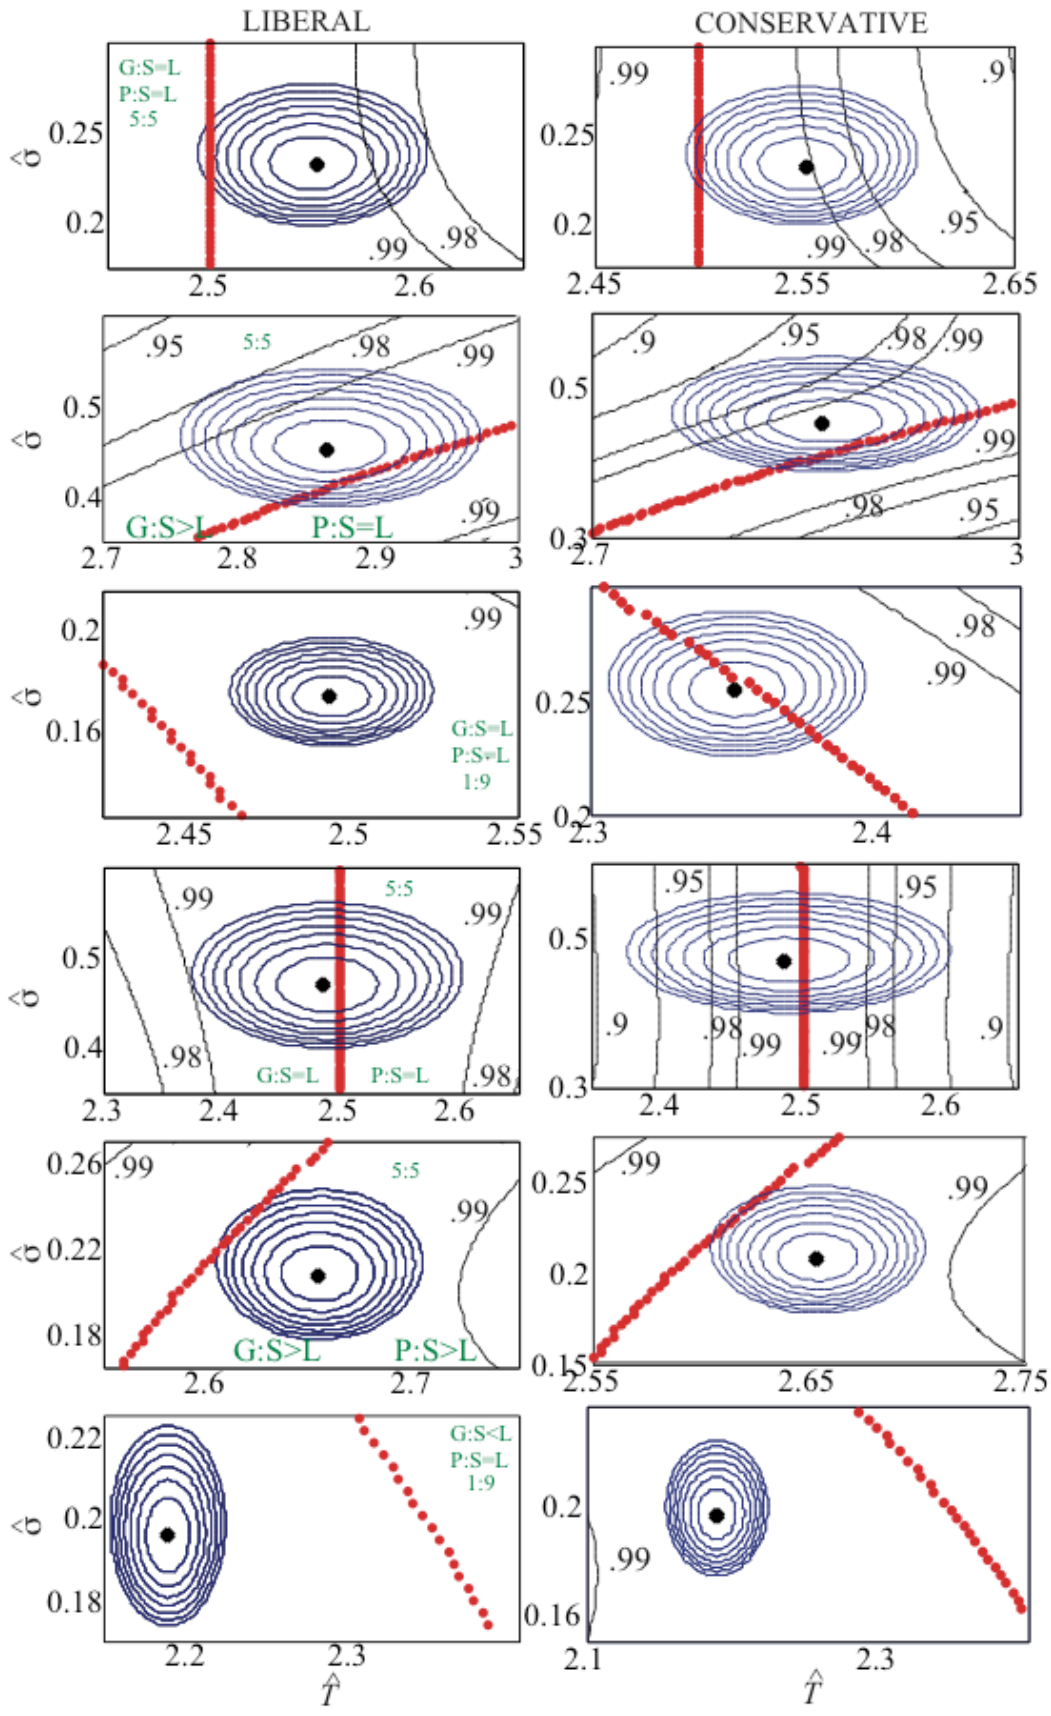
LIBERAL
CONSERVATIVE
G:S=L
P:S=L
5:5
G:S>L
P:S=L
G:S=L
P:S=L
1:9
G:S=L
P:S=L
5:5
G:S>L
P:S>L
5:5
G:S<L
P:S=L
1:9
.99
.98
.95
.9
$\hat{\sigma}$
$\hat{T}$

Figure 6.9.Superposed relative likelihood contours (concentric blue ovals) and relative gain contours (light black curves). The heavy red dotted curve is the ridge of the expected gain function (the locus of target latencies that maximize gain at a given level of variability). Left-panels illustrate the expected gain function defined in terms of simple proportions of MPEG. Right-panels (for the same sessions presented on the left-panel) illustrate the expected gain function defined in terms of the expected gains in between 2 and 3 seconds (see Figure 5.1 for clarification of this distinction).

Proportion of Maximum Possible Expected Gain:

In this section, we report the proportion of maximum possible expected gain - MPEG (expected gain at $\hat{T}_o$) at $\hat{T}$. In calculating these proportions we normalized MPEGs for different levels of $\hat{\sigma}$ (see Figure 6.9). The proportion of MPEG was calculated in two different ways (see Figure 5.1):

1- For each session, expected gain at maximally likely $\hat{T}$ was divided by MPEG given $\hat{\sigma}$ (either individual or session-based). Thus, this proportion took a negative value (e.g. -5 / 10 = -0.5) when the expected gain at $\hat{\sigma}$ was negative and MPEG was positive. In most conditions, the output of this calculation constituted a liberal model for optimality (see Figure 5.1, bottom panel).

2- For each session, the proportion of MPEG was computed as a function of the range of gains one could gather within the 2 and 3 s, which constituted the short and long reference durations of the switch task. In other words, the proportions were computed in terms of the minimum and maximum expected gains one could earn if s/he would respond between short and long referents only (see Figure 5.1, bottom panel).

In most cases, the output of this calculation constituted a more conservative model of optimality (compare the left and right panels of Figure 6.9). In this section, we report proportions of MPEG that were computed in both ways. Further, for each type of calculation, we report the proportion of MPEG for both individual and session-specific temporal uncertainty. Figure 6.10 depicts the histogram of conservative and liberal mean proportion of MPEG assuming both session and subject-specific temporal uncertainty.

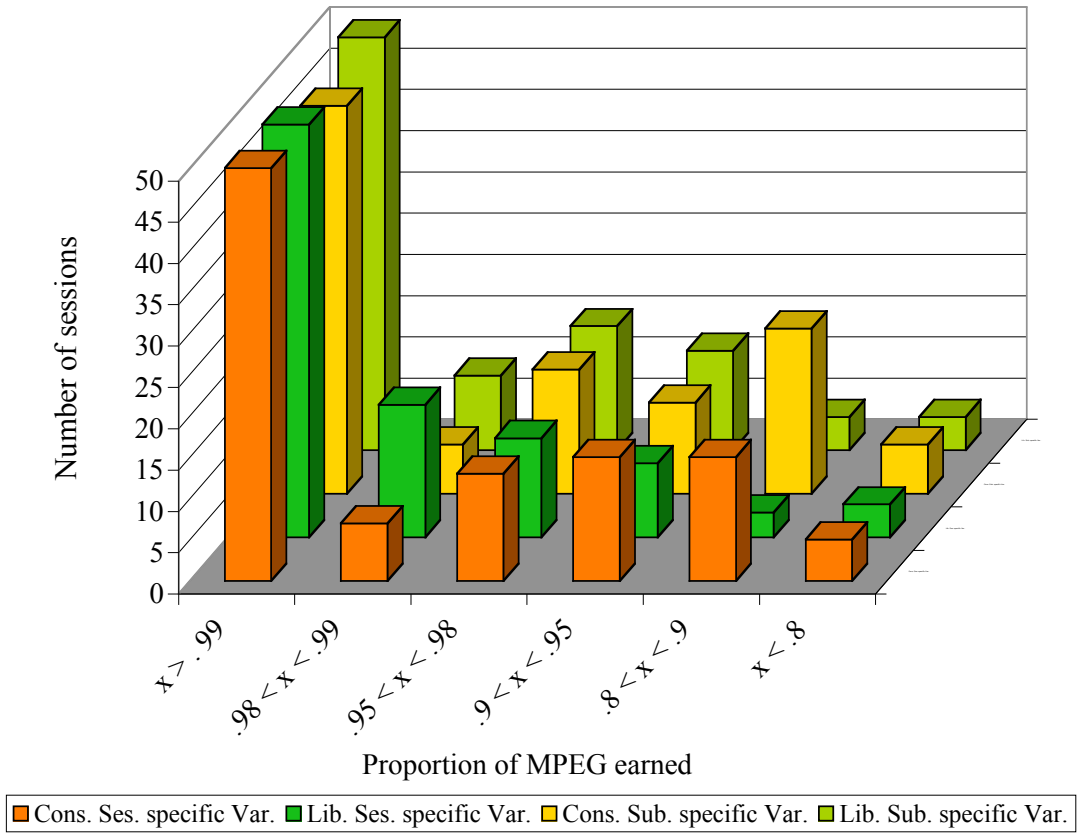

Figure 6.10. Frequency of sessions under which different proportions of MPEG (e.g. x > .99) were observed. Different colors stand for different ways of calculating the proportions of MPEG. Orange and yellow bars depict the conservative estimates of proportion of MPEG assuming session-specific and subject-specific temporal uncertainty, respectively. Green and lime bars depict the liberal estimates of proportion of MPEG assuming session-based and subject-based temporal uncertainty.

Figure 6.10 suggests that different ways of calculating the proportions of MPEG did not change these scores much. We tested if there was a difference in the overall proportions of MPEGs when they were calculated in four different ways: 1- *Conservative with session-specific uncertainty*, 2- *Conservative with subject-specific uncertainty*, 3- *Liberal with session-specific uncertainty*, and 4- *Liberal with subject-specific uncertainty*. One-way ANOVA did not reveal significant differences across proportions of MPEG calculated in these four ways (see Figure 6.10). These proportions when calculated for each individual subject and averaged across them were .95 ± .017 (session-specific uncertainty & conservative), .96 ± .015 (subject-specific uncertainty & conservative), .94 ± .019 (session-specific uncertainty & liberal), and .96 ± .018 (subject-specific uncertainty & liberal). In making these calculations, we excluded those sessions in which the MPEG across the whole parameter plane was 0 from this analysis, since there is no straightforward way of calculating "liberal proportions" in those cases. For these conditions, we have calculated the proportion of MPEG for individual subjects assuming both session-specific and subject-specific temporal uncertainty. When they were averaged across all subjects, these proportions were .86 ± .09 and .84 ± .09 assuming session and subject-specific uncertainty, respectively. In other words, when a subject faces only losses and can therefore only hope to minimize them, subjects do worse than when they can actually gain. Subjects found the conditions where they could only lose frustrating and unpleasant.

Improvement in the course of sessions:

We checked if subjects got closer to the optimal switch latency over the course of session. In order to answer this question, we evaluated the F-ratio of variances of $|\hat{T} - T_o|$

during the first and last quarters and deciles. Figure 6.11 depicts the frequency of different F-ratios presented on a logarithmic scale for both first and last quarters (top panel) and deciles (bottom panel). As this figure suggests in a large proportion of the sessions, F-ratio was around 1 suggesting no change in the absolute temporal distance to the optimal switch latency across.

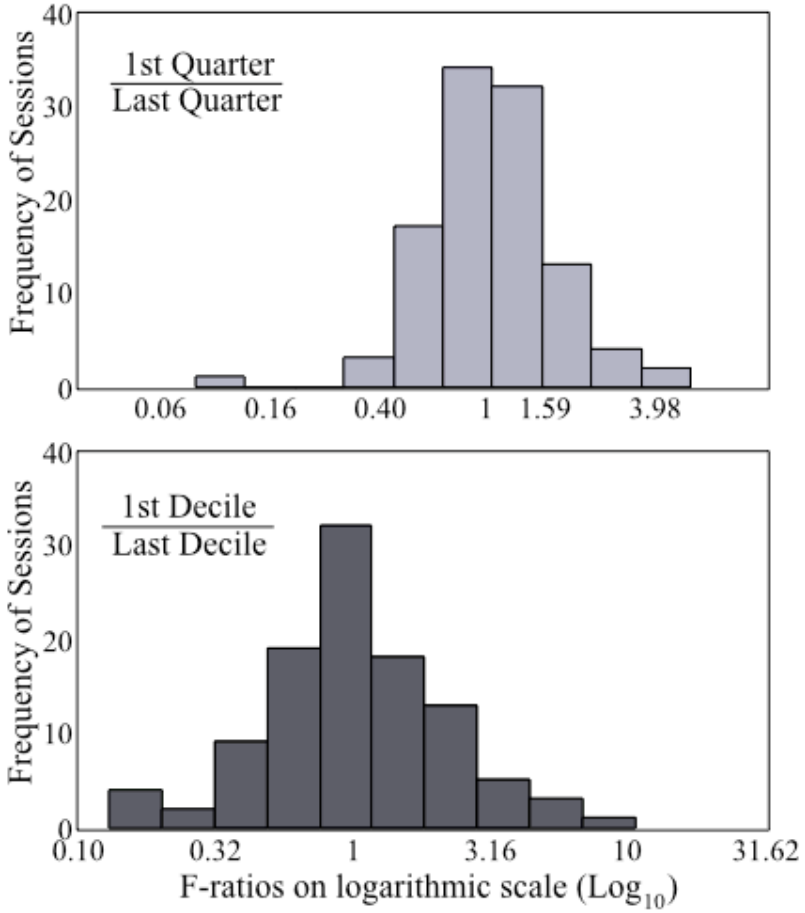

Figure 6.11. Frequency of sessions in which different F-ratios were observed. In the top panel we present the data gathered from the comparison of first and last quarters and in the bottom panel we present the data gathered from the comparison of first and last deciles.

Figure 6.12 depicts the proportion of sessions in which there was a significant difference in $|\hat{T} - T_o|$ across first and last portions of the sessions. It suggests in only a small proportion of the sessions, we observed a significant change in the absolute

distance to the optimal switch latency. When there was a change however, this was mostly in the direction of improvement rather than disruption.

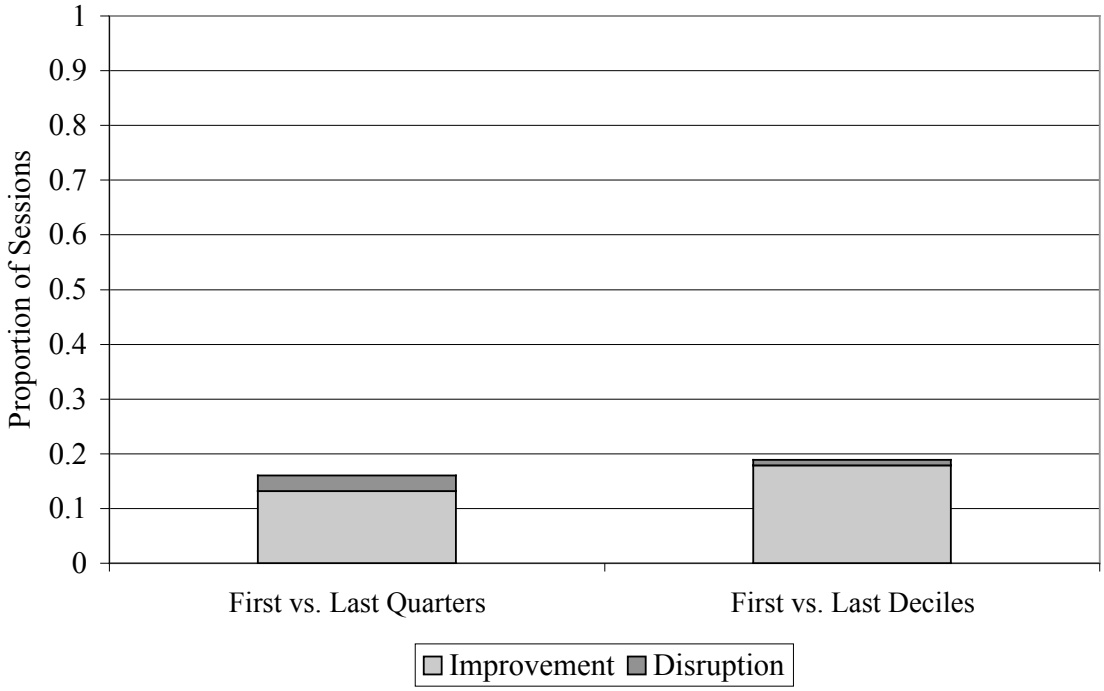

Figure 6.12. The proportion of sessions in which there was a significant change in absolute temporal distance to the optimal switch latency. Light gray bars represent the proportion of those sessions in which there was improvement, that is subjects got closer to the optimal switch latency. Dark gray bars represent the proportion of those sessions in which there was disruption, that is subjects got farther away from the optimal switch latency.

Further, we have conducted a paired sample t-test across the mean absolute temporal distance in the initial and latter portion of the sessions both in terms of quarters and deciles. We did not find any close to significant differences in the mean temporal distances between first and last quarters and deciles of sessions. Based on these analyses we concluded that there was no general improvement or disruption throughout the session.

Numerical vs. Graphical Presentation of Reward and Penalty Magnitudes:

Some of the task parameters have been tested both with numerical and graphical presentations of rewards and penalties within the same subject. The order of these numerical and graphical presentations were counter-balanced across subjects. We investigated, everything else kept constant, if numerical vs. graphical presentations of reward and penalty magnitudes resulted in differences in proportion of MPEG (measured in four different ways), absolute temporal distance from $\hat{T}_o$ (for session-specific and subject-specific temporal uncertainty), and temporal uncertainty. We compared these measures using paired-sample t-test. Our analyses did not reveal any significant effect of type of reward/penalty presentation on any of these measures. CV's tended to be higher for graphical sessions (assuming session-specific temporal variability) compared numerical sessions, however, this difference did not reach statistical significance, $t(35) = -1.883$, $p = .07$. The proportions of MPEGs were higher for numerical sessions in all four different calculations of these proportions, however none of these differences reached statistical significance. Finally, as can be deducted from these findings, the absolute temporal distance tended to be larger for graphical sessions compared to numerical ones.

Effect of session order:

Finally, we compared if there was any effect of order of session presentation using repeated measures t-test. For these analyses, we used only those task conditions that were tested twice within the same subject, in order to control for all other aspects. Our analyses did not reveal any effect of session order on any of the measures listed in the previous section. This shows that as there was no general improvement within a session, there was also no general improvement across sessions.

SECTION 7

RESULTS OF MICE EXPERIMENTS

General Results:

In all the analyses conducted with mouse data, we only used the last 10 sessions of each phase in order to ensure that analyses involved the steady state performance data. In order to determine the stability of CV, for each individual mouse we conducted the regression of CV on $\hat{T}$ (see Figure 7.1). The slopes of regressions averaged across subjects was found to be significantly lower than the slope of 0, $t(11) = -4.22$, $p < .01$ (mean slope = –0.12). We further conducted quadratic regression of each subject's CV on $\hat{T}$ during Phase 1, in which involved equal probabilities and was run with all subjects. Our analyses revealed a significant quadratic regression of CV on $\hat{T}$ ($p < .001$), which reflects inter-subject differences in terms of temporal variability. Based on this finding, we decided to use phase-specific temporal uncertainty rather than subject-specific temporal uncertainty in all of the analyses reported in this section.

As in the analyses of data gathered from human subjects, in analyzing mice data for each phase, we estimated the mean (target) switch latency ($\hat{T}$) and the optimal switch latency ($\hat{T}_o$), given the estimate of $\hat{\sigma}$ for that phase. To allow visual comparison of the $\hat{T}$s and $\hat{T}_o$s, Figure 7.2 plots both.

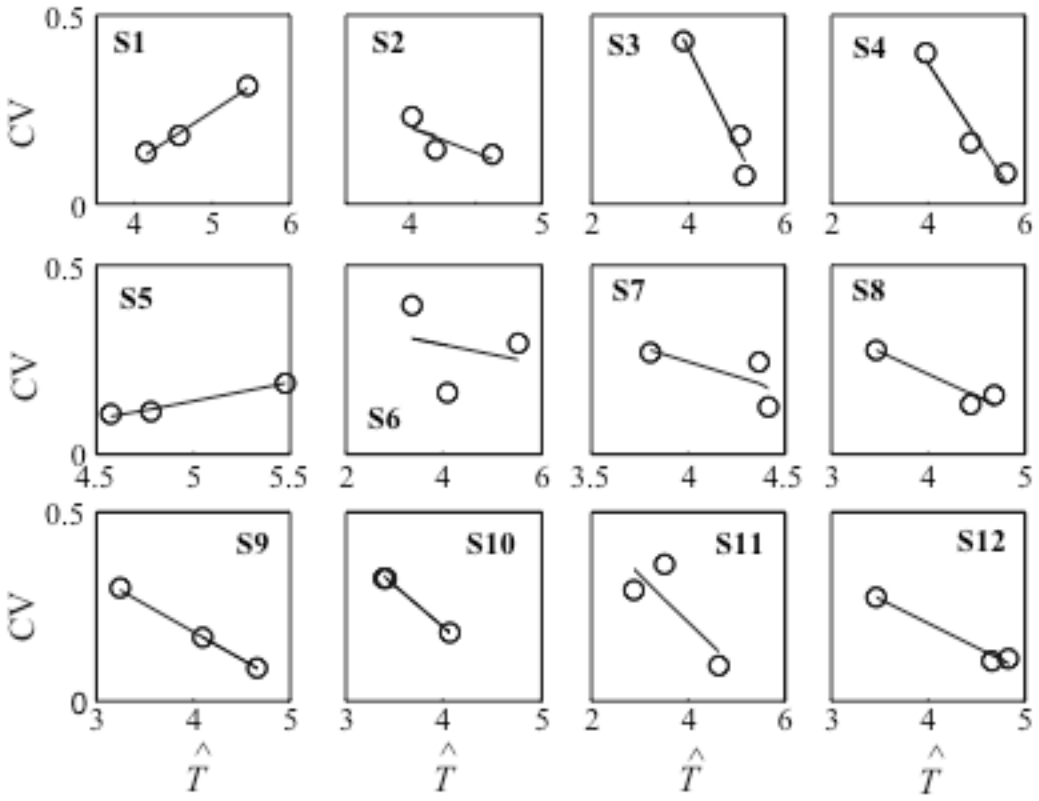

Figure 7.1. Linear regression lines fitted on the CVs observed in individual phases by individual subjects. CVs were plotted as a function of ranked $\hat{T}$s.

Figure 7.2 suggested that, $\hat{T}$s are in general close to $\hat{T}_o$s. It is also clear that as $\hat{T}_o$ (estimated optimal switch latency) increases so does $\hat{T}$, the estimate of the subject's target latency. We have conducted mutual regression of $\hat{T}$ and $\hat{T}_o$ for each subject and compared the average slope to 0 and 1. We found that the average slope was significantly higher than the slope of 0, $t(11) = 6.15$, $p < .0001$ and not significantly different from the slope of 1 ($p = .18$). This shows that as human subjects also mice were sensitive to the probabilities that determine $\hat{T}_o$ and adjusted their target latency ($\hat{T}$) appropriately.

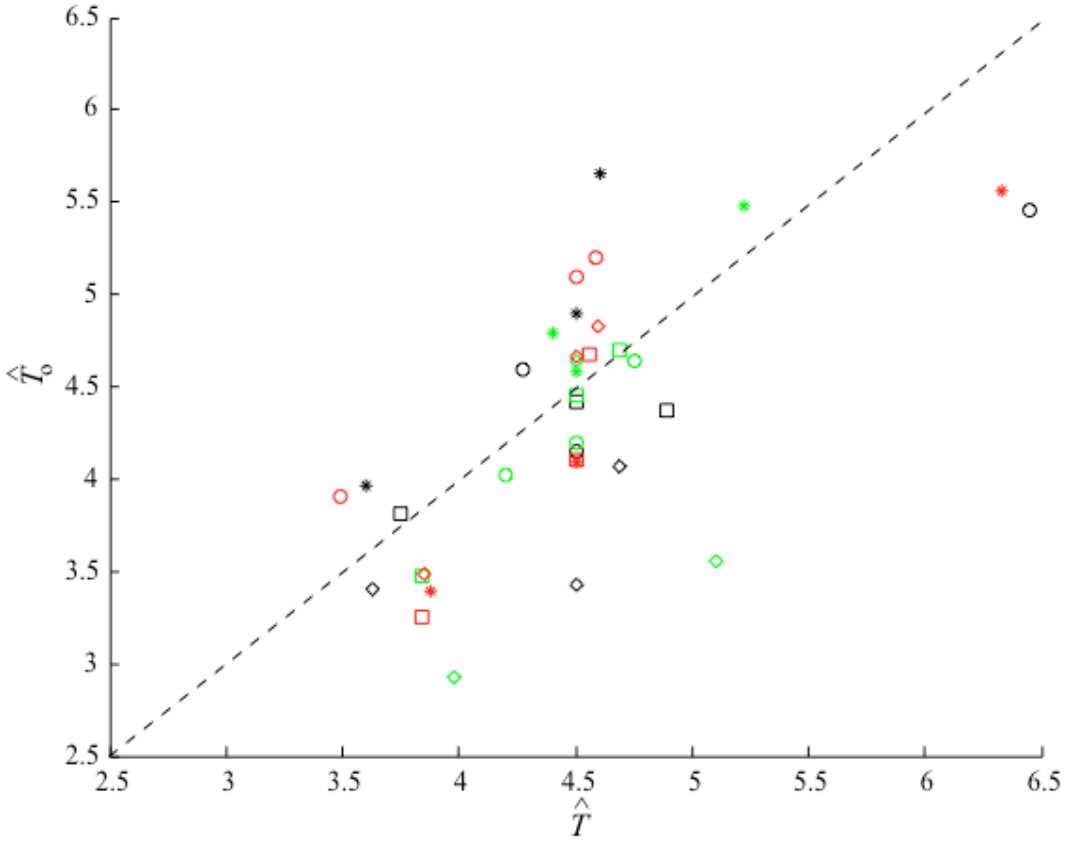

Figure 7.2. T_o as a function of $\hat{T}$ for mice subjects. Each symbol represents T_o for a given phase and plotted as a function corresponding $\hat{T}$ in that particular phase. Each color/shape corresponds to a different subject. For the purposes of comparison, we have drawn the line (diagonal dashed line) of slope 1 through the plot of $\hat{T}$.

We further confirmed the similar effect of the probabilistic conditions on $\hat{T}_o$s and $\hat{T}$s by running a two-way ANOVA with phase and the T_o vs. $\hat{T}$as factors. We conducted two-way ANOVA separately for different groups of mice that experienced probability conditions that differed across groups. In both groups, two-way ANOVA revealed a significant effect of probability condition, $F(2,20) = 39.37$, $p < .001$ and $F(2,20) = 24.03$, $p < .001$ for both groups, respectively. In one of the groups we found a significant difference between $\hat{T}$ and $\hat{T}_o$, $F(1,10) = 7.88$, $p < .05$ while this difference did not

approach significance in the second group. We did not find any significant interaction in either group.

Figure 7.3 depicts the relation between the probability conditions and $\hat{T}$ and $\hat{T}_o$s. In line with the statistical outputs, this plot suggest the effect of probability conditions on $\hat{T}$ and $\hat{T}_o$ and the strong relation between these two measures. Figure 7.3 is derived from data gathered from two different experimental groups, which received different asymmetrical probability conditions.

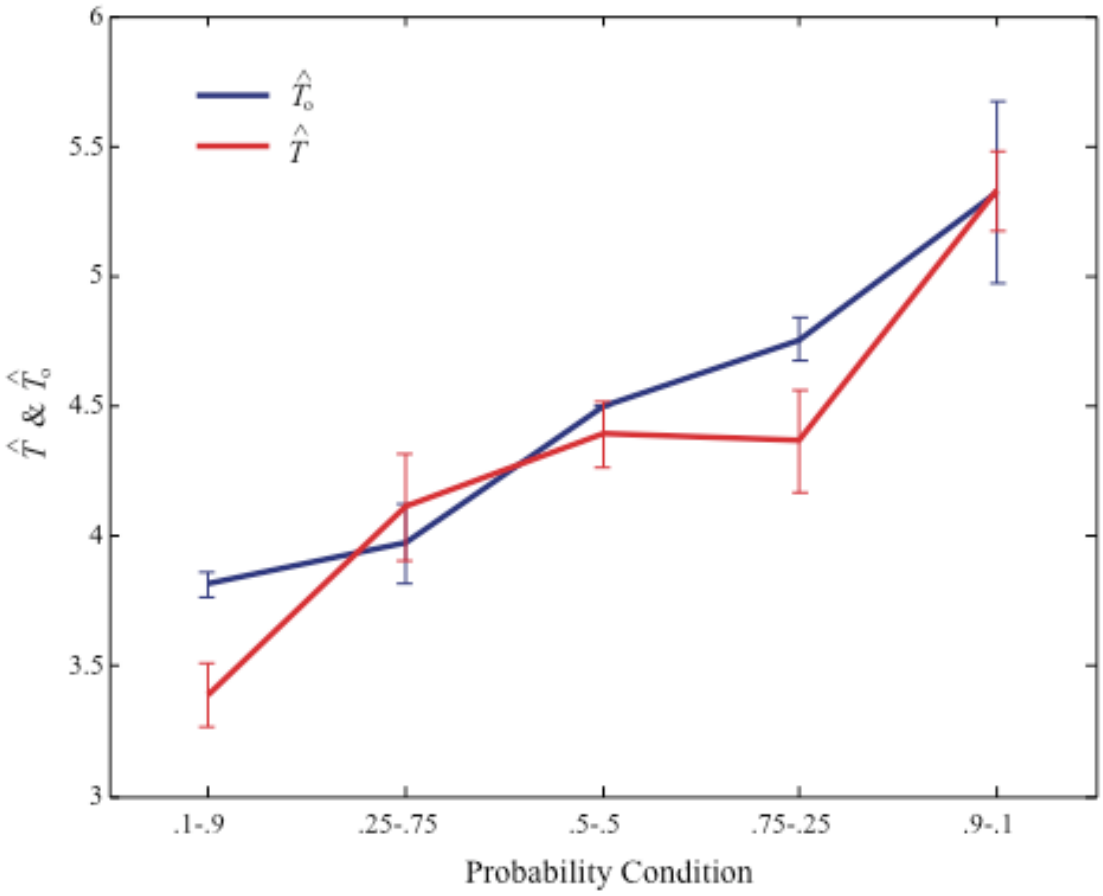

Figure 7.3. Average $\hat{T}$s to $\hat{T}_o$s as a function of different probabilistic conditions. Error bars depicts the standard errors.

We quantified the closeness of $\hat{T}$s to $\hat{T}_o$s with directional and absolute temporal distances. When the temporal distances were considered directionally, subjects' switch latencies were on average 150 ms earlier than the optimal switch latency (147 ± s.e. 91 ms). When the temporal distances were considered in terms of their absolute values, subjects' switch latencies were on average 440 ms farther away (non-directional) from the optimal switch latency (436 ± s.e. 60 ms). These temporal distances between $\hat{T}$ and $\hat{T}_o$ correspond to around 2.5% and 7% of 6 s range of possible values that $\hat{T}$ could take in our experiments. The histograms of both absolute and directional temporal discrepancies from the optimal switch latency ($\hat{T}_o$) are depicted at Figure 7.4.

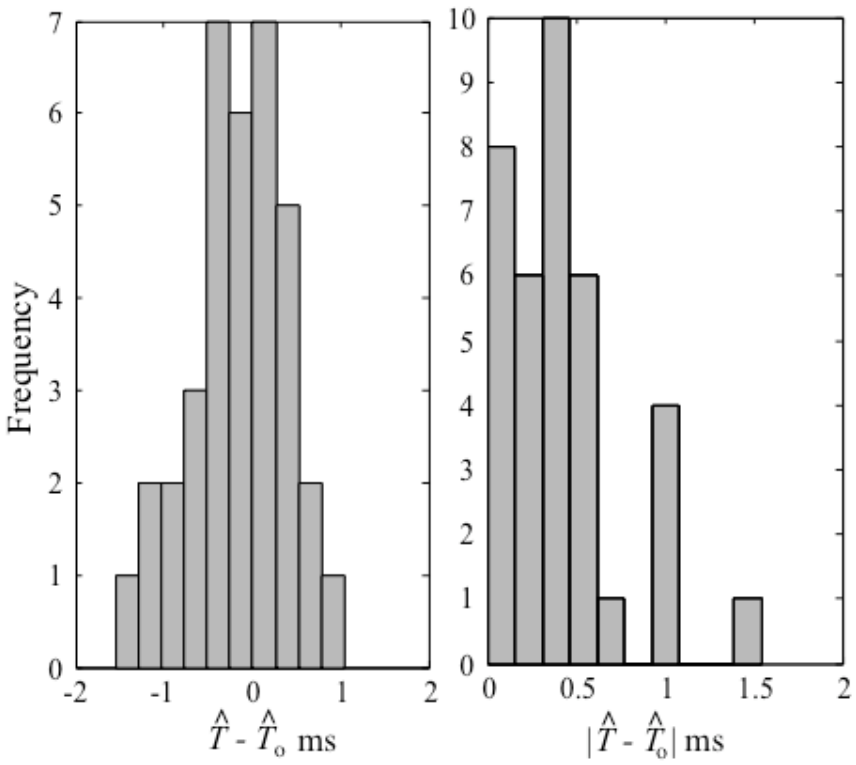

Figure 7.4. Histogram of directional (left panel) and absolute (right panel) temporal distances from $\hat{T}_o$.

Unlike in the human case, the between-subject differences in temporal uncertainty

(CV) were not found to be significant. Finally, we have looked at the absolute temporal distance from the $\hat{T}_o$ (optimal switch latency) as a function of $\hat{T}_o$. Figure 7.5 depicts the absolute temporal distances from the optimal switch latency for individual subjects and phases. We have run regressions for each individual subject and found that the mean slope (across different subjects) did not differ significantly the slope of 0. Confirming this finding, Figure 7.5 suggests that the absolute temporal distance to the optimal switch latency was quite constant across a wide range of $\hat{T}_o$ values. On the other hand, when the optimal switch latency approached to the referent durations or fell out of the referent duration range, the absolute temporal distances to $\hat{T}_o$ tended to increase. On the other hand, quadratic regression revealed a significant relation only in 1 out of 12 subjects.

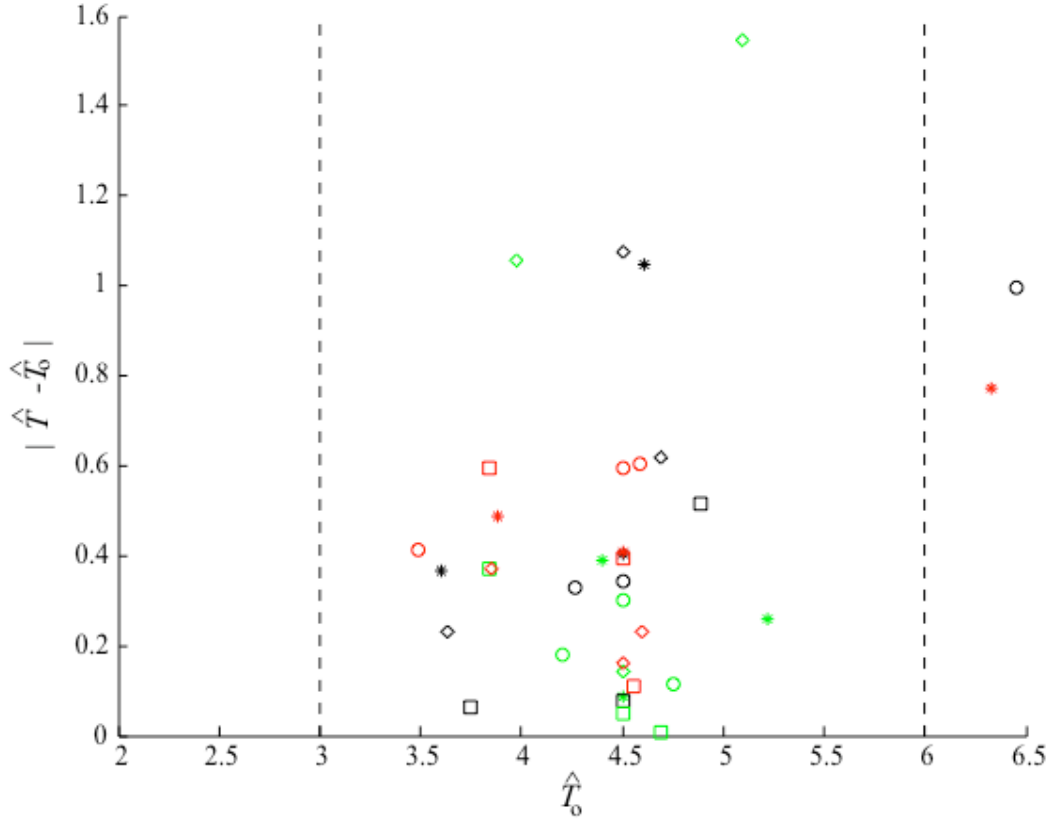

Figure 7.5. Absolute temporal distances from $\hat{T}_o$ as a function of $\hat{T}_o$ for mice

subjects. Each color/shape correspond to a different subject. For each subject, we had three values each of which was derived from three phases they were run in.

Finally, we have looked at the histogram of the statistical distances between the estimated optimal switch latency and the estimated target latency for a given subject in a given phase (see Figure 7.6). As in the human case, Figure 7.6 suggests bimodality in the frequency of sessions as a function of different likelihood ratios. This can be worded as in a given phase, the optimal switch latency was either very likely under the empirical data or it was more than 1000 times less likely. The high frequency of the latter cases is very likely to be caused because of very precise estimates in the model parameters.

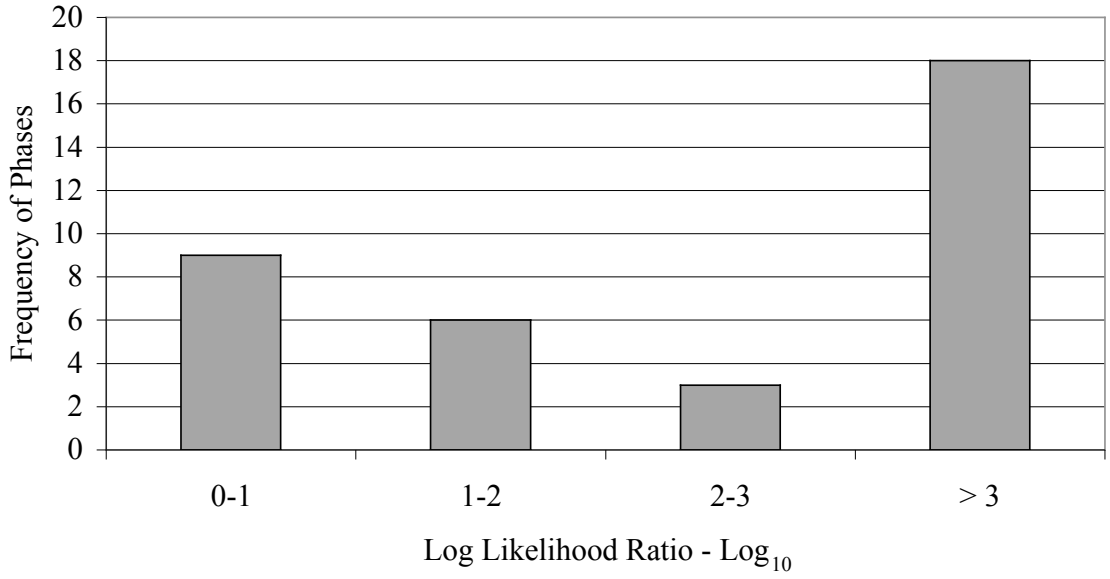

Figure 7.6. Frequency of phases as a function of likelihood ratios for mice subjects.

Independence of Successive Switch Latencies:

As a check on the independence of successive switch latencies, we conducted autocorrelations on the switch vectors (observed in a given phase) at different signed

lags. Prior to running the autocorrelations, we detrended the switch vectors. Autocorrelation was applied on a randomly selected five different subjects' phases. Figure 7.7 shows the output of our analyses of 5 subject/phases. In none of these 5 cases, there was considerable autocorrelation between successive switch latencies and thus one can conclude that the decisions made from trial to trial were not dependent on the decision made prior to them.

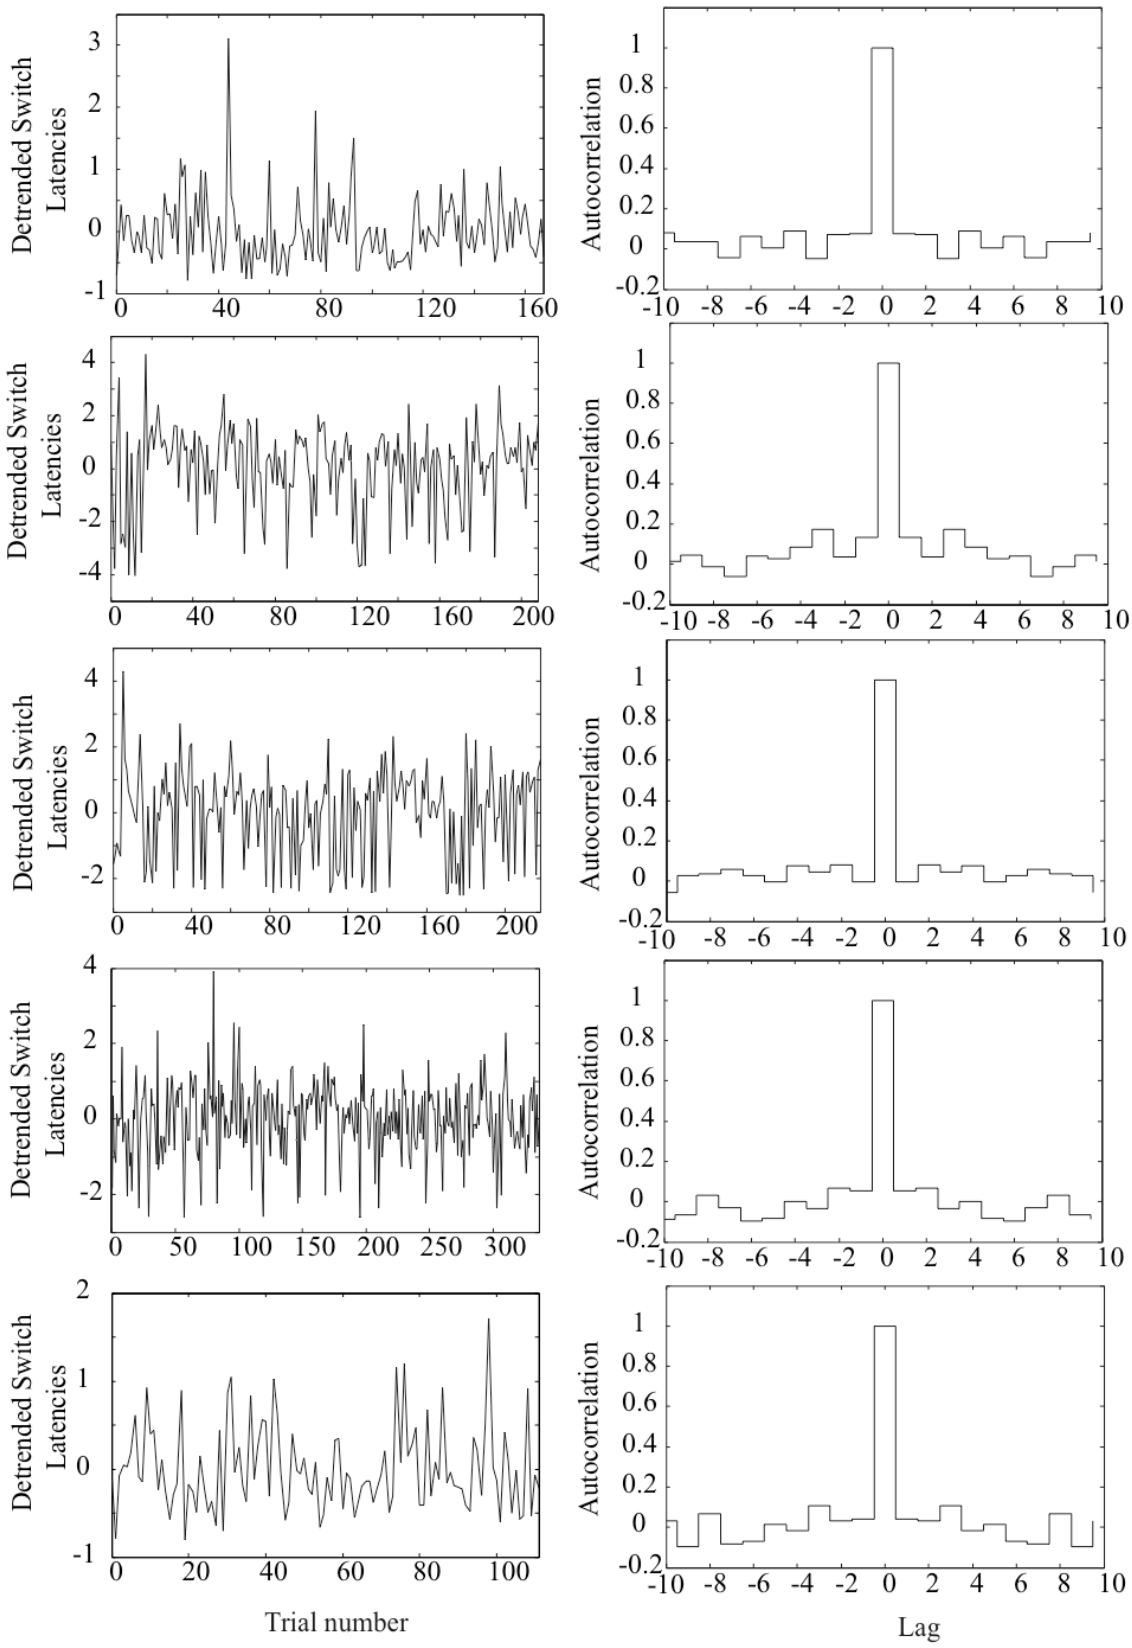
Detrended Switch Latencies
Autocorrelation
Trial number
Lag

Figure 7.7. Independence of successive switch latencies for mice subjects. Top Panel: Detrended trial-by-trial switch latencies. Bottom Panel: Output of autocorrelation run on detrended switch vector.

Proportion of Maximum Expected Gain:

Figure 7.8 depicts the expected gain functions and where the likely models of subject's data falls in this expected gain function. These figures can be read in the same way as those presented in human experiment results. In 7.8 we depict 4 individual subjects' data across all probabilistic conditions under which they were run.

Figure 7.8 shows that $\hat{T}$ changed considerably as a function of probabilistic conditions (across phases). Figure 7.8 further shows that mouse subjects' $\hat{T}$ were often very close to $\hat{T}_o$ across different phases. Moreover, in those cases where the subject diverged considerably from $\hat{T}_o$ (that is, where the statistical distance from optimality was large), there was not a large loss in the expected gain relative to what a mouse with that variability could in principle obtain (that is, the relative loss of gain was small). For example, in 3rd panel (.1:.9 condition) of Mouse 8, the relative gain ridge does not intersect the outermost likelihood contour, meaning that the statistical distance from optimality is large. However, the expected gain at the maximum likelihood point (the filled circle) is greater than 99% of the gain that a mouse with that variability could attain.

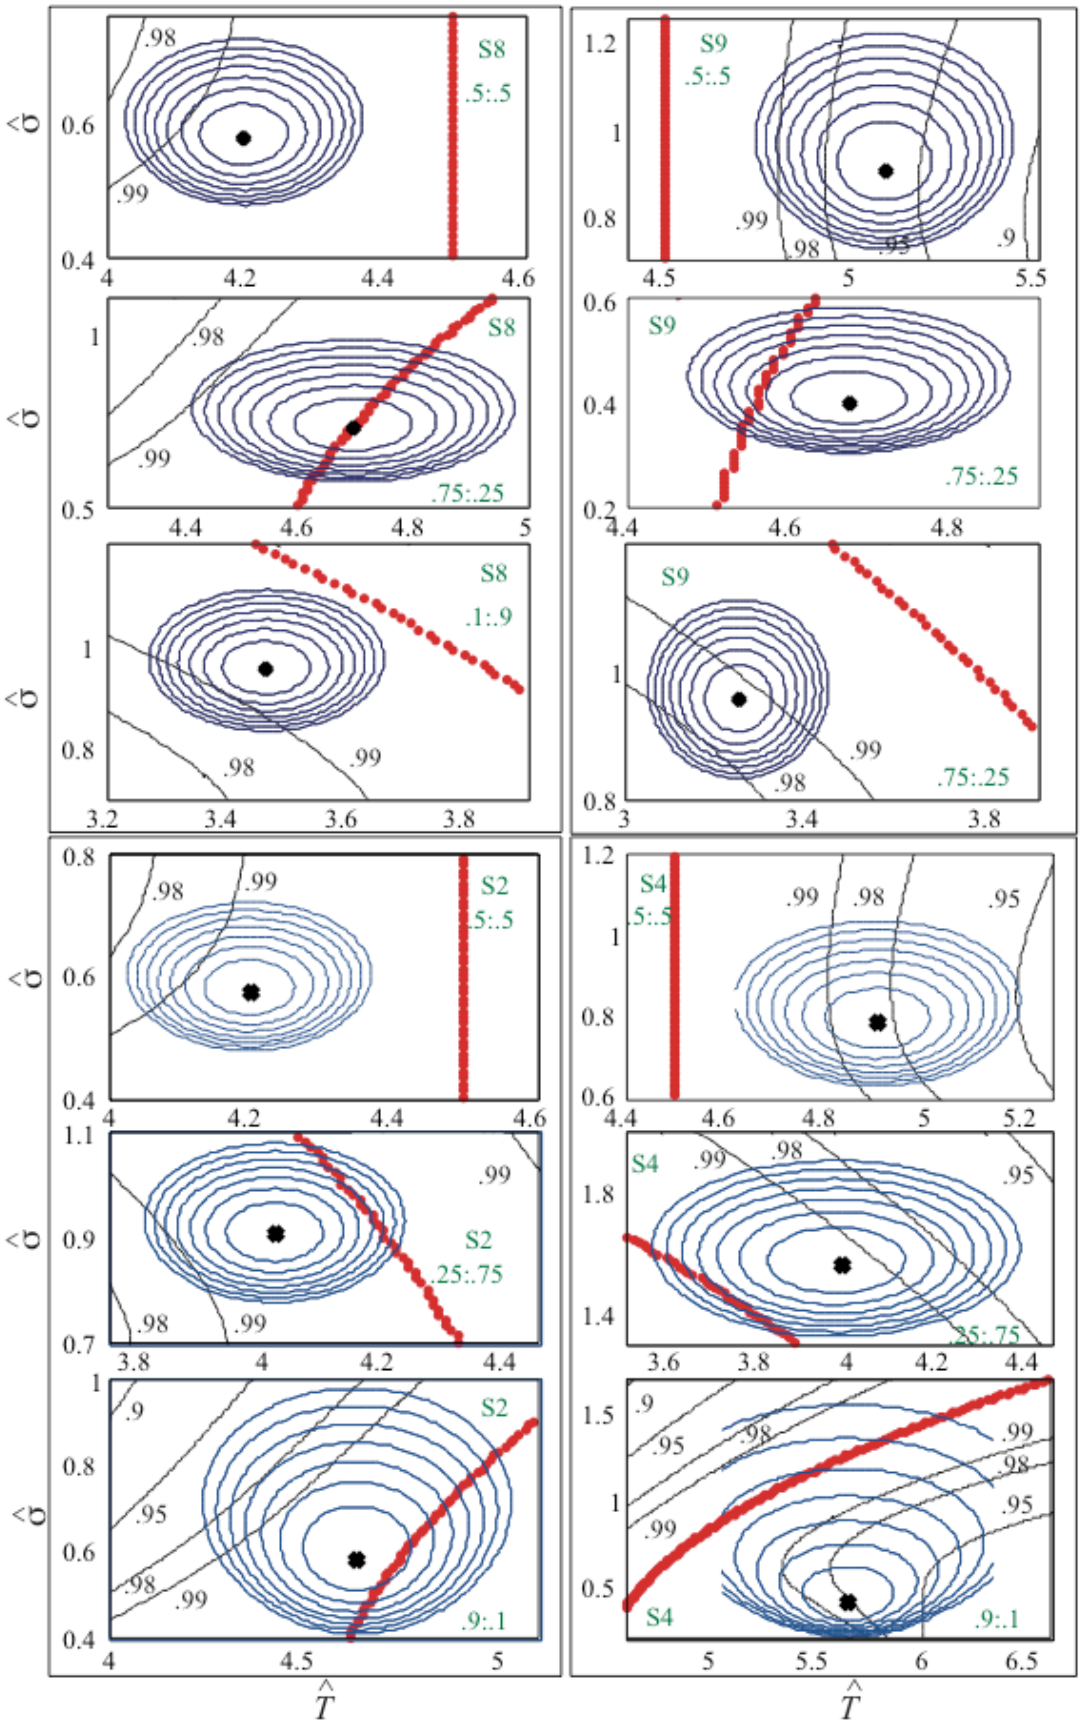
S8
.5:.5
S8
.75:.25
S8
.1:.9
S9
.5:.5
S9
.75:.25
S9
.75:.25
S2
.5:.5
S2
.25:.75
S2
.9:.1
S4
.5:.5
S4
.25:.75
S4
.9:.1
$\hat{\sigma}$
$\hat{T}$

Figure 7.8. Likelihood functions superimposed on the relative expected gain function. Here we present data from 4 individual subjects. For each subject we present 3 plots, each corresponding to a different probability condition that particular subject was tested in. Subject numbers and probability conditions are stated by green text within each plot.

We have compared the proportion of MPEG across different probabilistic conditions separately for two groups of mice. For comparison, we used repeated measure one-way ANOVA. Our analysis did not reveal a significant difference in proportion of MPEG across different probabilistic conditions for either group. The average proportion (averaged across subjects) was .98 ± .01. Figure 7.9 depicts the histogram of proportions of MPEG.

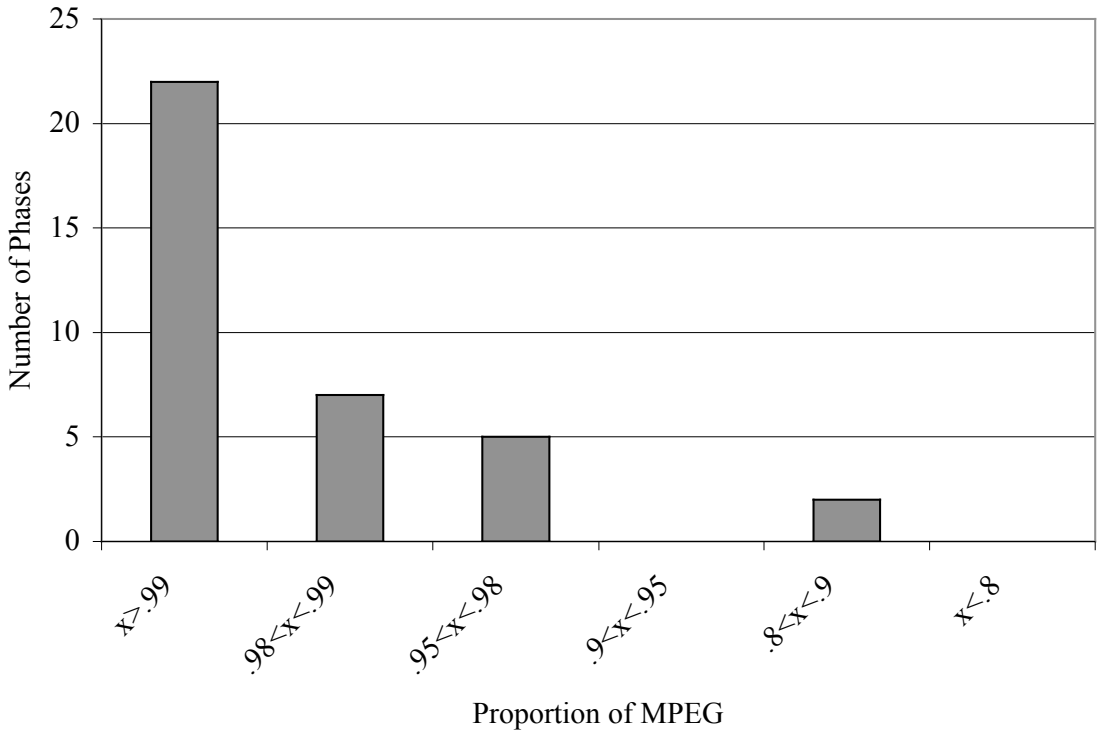

Figure 7.9. Observed proportion of MPEG in each phase.

Improvement in the course of sessions:

Finally, we checked if mice subjects got closer to the optimal switch latency over the course of last 10 sessions of the phase. In order to answer this question, we evaluated the F-ratio of variances of $|\hat{T}-T_o|$ during the first and last quarters and deciles of each subject/phase. Figure 7.10 depicts the frequency of different F-ratios presented on a logarithmic scale for both first and last quarters (top panel) and deciles (bottom panel). As this figure suggests in a large proportion of the phases, F-ratio was around 1 suggesting no change in the absolute temporal distance to the optimal switch latency across.

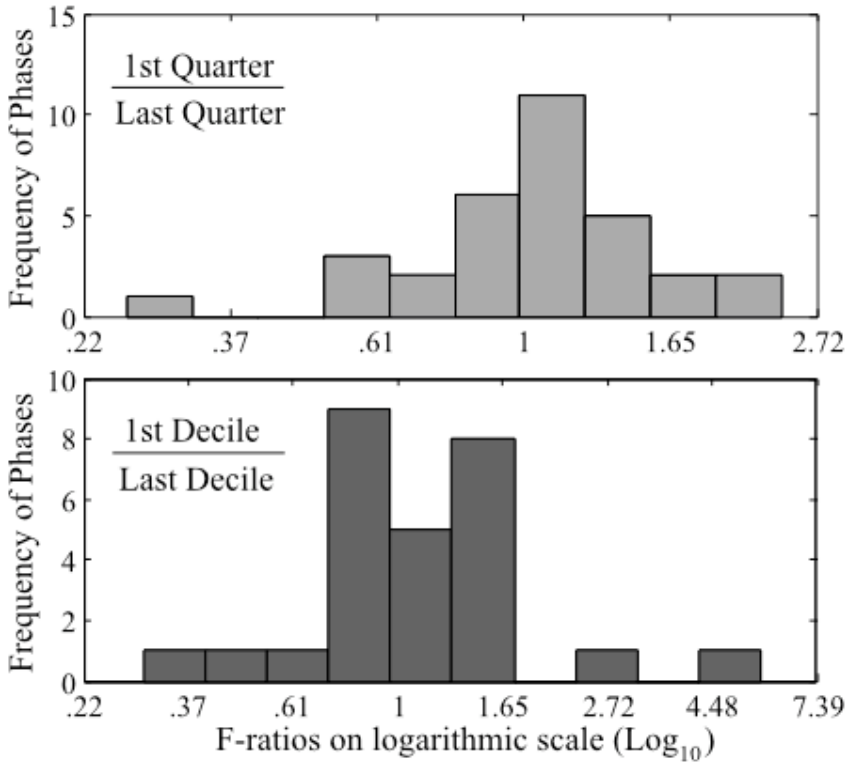

Figure 7.10. Frequency of phases in which different F-ratios were observed with mice subjects. In the top panel we present the data gathered from the comparison of first and last quarters and in the bottom panel we present the data gathered from the comparison of first and last deciles.

Figure 7.11 depicts the proportion of phases in which there was a significant difference in $|\hat{T} - T_o|$ across first and last portions of the phases. It suggests in only a small proportion of the phases, we observed a significant change in the absolute distance to the optimal switch latency. When there was a change however, this was mostly in the direction of improvement rather than disruption.

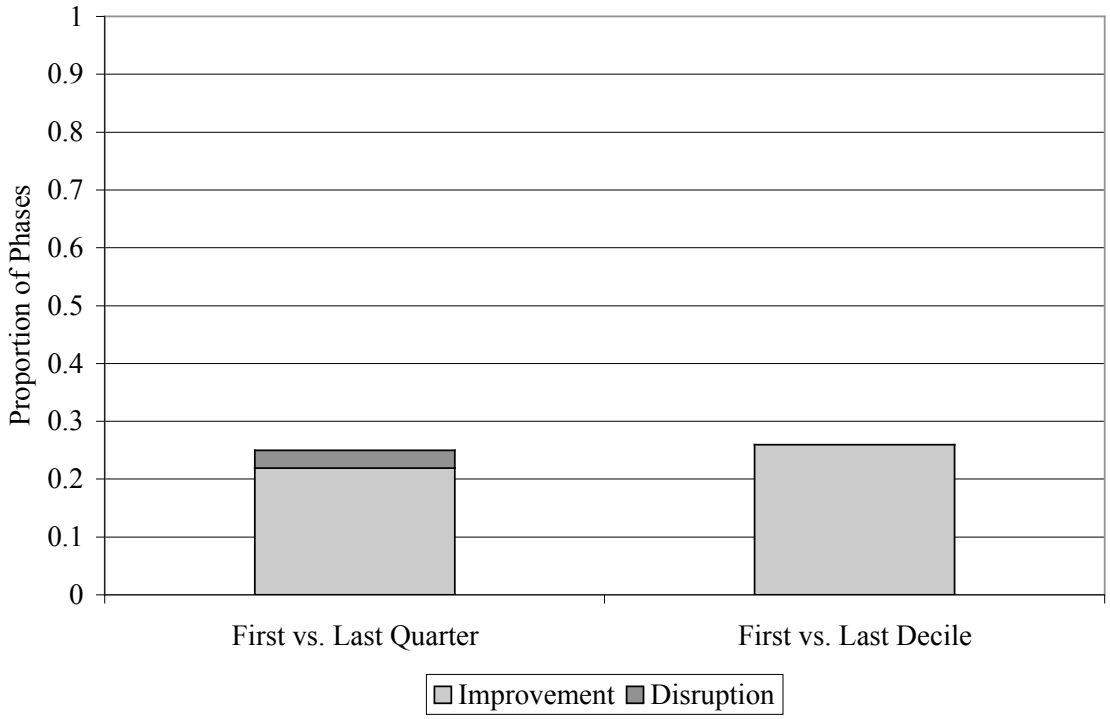

Figure 7.11. The proportion of phases in which there was a significant change in absolute temporal distance to the optimal switch latency. Light gray bars represent the proportion of those phases in which there was improvement, that is subjects got closer to the optimal switch latency. Dark gray bars represent the proportion of those phases in which there was disruption, that is subjects got farther away from the optimal switch latency.

Further, we have conducted a paired sample t-test across the mean absolute temporal distance in the initial and latter portion of the sessions both in terms of quarters and deciles. We did not find any significant differences in the mean temporal distances between first and last quarters and deciles of phases. In the case of deciles, the value of

$|\hat{T} - T_o|$ tended to be smaller during the latter portion of the phase however, this difference did not reach statistical significance. Based on these analyses we concluded that there was no general improvement or disruption throughout the session.

SECTION 8

SIMULATIONS

When considered as a model of subjects' decision-making processes, our optimal temporal decision-making model, similar to optimal movement planner, assumes that subjects can represent the task relevant magnitudes and uncertainties and apply combinatorial operations on them. Specifically, under our model subjects are assumed to represent the payoff matrix and extrinsic and intrinsic uncertainty related to the occurrence of the trials and the representation of temporal intervals, respectively. Given the information about these task relevant parameters and computational power, we propose that subjects planned the optimum switch latency from the very start. This approach has specific quantitative and qualitative predictions about the distribution of switch points over trials. To summarize, under an accurate representation of the task relevant parameters and combinatorial operations applied on these representations, one would expect the central tendency of the switch points to be around the optimum aim point predicted by the model. Moreover the uncertainty around this aim point would be expected to be random/not systematic as a function of the trials. Thus, according to the approach assumed by our model, one would not expect any form of convergence to the optimum point of responding over trials. Our data have suggested that in general convergence is not a quantitative property of the successive switch latencies. Given these assumptions about the decision-making process and the form of its inputs, our approach represents an information processing view of cognition.

On the other hand, a critical question one should ask before drawing such conclusion is how much representational structure and computational power is needed in

order to account for the empirical findings. Before attributing the power of our model to the information-processing (representational and computational) framework that it adopts, one should test the possibility that the empirical data can also be explained by non-representational accounts. Neural timing model of Grossberg & Schmajuk (1989) is an example of a model with only modest representational assumptions. Their model is called spectral timing model, because it assumes an array of stimulus-activated neurons with a spectrum of activation dynamics. The onset of a temporal warning stimulus innately activates a particular array unique to that stimulus. The dynamics of activation within the array varies systematically: Activity in the neuron with the most rapid dynamics rises and decays with short time constants. These time constants become progressively slower as one moves from neuron to neuron within the array. Thus, the locus of maximum activity within the array indicates the time elapsed since the activating stimulus.

When an unconditioned stimulus (US) is experienced, associative connections between "US neurons" and the neurons in the timing array are incremented. The increments are a monotonically increasing function of the momentary level of activity of a given neuron within the array. Thus, early in the development of these associations—after relatively few trials—the neurons whose activity peaks around the time of US occurrence, have stronger associations with the US neurons. Appropriate thresholding leads to an appropriately timed response to the US. We presented this model since it constitutes a representative alternative model of timing that is in the spirit of associative learning. On the other hand, we did not consider this particular associative model in itself since it has a well-known problem of simulating temporally controlled

responses only for the initial few number of trials. After that point, each neuron in the array irrespective of their activation has an asymptotically strong association with the US. This results in the disappearance of timed anticipatory responding with prolonged training, which is, of course, not what is in fact observed.

The existence of this and other associative/neural models that attempt to minimize the amount of representational and computational power that is assumed lead us to ask, "How much representational-computational power is required to account for the nearly optimal risk assessment that we see in both the human and the mouse data?" A formal way of answering this question is running simulations of models that assume different degrees of representational structure underlying it and evaluate their performances with the quantitative characteristics of our data.

Models that make the fewest representational assumptions are trial-and-error models with an array of response-latency options. Because there is no representational structure assumed in these models, the system does not know any structural/metric relation between these response alternatives. For instance, in the temporal domain, for a given pair of response alternatives/latencies the system does not know one is larger than the other one, namely the ordinal relation. So, in sampling latencies from trial to trial, it cannot move systematically in the direction of better latencies. Lacking the representational power required to direct sampling systematically in the right temporal direction, these models should use sampling rules that do not require any representation of the temporal ordering of intervals.

One of the common implementation of such a rule is adjusting the weights of latencies as a function of failure and success and next time sampling across the latencies

by utilizing these weights. Under the assumptions of these models, this is the only information about how good a latency is. In the following sub-section we will describe the details of the associative weight updating and sampling rules we used in our simulations.

8.1. Simulation Procedure:

There are two major components of our simulations:

1) Associative weight adjustment rule
2) Sampling rule utilizing associative weights

8.1.1. Associative Weight Updating Rule:

In adjusting the associative weights we used an algorithm that closely represents the well-accepted Rescorla-Wagner model for adjusting the associative weight of CS (will be referred to as "latency" after this point). Rescorla-Wagner model has three parameters: α , β, and λ. λ takes the value of 1 when the latency is reinforced in a discrete trial and it takes the value of 0 when it is not reinforced or penalized in a discrete trial. α is the factor used to increment the weight of a latency, when it is reinforced:

$$V + \alpha(\lambda - V) = V + \alpha(1 - V)$$

and β is the factor used to decrement the weight of a latency when it is not reinforced (or penalized).

$$V + \beta(0 - V) = V - \beta V$$

In these formulations V represents the current associative weight.

Let's assume that latency's current associative weight (V) is .2 and the values of α and β constants are .01 and .02, respectively. Further, assume that in Trial 1 latency is

reinforced. At the end of this reinforced trial the associative weight of the latency will be incremented in the following fashion:

$$Trial_1: 0.2 + 0.01 \times (1 - 0.2) = 0.208.$$

Let's assume that in Trial 2 latency is reinforced again. At the end of this trial the associative weight of the latency will again be incremented in the following fashion:

$$Trial_2: 0.208 + 0.01 \times (1 - 0.208) = 0.216$$

Now assume that in Trial 3 latency is not reinforced. At the end of this trial the associative weight of the latency will be decremented in the following fashion:

$$Trial_3: 0.216 + 0.02 \times (0 - 0.216) = 0.212$$

When one assumes a single latency (or CS) that can have associative connection with US, the associative weight might be reflected in the vigor of the conditioned response (CR) at that latency (CS) elicits or the probability of the CS eliciting a response at that latency. Thus, on the face value this process well fits in the simple classical and operant conditioning protocols, in which the organism learns the relation between two stimuli (classical conditioning) or a response and a stimulus (operant conditioning).

In cases that involve more than a single latency the process of updating associative weights and mapping these associative weights to the responses takes a more complex form. Our simulations will involve multiple possible response latencies. Regarding the updating of associative weights, we will consider two main possibilities. First, one can treat each latency's associative weight independent of others. In these cases, in a given trial only the associative weight of the latency that was implemented in that particular trial is updated. The associative weight of other latencies that were not implemented in that particular trial stays unchanged.

Another way of treating the updating of associative weights in these situations is defining a single asymptotic associative weight that is distributed across all possible latencies. In this case, at the end of a trial the associative weights of all latencies are updated. Specifically, the associative weight of the reinforced latency is incremented and this incremented amount is decremented from all other latencies. This way the reinforced latency inhibits all other latencies. In the case of an unreinforced (penalized) latency, the associative weight of the relevant latency is decremented and that decremented amount is added to the associative weight of all other latencies. This way the not reinforcing a latency facilitates all other latencies.

8.1.2. Sampling Rule:

In cases that involve a single latency, the effect of the associative weight is usually modeled as being reflected in the vigor or probability of responding occurring in a discrete trial or not. The cases that involve more than a single latency however the problem again takes a more complex form. Since there are multiple latencies in play, the mapping problem takes the form of deciding which latency is going to occur in a given trial as function of the underlying associative weights. There are multiple ways of mapping associative weights to the responses. In other words, given there are multiple latencies with different associative weights, there are multiple ways of determining a response as a function of these associative weights. Some of these sampling rules can be listed as the following.

a) Winner-Takes-All

b) Sampling randomly across latencies with maximal weights

c) Probabilistic sampling

Before going into the details of different sampling rules, we would like to describe how a temporal task (such as ours) could be treated in associative terms. As in the spectral timing model, one can assign a different latency for each different temporal interval that span over a range (e.g. L_1 = 1 s, L_2 = 2 s, L_3 = 3 s, L_4 = 4 s, …, $L_n = n^{th}$ s). Keeping this structure in mind, assume that a subject is reinforced for its first response 10 seconds after the onset of a signal, namely fixed-interval (FI-10 s) schedule. In such protocols, after a number of trials subjects exhibit anticipatory responses around 10 s. The probability of anticipatory responding decreases as one gets farther away from 10 s in both directions (e.g. in a peak-interval trial). We can assign a different latency for each second up to 20 s. In associative terms one would explain the higher expectancy of reinforcement around 10 s by the higher associative weight that L_{10} has gained compared to other latencies in the set. In other words, experiencing reinforcement around 10 s resulted in higher associative weight for L_{10}, which is the latency associated with that duration. Not being reinforced at durations other than 10 seconds, results in decrements in the associative weight of other Ls that are associated with other temporal intervals. Different sampling rules that will be presented in the next subsection will be explained in these terms.

1. Winner-Takes-All:

With winner-takes-all function, in a given trial latency with the highest associative weight will elicit the response at the temporal interval it stands for. For instance in the following set of latencies: [L_1 = .16, L_2 = .1, L_3 = .01, L_4 = .17], where each latency stands for a second-long temporal intervals L_4 will be the one that will elicit the response at the 4^{th} second of the trial since it has the highest associative weight

compared to other latencies in the set. The associative weight of this latency will be updated based on its consequence in that trial. It still being the latency with the highest associative weight by the end of this trial, it will again elicit the response at the 4^{th} second of the next trial. If on the other hand, its associative weight is decremented (as a result of non-reinforcement or penalty) such that it does not carry the highest associative weight by the end of the trial, in the next trial the response will be elicited by another latency with the highest associative weight at the temporal interval it stands for. As one can imagine this sampling rule would not result in much variation in terms of the delays of responses across trials. This is because once a latency that is likely to be reinforced is sampled, the algorithm will keep sampling the same latency over and over again.

In our simulations we induced variation in sampled response latencies by introducing noise in the sampling process of the latencies. In order to achieve this derived from the variability signature of the empirical data, we formed an error distribution around each latency. Although the simulation would aim at the latency with highest associative weight, it could end up testing other latencies, the probability of which was defined by the empirically derived error distribution formed around the aimed latency. For instance, even though the simulation was supposed to pick L_4 because of its highest associative weight, it could end up testing L_5. In such cases, it would still update the associative weight of CS_4 as a consequence of testing L_5. This manipulation induced a level of exploration to this sampling rule.

2. Sampling randomly across latencies with the highest weights:

Sampling randomly across latencies with the highest associative weights as a

process resembles the winner-takes-all sampling rule. In this rule, the algorithm ranks the latencies based on their associative weights in a descending order and chooses randomly across latencies that constitute the top portion of the rank. For instance, for the following latencies set: [L_1 = .1, L_2 = .09, L_3 = .01, L_4 = .12, L_5 = .2, L_6 = .45, L_7 = .12, L_8 = .03, L_9 = .04, L_{10} = .17] if the algorithm is asked to sample randomly across 5 latencies with highest associative weight, it will randomly choose across [L_6, L_{10}, L_4, L_7, and L_1]. In this subset of the latencies, each latency will have equal probability of being sampled. We will refer to this rule as "Max 5" at the following paragraphs.

Assume that L_4 is randomly sampled in the current trial. If it is reinforced its associative weight will be incremented and if it is not reinforced (or penalized) its associative weight will be decremented, accordingly. If the former case, that is when the latency is reinforced, the latencies subset (n=5) will stay the same for the next trial. In the latter case, that is when the latency is not reinforced, the latencies subset for the next trial would change if the current associative weight of L_4 will not the highest compared to latencies that were not included in the previous latencies subset. This sampling rule aimed to establish a higher level of exploration across a wider range of latencies compared to winner-takes-all functions. Thus, it prevents the simulation will be stuck with a latency when it is reinforced, which was generally the case in the winner-takes-all function.

3. Probabilistic sampling rule:

Probabilistic sampling rule has a more sophisticated utilization of the associative weights in sampling across latencies compared to the first two sampling rules presented above. The sampling occurs probabilistically as a function of the associative weights of

the latencies. For instance, in the following set of latencies: [L_1 = .15, L_2 = .3, L_3 = .75, L_4 = .06, L_5 = .57, L_6 = .36, L_7 = .06, L_8 =.6, L_9 = .12, L_{10} = .03], the probability of L_1 being sampled in the current trial is .15/3 → p (L_1) = .05, the probability of L_2 being sampled is .3/3 → p(L_2) = .1, the probability of L_3 being sampled is .75/3 → p(L_3) = .25, etc. The denominator (3) is the sum of all associative weights in the latencies set.

When the sample latency is reinforced in the current trial, its associative weight is incremented, which increases the probability of that latency to be sampled in the next trial. Let's assume that L_2 was sampled and reinforced in the current trial. For α value of .05, the associative weight of L_2 at the end of current trial will be .3 + ((1-.3)*.05) = .335 and thus the probability of L_2 to be sampled in the next trial will increase to .335/3.035 → p(L_2) = .11 from p(L_2) = .1. In the case of non-reinforcement or penalty, the associative weight of the sample latency will be decremented which will decrease the probability of that latency being sampled in the next trial.

All the examples presented under different sampling rules used the first type of associative updating that is associative updating assuming independent asymptotic associative weights for each latency.

8.1.3. Associative Weight Updating Rules with Different Sampling Rules:

For the simulations that used winner-takes-all and sampling from latencies with highest associative weights, we used an associative weight updating rule that assumed independent asymptotic associative weight for each latency in the set. In other words, when the associative weight of a latency was incremented or decremented it did not have any effects on the associative weights of other latencies. The associative weight of each

latency was treated independent of the associative weight of other latencies. Examples of associative weight updating presented under these sampling rules well represent the updating rule used.

In the case of probabilistic sampling rule, however, we implemented a different kind of updating rule. In this case, when the associative weight of a latency was incremented, the associative weight of all other latencies was decremented. The absolute value of sum of the decrementation was equal to the amount of incrementation applied to the associative weight of the reinforced latency. In other words, the associative weight of reinforced latency would be augmented at the expense of all others. To illustrate this process let's take the latency set presented under the probabilistic sampling subsection: [L_1 = .15, L_2 = .3, L_3 = .75, L_4 = .06, L_5 = .57, L_6 = .36, L_7 = .06, L_8 =.6, L_9 = .12, L_{10} = .03]. Again let's assume that L_2 was sampled and reinforced in the current trial. For α value of .05, the associative weight of L_2 at the end of current trial will be .3 + ((1-.3)*.05) = .335 and the associative weights of each other latency will be decremented by .035/9 = .0039, denominator being n-1. Thus the probability of L_2 to be sampled in the next trial will be .335/3 → $p(L_2)$ = .112. This is slightly higher than the probability that was determined in the previous calculation presented under sub-section of "*Probabilistic Sampling*" ($p(L_2)$ = .11) in which we assumed independent asymptotic associative weights. As mentioned earlier, with such an updating rule the latency that was reinforced actively inhibited all other latencies by taking over their associative weight. When the associative weight a latency reached to 0, we excluded that latency from our set. This ensured that the associative weight of each latency in the set was positive and their sum added to the same value, namely the asymptotic associative weight.

On the other hand, when the sampled latency would not be reinforced in the current trial its associative weight would be decremented by $V \times \beta$. This associative weight would be added to the associative weight of all other latencies in the set at equal amounts. This again ensured a constant sum of associative weight across all latencies in the set. The constant sum of associative weights across all latencies can be treated as the additivity of all probabilities to 1.

Up to this point, we have presented the associative updating and sampling rules that we have implemented. In our simulations, we further varied both the associative parameters and the degree of representational power/substitution. In the following subsection, we present the parameters that were varied.

8.1.4. Associative and Representational Manipulations:

Derived from our original motivation, we substituted for the representation power (the power that optimal model had) in the case of simulations by varying simulation parameters. Through these variations we intended to test the effect of different levels of representational substitution on the performance of the simulations in explaining the empirical data. Representational manipulations included the restriction of the temporal response possibilities and reward/penalty magnitudes. In addition to representational variations, we also tested the effect of associative factors α and β (learning rate parameters) on the performance of the simulations. For this we ran simulations under different combinations of values for α and β. Finally, we tested the effect of semi-representational factor on the performance of simulations. This last factor was the

number of latencies that the simulations could sample from. The details about these parameters and the nature of their integration is presented in the following subsections.

1. Alpha:

Alpha is the learning rate that was used to update the associative weight of a latency after its reinforcement. Higher the value of α, faster is the learning in the course of reinforced trials. Alpha was assigned values of .01 or .05.

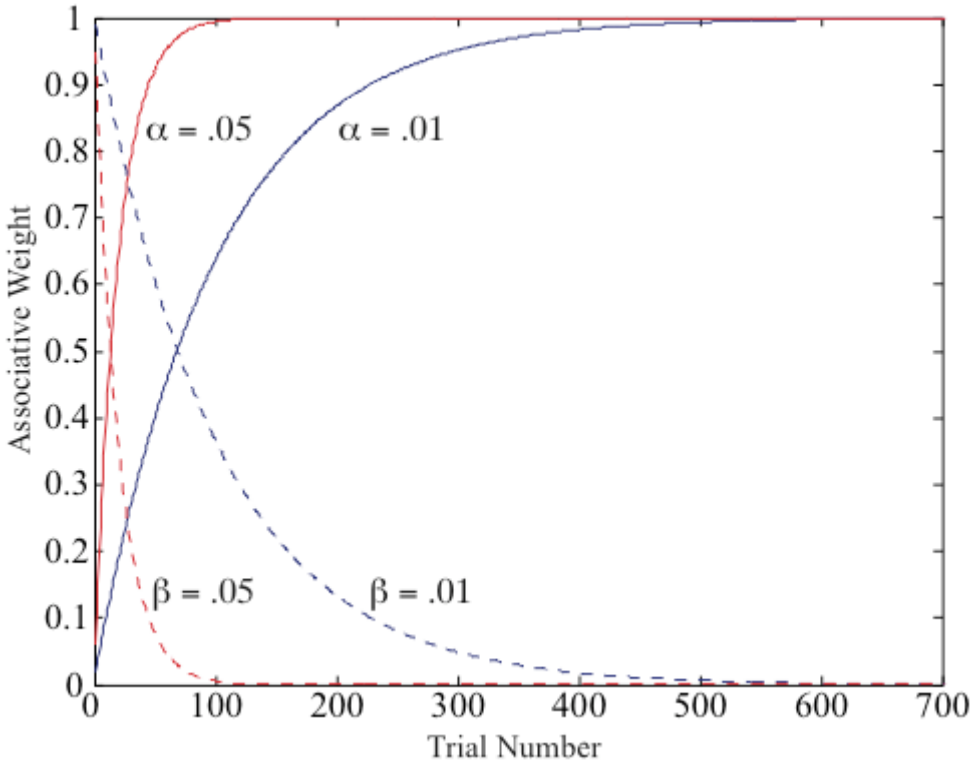

Figure 8.1. Effects of different values of α and β on the rate of acquisition.

2. Beta:

Beta is the learning rate that was used to update the associative weight of a latency after its non-reinforcement (or penalty). Higher the value of beta, faster is the learning due to non-reinforcement. Beta was assigned values of .01 or .05. The ratio between α and β determines the relative effects of reward and non-reinforcement/penalty on the

updating of the associative weight. For illustrative purposes, we have simulated the effects of different values of α and β on the rate of acquisition and extinction, respectively. The results of these simulations are demonstrated in Figure 8.1. As this figure suggests different values of these learning parameters have clear effects on the rate of acquisition.

3. Test Range:

Test range constituted the main representational substitution in our simulations. The test range was simply the temporal interval that simulations were allowed to sample latencies from. It can be considered as compensation for the lack of representational power by an add-hoc restriction of temporal response possibilities. For instance, if the test range was determined to be 1-2 s, the simulation could sample only those latencies that fell in between 1 and 2 seconds. Three test ranges that we supplied to simulations for human data were 1-4 s, 0.1-3 s, and 0.1-60 s (short and long referents for humans were 2 and 3 s). Three test ranges that we supplied to simulation for mice data were 1-8 s, 2-7 s, and 0.1-60 s (short and long referents for mice were 3 and 6 s).

Note that all the test ranges we supplied as representational substitutions to the simulations were concentrated around the critical temporal intervals. In this sense, they forced the simulations to sample latencies from a temporal range that was likely to lead to a point within the critical range. Representational substitution got stronger as the test range was set tighter around the critical temporal intervals. For instance, 1-4 s test range for humans and 1-8 s & 2-7 s test range for mice greatly diminish the opportunities to sample the latencies outside the critical temporal intervals. In other words, with these test ranges the simulations are doomed to sample latencies from the critical temporal range.

Although, similar levels of representational substitutions were supplied under other test ranges, they were not as representationally constraining as these ranges. This particular manipulation would be considered more of a compensation/rather representational aid since with this manipulation we did not introduce a metric or ordinal structure, but only restricted the response possibilities. In other words, the manipulation of the test range does not tell the simulation which direction to go given a consequence of an implementation.

4. Magnitude Effect:

In our human experiments, the pay-off matrix varied, that is, the gains and losses associated with the four different outcomes of a trial. And the pay-off metric appropriately affected subjects' target switch latency. In running simulations that vary the power of the underlying representation, the question arises how to factor in the pay-off matrix. In our simulations, the magnitudes were introduced as multiplicative factors that affected the learning rate parameters of α and β. In other words, in the case of reinforcement α was multiplied by the magnitude of reward

$$V + ((1 - V) \times \alpha \times R^{+})$$

where R^{+} is the reward magnitude earned. In the case of non-reinforcement or penalty β was multiplied by the magnitude of penalty

$$V - (V \times \beta \times R^{-})$$

where R^{-} is the penalty magnitude received (in the case of non-reinforcement β was multiplied by 1). If the magnitudes were determined not to have any effect on the updating process, no matter what the difference was between the payoff of successes in short and long trials, they resulted in the same amount of increment/decrement in the

associative weight. The representational aid of magnitudes were expected to help simulations better approximate the empirical data since they have a critical role in the solution to the optimization problem. The magnitude effect was not a parameter when we simulated in mouse data since in all task conditions of mice experiments the gains were symmetrical and there was no penalty.

5. Number of Latencies within the Critical Response Range

The critical temporal range was 2-3 s for human subjects and 3-6 s for mice subjects. With this parameter we determined the number of latencies that fell in between these critical temporal ranges. The number of latencies was set either to 5 or 25. These latencies were set at equally distant intervals to each other on a logarithmic scale. Then, taking into account the specified number of latencies that we determined to fall within the critical range, the simulation also formed a number of latencies values within the determined test range (e.g. 0.01-60 seconds) that were also set at logarithmically equal intervals to each other.

In assuming latencies evenly spaced on a logarithmic scale, we are, of course, introducing some representational structure. Without some minimal representation of the duration of an elapsing interval, there does not appear to be any way of simulating timed behavior. The number of latencies used can also be treated as a measure of temporal precision, namely higher the number of latencies higher the resolution of temporal intervals to sample. Table 8.1 summarizes all the associative and representational parameters used in our simulations of human and mice data separately.

Summary of the procedure:

At the beginning of each simulation, each latency was assigned a random weight. As mentioned earlier, in running these simulations, a latencies is sampled from the test range according to a specific sampling rule and the program checks if that latency is successful or not. In our case, a latency was considered to be successful when it fell in between the critical durations (e.g. 2 =< x <= 3 for humans) and unsuccessful when it fell out of this range (*x < short referent or x > long referent*). Given the variability in the empirical data, sampling also in these simulations was not defined to be a noiseless process. We assumed a level of noise for these latency values, which are sampled for testing. We have used empirical temporal variability in modeling noise around each latency. This error distribution was formed using the CV ($\hat{\sigma}/\hat{T}$) gathered from the session to be simulated.

Table 8.1 Simulation Parameters

	HUMAN	MICE
Short Duration:	2 seconds	3 seconds
Long Duration:	3 seconds	6 seconds
Number of Latencies	5, 25	5, 25
The range of latencies / Test Range:	[1-4 s], [0.1-3 s], [0.1-60 s]	[1-8 s], [2-7 s], [0.1-60 s]
Magnitude Effect	[YES], [NO]	[NO]
Alpha (c_1)	[0.01], [0.05]	[0.01], [0.05]
Beta (c_2)	[0.01], [0.05]	[0.01], [0.05]

In other words, for a given latency value (strategy), the actual value tested would be off by some degree. Therefore, even the latency to be tested (strategy) would be within the critical range such as 2.1, the real value that was tested could be 1.9, which would be out of range (for a short referent of 2 s). In this case, the original latency (2.1) would be treated as unsuccessful. Similarly, when the selected latency (strategy) would be out of critical range such as 1.9, the real value that was tested could be 2.01, which would be within the critical range (for a short referent of 2 s). In this case, the original latency (1.9) would be treated as successful. Briefly, for a strategy that was actually within the critical range, one could end up with failure and this would be more likely to happen for the latencies that are closer to end points of the critical range. Finally, there was no ordinal relation assumed during sampling.

8.2. Comparison with Empirical Data:

We ran the simulations for each subject's each session. For each of these simulations, we estimated the mean response latency, $\hat{T}_s$ and computed the statistical distance between $\hat{T}_s$ and $\hat{T}$ at the level of $\hat{\sigma}$. We then compared the statistical distance between $\hat{T}_s$ and $\hat{T}$ to the statistical distance between $\hat{T}_o$ and $\hat{T}$. Figure 8.2 depicts a sample comparison of these statistical distances, in which simulation does a worse job in explaining the empirical data compared to the optimal decision-making model, when performance is operationalized in terms of statistical distance. The output of the simulation was considered to be successful when the statistical distance between $\hat{T}_s$ and $\hat{T}$ was equal to or smaller than the statistical distance between $\hat{T}_o$ and $\hat{T}$ at the level of $\hat{\sigma}$.

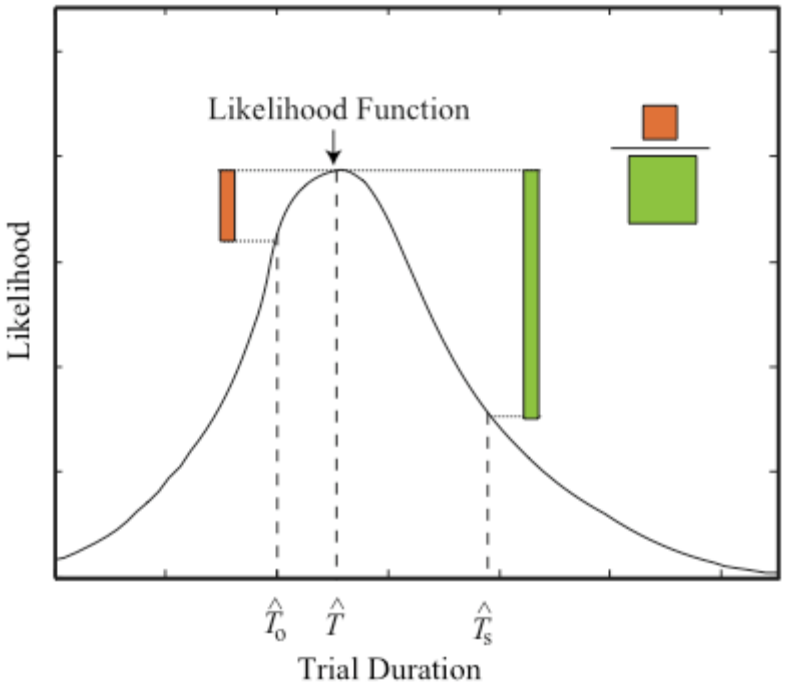

Figure 8.2. Illustration of comparison of statistical distance between $\hat{T}_s$ and $\hat{T}$ to statistical distance between $\hat{T}_o$ and $\hat{T}$. Orange block corresponds to the statistical distance between $\hat{T}_s$ and $\hat{T}$ and the lime block corresponds to the statistical distance between $\hat{T}_o$ and $\hat{T}$. These comparisons were conducted with the likelihood function at the level of $\hat{\sigma}$.

For each simulation parameter combination, we have run 5 separate simulations. From these 5 simulations, we evaluated those that simulations that had the median statistical distance to $\hat{T}$. We used medians instead mean as central tendency of statistical distance because simulation results with very large distances were not uncommon and they dominate the means $\hat{T}_s$ and $\hat{T}$. Indeed, this is a reason all by itself to reject these models, since we did not observe such large statistical distances in our empirical data set.

SECTION 9

RESULTS OF SIMULATIONS

Here, we report the output of these simulations for each different representational and associative simulation parameter combinations. Irrespective of sampling rule, the simulations ran in human data set did a very bad job in approximating the empirical data compared to the performance of our model. With the test ranges of 0.1-3 and 0.1-60 s in only around % 1 of instances the simulation did as good or better job compared to our model. Note that this level of performance was observed even when the simulations were aided by constraining their response to the actual response interval (e.g. 0.1-3 s). The very small number of successes was mostly observed when the simulations were aided by the representation of the reward/penalty magnitudes.

Figure 9.1 depicts the data for mice subjects. In the case of mice when the we substituted for the representation of temporal intervals, by constraining the sampling range 2 seconds below the short duration and above the long duration, the percentage of successful instances ranged between % 5 – % 25. The percentage of successful instances increased as a function of the complexity of the sampling rule.

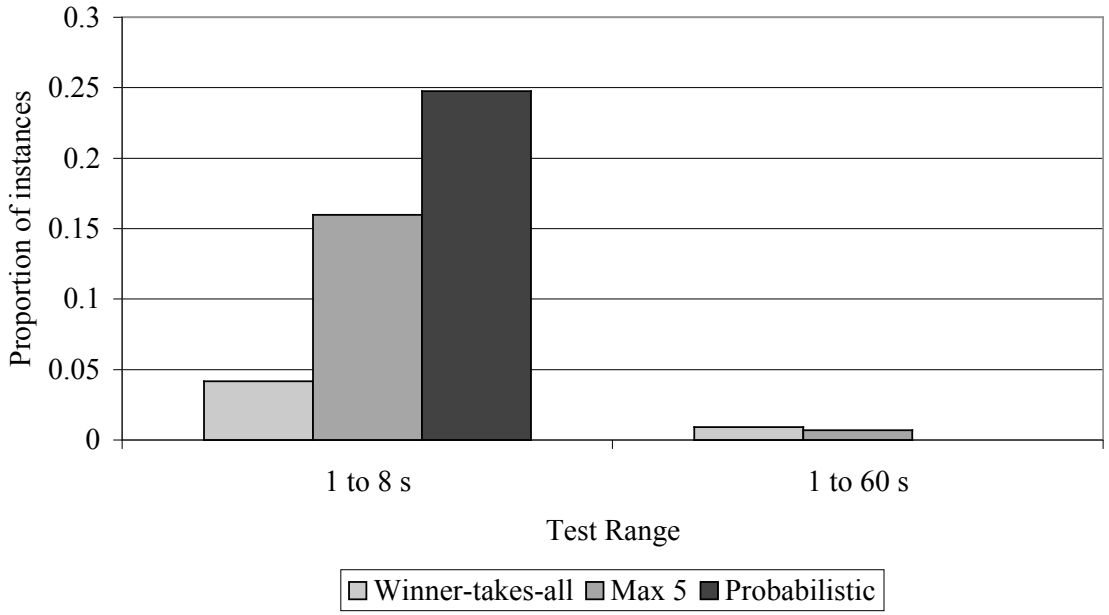

Figure 9.1 The proportion of simulations the exhibited successful predictions for mice data set. Here we present the results from two test-ranges for three different sampling rules. On the same graph, we differentiate between different sampling rules by different shades of gray.

We further tested how well the simulations performed when the opportunity of sampling latencies outside the critical range was dramatically diminished further by constraining the test range 1 second below the short duration and 1 second above the long duration. Figure 9.2 depicts this data for human (top panel) and mice (bottom panel) subjects. Figure 9.2 suggests that even when the test range was tightly constrained around the critical temporal range, for human data the proportion of successful instances was at most .15. The effects of sampling rules became clearer when simulations were not allowed to sample latencies that fell outside the range delimited roughly by the short and long referent durations. Specifically, the proportion of successful instances increased monotonically as a function of the complexity of the sampling rule.

A very similar effect of the sampling rules was also observed for mice data,

namely the proportion of successful instances increased dramatically with the complexity of the sampling rule. Overall, the proportion of successful instances was higher for simulation ran for mice data compared to human data. This is very likely to have occurred due to the relative tightness of the test range around the critical range. In simulations ran for human data, this test range was 1-4 seconds. Considering that the critical response interval was in between 2-3 seconds the total duration that falls out of the critical range; ((2-1)+(4-3))/(3-2) = 2, was % 200 of the response range. In the case of mice data, this duration was ((3-2)+(7-6))/(6-3) = .67, % 67 of the response range. In that sense, the simulations ran for human data with 1-4 s test range is more comparable to simulations ran for mice data with 1-8 s test range. In 1-8 s test range, the total duration that falls out of the critical interval ((3-1)+(8-6))/(6-3) = 1.33, was % 133 of the response range. This very relation suggests that the substitution for the representation of temporal intervals was a good predictor of the success rate of simulations.

Finally, we tested the effects of different simulation parameters on the absolute temporal distance between $\hat{T}_s$ and $\hat{T}$ as a function of all simulation parameters for different sampling rules.

1. Effects of associative parameters:

In the simulations ran for human data set, we found a significant effect of α on the temporal distance between $\hat{T}_s$ and $\hat{T}$ only during probabilistic sampling, $t(3778) = -1.96$, $p = .05$. In the simulations ran for mice data set, the significant effect of α was observe only during winner-takes all function, $t(1294) = -2.84$, $p < .01$. In both cases, a higher

value of α resulted in predictions that were farther from $\hat{T}$ under this sampling rule. For either data set we did not find a significant effect of β and number of latencies that fell within the critical range under any sampling rule.

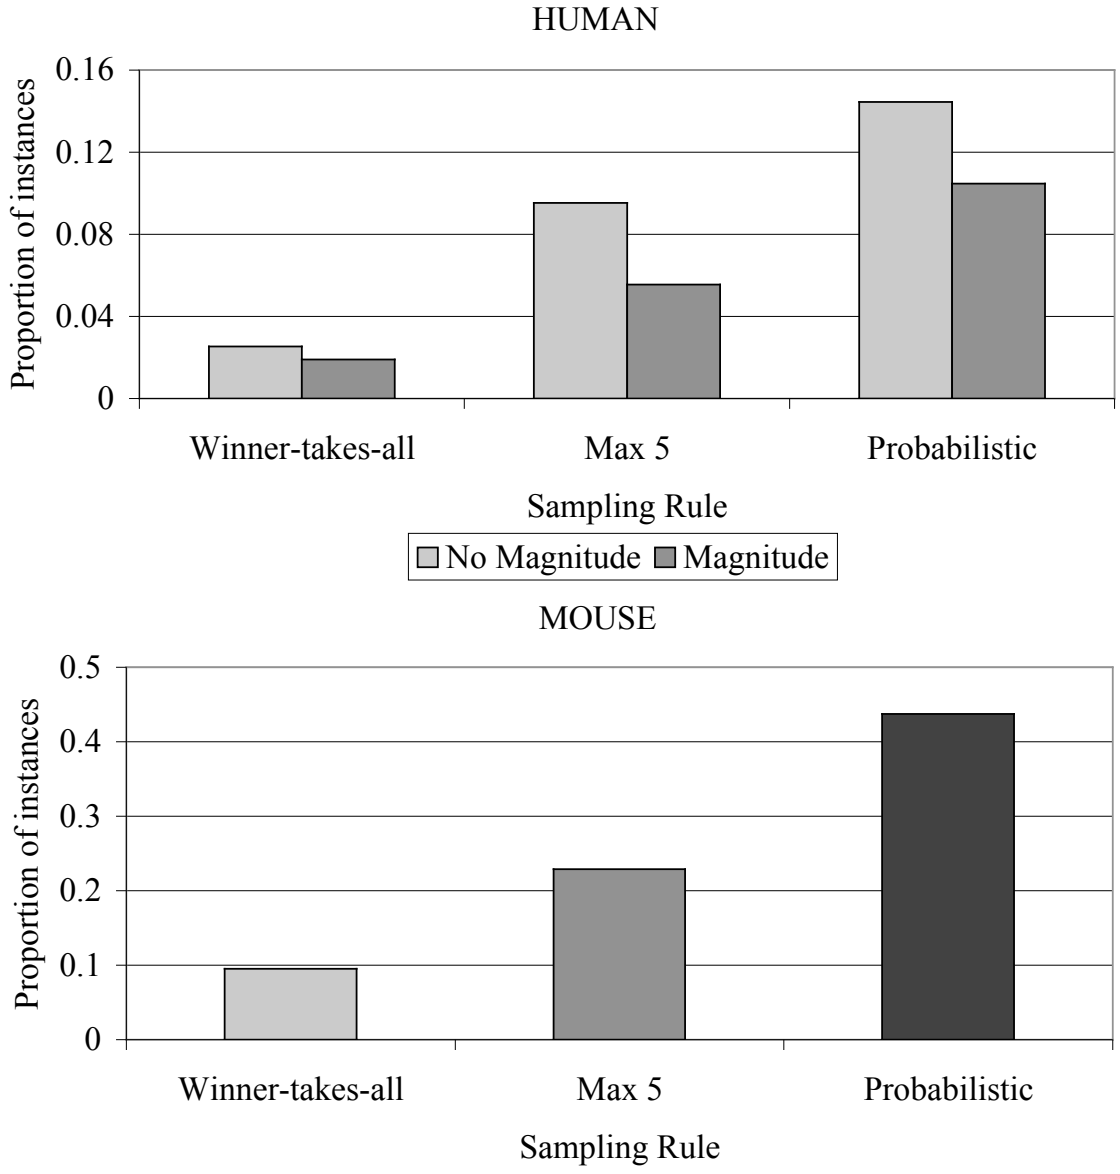

Figure 9.2. Proportion of successful instances as a function of different sampling rules for human (top panel) and mice (bottom panel) under conditions in which the simulations were not allowed to sample from latencies outside the critical response range.

2. Effects of representational parameters:

In the simulations ran with winner-takes-all function for human data set, we found a significant main effect of test range on the temporal distance, $F(2,3774) = 267.74$, $p < .0001$, no significant main effect of magnitude or its interaction effect with test range. In the simulations ran with winner-takes-all function for mice data set, we also found a significant effect of test range on the temporal distance, $F(2,1293) = 163.48$, $p < .0001$. In both cases, tighter test range resulted in closer to empirical estimations.

In the simulations ran with "Max 5" rule for human data set, we found a significant main effect of test range, $F(2,3774) = 1036.51$, $p < .0001$, a significant main effect of magnitude $(1,3774) = 3.92$, $p < .05$, and a significant interaction effect, $F(2,3774) = 6.51$, $p < .01$. We also observed a significant effect of temporal range in the case of mice data set, $F(2,1293) = 1167.5$, $p < .0001$. In both cases, a tighter test range again resulted in estimates that were closer to empirical estimations. The effect of magnitude was in the expected direction, that is the introduction of magnitude representation resulted in temporal distances that were closer to empirical estimates. This effect was pronounced during a wider test range.

In the simulations ran with probabilistic sampling rule for human data set, we found a significant main effect of test range, $F(2,3774) = 1760.49$, $p < .0001$, a significant main effect of magnitude representation, $F(1,3774) = 14.8$, $p < .001$, and significant interaction effect, $F(2,3774) = 10.16$, $p < .0001$. We also found a significant effect of test range for mouse data set, $F(2,1293) = 2587.82$, $p < .0001$. For both data set as tighter test

range again resulted in estimates that were closer to empirical estimations. For the probabilistic sampling rule, the effect of magnitude was in opposite direction with what we observed with "Max 5" rule. The introduction of magnitude representation resulted in estimates that were farther from the empirical estimations. This observation was pronounced during a wide test range.

Finally, we present the average temporal distance between $\hat{T}_s$ and $\hat{T}$ as a function of sampling rule and test range for both human and mice data set. The mean $\hat{T}_s$ and $\hat{T}$s are presented in Table 9.1 separately for mice and human data set. A quick look to Table 9.1 suggests that in general the temporal distance of simulation outputs from the empirical estimations are quite different from the temporal distance between the prediction of the optimal model and the empirical estimations. The only condition in which the simulation outputs approximate the model's performance is when the test range is forced to be tightly clustered around the critical range and the probabilistic sampling rule is used. However, note that these are the conditions in which the simulations are not allowed to sample latencies from outside the critical temporal interval.

Overall, the performance of the simulations compared to the prediction of the optimal decision-maker was not uniform across different task parameters. In other words, successful cases were specific to a number of task parameters, which makes it difficult to attribute a general success rate to the trial-error learning algorithms.

Table 9.1. Absolute temporal distance between simulation and empirical estimations as a function of sampling rule for human and mice data set.

Test Range	Winner-takes-all $\lvert\hat{T}_s - \hat{T}\rvert$	Max 5 $\lvert\hat{T}_s - \hat{T}\rvert$	Probabilistic $\lvert\hat{T}_s - \hat{T}\rvert$	Empirical $\lvert\hat{T}_o - \hat{T}\rvert$
HUMAN DATA SET				
1-4 s	.78 ± .009 s	.46 ± .006 s	.38 ± .006 s	.158 ± .014 s
.1-3 s	1.75 ± .016 s	1.65 ± .008 s	1.55 ± .006 s	
.1-60 s	5.99 ± .293 s	6.36 ± .168 s	5.7 ± .116 s	
MOUSE DATA SET				
2-7 s	1.33 ± .026 s	.72 ± .02 s	.50 ± .016 s	.436 ± .058
1-8 s	1.94 ± .04 s	1.08 ± .026 s	.73 ± .019 s	
.1-60 s	9.6 ± .62 s	10.88 ± .29 s	8.23 ± .148 s	

Discussion

The only cases in which the simulations achieved non-negligible rates of success when pitted against a model that computed the optimal switch latency were those in which the pre-determined test range greatly diminished the opportunity of sampling latencies outside critical range. These cases were when the simulations were provided with a test range of 2-7 s, 1-8s (for 3-6 s critical range) and a test range of 1-4 s (for 2-3 s critical range). The tighter the test range was (relative to the critical interval), the higher was the chances that simulation using it would succeed in approximating the performance

of optimal model. In all other test intervals, they performed considerably badly compared to the performance of optimal model. The effect of the test range was also reflected on the temporal distance between $\hat{T}_s$ and $\hat{T}$ in the same fashion. Restricting the range for possible response latencies, of course, amounts to whispering in the model's ear where it must focus. It in effect compensates for the models' (assumed) inability to represent the structure of temporal intervals.

Unlike, the representational substitution for temporal intervals through constraining the testing range, the representational aid of reward/penalty magnitudes was not observed to have a consistent and strong effect in determining the rate of successes. The representation of magnitudes was found to help simulations only in those cases where the models were not aided by strongly restricting the possible latencies. When the possible latencies were strongly restricted (e.g. 1-4 s for 2-3 s critical interval), the representation of magnitudes was found to decrease the rate of success. Given the critical role of reward magnitudes in explaining the empirical data, this finding suggests that the representation of reward/penalty magnitudes were not integrated in the learning rule in a such a way that the information it provided was utilized in the same way that it was utilized in the optimal model. This further suggests that the combinatorial properties of the representations are also an important aspect of information-processing.

Associative parameters were not found to have any major and consistent effects on the success rate of the simulations. Both a and b are the core parameters of the associative learning models (l was kept constant – 1 for reinforcement and 0 for non-

reinforcement/penalty). Despite their critical role in the output of these models, these parameters did not predict the performance of the simulations.

Finally, sampling rules were found to affect the success rate of simulations only when the test range was set tightly around the critical intervals. In these cases, probabilistic sampling rule was found to increase the success rate of simulations. A closer look at the probabilistic sampling rule reveals that this rule implicitly uses the knowledge about the properties of probabilities. For instance, the additivity of probability to 1 is a major property of this sampling rule.

SECTION 10

CONCLUSION

Decision-making under uncertainty is a cognitive process that has been traditionally studied using descriptive paper-pencil tasks. The results of these studies have been interpreted as reflecting the non-normativeness of human decision-makers. The non-normativeness of the decisions made in these tasks has been attributed either to the biased representation of information that goes into the decision-making process (e.g. probability) and/or non-normative decision-making processes themselves (e.g. heuristics). More recent research, on the other hand, revealed that these non-normative decisions might be particular to problems in which information about the uncertainty is described verbally (explicitly) to the decision-maker by the experimenter. These researches basically reported closer to normative decisions when human decision-makers experienced the uncertainty and the consequences of their decisions over time. Further research also revealed that humans are optimal decision-makers in tasks that require motor action when the uncertainty is an intrinsic property of the sensori-motor system. The same decision-making tasks when involved also extrinsic uncertainty (imposed by the experimenter) resulted in divergences from the optimality.

We investigated decision-making under uncertainty using a temporal discrimination task. The consequences of different responses (payoff matrix) and extrinsic probabilities constituted two critical experimental manipulations. The third critical constituent of the temporal decision-making task was the intrinsic uncertainty, which could not be controlled either by the experimenter or the subject. This form of

uncertainty was an intrinsic property of the cognitive system, a limit on the precision with which it can estimate elapsed intervals.

Given the temporal nature of our task, the dominant source of intrinsic uncertainty was assumed to originate from the representation of temporal intervals in the decision-maker's mind. The representation of temporal intervals and their psychophysical properties are known to be quite similar across different species (e.g. mice, rats, pigeons, human). Use of a temporal discrimination paradigm (thus temporal information processing) as the context for decision-making under uncertainty added a translational character to our decision-making task and models developed based on the findings. This allowed us to study decision-making under uncertainty in mice using the same.

The translational character of this complex and experiential decision-making task further motivated us to assess the explanatory power of representational/computational versus associative processes, the latter of which represents the dominant models of animal learning. In the following sections, we summarize the results from human experiments, mice experiments and simulations.

Hill-climbing or Parachute Landing?

More than half of the time both human and mice subjects responded such that they ensured more that 95 % of the maximum possible expected gain. They were doing so by positioning their switch latency within 5-7 % of the available response range. The question of interest given this finding is "How did they get to the top of the hill?". There are two possible accounts for this finding. The first account that is represented by our optimal decision-making model suggests that early on in the session, subjects formed an accurate representation of the task parameters (reward/penalty magnitudes and extrinsic

and intrinsic uncertainty) and they computed the optimal point switch latency given this information. The alternative account that is more consistent with models developed in the area of psychology and computer science (e.g. reinforcement learning), would suggest that subjects found the top of the gain hill by using hill-climbing algorithm (which actually assumes some representation of gain and metric of the response dimension). The critical difference between these accounts is that the first account assumes the calculation of the locus of the optimum while the latter account assumes the discovery of optimum through trial-error.

Investigation of our empirical data set and the performance of the simulations conducted in this data set provided a ground for assessing the validity of these accounts and set a very plausible answer to the question of "How did subjects get to the top of the hill?" A close look to Figure 6.9 and 7.8 reveals that in general even though the expected gain plane was shallow around the top of the hill, the subjects' aimed switch latencies were closer to the optimal switch latency (top of the hill) than the contour delimiting 99 % of the MPEG. This observation was generalizable to other data sets not presented in these two figures. Further, our investigation of performance in the course of session did not reveal a general improvement over trials in either human or mice data sets.

These two basic observations clearly favor the first account over the second one. In general hill-climbing algorithms take a very long time to converge on the top of the hill. This fact gets more pronounceable particularly when the plane around the top of the hill is shallow as in our data set. On the other hand, findings reported above suggest that subjects located their responses at the top of the gain hill and they did so early on in the session. They did not exhibit any form of improvement in the course of session, which

would be expected to occur if they were using a hill-climbing like algorithm to discover the optimum. We conclude that our subjects landed on the top of the hill aiming for it from the very start rather than blindly climbing the hill trying to find its peak.

Human Experiments

Our results of human experiments have shown that, more than half of the time human subjects planned temporal strategies that ensured well above % 95 of the maximal possible expected gain given asymmetrical payoff matrices and under uncertainty that originated from both intrinsic (due to imprecision in their representation of elapsed time) and extrinsic stochastic processes. Through these temporal strategies, human subjects on average were found to constrain their temporally strategies within less than % 5 of the total temporal response range (3 seconds). This corresponded to approximately 150 ms of temporal window around the optimal switch latency that would maximize the gain.

The close correspondence between the empirical target switch latencies and the optimal switch latencies suggests that human subjects combined information about the reward/penalty magnitudes (payoff matrix), extrinsic uncertainty, and intrinsic uncertainty in a normative fashion. This further suggests an unbiased estimate of extrinsic and intrinsic uncertainties by decision-makers, which was shown not to be the case when extrinsic uncertainty was presented in a verbal form (described). Briefly, our results characterized human subjects as normative decision-makers under uncertainty with multiple sources, at least in the context of temporal information-processing.

Our trial-by-trial analyses of the responses suggested that subjects planned and implemented the optimal temporal strategy starting early in the session and their responses varied around this aim point due to the stochastic processes underlying timing

and/or memory for temporal intervals. There was not a systematic improvement during a session (that is, with a given set of relevant pay-off and probability parameters). This is in line with the results and interpretation of Trommershäuser et al. (2003a, 2003b).

Performance of human subjects achieved in our tasks combined with recent findings from experiential decision-making tasks suggest that non-normative solutions to the decision-making problems in classical paper-pencil tasks might be mainly due to the non-experiential/descriptive means of information gathering. On the other hand, when the uncertainty is an intrinsic property of an input used in decision-making processes, the decision-making process might very well take the uncertainty information into account in a normative manner. Our findings revealed that human subjects are optimal decision-makers under uncertainty that originates from their very representation of the task attribute that is temporal intervals. This finding has a clear parallelism with the findings of Trommershäuser et al. They too showed that human subjects were optimal decision-makers in motor planning tasks under uncertainty that was intrinsic to their motor system.

Our results allowed us to extend the claim of normative processing of intrinsic uncertainty to decision-making cases, in which the intrinsic uncertainty originates from the representation of temporal intervals. This suggests that our subjects should have represented the uncertainty in their representation of temporal intervals and used it as a critical input in planning a temporal strategy for maximizing their gain. Despite the fact that this might be characterized as a meta-cognitive process, the decision-making process that uses this information seems to be an automatic one, which does not require conscious effort on the part of the decision-maker.

Our results also constitute a divergence from some of the findings of the same researchers (e.g. Maloney, Trommershäuser, & Landy, 2006). Maloney et al. showed that when extrinsic uncertainty was introduced to their motor planning tasks, subjects were suboptimal at planning the optimal aim point. Their interpretation of this finding was that human mind might be normative in dealing with intrinsic uncertainty but not in dealing with extrinsic uncertainty. Our findings on the other hand revealed that subjects could exhibit optimal decision-making under conditions in which the uncertainty originated from both intrinsic and extrinsic stochastic processes. This is indeed consistent with the findings of Barron et al., Hertwig et al., and Weber et al.

A simple comparison of the spatial task of Maloney et al. with our temporal task might shed light on the reasons of this discrepancy. In their spatial task, in each trial, two identical targets were presented simultaneously. The penalty regions were color-coded and the subjects were verbally informed about the probability condition at the beginning of each block. The probability condition was always defined by %50 which stayed constant across each block (36 trials each). The outcome at each trial not only depended on the subjects' movement end point (intrinsic) but also to a chance element that was independent of subject's response, where rewards and penalties were stochastic. Maloney et al. expected that a normative observer would be expected to give the same responses in cases when the penalty was $ 100 at 50 % of the time and a certain penalty of $ 50.

Maloney et al. allowed their subjects 36 trials in each block to gather information about the extrinsic probability. This might not constitute enough number of experiences for one to gather an accurate representation of extrinsic uncertainty. In our tasks on the other hand, the extrinsic uncertainty was experienced over 500 trials that constituted each

session. This constitutes almost 15 times more number of trials than the ones that were used in each block in Maloney et al. In other words, we might have observed optimal decision-making under extrinsic uncertainty solely because our subjects were allowed to gather and use this information over the course of large number of trials. This discrepancy also might have occurred simply due to the fact that extrinsic uncertainty (although also experienced) was verbally communicated to their subjects but they were solely experienced by our subjects. In this sense, our introduction of extrinsic uncertainty represents more of Barron et al., Hertwig et al., and Weber et al's.

Given this state of affairs, our experiments constitute the first case in which the human subjects are shown to be optimal under uncertainty that originates from both intrinsic and extrinsic processes.

The optimal performance in our task requires one to integrate the magnitudes of rewards and penalties with the intrinsic and extrinsic uncertainty. Thus, our results suggest that such an integration of information about critical task parameters happens in subjects' mind from early on in the session. Briefly, in some respects our results confirm the optimal decision-making with intrinsic uncertainty and different payoff matrices and in other respects they confirm optimal decision-making under extrinsic uncertainties.

A closer look at the data showed that intrinsic uncertainty in our task originated from the representation of temporal intervals and was independent of the uncertainty that was due to planning (at least in the case of humans). This claim is based on our findings that variability measures of timing distributions did not change as a function of task conditions. This finding is consistent with the findings of Trommershäuser et al. (2003b) who reported no change in the motor variability as a function of their motor task

parameters. Their intrinsic uncertainty was dominantly due to the noise in the motor system and did not originate from the motor planning. Based on these findings, the intrinsic uncertainty in our tasks might be modeled solely as uncertainty in the timing or temporal memory. This finding was further supported by the fact that in almost all cases this timing uncertainly did not get better or worse over the course of sessions.

Finally, we found slight differences in the processing of reward/penalty magnitudes when they were numerical vs. graphical. For instance, the analyses of absolute temporal distances from the optimal point (both earlier and later) revealed that with graphical presentation of gains, subjects tended switch farther away from the optimal switch latency. These differences might be due to the different mappings from numerical and graphical (analogue) external representations of quantities to their representations in mind.

Our findings from human experiments offer a new interpretation for the results of classical decision-making tasks, which were presented as empirical evidence to non-normative information-processing/decision-making under uncertainty in humans. Our results signify the importance of the experiential nature of uncertainty for the expression of normative decision-making processes. This very fact motivated us to investigate the same processes in animal subjects.

Mice Experiments

In line with the optimal decision-making processes observed with humans, decisions made under both intrinsic and extrinsic uncertainty by mice were also found to be quite close to the optimal decision criterion. The aimed switch points of the mice on average ensured above % 95 of the MPEG. These aim points corresponded to 2.5-7 % of

all the values that the switch latencies could take. These measures suggest that mouse decision-making processes that take place under different sources of uncertainty are optimal.

Comparison of the human and mouse performance in these tasks further reveals a possibly better performance of mice compared to humans. On the other hand, our mouse tasks did not involve any asymmetry in terms of the reward magnitudes and any penalty. Lack of the penalty variables and/or the lack of asymmetries in the reward magnitudes in the payoff matrix might have mediated the slightly better of performance of mice compared to humans.

Briefly, the level of performance observed with mice suggests that mice subjects were sensitive to these decision parameters. This further suggests that they represented temporal intervals, extrinsic and intrinsic uncertainties and combined them in a normative fashion such that its output predicts the optimal switch latency. Thus, one can assume that in non-verbal decision-making tasks such as ours, the underlying decision-making processes in human and mouse mind resemble each other.

Representational/Computational versus Associative Processes

Optimal temporal decision-making model, assumes that representations of temporal intervals, extrinsic and intrinsic uncertainty are used as inputs in computing the optimal switch latency. Thus optimal temporal decision-making model assumes a representational and computational power underlying the decision-making process.

Given the long-standing controversies over representational explanations in psychology and frequent suggestions by, for example, the connectionist modeling community, of non-representational associative models as alternatives to information

processing models, we were also led to consider how good an account could be obtained from some models of this kind. The basic assumption of these models is that responses (or the probability of their occurrence) are determined as some function of their associative weight. The associative weights are not predetermined. They are updated through course of discrete experiences of consequences that are contingent on some stimulus or response. This process is assumed to abide by certain updating rules, which characterize the associative learning models. In our simulations, we used a widely accepted associative learning rule, namely Rescorla-Wagner model. We considered three different functions for mapping the probability of a response occurring on a given trial (or, at least, having its associative strength updated) to its associative weight.

Our simulations had five parameters. We varied rates of learning (a) and extinction (b), the density and range of possible response latencies whose associative weights could be incremented or decremented by success or failure on a given trial, and the effect or lack of effect of magnitude of the experienced gain or loss.

Remaining two parameters were manipulated to provide different degrees of representational substitution or power to the simulations. One of these parameters was the test range. When this range is constrained by the critical temporal parameters of the task, it provides a strong substitute for the representation of temporal intervals in the simulation, because it constrains the range during which the responses should occur. When this range is determined to be larger, the simulations entertain a weaker substitute for representational power since they are allowed to sample from outside the critical temporal range. As noted earlier, this manipulation can only be characterized as a substitute for the representation rather than the introduction of the representation of

temporal intervals itself, since it did not provide a metric or ordinal structure to the simulations.

Another parameter that determined representational power involved the representation of magnitude of consequences. In human experiments there were different magnitudes of rewards and penalties across different sessions. For a system that cannot represent magnitudes, the effect of any experienced reward/penalty on the updating of the associative weight would be equal. For a system that has a representation of magnitudes, the experience would result in an associative updating factor that are proportional to the magnitude of the consequences. In half our simulations we allowed the reward/penalty magnitudes to determine the increment and decrement factors while in the other half all experienced rewards/penalties had equal effects on the updating of associative weights. The number of samples can also be interpreted as relating to the representational substitution. For instance, in cases when the test range is constrained by the critical temporal intervals, higher number of samples would provide a higher "temporal resolution" for the sample of latencies.

Our analyses of the performance of simulations revealed that compared to optimal model overall simulations using an associative learning rule using different combinations of five parameters did not do a good job in accounting for the empirical data. Further analyses of the performance of simulations as a function of different simulation parameters and sampling rules revealed a clear role of assuming representations in accounting for the empirical data. For instance across all sampling rules used by the simulations, almost for all of the instances in which the simulation did as good as or better job than the optimal model in explaining the data, the test range was constrained

around the critical temporal intervals, (e.g. 1-4 s in human case). When this constraint was defined even in the actual response range (e.g. 0-3 s in human case), the trial duration during which one could respond, the simulations almost never succeeded in explaining the empirical data. This result suggests the critical role of representation of temporal intervals in explaining the empirical data. When this representational substitution is reduced even by a small degree associative updating rule fails to explain the empirical data. However, it is important to note that even when these simulations were provided with a substitution for the representation of temporal intervals, in only a small portion of the instances they explain the empirical data as good as or better than the optimal model. This suggests that the metric structure that underlies the representation of temporal intervals is a main predictor of temporally controlled responses observed in our data set.

Another parameter that related to the representational power was the representation of reward/penalty magnitudes. The representation of magnitudes and their proportional effect on the associative updating rule was found to increase the number of correct instances when the simulations had the weaker substitution of temporal representational power (wider test-range). On the other hand, when they were provided with the substitution temporal representation, the representation of the magnitudes decreased the number of instances of success in explaining the empirical data. This is very likely to be a function of the associative adjusting rule, itself that used the magnitudes in updating the associative weights.

Although increased precision of latencies seemed to increase the number of successful instances, their effect was not pronounced in the form of statistical

significance. Finally, the important role of representations in account for empirical data was reflected by the differential success rates of different sampling rules used by the simulations. Probabilistic sampling rule resulted in higher number of instances in which the simulations were as good as or better than the optimal model in explaining the empirical data compared to two other sampling rules. Within these two other sampling rules, the one that used a semi-probabilistic sampling scheme (sampling randomly across five criteria with highest associative weights) did a better job than the one that did not use a probabilistic approach at all. When one looks closer at the sampling rules used in the probabilistic case, the knowledge that probabilities add up to 1 was implemented in the functioning of these simulations, which reflects another kind of representation: the knowledge about the compositional properties of probabilities. When the implementation of this knowledge was aided further by the substitution for the representation of temporal intervals, the simulations were most likely to succeed in accounting for the empirical data.

Unlike the parameters that determined the representational power/state of simulations such as temporal test range and magnitude effect, the parameters that were directly related to associative updating process itself such as a and b were not found to be the simulation parameters that reliably affected the number of successful instances. The only simulations that these traditional associative parameters were found to have an effect on (in addition to other parameters) were the ones that used winner-takes-all sampling rule. This was the sampling rule that resulted in the least number of successful instances. As mentioned in the results section, small a values tended to help the simulations under this sampling rule. Small a values would mediate this effect by inducing exploration by

incrementing the associative weight in only small amounts in the case of success. In probabilistic sampling rule, which was found to be far more successful compared to two other sampling rules, a and b values were not found to have any effect on the number of successful instances. This very finding also signifies the far more important role of representations than the components of associative learning in explaining observed behavior.

Briefly, our attempts to account for the empirical data through simulations using associative learning rules with different sampling rules and levels of representational manipulation revealed that these models best explain the data when they are compensated for the lack representational power. However, the rate of these successes was still very low. This suggests the necessity of assuming metric structure underlying the representation of temporal intervals for modeling decision-making in our task (see Montemayor & Balci, under review, for an in-depth discussion of this issue). This fact also suggests that as much as (if not more) the representations have critical role in explaining subject's data in these complex decision-making tasks, also their combinatorialness is a crucial condition to account for the empirical data. These results can be generalized to the validity of associative vs. information-processing accounts of decision-making processes.

Contribution to the study of interval timing:

Several widely used protocols have been developed in order to study interval timing in animals, including the peak procedure (e.g. Catania, 1970), the bisection method (e.g. Meck & Church, 1983), and the temporal generalization (Church & Gibbon, 1982), and the time-left protocol (Gibbon & Church, 1981). Most of these protocols have

further also been adapted to study interval timing in human subjects (Rakitin, Gibbon, Penney, Hinton, Malapani, & Meck, 1998 for peak procedure, Balci & Gallistel, 2005 for bisection method, Wearden & Bray, 2001 for temporal generalization, for review see Balci, Moore, & Brunner, 2007a). The switch paradigm used here is a new paradigm of studying interval timing in both mice (but see Fetterman & Killeen, 1995 and Platt & Davis, 1983 for a variant conducted with pigeons) and human subjects.

It has several advantages over the traditional protocols. For instance, the free-operant environment and the observed response type allows the experimenter to observe the real-time comparison and discrimination of durations (since the response could be observed at any point during the trial), which cannot be done in most temporal discrimination tasks. For instance, in bisection method, which is a widely used technique of studying temporal discrimination, the subjects' decision regarding the similarity of a probe to the referents takes place after the offset of the probe duration. Because of this procedural constraint along with the way the psychometric function is constructed, the subjects are presented with a limited sample of probe durations; each probe being presented for a number of times (e.g. 10 probe durations each presented for 10 times). Consequently, the subjects' responses are observed only for a limited number of probes determined by the experimenter.

On the other hand, with our paradigm, we introduced a new way of demonstrating and quantifying temporal discrimination (also see Fetterman et al. 1995). In this new way of demonstrating phenomenon, the cumulative distribution of the switch points –time when the subject leaves the short hole for the long hole- (the raw data) is the psychometric function of the duration discrimination, itself. Therefore, one does not need

to fit a function to a limited number of data points for gathering the psychometric function. This way there are no free parameters, which in turn increases the certainty about what is represented by the shape and location of the curve. In bisection method on the other hand, as stated above the psychometric function is constructed by fitting a function to a few number of points. This function has a number of free parameters and the best fitting function explains only a portion of the variance (R^2).

The real-time responding allowed by free-operant environment, further eliminated the potential effect of the nature of spacing between the probes on the psychometric function, which is an issue again in the case of bisection method. For instance, Allan (2002), Wearden and Ferrara (1995, 1996) and Wearden, Rogers, and Thomas (1997) reported that point of subjective equality was larger for linearly spaced probe durations compared to logarithmically spaced probes in humans. The disadvantages of bisection listed above largely constitute the disadvantages of the temporal generalization protocol. Unlike in the case of these protocols, in the switch paradigm, there are no probes determined by the experimenter.

The switch paradigm also has advantages over the peak procedure. Analyses of peak procedure requires the researcher to use a complex data-parsing algorithm to recognize the starts and stops of temporally controlled responses on a trial-by-trial basis. The resultant units of analysis might not coincide with an actual response. More importantly, different data-parsing algorithms are likely to detect different starts and stops (see Taylor, Horvitz, & Balsam, 2007). The output of these algorithms are further argued to be response-rate dependent (e.g. Odum, Lieving, and Schall, 2002). The switch paradigm does not require use of data-parsing algorithms in order to gather the behavioral

measures. It uses a clear-cut unit of analysis that is the latency with which the subject leaves the short-latency alternative for the long one, which does not have any response-rate dependency. Unlike other timing protocols, it does not involve any probe trials in order to gather timing measures. Finally, it allows the researchers to investigate the processing of different forms of information (e.g. probabilities, reward magnitudes, etc.). Procedural and analytical integration of these attributes is not as straight forward in the case of other interval timing protocols.

Briefly, switch paradigm provides an alternative way of conducting translational research on the process and memory of interval timing. Its advantages over traditional paradigms characterizes this paradigm as a strong alterative over traditional techniques of studying interval timing.

Lastly, our experiments have an empirical contribution to the study of interval timing. Our results have shown that mice can discriminate 1:2 ratio of temporal intervals. To our knowledge this is the first study that reports successful discrimination of 1:2 ratio temporal intervals by animals. Our earlier work has shown that when mice are trained originally with 1:2 ratio of temporal intervals, they cannot discriminate the referents (Balci et al., 2007b). In this work, however mice were initially trained with 1:3 ratio of temporal intervals and then introduced 1:2 ratio. This suggests that the maximum temporal precision reported in earlier studies is not due to representational/cognitive constraints of the subjects but procedural constraints such as the use/order of training protocols.

Contribution to behavioral/cognitive phenotyping:

In addition to its efficiency and simplicity compared to other interval timing paradigms, the protocols that we used in this experiment also proved to be work with mice. These tasks involve higher-level cognitive processes. Thus using these protocols one can conduct higher level cognitive phenotyping of genetically modified mice (see Gallistel, King, Gottlieb, Balci, Papachristos, Szalecki, & Carbone, 2007 for a similar approach).

Further, the switch task allows the pharmacology and genetics of impulsivity which characterizes the symptoms of certain psychopathologies such as schizophrenia. Impulsivity in our paradigm can be well quantified in terms of the proportion of the response distribution that is best modeled by a Weibull distribution (Balci & Freestone, in preparation).

Future Research:

In current research, we have used a wide range of probabilities (.1-.9) and reward/penalty magnitudes (0-50) in order to predict a large range of optimal points. On the other hand, fixed interval pairs that were to be discriminated in the human task was found to be easy for the subjects given their temporal uncertainty. For this very reason, the optimal switch latencies did not range as large as they would if the temporal intervals would be more difficult to discriminate. In order to ensure a larger range of optimal switch latencies, one can design harder temporal discrimination task. Derived from Weber's Law, this can be achieved by increasing the value of short and long temporal intervals while keeping 1 second temporal distance across them (e.g. 5 vs. 6 s. instead of 2 vs. 3 s.). This particular sort of manipulation would make temporal discrimination

harder while ensuring a wide enough temporal interval (long duration minus short duration) during which one can observe different target latencies.

Another aspect of the task that future research might aim at establishing is testing subjects under steeper expected gain functions. This could be achieved by increasing the relative magnitude of penalties compared to reward magnitudes. We have tried to achieve this by introducing a penalty of 50 points in certain sessions in which the reward magnitudes ranged between 1-9 points. These two manipulations, namely decreasing the discriminability and increasing the relative penalty magnitude for misses would certainly make the gain maximization problem harder than it was for our human subjects, which in turn would result in a more sensitive test of optimality.

Future research might focus on the adjustment of temporal decision-criterion as a function of changing probabilities of short and long intervals. The question would be if subjects would adjust their decision criterion as a function of the probability of the targets calculated over all trials or would they detect the change in the probabilities of short and long targets and shift their temporal strategy drastically.

An independent question that further research can investigate is if and how subjects acquire and integrate the temporal uncertainty that is introduced by the experimenter. This question can be addressed by introducing a level of noise to the presentation of short and long durations. The noise introduced by the experimenter can be modeled in different ways. It might be consistent with the model of temporal error (e.g. Gaussian) or it might be inconsistent with the model of temporal error (e.g. Exponential). Through these experiments, one can investigate if the acquisition of experimentally introduced temporal noise is facilitated by its consistency with the nature of temporal

representation in the mind. Faster acquisition of introduced temporal uncertainty as a function of its similarity to uncertainty in its representation would have important suggestions about the nature of information flow in mind: Is it belief-driven or not? Here what is meant with belief is the knowledge/priors about one's own representations.

Another line of research can also investigate the nature of integrating travel durations in between two different targets by experimentally manipulating the travel durations across different sessions, keeping other task parameters stable. There are multiple possibilities that might explain the integration of travel duration in decision-making process. The subject might be aiming for the actual switch duration or the actual switch duration plus the travel duration or any where in between these two points. This would help us clarify the nature of mapping between the switch point as a temporal decision criterion and the actual task with its possible outcomes.

Finally, further research can investigate the effect of reward and penalty magnitudes in mouse subjects. The reward magnitude can be manipulated by assigning different amounts of reinforcement (either directly or through increased duration of access to reward) for different options. The penalty magnitude can be manipulated by assigning different time-out periods (ITI) for different kinds of errors.

LITERATURE CITED

Allan, L. G. (2002). The location and interpretation of the bisection point. *Quarterly Journal of Experimental Psychology B, 55*, 43-60.

Balci, F. & Gallistel, C. R. (2006). Cross-domain transfer of quantitative discriminations: Is it all a matter of proportion? *Psychonomic Bulletin & Review, 13(4)*, 636-642.

Balci, F., Moore, H. and Brunner, D. (2007a, in press). Timing Deficits in Aging and Neuropathology. In Jennifer L. Bizon & Alisa G. Woods (Eds.), *Animal Models of Human Cognitive Aging*, Humana Press.

Balci, F., Papachristos, E. B., Gallistel, C.R., Brunner, D., Gibson, J., & Shumyatsky, G. P. (2007b, in press). Interval-Timing in the Genetically Modified Mouse: A Simple Paradigm. *Genes, Brain, and Behavior*.

Barron, G. & Erev, I. (2003). Small feedback-based decisions and their limited correspondence description based decisions. *Journal of Behavioral Decision Making*, *16*, 215-233.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433-436.

Catania, A.C. (1970) Reinforcement schedules and psychophysical judgments: A study of some temporal properties of behavior. In: Schoenfeld W.N. , editor. *The theory of reinforcement schedules.* New York: Appleton-Century-Crofts; pp. 1–42.

Cordes, S., Gelman, R., Gallistel, C.R., & Whalen, J. (2001). Variability signatures distinguish verbal from nonverbal counting for both large and small numbers. *Psychonomic Bulletin & Review, 8*, 698–707.

Fanselow, M. S. (1993). Associations and memories: The role of NMDA receptors and long-term potentiation. *Current Directions in Psychological Science, 2(5)*, 152-156.

Fetterman, J.G., & Killeen, P.R. (1995). Categorical scaling of time: Implications for clock-counter models. *Journal of Experimental Psychology: Animal Behavior Processes*, *21*, 43-63.

Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, *47*, 381-391.

Fitts, P. M. & Petersen, J. R. (1964). Information capacity of discrete motor responses. *J. Exp. Psychol., 67,* 103-112.

Gallistel, C. R. (2002). Frequency, contingency and the information processing theory of conditioning. In P. Sedelmeir & T. Betsch (Eds.), *Frequency processing and cognition*. Oxford, UK: Oxford University Press, pp. 153-171.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. *Psychological Review, 107*, 289-344.

Gallistel, C. R., & Gibbon, J. (2001). Computational Versus Associative Models of Simple Conditioning. *Current Directions in Psychological Science, 10*, 146-150.

Gallistel, C. R., King, A. P., Gottlieb, D., Balci, F., Papachristos, E. B., Szalecki, M., & Carbone, K. S. (2007). Is Matching Innate? *JEAB, 87*, 161-199.

Gibbon, J. (1977). Scalar expectancy theory and Weber's Law in animal timing. *Psychological Review*, *84*, 279–335.

Gibbon, J., & Church, R.M., 1981. Time left: linear versus logarithmic subjective time. *J. Exp. Anal. Behav. 7*, 87–107.

Gluck, M. A., & Thompson, R. F. (1987). Modeling the neural substrates of associative learning and memory: a computational approach. *Psychological Review, 94(2)*, 176-191.

Grossberg, S., & Schmajuk, S. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks, 2*, 79-102.

Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature, 394*, 780--4.

Hawkins, R. D., & Kandel, E. R. (1984). Is there a cell-biological alphabet for simple forms of learning? *Psych. Rev., 91*, 375-391.

Hertwig, R., Barron, G., Weber, E. U. & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534-539.

Hudson, T. E., Maloney, L. T. & Landy, M. S. (under review). Optimal movement timing with temporally asymmetric penalties and rewards.

Kacelnik A & Bateson M (1996) Risky theories – the effects of variance on foraging decisions. *American Zoologist, 36*, 402–434.

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

Kahneman, D. & Tversky, A. (1979). Prospect Theory: An analysis of decision under risk. *Econometrica*, *47*, 263-291.

Maloney, L. T. (2002). Statistical decision theory and biological vision. In Heyer, D. & Mausfeld, R. (Eds.), *Perception and the Physical World: Psychological and Philosophical Issues in Perception.* New York: Wiley, pp. 145-189.

Maloney, L. T., Trommershäuser, J. & Landy, M. S. (2006), Questions without words: A comparison between decision making under risk and movement planning under risk. In Gray, W. (Ed), *Integrated Models of Cognitive Systems*. New York , NY : Oxford University Press.

Mamassian, P., Landy, M. S., & Maloney, L. T. (2002), Bayesian modeling of visual perception. In Rao, R., Lewicki, M., & Olshausen, B. (Eds), *Probabilistic Models of the Brain; Perception and Neural Function*. Cambridge, MA: MIT Press, 13-36.

Marsh, B & Kacelnik, A (2002). Framing effects and risky decisions in starlings. *Proceedings of the National Academy of Sciences, 99 (5)*, 3352-3355.

Mechner, F. (1958). Probability relations within response sequences under ratio reinforcement. *Journal of the Experimental Analysis of Behavior, 1*, 109-122.

Montemayor, C. & Balci, F. (under review). Compositionality in language and arithmetic.

von Neumann, J., and Morgenstern, O. (1947). *Theory of games and economic behavior*, Princeton University Press.

Odum, A. L., Lieving, L. M., & Schaal, D. W. (2002). Effects of d-amphetamine in a temporal discrimination procedure: Selective changes in timing or rate dependency? *Journal of the Experimental Analysis of Behavior, 78*, 195–214

Pavlov, I. P. (1927). *Conditioned Reflexes* translated by G V Anrep. Oxford University Press, London, England

Pavlov, I. P. (1928). *Lectures on Conditioned Reflexes. Twenty-Five Years of Objective Study of the Higher Nervous Activity (Behaviour) of Animals.* Trans. by W. H. Gantt & G. Volborth. Introduction by Walter B. Cannon. New York: International Publishers.

Papachristos, E. B., & Gallistel, C.R. (2006) Autoshaped Head Poking in the Mouse: A Quantitative Analysis of the Learning Curve. *Journal of the Experimental Analysis of Behavior,85*, pp. 293-308.

Platt, J. R., & Davis, E. R. (1983). Bisection of temporal intervals by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 9,* 160-170.

Rakitin, B. C., Gibbon, J., Penney, T. B., Hinton, S. C., Malapani, C. & Meck, W. H. (1998). Scalar Expectancy: Theory and peak interval timing in humans. *Journal of Experimental Psychology: Animal Behavior Processes, 24*, 1-19.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In Black, A. H., & Prokasy, W. F. (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.

Samuels, R., Stich, S. & Bishop, M. (2002). Ending the Rationality Wars: How to Make Disputes About Human Rationality Disappear. In Renée Elio, (Ed), *Common Sense, Reasoning, and Rationality.* Oxford: Oxford University Press, pp. 236-268.

Shafir, S. (2000). Risk-sensitive foraging: The effect of relative variability. *Oikos, 88,* 663–669.

Smyrnis, N., Evdokimidis, I., Constantinidis, T. S., & Kastrinakis, G. (2000). Speed-accuracy trade-off in the performance of pointing movements in different directions in two-dimensional space. *Experimental Brain Research, 134*, 21–31.

Taylor, K.M., Horvitz, J.C., Balsam, P.D. (2007). Amphetamine affects the start of responding in the peak interval timing task. *Behav Processes, 74*, 168–175.

Thaler, R. H., Tversky, A. Kahneman, D. & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *Quarterly Journal of Economics*, *112*, 647-661.

Thorndike, E.L. (1911). *Animal intelligence: Experimental studies*. New York: Macmillan.

Trommershäuser, J., Gepshtein, S., Maloney, L. T., Landy, M. S. & Banks, M. S. (2005). Optimal compensation for changes in task-relevant movement variability. *Journal of Neuroscience*, *25*, 7169-7178.

Trommershäuser, J., Landy, M. S. & Maloney, L. T. (2006), Humans rapidly estimate expected gain in movement planning. *Psychological Science, 17*, 981-988.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision, 16*, 255-275.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A, 20*, 1419-1432.

Trommershäuser, J., Mattis, J., Landy, M. S. & Maloney, L. T. (2006), Limits to human movement planning with delayed and unpredictable onset of needed information. *Experimental Brain Research, 175*, 276-284.

Tversky, A. & Fox, C. R. (1995), Weighing Risk and Uncertainty. *Psychological Review, 102*, 269-283.

Tversky, A. & Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Risk and Uncertainty*, *5*, 297-323.

Wagner, A.R., & Rescorla, R. A. (1972) . Inhibition in Pavlovian conditioning: Application of a theory. In: Boakes RA, Halliday MS., editors. *Inhibition and learning.* London: Academic Press; pp. 301–336.

Wearden, J.H., & Bray, S. (2001). Scalar timing without reference memory: Episodic temporal generalization and bisection in humans. *Quarterly Journal of Experimental Psychology, 54B*, 289-310.

Wearden, J. H., & Ferrara, A. (1995). Stimulus spacing effects in temporal bisection by humans. *Quarterly Journal of Experimental Psychology, 48B*, 289–310.

Wearden, J. H., & Ferrara, A. (1996). Stimulus range effects in temporal bisection by humans. *Quarterly Journal of Experimental Psychology, 49B*, 24–44.

Wearden, J. H., Rogers, P., & Thomas, R. (1997). Temporal bisection in humans with longer stimulus durations. *Quarterly Journal of Experimental Psychology, 50B*, 79–94.

Weber, E. U., Shafir S., & Blais, A.-R. (2004). Predicting risk-sensitivity in humans and Lower animals: Risk as variance or coefficient of variation. *Psychological Review, 111*, 430-445.

Whalen, J., Gallistel, C.R. & Gelman, R. (1999). Non-verbal counting in. humans: The psychophysics of number representation. *Psychol. Sci., 10*, 130-137.

Wu, S.-W., Trommershäuser, J., Maloney, L. T. & Landy, M. S. (2006), Limits to human movement planning in tasks with asymmetric gain landscapes. *Journal of Vision, 6*, 53-63.

CURRICULUM VITA

Fuat Balci

Education

2002 B.A. in Psychology
Koc University, Istanbul, Turkey

2004 M.S. in Cognitive Psychology
Rutgers, The State University of New Jersey, New Brunswick, NJ

Employment

2004-2006 Teaching Assistant
Rutgers, The State University of New Jersey, New Brunswick, NJ

Publications

Balci, F. & Gallistel, C. R. (2006). Cross-domain transfer of quantitative discriminations: Is it all a matter of proportion? *Psychonomic Bulletin & Review, 13(4)*, 636-642.

Balci, F., Moore, H. and Brunner, D. (2007a, in press). Timing Deficits in Aging and Neuropathology. In Jennifer L. Bizon & Alisa G. Woods (Eds.), *Animal Models of Human Cognitive Aging*, Humana Press: Totowa, NJ.

Balci, F., Papachristos, E. B., Gallistel, C.R., Brunner, D., Gibson, J., & Shumyatsky, G. P. (2007b, in press). Interval-Timing in the Genetically Modified Mouse: A Simple Paradigm. *Genes, Brain, and Behavior*.

Gallistel, C. R., King, A. P., Gottlieb, D., Balci, F., Papachristos, E. B., Szalecki, M., & Carbone, K. S. (2007). Is Matching Innate? *JEAB, 87*, 161-199.