Three Experimental Studies of Reward and Decision Making

Thesis by

Min Jeong Kang

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2010

(Defended August 20, 2009)

To the memory of my father

Acknowledgments

I would like to thank my advisor and mentor Colin Camerer. I am very much indebted to him as I would not have been able to complete this long and demanding journey towards the PhD without his guidance.

He has been a ceaseless source of insights, ideas, and energy (and humor, of course). I always found him intellectually stimulating and much enjoyed his jokes in and outside the lab. Also, I admire his "mastery" of the English language and flamboyant writing style—I wish I could write like him! Despite his busy schedule, he is always easily accessible whenever I need his input academically or personally.

It is very easy for graduate students to drift away without constant reminders, checking in, and prodding. Through all these years, he has been diligent about paying attention to my progress and keeping me on track. I always felt like he cared about my work. I very much appreciate his devotion of his energy and time to advise and train me.

Of course, no thesis would be complete without a thesis committee. In this regard, I am deeply grateful to Antonio Rangel, Ralph Adolphs, and Peter Bossaerts for taking time out of their busy schedules to generously serve on my committee and for their valuable comments on my dissertation. Further thanks are due to Antonio for his contributions and guidance in the second study, on hypothetical and real choices.

Further, I would like to thank all the professors in HSS that taught me at various points in the past five years. They set a foundation for understanding and conducting research in social sciences, which was necessary and crucial for this thesis. In this regard, I gratefully acknowledge my undergraduate advisor, Youngse Kim at Yonsei University. He inspired and encouraged me to pursue a PhD in the U.S. and I am very thankful for what I learned from him in college and for all of his support. Without his guidance and encouragement, I would not have come here and this thesis would not have existed.

Throughout my graduate research, I had the good fortune of working with many skilled collaborators, to whom much thanks is due. My second chapter was coauthored with Ming Hsu, Ian Krajbich, Joseph Wang, Samuel McClure, George Lowenstein and Colin Camerer. My third chapter on hypothetical and real decisions was a collaboration with Mickael Camus and Colin Camerer. My fourth chapter on optimal timing decisions is an ongoing effort with Colin. Furthermore, Meghana Bhatt, Carmina Caringal, Noah Myung and Cary Frydman deserve acknowledgement for their help. Colin financially supported all of the research; his financial support is greatly appreciated. I am also very grateful to Mike Lipper for his generous financial support; my final year at Caltech was funded from the Lipper Family Fellowship in Neuroeconomics.

I would also like to thank the many other people who helped me over the years to complete this thesis. I thank Ming Hsu for being such a wonderful coach to get me started on fMRI analysis. I also thank Hilke Plassmann, Todd Hare, Cendri Hutcherson and Camerer and Rangel lab members for useful comments and help. Especially, I appreciate Hilke's warm friendship and willingness to help anytime.

Many thanks go to HSS, SSEL, and CASSEL staff for administrative assistance to run experiments. I especially appreciate help from Tiffany Kim and Walter Yuan. I should not forget to mention Laurel Auchampaugh; she has been of great help all the time and I am very grateful for that.

Personally, I would like to give warm thanks to friends at the Caltech dance clubs and CKGSA for all the parties, dancing and good times. I thank Oleg Kogan and David Young for fun times dancing together and listening to my concerns. I will not forget the countless chit-chats with my classmates, Sera Linardi and Noah Myung. They have been great friends and counselors.

I would especially like to thank Daniel Holland. He deserves a big hug for his constant support. He has been of great help, both academically and personally. I will not forget his patience and support throughout this difficult time. In particular, I am grateful to him for his love, his mastery of the English language (though slightly less sophisticated than Colin's for now), and keeping my clothes clean when I was too busy with this thesis to do it myself.

I would like to reserve my last, but biggest, thanks for my family. I appreciate all their support and love from afar. I dedicate this thesis to my late father.

Abstract

This thesis investigates reward-driven decision making using a combination of techniques such as functional neuroimaging, behavioral experiments, and pupilometry. Different aspects of reward and decision making were examined in three different studies, including the nature of curiosity, similarities and differences in hypothetical and real decisions, and optimal-timing decisions. In the study of curiosity, my colleagues and I conceptualized information as a rewarding object and curiosity as a form of reward anticipation of the rewarding information. We explored this hypothesis using a combination of functional neuroimaging, pupillometry, behavioral experiments, and memory-retrieval experiments. In the study of hypothetical and real decisions, the neural differences and similarities underlying these types of decisions were explored. We discuss potential implications of the findings on scientific practices and suggest the possibility of a new use of fMRI to improve the prediction of real choices based on hypothetical choice data. In the third study, we explore how people make timing decisions when motivation to delay trades off against the motivation to take an action immediately. We experimentally test current theory and examine how strategic decisions become sophisticated over time. We further hypothesize about psychological processes that could guide decision making under uncertainty and time pressure.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis contains three independent experimental studies concerned with the general theme of human reward processing and decision making. In recent years, the development of brain imaging technology has enabled researchers in various disciplines to expand our understanding of decision making and valuation. This thesis continues the trend using a combination of methods such as functional imaging, behavioral experiments, and pupilometry. The studies herein contribute to an understanding of reward-driven decision making: (1) by extending a dimension of reward to include more abstract objects of higher complexity (information, consumer products); (2) by comparing different modalities of decision making (hypothetical and real) on a neural level that had not been explored previously; (3) testing decision making when different motives conflict in a complex social environment.

In the following chapter is an experimental study of curiosity itself, which has long intrigued philosophers and social scientists and has been described as an important motivator of learning. Despite its importance, the psychological and neural substrates of curiosity remain poorly understood. In the study, information has been conceptualized as a rewarding object and curiosity as a form of reward anticipation of the rewarding information. This hypothesis is explored using a combination of functional neuroimaging, pupillometry, behavioral experiments, and memory-retrieval experiments.

The second study examines the neural differences and similarities underlying hypothetical and real consumer decision making. Many studies in social sciences rely on hypothetical choice data to infer actual choices, which are often difficult to obtain. However, studies in economics have found that hypothetical valuations of goods reflect a systematic bias compared to actual preferences. Based on this behavioral difference, we further ask if there are different neural circuitries engaged in hypothetical and real decisions, and different neural computations are carried out in the brain for hypothetical and for real decisions. We discuss potential implications of the findings on practices in imaging neuroscience and psychology and further suggest the possibility of a new use of fMRI to neurally predict real choices based on hypothetical choices.

In the last chapter, optimal timing decisions are investigated when motivation to delay trades off against motivation to take an action immediately. Current theory is experimentally tested in two different environments, and we examine how strategic decisions become sophisticated over time. We further hypothesize psychological processes that could guide decision making under uncertainty and time pressure.

Chapter 2

Epistemic Curiosity Activates Reward Circuitry and Enhances Memory[1]

"The important thing is not to stop questioning. Curiosity has its own reason for existing. One cannot help but be in awe when he contemplates the mysteries of eternity, of life, of the marvelous structure of reality. It is enough if one tries merely to comprehend a little of his mystery every day. Never lose a holy curiosity."

−Albert Einstein

2.1 Introduction

Curiosity is the complex feeling and cognition accompanying the desire to learn what is unknown. Much as curiosity motivates other animal species to explore their environment and find sources of food, it plays a critical role in motivating learning and discovery in humans, especially by creative professionals, and is necessary for increasing the world's store of knowledge. Apples had fallen from a tree before Newton and have done so after him, but it was his curiosity that inspired him to formulate the theory of universal gravitation out of such a seemingly inconsequential phenomenon. Einstein, once said, "I have no special talents. I am only passionately curious (Hoffmann 1972)."

Although curiosity has made great contributions to advance knowledge and technologies, curiosity has a dangerous side in its association with exploratory behaviors having harmful consequences. An ancient example is the mythical Pandora, who opened a box that unleashed misfortunes on the world. Curiosity is often accused of causing teenagers'

[1] This chapter has been previously published in Kang et al. (2009) and the definitive version is available at www.blackwell-synergy.com.

thrill seeking, drug and alcohol use, and sexual experimentation. Moreover, pathological curiosity leads to behavioral disorders such as perverted voyeurism. Technology such as the Internet augments both good and bad effects of curiosity, by putting both enormous amounts of information and potentially dangerous social encounters a mouse click away.

Despite its importance, the psychological and neural underpinnings of human curiosity remain poorly understood. Philosophers and psychologists alike have described curiosity as an appetite for knowledge, a drive like hunger and thirst (Loewenstein 1994), the hunger pang of an "'info-vore" (Biederman & Vessel 2006), and "the wick in the candle of learning" (William Arthur Ward). In reinforcement learning, a *novelty bonus* is used to motivate the choice of unexplored strategies (Kakade & Dayan 2002); curiosity can be thought of as the psychological manifestation of such a novelty bonus.

Our research is guided, in part, by the theory that curiosity arises from an incongruity or "information gap"—a discrepancy between what one knows and what one wants to know (Loewenstein 1994). The theory assumes that the aspired level of knowledge increases sharply with a small increase in knowledge, so that the information gap grows with initial learning. When one is sufficiently knowledgeable, however, the gap shrinks and curiosity falls. If curiosity is like a hunger for knowledge, then a small "priming dose" of information increases the hunger, and the decrease in curiosity from knowing a lot is like being satiated by information.

In the information-gap theory, the object of curiosity is an unconditioned rewarding stimulus: unknown information that is anticipated to be rewarding. Humans (and other species, such as cats and monkeys) will expend resources to find out information they are

curious about, much as rats will work for a food reward (Loewenstein 1994). Based on this observation, we hypothesized that the striatum would be linked to curiosity, as a growing body of evidence suggests that activity in the human striatum is correlated with the level of reward signals (Hare et al. 2008; Knutson et al. 2000; McClure et al. 2004; O'Doherty 2004).

Guided by the ideas mentioned above, we explored the neural correlates of curiosity in one study and tested the hypotheses derived from its findings in two separate studies. In all studies, subjects were presented with series of trivia questions chosen to create a mixture of high and low "epistemic" curiosity (figure 2.1).[2] Subjects were instructed to read each question, guess the answer, and rate both their curiosity and how confident they were that they knew the answer (*P*). They were then shown the question again along with the correct answer (figure 2.2).

In the first experiment subjects read the questions during fMRI. In the second experiment they performed the same task without scanning, and their memory for answers was tested in a follow-up session 1−2 weeks later. In the third experiment, we behaviorally tested whether curiosity is indeed a form of reward anticipation.

2.2 Study 1: Neural Correlates of Curiosity

Experimental Design and Methods

Participants and Stimuli

[2] Epistemic curiosity refers to a desire to acquire knowledge and applies mainly to humans (Loewenstein, 1994).

Nineteen Caltech students were scanned (average age: 21.7 ± 3.5 years; 14 males, 5 females; 18 right-handed)[3]. They earned $20 for participation. Informed consent was obtained for all three experiments, using a consent form approved by the Internal Review Board at Caltech.

The stimuli consisted of 40 trivia-type questions, regarding various topics, chosen to measure curiosity about semantic knowledge. The questions were pre-tested to reliably evoke a range of curiosity levels across individuals (for sample questions, see figure 2.1; for the complete list of questions, see appendix 2.7.1).

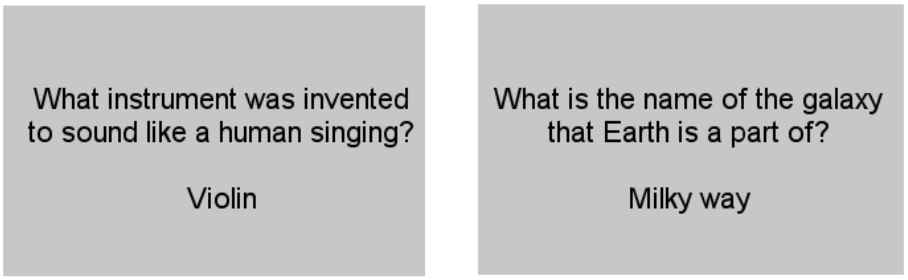

Figure 2.1 Sample questions with relatively high (*left*: average 5.72 out of 7) and low (*right*: 2.28 out of 7) curiosity ratings.

Experimental Procedure

Written instructions were provided outside the scanner (appendix 2.7.2). Once subjects understood the experimental procedure, they were put in the scanner for the task. Each experimental session consisted of 4 runs, with each run containing a set of 10 questions, which were randomly presented within each run. There was a one-minute break between runs, due to physical restrictions on the scanner. Each trial consisted of 5 epochs: (1) the

[3] Initially, 20 subjects participated, but one subject was discarded because he received instructions in the scanner instead of outside, and showed odd behavioral data—his mean curiosity rating was "7," the maximum allowed, not creating enough variation in curiosity across questions.

first presentation of a question, (2) selection of curiosity rating, (3) confidence rating, (4) second presentation of the question, and (5) display of the answer. The curiosity and confidence rating epochs were self-paced; the durations of the three presentation/display epochs were independently randomized within a set range, which is explained below and shown in figure 2.2. The task was presented to subjects through MRI compatible goggles. Subjects were given 12 to 15 seconds to read the question, followed by a fixation screen displayed for 4 to 6 seconds, and they were instructed to silently guess the answer while reading a question. The subjects were then asked to indicate the extent of their curiosity about the correct answer as well as the level of confidence they had in their guess. The curiosity and confidence rating epochs were self-paced; the subjects moved on to the next screen by making their selection with an MRI-compatible button box. These rating screens were also followed by fixation-cross screens. After the confidence level was entered, the question was presented again for 3 to 5 seconds, and then the answer was revealed, below the question, for 4 to 6 seconds. To keep the motor requirement of the task minimal, the presentation of questions and answers was designed to not be self-paced. Each complete cycle, or trial, took about a minute, with the entire experiment lasting approximately 45 minutes. Following the scanner part, the subjects were asked to self-report their initial guesses at the correct answers to the questions and to fill out a questionnaire. This self-reporting part was conducted outside of the scanner, due to the difficulty of collecting verbal, typed, or written responses while the subject is inside the scanner tube.

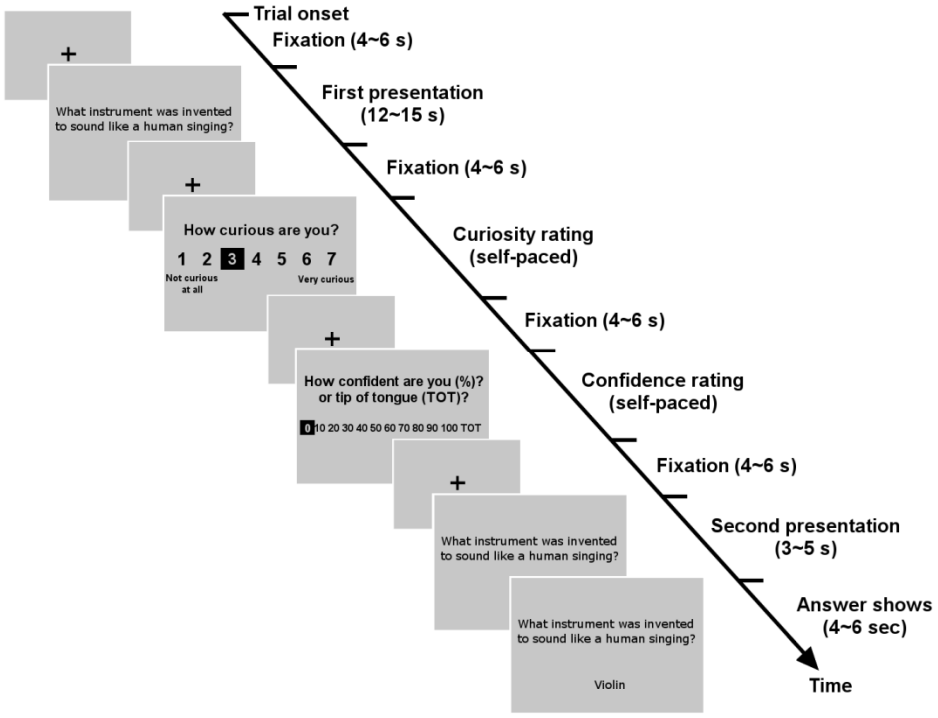

Figure 2.2 Experimental timeline.

fMRI Data Acquisition and Preprocessing

Data were acquired using a 3-T Siemens (Erlangen, Germany) Trio scanner at Caltech. A set of high-resolution ($0.5 \times 0.5 \times 1.0$ mm^3) T1-weighted anatomical images was first acquired to enable localization of functional images. Whole-brain T2*-weighted echoplanar images with blood-oxygenation-level-dependent (BOLD) contrast were acquired in 32 axial slices (64×64 voxels; 3 mm thickness and 3 mm in-plane resolution) at a repetition time of 2000 ms and echo time of 30 ms. The scan sequences were axial slices approximately parallel to the AC-PC axis. The fMRI data were preprocessed using SPM2 (Welcome Department of Imaging Neuroscience, Institute of Neurology, London, UK). Functional scans were first corrected for slice timing via linear interpolation. Motion correction of images was performed using a 6-parameter affine

transformation followed by nonlinear warping using basis functions (Ashburner & Friston 1999). Finally, images were smoothed with a Gaussian kernel of 8 mm FWHM.

fMRI Data Analysis

The data analysis was conducted using the random effects general linear models (GLM) for event-related designs in SPM2. In all imaging data analyses, the regressors were modeled using box-car functions, convolved with a canonical hemodynamic response function. All four runs were concatenated and treated as one run and to control for the variance between sessions from concatenation, session dummy variables were included as separate regressors. Images were adjusted for both global intensity, using proportional scaling; and for low-frequency physiological drifts, using a high-pass filter of 128 seconds. Autocorrelation of the hemodynamic responses was modeled as an AR(1) process. Parameters were estimates from a ReML procedure. We initially performed 19 separate single-subject analyses: voxel-by-voxel statistical parametric maps of the t-statistic for each contrast of interest were defined for each subject. These contrast maps were then integrated to derive contrast images for second-level group T-tests and ANOVA (Ashburner et al. 1997; Friston et al. 1995; Genovese et al. 2002).

1. Curiosity Median Split Analysis. Each subject's trials were split into two conditions (high or low) according to where they fell relative to that individual's median curiosity level. Then all five epochs in each trial (first presentation, curiosity rating, confidence rating, second presentation, and answer display) were classified as being in the high- or low-curiosity condition according to the condition to which the whole trial had been assigned. Thus, there were two curiosity conditions for each epoch, resulting in a total of

10 separate regressors of interest. Each regressor was time locked to stimulus presentation. A GLM including these 10 regressors plus regressors of no interest was estimated. The 10 regressors of interest were modeled using box-car functions with the length of each epoch (e.g., the presentation time for the first answer) as the corresponding box-car duration. We then calculated contrasts to compare the effects of high versus low curiosity.

2. Curiosity interaction analysis. We also examined whether the brain activations identified in the median-split analysis increased linearly with curiosity level, rather than being associated with two levels (high or low). We estimated a GLM in which normalized curiosity was a parametric modulator for each of the five epochs.

3. Residual curiosity analysis. This analysis was performed to investigate the effect of curiosity that is dissociated from confidence level, P, and uncertainty, $P(1-P)$.[4] To do this, we first regressed curiosity on P and $P(1-P)$ [5] (with a constant) and then took the residuals from this regression to construct a new variable, called the "residual curiosity." Further, to study the interaction between residual curiosity and correctness of guess, we divided the answer-display epochs into correct- and incorrect-guess conditions. This procedure resulted in a total of 6 conditions of interest: first presentation, curiosity rating, confidence rating, second presentation, answer display preceded by a correct guess, and

[4] Uncertainty, $P(1-P)$, measures a subject's uncertainty about a guess. A guess is a random variable (more specifically, a Bernoulli random variable) with two outcomes, correct or incorrect; the two outcomes have probabilities of P (confidence level) and $(1-P)$, respectively. The uncertainty associated with the random variable (or uncertainty about which outcome will occur) is measured by the entropy, which in this case is a monotonic function of the variance $P(1-P)$. $P(1-P)$ is a quadratic function with a maximum at $P = 0.5$ and minima at $P = 0$ and 1. For example, suppose that you are 100% (or 0%) sure about your guess. In this case, your confidence level, P, is 1 (or 0). Your uncertainty about the outcome will be minimal. In contrast, if your confidence level about your guess is .5, then you are most uncertain about which outcome will occur because the two outcomes are equally likely.

[5] Raw curiosity level was normalized at an individual level. See **behavioral results**.

answer display preceded by an incorrect guess. We then estimated a GLM in which residual curiosity, $P(1 - P)$, and P were parametric modulators for each of these six conditions.

Behavioral Results

Curiosity is Correlated with Uncertainty P(1 − P) and Peaks around Confidence P = 0.5

Prior to all the analysis herein, raw curiosity ratings, ranging in value from 1 to 7, were normalized for each individual (subtracting each individual's mean and dividing by each individual's standard deviation). The information-gap theory predicts that curiosity should increase with statistical uncertainty, $P(1 - P)$, since people who know very little have not had their curiosity piqued, and those who know a lot are satiated. Reported curiosity was indeed an inverted U-shaped function of P, reaching its maximum when P was around .50 (figure 2.3). Curiosity correlated positively with $P(1 - P)$, $r = .44$, $p < 0.0005$. Most subjects showed this relation; estimated peak curiosity was at values of P between .45 and .55 in three-quarters of the subjects (table 2.1).

Further, at an individual level, most subjects show a quadratic relationship between curiosity and confidence level. Table 2.1 includes results from regressing fitted curiosity against a constant, P, and $P(1 - P)$ and the confidence level that has the maximum curiosity (based on the regression estimates). Note that 12 of 19 subjects have estimated probabilities with maximum curiosity between .40 and .60.

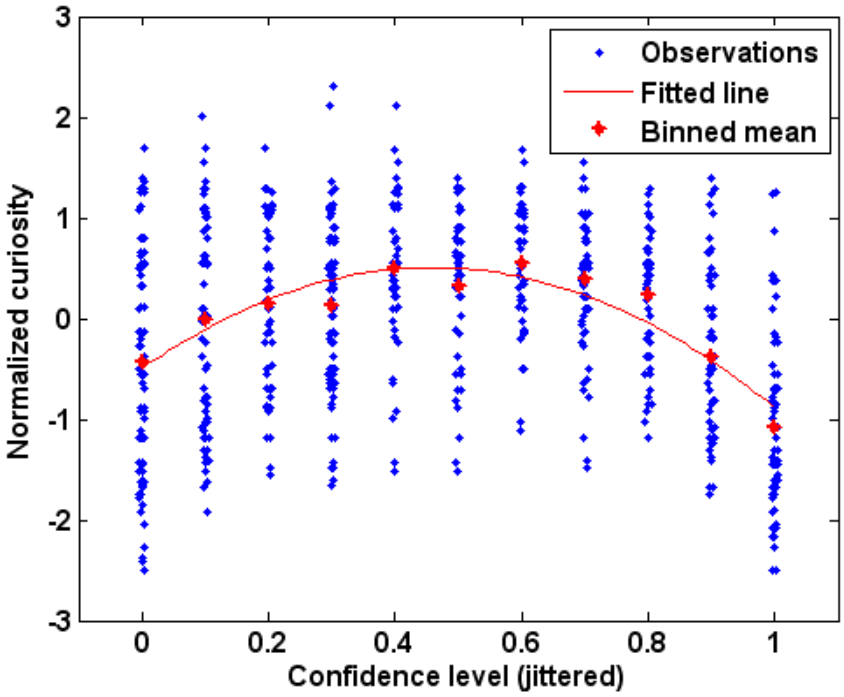

Figure 2.3 Distribution of curiosity against confidence: The confidence scale ranged from 0% to 100% but was rescaled to range from 0 to 1. There was also a tip of the tongue (TOT) response option (Maril et al. 2001) but there were too few of these responses to analyze so they were excluded. All confidence ratings are jittered by adding a uniform random number over the interval [-0.005, 0.005] to convey data density. Red stars indicate mean curiosity at each confidence level. The solid curve line is the regression line of curiosity against confidence P and uncertainty $P(1 - P)$. The estimated regression is Curiosity = -0.49 – 0.39P + 4.77$P(1 - P)$ + Residual curiosity.

Table 2.1 Individual regression analysis of normalized curiosity level against *P* and *P(1 − P)*

Subject ID	Coefficients			R^2	*P with maximum curiosity*
	Intercept	*P*	*P(1 − P)*		
1	-0.06	-1.51	4.00	0.19†	0.32
2	-0.4	-0.32	4.16	0.13	0.47
3	-0.17	-0.39	2.47	0.08	0.43
4	0.26	-1.38	3.18	0.33**	0.29
5	-0.88	0.33	4.77	0.19*	0.54
6	-0.49	-0.41	5.24	0.30*	0.47
7	-0.21	-1.20	6.40	0.56**	0.42
8	0.38	-1.68	2.51	0.49**	0.17
9	-1.73	0.74	8.25	0.44**	0.55
10	-0.61	0.86	1.71	0.12	0.76
11	-0.77	0.87	2.07	0.11	0.72
12	-1.58	-0.02	9.86	0.53**	0.51
13	0.1	-1.60	6.11	0.66**	0.38
14	-1.14	-0.12	7.84	0.69**	0.50
15	-0.43	-0.34	4.16	0.13	0.47
16	-0.48	0.07	3.68	0.13	0.52
17	-2.16	0.77	9.89	0.68**	0.55
18	-0.92	1.07	2.40	0.12	0.73
19	-0.85	-0.67	8.37	0.54**	0.47

Note: Confidence ratings (0% ~ 100%) are rescaled to range from 0 to 1.
† significant at $p<0.05$
* $p<0.01$
** $p<0.001$

fMRI Results

In this section, we first focus on brain activity when questions were initially presented and then discuss brain activity when answers were presented. Results are reported for brain areas with significant activity at an uncorrected *p*-value of 0.001 and cluster size $k \geq 5$ unless noted otherwise.

Curiosity is Correlated with Activity in Reward Regions

The first question-presentation epoch was associated with the high- or low- curiosity condition according to the individualized median curiosity level. We created a contrast that identified regions whose activity was greater in response to high curiosity than in response to low curiosity. Significantly activated regions include the left caudate; bilateral prefrontal cortex (PFC), including inferior frontal gyrus (IFG) (figure 2.4a); and parahippocampal gyri (PHG) (table 2.2). Activations in the putamen (*x, y, z* = 21, 9, 9), $t(18) = 3.15$, and the globus pallidus (*x, y, z* = 12, -6, 0), $t(18) = 3.94$, were significant at a more lenient *p*-value of 0.005 (uncorrected), but no activation was found in the nucleus accumbens. The area of significant activation in the left caudate overlaps with areas of significant activity identified by the models using subject-normalized linear curiosity and residual curiosity as the regressors (figure 2.4b). This finding is consistent with the view of curiosity as anticipation of rewarding information.

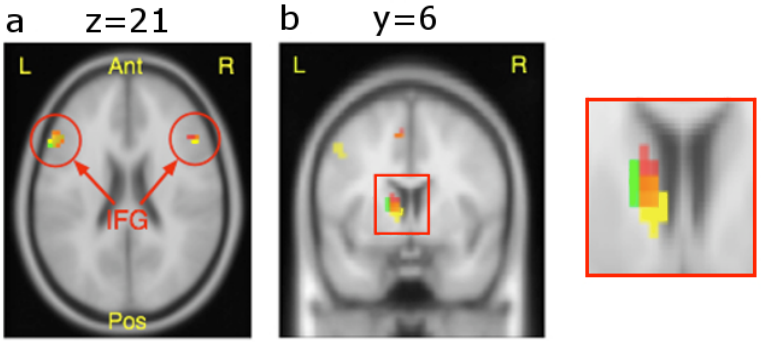

Figure 2.4 Differential brain activity in high and low curiosity trials during first question presentation ($p<.001$ uncorrected, extent threshold ≥ 5). Overlapped regions of activation in bilateral PFC (a) and the caudate (b) by curiosity from three different dependent variable models: High > Low median-split curiosity (red); linear in curiosity (yellow); linear in residual curiosity from the figure 2.3 regression (green). Note that we did not find activation in OFC regions as one might expect, but this is not surprising because our sequence was not optimized to detect OFC activations. (*right*) Overlapped close-up of caudate activations from the same three different models in (b).

Table 2.2 Brain regions associated with high curiosity, relative to low curiosity, during the first question presentation

Region	L/R*	Coordinates			Spatial Extent (voxels)	t
		x	y	z		
Caudate	L	-9	3	3	10	4.04
IFG/BA45	L	-54	24	21	112	5.71
	R	48	24	21	5	4.01
PHG	L	-33	-39	-12	21	4.04
	R	36	-30	-18	5	4.46
Medial Frontal Gyrus	L	-12	36	48	26	4.49
MFG, Pre-motor Cortex	L	-27	15	57	70	5.71
Lingual gyrus	R	18	-63	-3	11	4.57
Cerebellum	R	36	-69	-36	34	4.67

All locations are reported in MNI coordinates. * Laterality

Curiosity is Correlated with Memory-Related Regions when Incorrectly Guessed Answers are Revealed

When the answers were revealed, activations in areas linked to learning and memory were much stronger if the subject's prior guess had been incorrect, rather than correct. Areas differentially activated when subjects guessed incorrectly included the bilateral putamen and left IFG (Broca's area) (figure 2.5).

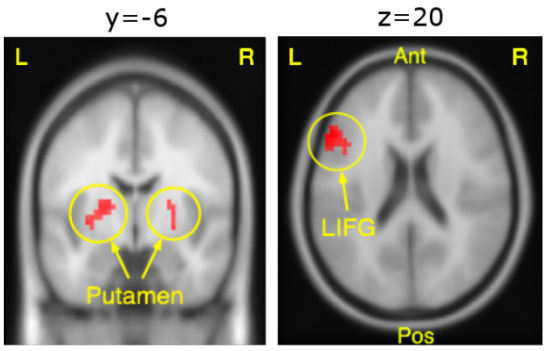

Figure 2.5 Regions which are more active in response to answers after making incorrect guesses than to correct guesses: Bilateral putamen (right: *x, y, z* = -24,-9,6, *t*(18)=4.63; left: *x, y, z* = 24,-9,12, *t*(18)=4.77), left IFG (Broca's area, BA 44/45).

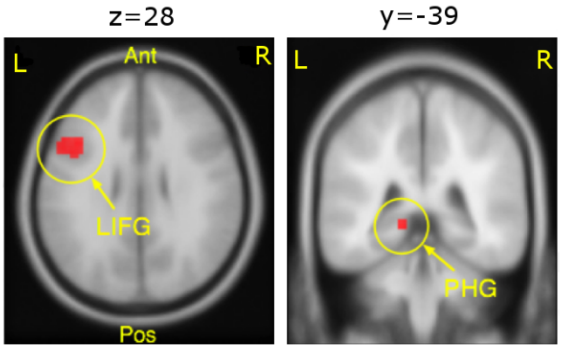

Figure 2.6 Regions having activity correlated with novel information (wrong answers × curiosity). Left IFG (*left*), Left PHG (*right*).

Table 2.3 Brain regions linearly correlated with residual curiosity during answer display when subjects initially guessed incorrectly

Region		MNI			Spatial Extent	T statistic
	L/R	x	y	z	(voxels)	
Parahippocampal Gyrus	L	-24	-27	-6	19	4.69
Inferior Frontal Gyrus	L	-54	9	24	76	4.48
DLPFC/BA 9	L	-51	15	30		4.23
Inferior Frontal Gyrus	L	-45	30	3	88	5.98
Lingual Gyrus	L	-12	-63	-6	40	5.31
Superior Temporal Gyrus	L	-60	-57	12	5	4.6
Superior Frontal Gyrus	L	-21	48	12	9	4.36
Medial Frontal Gyrus	L	-6	15	51	8	4.07
Cerebellum	R	9	-72	-30	125	6.12

Furthermore, curiosity level modulated the activations during the answer display. After an incorrect guess, left PHG and left IFG activations during the answer display were positively correlated with residual curiosity (figure 2.6, table 2.3). Bilateral midbrain regions (left: *x, y, z* = -12,-24,-6, $t(18) = 3.37$; right: *x, y, z* = 12,-21,-18, $t(18) = 3.97$) were also activated at $p<0.005$ (uncorrected), as was the hippocampus (*x, y, z* = -27,-33,-6, $t(18)=3.2$) (not shown). The identified area in left IFG was dorsal to areas identified in the analyses of the question epoch (figure 2.4a) and was part of Broca's area, which is

important for language comprehension (Bookheimer 2002). When subjects guessed correctly, residual curiosity did not correlate with any of the identified regions.

Because memory-related regions were differentially activated in response to answers presented after incorrect guesses, and the activity was modulated by curiosity, we hypothesized that curiosity would be associated with "memory enhancement" for new information (in this paradigm, a correct answer is new information if it follows an incorrect guess). That is, we hypothesized that after guessing incorrectly, people would be more likely to remember the answer to a question if they were curious to know it.

The findings from the fMRI study suggested that curiosity is anticipation of rewarding information and that it may enhance learning of new information. We tested these hypotheses in separate experiments. We first describe the experiment that tested the memory-enhancement hypothesis and then report the experiment that tested the reward-anticipation hypothesis.

2.3 Study 2: Curiosity Enhances Memory

Experimental Design and Methods

Participants and Task

Sixteen Caltech students (11 males, 5 females) participated in the first task. Fifteen subjects returned for a follow-up memory test. Two returned within 1 week and exhibited too high recall rate (over 90%) and one returned in 3 weeks—the data from these three subjects were not used for the memory analysis reported below.

The experimental procedure was very similar to that for the functional imaging study, having a few modifications: (a) the order of questions was randomized across all 40 of the question trials, (b) fixation screens were removed (these were necessary in fMRI to allow the BOLD signal to dissipate between decision epochs, but are unnecessary in a behavioral study), (c) the first presentation screen time was fixed at 10 seconds of exposure (rather than 12~15 seconds in the fMRI study), (d) a 'count-down' of five seconds was presented onscreen immediately before the answer was displayed in order to attract subjects' attention and precisely quantify their pre-answer anticipation in terms of pupil size, (e) since there were few tip-of-the-tongue responses in the fMRI study, the TOT option was removed from the confidence scale, (f) between the curiosity and the confidence rating screens, a "give answer" screen was presented and subjects were asked to state their guess out loud so that an experimenter could record it.

Measures

Behavioral measures in this study included those of Study 1 (curiosity and confidence levels, guesses) with the change that subjects' guesses were recorded during the task (as a check on post-scanner over-reporting of correct guesses in Study 1, which was minor; see table 2.4). This study included two additional behavior measures, including pupil dilation response and whether the subjects recalled the correct answers in a follow-up session. Upon completing the (initial) task, subjects were unexpectedly asked to return within 11-16 days for a follow-up study. Twelve returned in about 2 weeks and were used in the analysis. In the follow-up session, subjects were shown the same questions and asked to recall the correct answers (earning $0.25 for each correct answer), in addition to $15 for participation.

Pupil dilation response (PDR) was measured before and after the answer display using a Mobile Eyelink II eyetracker (SR Research, Osgoode, Ontario) at 250 Hz. Experiments were conducted and analyzed in Matlab (Mathworks, Natick, MA) using the Psychophysics Toolbox (Brainard 1997) and the Eyelink Toolbox (Cornelissen et al. 2002). In regards to the pupil dilation data, blinks were treated as missing data and removed. We focused on the time interval from 4.8 seconds before to 4.8 seconds after the onset of the answer display. After normalization, we split the pupillary data collected over this interval into groups of high-, middle-, and low-curiosity level. The data were then averaged every 400 msec across subjects (figure 2.7).

Table 2.4 Two-sample t-test of the accuracy rate of fMRI subjects' post-scan guesses and Study 2 subjects' online guesses (equal variances assumed)[6]

	Study 1 fMRI subjects	Study 2 Behavioral subjects
Mean	0.31	0.27
Variance	0.01	0.01
Observations	19	15
Pooled Variance	0.01	
p-value	.11	

Results

Pupils Dilate in Response to Curiosity-Piquing Questions

Trials were divided into three terciles based on curiosity as measured in the initial session. On high-curiosity trials, PDR responses ramped up 1 to 2 s before the answer

[6] One outlier is excluded from the subject group in Study 2. The outlier subject was a visiting international student for the summer, who was not familiar with American cultural trivia and showed strong underperformance (accuracy rate = 0.075). The fMRI-behavioral group difference is slightly larger (.31 vs. .26) and slightly more significant if the outlier is included.

onset, peaked 800 ms after, and then dropped back to baseline around 2 seconds afterward (figure 2.7). Average PDR during anticipation (1 second before the answer onset) was significantly higher for high curiosity items as compared to middle curiosity items ($p < 0.03$, one-tailed t-test), and modestly different for middle as compared to low curiosity items ($p = 0.13$, one-tailed t-test). When the answer appeared (0~1000 ms after the onset), the average PDR was significantly different among all three groups ($p < .03$ or lower, one-tailed t-test).

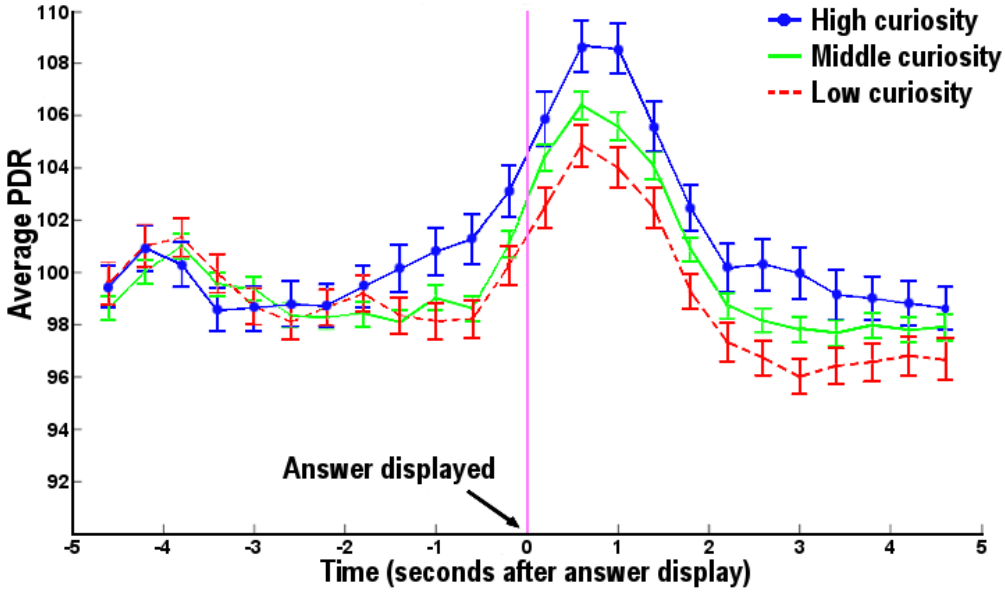

Figure 2.7 Pupillary response. Curiosity correlates with pupil dilation before and just after answers are revealed. The y-axis shows individually normalized pupil dilation (n=16) around the time of answer display, for different levels of curiosity: high (above the 67th individual percentile) in blue, low (below the 33rd individual percentile) in red, and middle in green. The average pupil dilation in this time interval for each subject was normalized to 100.

Ordinary least squares (OLS) regression analysis was performed to confirm the effects of curiosity on PDR, while controlling for individual fixed effects and a quadratic time trend. The regression analysis shows that an increase of one standard deviation in

curiosity level results in a 0.76% (standard error 0.34%) and 0.86% (0.35%) increase in average PDR during the anticipation and answer viewing, respectively (table 2.5).

Table 2.5 Pupil dilation response (PDR) regressions. Regress average PDR (individual mean=100) with curiosity level (CURIO), controlling for individual fixed effects (results not shown) and a quadratic time (QUESTION) trend; standard errors in parentheses. N<640 since some PDR are missing (blinks, etc.)

Period of Interest	Anticipation (-1~0 s)	Answer viewing (0~1 s)	Drop-off (1~2 s)
Constant	106.741** (1.745)	107.941** (1.721)	102.200** (1.607)
CURIO	0.740† (0.358)	0.738† (0.354)	0.384 (0.329)
confidence P	0.089† (0.036)	0.254** (0.036)	0.247** (0.033)
uncertainty P(1 – P)	0.0008† (0.0004)	0.0024** (0.0004)	0.0027** (0.0004)
QUESTION	-0.398** (0.120)	-0.433** (0.118)	-0.392** (0.110)
$QUESTION^2$	0.006† (0.003)	0.006† (0.003)	0.005 (0.003)
N	632	639	636
F	6.79	9.809	9.476
F-test p-value	0	0	0
R^2	0.155	0.216	0.211

Note: † t-test significant at $p<0.05$; ** at $p<0.001$
QUESTION denotes the number in which the question appeared in the temporal order to capture adaptation effects. Note that curiosity reliably increases PDR just before and after viewing the answer (second and third column results) but is insignificant in the 1-2 s after the answer presentation, while the effects of confidence and uncertainty persist.

Initial Curiosity Enhances Subsequent Memory for Incorrectly Guessed Answers

Curiosity as expressed in the initial session had a strong effect on subsequent recall of the answers to the questions that were initially guessed incorrectly (figure 2.8). The accuracy rates differed significantly between high- and middle-curiosity items; middle- and low-curiosity items; and high- and low-curiosity items (all p's <0.05, paired one-tailed t-tests).

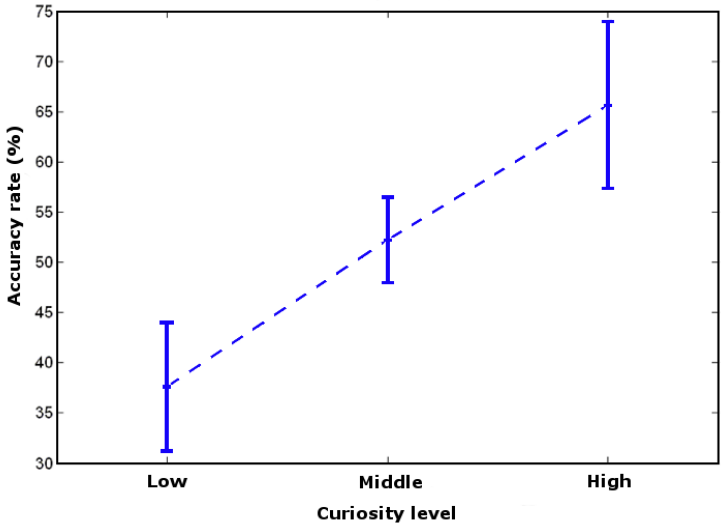

Figure 2.8 Percentage accuracy on the memory test for trials in which subjects initially guessed incorrectly. Focusing only on the answers initially guessed incorrectly, we divided the answers into three (high/middle/low) curiosity groups. Then the average accuracy rate was computed for each group across subjects, by dividing the number of correct recalls by the number of total answers to be recalled.

We repeated the same procedure with the residual curiosity (as defined in Study 1). Average accuracy rates are 56% (5.14%) for high, 51% (5.81%) for middle, and 38% (6.40%) for low curiosity questions (standard error in parenthesis). The difference in accuracy rate between high vs. middle was in the right direction but insignificant ($p=0.17$, paired one-tailed t-test). However, the differences between middle vs. low and also between high vs. low were significant ($p<0.04$ and $p<0.01$, respectively, paired one-tailed t-tests). The result is a little bit weaker, but still consistent with the previous results. We also show a main effect of curiosity on recall when we repeated the analysis, including control variables P and $P(1-P)$. Consistent with the fMRI findings, these results support the hypothesis that curiosity activates memory regions differentially in response to surprising (incorrectly guessed) answers, resulting in greater accuracy of subsequent recall of the correct answers.

OLS and logistic regression analyses were also performed to confirm the effects of curiosity on memory enhancement. The dependent variable, correct recall in the memory test (coded as 1 for a correct recall and 0 for an incorrect recall), was regressed on curiosity level interacting with two dummy variables that indicate whether subjects initially guessed the answer correctly or not. The regression analyses also find consistent evidence that curiosity modulates the later recall rate for answers to questions that subjects initially guessed incorrectly (table 2.6).

Table 2.6 Memory test regressions: correct recall (1 = recalled the correct answer in memory test, 0 = otherwise) was regressed on confidence *P*, uncertainty *P(1 − P)*, initial correctness(INI-CORRECT, 1 = initially guessed correctly, 0 = otherwise), initial 'incorrectness'(INI-WRONG, 1 = initially guess incorrectly, 0 = otherwise), and curiosity level interacting with initial correctness (CURIO*INI-CORRECT, CURIO*INI-WRONG, respectively), controlling for individual fixed effects (fixed effect results not shown).

Regression Method	OLS (s. e.)	OLS (s. e.)	OLS (s. e.)	Logit (s. e.)	Logit (s. e.)
Constant (Last subject)	0.362** (0.075)	0.362** (0.075)	0.386** (0.070)	-0.788† (0.371)	-0.752† (0.375)
CONFIDENCE *P*	0.474** (0.072)	0.478** (0.087)	0.127 (0.091)	2.686** (0.444)	3.142** (0.676)
UNCERTAINTY *P(I-P)*	-	-0.017 (0.257)	-0.083 (0.240)	-	-1.573 (1.624)
INI-CORRECT(‡)	-	-	0.452** (0.052)	-	-
CURIO*INI-CORRECT	0.033 (0.040)	0.034 (0.042)	0.035 (0.039)	0.104 (0.232)	0.166 (0.247)
CURIO*INI-WRONG	0.072* (0.024)	0.072* (0.024)	0.078** (0.022)	0.345* (0.120)	0.350* (0.120)
N	520	520	520	520	520
F–statistic/ LR	6.952	6.505	11.540	99.763	100.735
F-statistic/LR-test p-value	0.000	0.000	0.000	0.000	0.000
R^2	0.147	0.145	0.257	0.187	0.189

Note: † t-test significant at $p<0.05$; * $p<0.01$;** $p<0.001$.

(‡) Logit analysis including initial correctness could not be performed due to multicollinearity.

2.4 Study 3: Curiosity as Reward Anticipation

Experimental Design and Methods

Participants

Participants in Study 3 were assigned to one of two conditions: a token condition (10 Caltech students, mean age = 23.4±3.3 years, 5 males, 5 females) and a time condition (20 Caltech students, mean age = 19.9±2.2 years; 12 males, 8 females). Initially, 21 subjects participated in the time condition, but one subject was excluded for further analysis because his data showed that he was not engaged in the task—first, he did not take enough time to comprehend a question (evidenced by the fact that he spent only 2.1 seconds, which is only a third of the time that other 20 subjects spent reading a question on average) and second, his curiosity and confidence ratings were highly correlated with anchor ratings (70% and 65%, respectively)[7], which means that he submitted anchor ratings instead of his own. Informed consent was obtained using a consent form approved by the internal review board at Caltech.

Task and Measures

The task and timeline were similar to the previous experiments except for two modifications (figure 2.9). First, in this experiment, subjects had to spend scarce resources—either an experimental token or waiting time—to learn answers. Second, 10 questions were added to the original 40 questions.[8]

[7] At the start of each rating screen, a computer randomly selected an initial rating which a subject started toggling buttons from. We call this starting-point rating an anchor rating.

[8] We were concerned that in the token condition some subjects might spend their tokens too quickly or too slowly, so that they would have either none left or many left before the last few questions. If so, their token choices in those later trials would not reflect their true desire for information. We therefore added 10 new questions at the end as padding and excluded them from the analysis, using only the same 40 questions

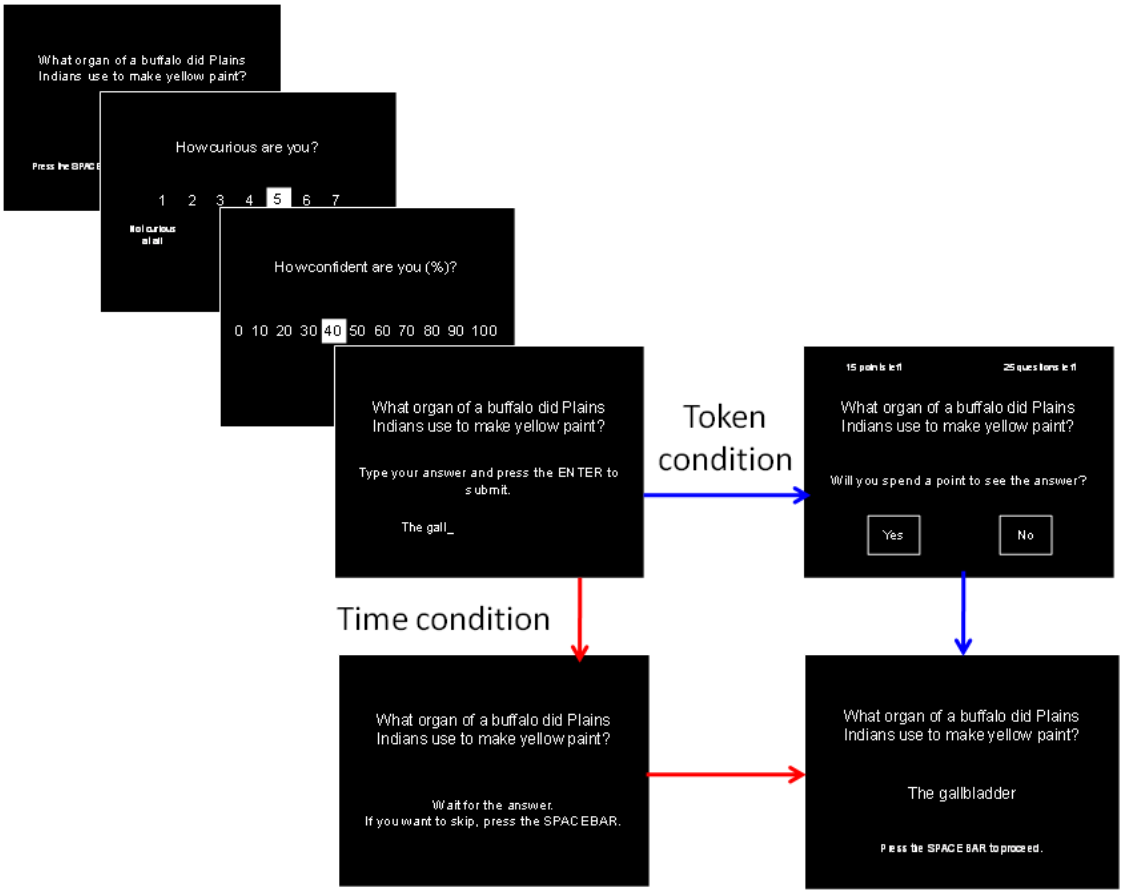

Figure 2.9 Timeline of Experiment in Study 3. Subjects read a question, rate curiosity and confidence level about the answer and type their guess. Then subjects can either spend a token (in the token condition) or wait for a certain amount of time (in the time condition) to see the answer if they wish. If they do not wish to see the answer, they can quickly move on to the next question.

A reward is an object or event that elicits approach and is worked for (Wise 2004). Requiring subjects to spend tokens or time measured their willingness to pay for information they were curious about and to evaluate curiosity-provoking information as a reward. The different conditions tested the robustness of the effect of curiosity to change in the type of resource that was spent.

---

used in the previous studies (including the last 10 questions does not change the result, however). Indeed, three of 10 subjects spent all their tokens and one of 10 subjects had more than sufficient tokens to see all 10 answers left by the time the last 10 padding questions started. There is no budgeting problem in the time condition, so all 50 questions were included in the analysis.

In the token condition, subjects had to spend 1 of their 25 experimental tokens to find out the answer to a question. Subjects read each question, reported their curiosity and confidence levels, and typed their guess. After guessing, they could pay 1 token to see the answer immediately. The tokens did not have cash value, but since there were 25 tokens and 50 questions, spending a token on an answer meant skipping another answer, or a lost opportunity to see another answer (an opportunity cost). Given the previous finding in Study 1 that high curiosity was correlated with activity in the striatal region, we hypothesized that when subjects were more curious, they would anticipate higher reward from learning information, so they would be more likely to spend a token. Other results were possible, however. Subjects could allocate tokens on the basis of their confidence in their guesses (to find out answers they are confident they know, or confident they do not know), impatiently use all their tokens in the beginning, alternate spending and saving, or exhibit some other pattern unrelated to curiosity. Any of these other token-spending patterns would be inconsistent with the reward anticipation hypothesis.

The second condition imposed a different cost: after guessing, if subjects wanted to see the answer, they had to wait until it appeared. They were told that the waiting time would vary randomly from trial to trial and that any amount of time from 5 to 25 s would be equally likely. Subjects could quit waiting and skip to the next question at any time; in this case, they would not get to see the correct answer. We hypothesized that subjects would be more likely to spend time, to wait longer, for the answers that they were more curious about.

Results

Logistic regressions showed that spending tokens or spending time were both strongly correlated ($p<0.001$) with curiosity (figure 2.10, table 2.7 and 2.8). The significance did not change when P and $P(1-P)$ were included, or when residual curiosity was used.

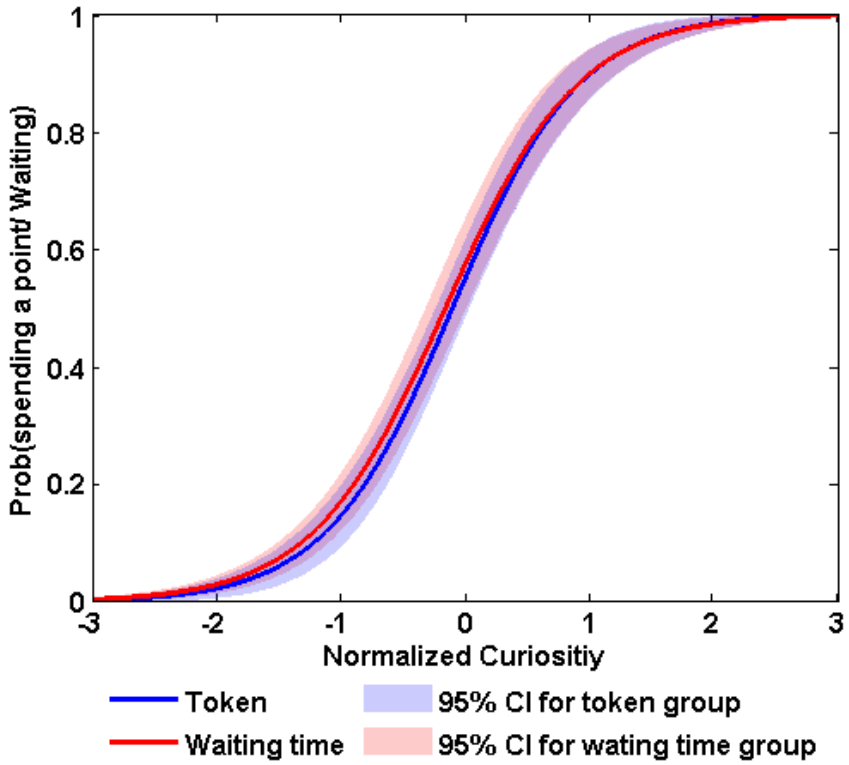

Figure 2.10 Curiosity correlates with willingness to pay the cost. Group logit curves relating normalized curiosity to the probability of spending a token (*blue*) or waiting time (*red*) to learn the answer to a trivia question. Logistic regression analyses were performed to test whether curiosity is correlated with spending tokens or time. For each condition, we pooled the data across subjects. CI=confidence interval.

Table 2.7 Group-level random-effects logistic regression of curiosity on decision to spend a token

Prob (spend a point)	Coefficient	S.E.	z	P>|z|	95% Confidence Interval	
normalized curiosity	1.96	0.18	10.96	0.000	1.61	2.31
constant	0.21	0.14	1.56	0.12	-0.05	0.48

Note: The dependent variable, decision to spend a token, was regressed on a subject-normalized curiosity rating and a constant. The dependent variable was coded as 1 if a subject spent a point and 0 otherwise.

Table 2.8 Group-level random-effects logistic regression of curiosity on decision to wait for an answer

Prob (wait)	Coefficient	S.E.	z	P>|z|	95% Confidence Interval	
normalized curiosity	1.88	0.13	14.47	0.000	1.63	2.14
constant	0.31	0.16	1.89	0.06	-0.01	0.63

Note: The dependent variable, decision to wait for the answer, was regressed on a subject-normalized curiosity rating and a constant. The dependent variable was coded as 1 if a subject waited and 0 otherwise.

Table 2.9 reports the results of individual logistic regressions for the effect of normalized curiosity on the decision to spend a token or time. At an individual-subject level, correlations between curiosity and spending were significant at $p<.01$ for most of the subjects. Figure 2.11 and 2.12 show the individually fitted probabilities of spending a token or time as a function of the normalized curiosity ratings, using the coefficients reported in table 2.9. Subjects waited for an additional 3.7 seconds as their normalized curiosity level increased by one standardized unit. In summary, for most of the subjects, curiosity is the strongest predictor of whether a subject will spend tokens or time to see the answer to a question.

Table 2.9 Summary of individual logistic regressions of the normalized curiosity on decision to spend cost

(a) Token condition

Subject ID	Coeff. for norm. curiosity	*z*-stat	*p*-value	pseudo-R^2
1	1.33	3.14	0.002	0.23
2	1.02	2.55	0.011	0.15
3	2.09	3.19	0.001	0.36
4	1.08	2.57	0.010	0.16
5	2.18	3.86	0.001	0.49
6	1.92	2.86	0.004	0.31
7	4.42	2.61	0.009	0.72
8	2.64	3.27	0.001	0.47
9	2.69	3.52	0.001	0.54

Note: (1) The coefficient for an implicit constant term is not reported here.

(2) For subject 10, the normalized curiosity level of -.524 predicts data perfectly—that is, the probability of waiting jumps at -.524 from 0 to 1.
(3) The coefficient for an implicit constant term is not reported here.

(b) Time condition

Subject ID	Coeff. for norm. curiosity	*z*-stat	*p*-value	pseudo-R^2
1	1.13	2.97	0.003	0.18
2	3.70	3.26	0.001	0.63
3	1.01	2.78	0.005	0.14
4	1.55	3.54	0.000	0.28
5	2.69	3.88	0.000	0.52
6	1.46	3.19	0.001	0.23
7	2.49	3.79	0.000	0.47
8	3.81	3.55	0.000	0.70
9	1.17	1.75	0.08	0.14
10	1.60	3.81	0.000	0.31
11	1.28	2.74	0.006	0.36
12	2.90	1.41	0.158	0.26
13	1.39	3.66	0.000	0.26
14	2.27	3.69	0.000	0.43
15	3.24	3.29	0.001	0.56
16	2.08	3.76	0.000	0.42
17	2.36	1.92	0.055	0.44
18	1.61	2.88	0.004	0.33

Note: (1) For subject 19, the normalized curiosity level of -.232 predicts data perfectly—that is, the probability of waiting jumps at -.232 from 0 to 1.
(2) For subject 20, the normalized curiosity level of -1.921 predicts data perfectly.
(3) The coefficient for an implicit constant term is not reported here.

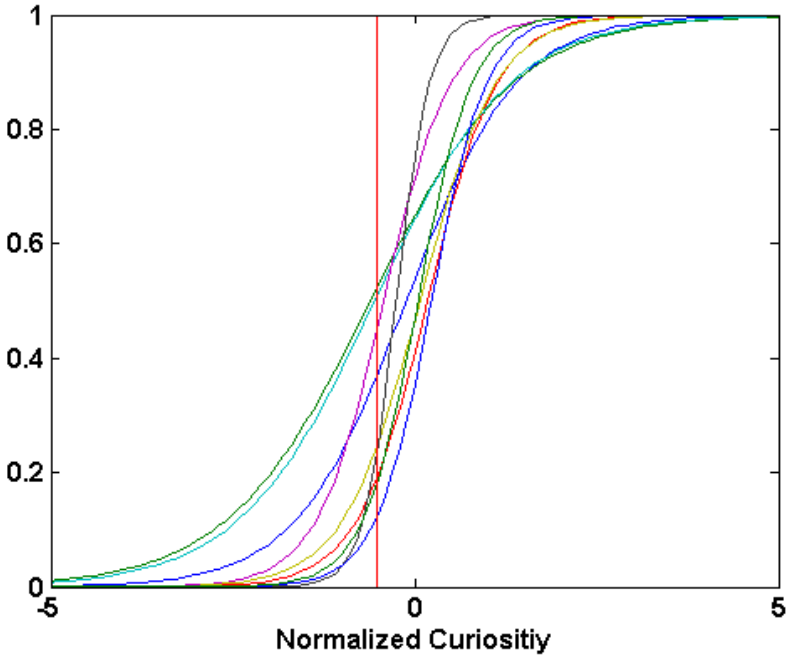

Figure 2.11 Individual choice probability curves in the token condition. Token-spending behavior depends on curiosity level. This pattern does not change when *P* and *P(1 − P)* are also included in the model. Y-axis: probability of spending a token.

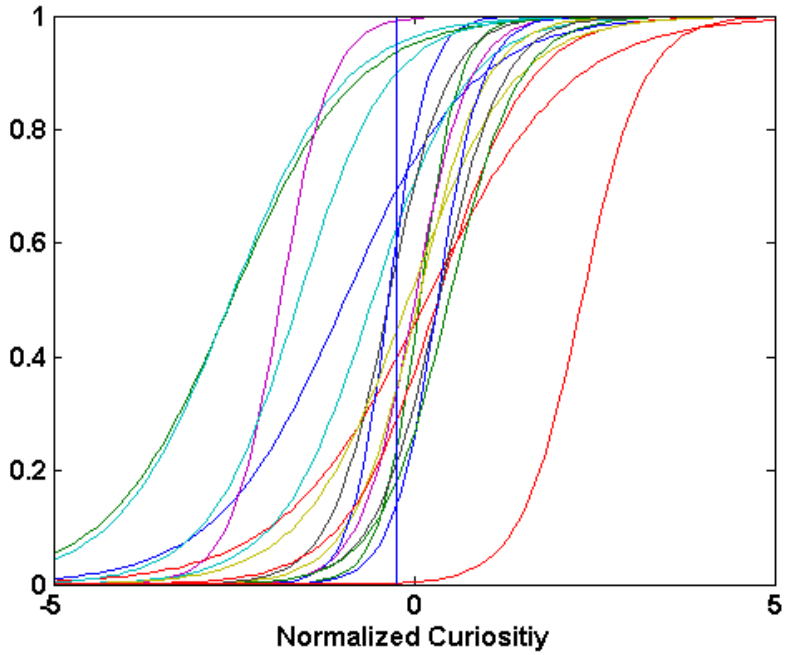

Figure 2.12 Individual choice probability logistic curves in the time condition. The decision to wait depends on curiosity level. For most of the subjects, this does not change when *P* and *P(1 − P)* are also included in the model. Y-axis: probability of waiting for the answer.

2.5 Discussion and Conclusions

Curiosity was considered an important drive by psychologists of a century ago, but research on it subsequently waned (Loewenstein 1994). This study was an attempt to revive interest in curiosity; measuring it by self-report, and using fMRI to study neural correlates of reported curiosity. The findings suggest hypotheses about memory and reward anticipation.

The correlations between reported curiosity and both lateral PFC and caudate activity are consistent with hypotheses that curiosity is linked to anticipation of information, and that information is a secondary reinforcer. Curiosity was correlated with activity in the caudate when a question was first presented. It is well established that the caudate is involved in reward anticipation and reward learning, across a wide variety of primary and secondary reinforcers (Delgado et al. 2003; Delgado et al. 2000), including social rewards such as benevolent reciprocity (Fehr & Camerer 2007; King-Casas et al. 2005), social cooperation (Rilling et al. 2002), altruistic punishment (De Quervain et al. 2004) and winning an auction (Delgado et al. 2008).

Previous studies have found that the expectation of feedback is sufficient for activation of the caudate (Aron et al. 2004). Our experimental design included feedback in the form of the correct answers presented to the subjects. If there is brain activity in anticipation of positive feedback, it should be modulated by the confidence level *P* (the more confident one is in being right, the more positive feedback one expects). The parametric design of the analysis, by correlating activity with curiosity levels and then with residual curiosity, precludes the possibility that the caudate activation was driven solely by expectation of

feedback from accurate guesses because residual curiosity and confidence are uncorrelated by construction. The finding that activity in the left PFC is correlated with curiosity is also consistent with the idea that curiosity is associated with an intrinsic value of learning, because neurons in the left PFC receive input from neurons in the substantia nigra via the dorsal striatum. The dorsal striatum responds to magnitude of primary rewards and reward prediction (Leon & Shadlen 1999; Rogers et al. 1999) and shows sustained phasic activations during reward expectation (Watanabe 1996)

There are also studies that have reported striatal activations in response to negatively valenced stimuli (Knutson et al. 2001) or during nonreward activity such as working memory and motor preparation (Cairo et al. 2004; Simon et al. 2002). Since our task did involve working memory and motor preparation, the striatal activation we found could have been due to increased attention, incentive salience, or other activities (as some studies suggest). Given that people tend to be more attentive to an object that is more rewarding, the fMRI evidence alone cannot rule out these interpretations from reward-anticipation interpretation.

Therefore, the reward anticipation interpretation was investigated further in a separate behavioral study. In Study 3, subjects were allowed to either spend a token or to wait to see the answer to a question. Both actions incurred a cost—a lost opportunity or lost time. People are generally willing to spend time and resources to obtain objects that they find rewarding. Thus, enhanced willingness to spend resources to find out the answers to more curiosity-provoking questions is consistent with the reward interpretation hypothesis and

not with the idea that the fMRI results indicate only effects of attention or incentive salience.

Recent computational neural network models suggest another compatible interpretation involving memory (Frank et al. 2001; O'Reilly & Frank 2006). Question stimuli trigger differing levels of curiosity, and the basal ganglia sends out a stronger signal to enable the lateral PFC to update, maintain, and internally represent questions eliciting higher curiosity, whose answer may be anticipated to be more rewarding. This internal representation in the lateral PFC, particularly in the left IFG, is a crucial component of long-term memory consolidation and is critical in the learning of new information (Paller & Wagner 2002).

The PHG and left IFG were activated in response to incorrect guesses, and the level of activation was correlated with the level of curiosity. These regions are thought to be involved in successful verbal memory encoding (Brewer et al. 1998; Paller & Wagner 2002; Wagner et al. 1998). In conjunction with the caudate and left PFC activations when questions were first presented, this activity suggests that curiosity strengthens memory of correct answers when subjects were initially wrong—that is, that curiosity is linked to the reward value of information and also enhances learning from new information.

This conjecture led to Study 2, in which we measured PDR and memory. Because pupil dilation is known to be linked to arousal, attention, interest, cognitive effort (Beatty 1982; Hess & Polt 1960; Kahneman 1973) and more efficient verbal learning (Kahneman & Peavler 1969), and because anticipatory pupillary responses increase following a stimulus that predicts rewards (O'Doherty et al. 2003), the observed correlation of curiosity with

pupil dilation is consistent with both reward anticipation and learning of novel information. The finding that curiosity enhanced later recall of novel information suggests that curiosity helps to consolidate novel information in memory. Having established that curiosity is a form of reward anticipation, we can also tie this research to the work of Adcock and her collegues (Adcock et al. 2006), who showed that anticipated monetary rewards modulate activations in the mesolimbic and parahippocampal regions and promotes memory formation prior to learning. Our results complement theirs by showing that *endogeneous* internal motivation manifested in curiosity recruits neural circuits similar to those that are recruited by *exogenous* incentives, and has a similar effect on learning.

We also found bilateral putamen activation during the answer display after incorrect guesses. Although no explicit reward or punishment was involved, subjects might have perceived guessing incorrectly as an unexpected inherent, or self-punishment and the severity of this punishment might have been based on confidence in the answer (e.g., subjects would have felt worse about an incorrect guess if they had been very confidence about that guess than if they had not been confident). The differential putamen activations we found might reflect this "internal punishment" aspect of guessing incorrectly. This interpretation is consistent with the recent finding that unexpected punishments and unexpected rewards produce very similar BOLD responses in the putamen at the time of outcomes (Knutson et al. 2000; Seymour et al. 2007).

The exploratory nature of our study does not allow us to examine all possible aspects of curiosity separately. It is certainly likely that curiosity works differently in different

sensory and knowledge domains. The trivia questions we used evoke what is often called "specific epistemic curiosity" (Berlyne 1954). This kind of curiosity is the desire for a particular piece of information and is often associated with motivations for academic achievement and scientific discovery. This type of curiosity is probably different from the sensation driven by stimulus novelty or the desire to avoid boredom or sensory deprivation. The latter type of curiosity is called diversive perceptual curiosity and can be found in various animals (Berlyne 1954). The curiosity we measured includes a desire to learn new information and anticipation of the rewarding information to be learned (because subjects received feedback). A study without feedback would isolate pure curiosity absent anticipation and learning, though it would not permit study of the response to correct and incorrect guesses. Further studies could also show whether curiosity is different for visual stimuli, semantic narratives (e.g., page-turner novels), social information (e.g., gossip), and subjects of "morbid curiosity."

Understanding the neural basis of curiosity has important substantive implications. Note that although information seeking is generally evolutionarily adaptive (Panksepp 1998), modern technologies magnify the amount of information available, and hence the potential effects of curiosity. Understanding curiosity is important for selecting and motivating knowledge workers who gather information (e.g., scientists, detectives, and journalists). The production of engaging news, advertising and entertainment is also, to some extent, an attempt to create curiosity. The fact that curiosity increases with uncertainty (up to a point) suggests that a small amount of knowledge can pique curiosity and prime the hunger for knowledge, much as an olfactory or visual stimulus can prime a

hunger for food; this observation might suggest ways for educators to ignite the wick in the candle of learning.

2.6 Appendix for Chapter 2

2.6.1 Questions, Answers, and Average Curiosity Ratings

The order of presentation of questions was randomized within each individual run.

Question **(Answer, Average curiosity rating across subject)**
Run 1
What rock and roll band performs "I want to Rock and Roll All Night"? (Kiss, 3.84)
What unfortunate handicap did Thomas Edison suffer from? (Deafness, 4.63)
What city was "Groundhog Day," starring Bill Murray, filmed in? (Pittsburgh, 3.53)
What book is the most shoplifted book in the world? (The Bible, 5.05)
What is the museum-surrounded space in Washington DC called? (The Mall, 3.89)
How long were Jerry Seinfeld and his pals sentenced in the series finale? (One year, 2.37)
What is the only type of animal besides a human that can get a sunburn? (Pig, 5.42)
Which school has the most students over age 25 according to US News? (University of Phoenix, 3.74)
What snack food is an ingredient in the explosive dynamite? (Peanuts, 5.63)
What is the most sober school according to The Princeton Review? (Brigham Young University, 5.10)
Run 2
What city is referred to as "Pittsburgh of the South"? (Birmingham, AL, 3.83)
What invention should make Ts'ai Lun, a 2nd century inventor, a household name? (Paper, 5)
What breed of dog is the only animal whose evidence is admissible in American courts? (Bloodhound, 4.94)
What animal can shed up to 30,000 teeth in its lifetime? (Shark, 4.33)
Who was the first host of the comedy show Saturday Night Live? (George Carlin, 3.5)

What is the only country in the world where women dominate the government? *
(Belgium, 5.89)
What type of political campaign is characterized by many stops in small towns?
(Whistle-stop campaign, 3.89)
What instrument was invented to sound like a human singing?
(Violin, 5.72)
From what city in the United States did Coca-Cola originate?
(Atlanta, GA, 3.78)
What animal's excrements are consumed as luxury food?
(Bats, 4.83)

Run 3

What everyday food will make a drug test show up positive?
(Poppy seeds, 3.61)
What industry uses 20% of China's harvested plants?
(Medicine, 4.72)
What electronic item is stolen most often on the NYC subways? *
(iPods, 4.11)
What famous person was Dolly the cloned sheep named after?
(Dolly Parton, 3.33)
What fictional character in Treasure Island lends its name to a fast food chain?
(Long John Silver, 4.11)
What is the name of the galaxy that Earth is a part of?
(Milky Way, 2.28)
What is the most abundant mineral in the human body?
(Calcium, 5.11)
What president has three 'A's in his first name where each has a different sound?
(Abraham Lincoln, 3.39)
What title was Catherine of Aragon known by after she divorced Henry VIII?
(Dowager Princess of Wales, 4.44)
What country has won the most Miss World beauty contests? *
(Venezuela, 3.89)

Run 4

What is the only country in the world that has a bill of rights for cows?
(India, 3.47)
What was the first animated film to win an Academy Award?
(Beauty and the Beast, 4.37)
What item on the McDonald's menu has the most calories? *
(Chicken Selects, 20 Piece, 4.84)
What city has the only drive-thru post office in the world?
(Chicago, 4.26)
What crime is punishable if attempted, but not if committed?
(Suicide, 4)
What secular philosopher's teaching influenced life in his country for 2000 years?
(Confucius, 4.63)

What Beatles song lasted the longest on the American charts?
(Hey Jude, 5)
What is the only type of lizard that has a voice?
(Gecko, 4.47)
Which sports athlete has appeared in McDonald's, Nike and Hanes advertisements?
(Michael Jordan, 2.89)
What was put in place by the Greeks before and during all the Ancient Olympic festivals?
(A truce, 4.47)

Note: The questions in asterisk (*) are updated or replaced by the following set:

What male body part did Mademoiselle magazine find to be the favorite of most women?	Eyes
What part of a woman's body were ancient Chinese artists forbidden to paint?	Her foot
What creature proved to be much faster than a horse in a 1927 race in Sydney, Australia?	The Kangaroo
What item on the McDonald's menu has the most calories?	Chocolate Triple Thick Shake (32 fl oz cup), 1160 kcal

Additional Questions and answers used in Experiment 3.

Questions	Answer
What did the girls in medieval Spain put in their mouths to avoid unwanted kisses?	Toothpicks
What drupaceous fruit were Hawaiian women once forbidden by law to eat?	The coconut
In parts of India, the oldest brother must marry first. If he cannot find a wife, what can he choose to marry?	A tree
How many years are in an eon (aeon)?	100 million
What fat substitute got FDA approval for use in snack foods, despite reports of diarrhea and cramps?	Olestra
What organ is found only in vertebrates (animals with a backbone)?	Liver
In 1875, who helped Daniel Peter invent "milk" chocolate?	Henry Nestle
What butterfly-shaped gland is located just in front of the windpipe?	The Thyroid
What is a shark’s skeleton made of?	Cartilage
Who was the first Christian emperor of Rome?	Constantine the Great

2.6.2 Instructions for Study 1

Instructions

Thank you for participating in this study on curiosity. During the experiment you will be asked a series of different trivia-type questions about things that you may or may not find yourself curious about. After presenting each question, we will ask you to rate (1) how curious you are to know the answer, and (2) how confident you are that you know the answer. The answer to the question will be revealed before you move on to the next question.

The questions are presented in a pre-programmed pace. So please wait for the next rating slide after you complete the task in the question slide. You will be given only 12 to 15 seconds to read each question. No matter how quickly you figure out the answer, the program will not let you continue until the full 12 to 15 seconds are up. The program will automatically move on to the next step when 12 to 15 seconds are over. Please respond quickly, but please do not speak/think aloud.

Here is the sample of a question.

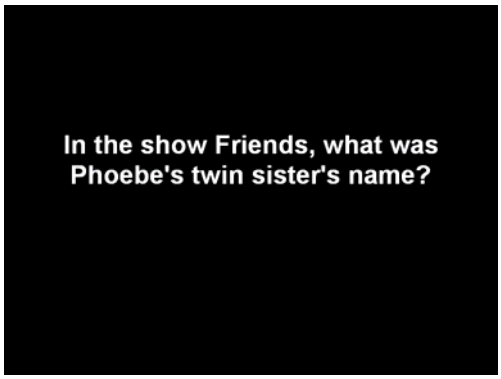

After a question is presented, a screen with a sliding scale from 1 (not at all curious) through 7 (very curious) will appear (see below for the sample). Select the number that represents how curious you are about the answer. When rating curiosity, use the right (1) and left (3) buttons to scroll along the response scale. To confirm your choice, press the top (4) button.

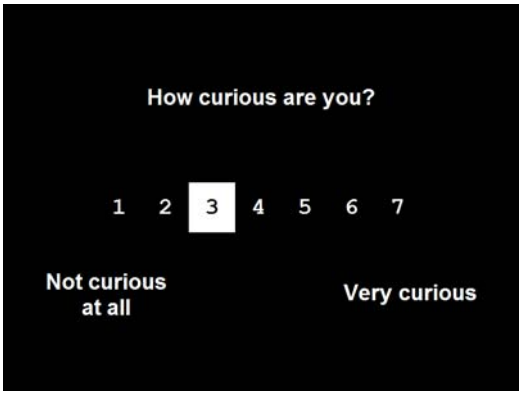

After making your selection, you will be asked to rate how confident you are of the answer or whether the answer is on the tip of the tongue (see below for the sample screen). You may respond with a sliding scale from 0% (not at all confident) through 100% (very confident) or the tip of the tongue (TOT) option. The 'tip-of-the-tongue' response indicates that you feel that you know the correct answer, even though you cannot remember the exact word that corresponds with the answer at the moment. If you feel that the correct answer is on the tip-of-your-tongue, please select "TOT" response. Before selecting your option, please silently say the word to see if you actually produce it. You may use the 1 (right) and 3 (left) buttons to scroll along the response scale. The rating steps are self-paced, so the program will not move on to the next step until you press the 4 (top) button to confirm your selection.

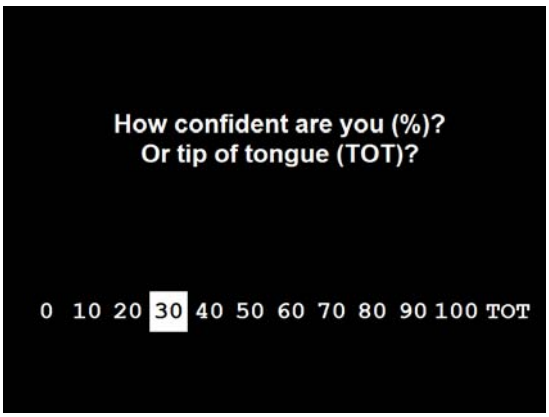

The fixation screen will appear between and after the rating slides and then the answer slide will follow.

Here are the fixation screen (left) and the answer slide (right).

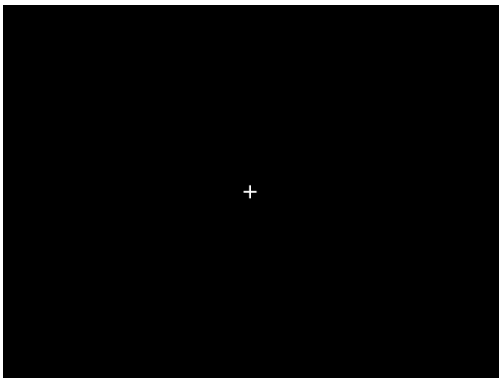

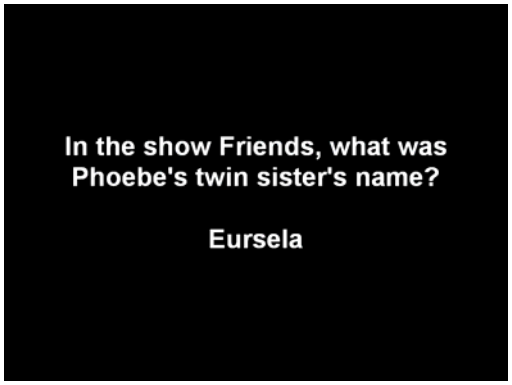

Before you move on to the next question, the fixation screen will appear again.

Before the experiment starts, you will be presented with 3 practice questions. After you answer and rate them, we will ask if you have any questions. At the end of the experiment, you will be asked to debrief whether your guess was right or not and to fill out a short questionnaire. You will be paid $20 for participating in this experiment. The information we obtain from the experiment will not be used for other commercial or non-academic purposes.

Bibliography

Adcock, R. A., A. Thangavel, S. Whitfield-Gabrieli, B. Knutson, & J. D. E. Gabrieli. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron,* **50** 507-517.

Aron, A. R., D. Shohamy, J. Clark, C. Myers, M. A. Gluck, & R. A. Poldrack. (2004). Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. *Journal of Neurophysiology,* **92** 1144-1152.

Ashburner, J., & K. Friston. (1999). Nonlinear spatial normalization using basis functions. *Human Brain Mapping,* **7** 254-266.

Ashburner, J., P. Neelin, D. L. Collins, A. Evans, & K. Friston. (1997). Incorporating prior knowledge into image registration. *Neuroimage,* **6** 344-354.

Beatty, J. (1982). Task-evoked pupillary responses, processing load, and the structure of processing resources. *Psychological Bulletin,* **91**(2) 276-292.

Berlyne, D. E. (1954). A theory of human curiosity. *British Journal of Psychology,* **45** 180-191.

Biederman, I., & E. Vessel. (2006). Perceptual pleasure and the brain. *American Scientist, 94,* 247-255.

Bookheimer, S. (2002). Functional mri of language: New approaches to understanding the cortical organization of semantic processing. *Annual Revies of Neuroscience,* **25** 151-188.

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

Brewer, J. B., Z. Zhao, J. E. Desmond, G. H. Glover, & J. D. E. Gabrieli. (1998). Making memories: Brain activity that predicts how well visual experience will be remembered. *Science,* **281** 1185-1187.

Cairo, T. A., P. F. Liddle, T. S. Woodward, & E. T. C. Ngan. (2004). The influence of working memory load on phase specific patterns of cortical activity *Cognitive Brain Research,* **21**(3) 377-387.

Cornelissen, F. W., E. M. Peters, & J. Palmer. (2002). The eyelink toolbox: Eye tracking with matlab and the psychophysics toolbox. *Behavior Research Methods, Instruments & Computers,* **34** 613-617.

De Quervain, D. J. F., U. Fischbacher, V. Treyer, M. Schellhammer, U. Schnyder, A. Buck, & E. Fehr. (2004). The neural basis of altruistic punishment. *Science,* **305** 1254-1258.

Delgado, M. R., H. M. Locke, V. A. Stenger, & J. A. Fiez. (2003). Dorsal striatum responses to reward and punishment: Effects of valence and magnitude manipulations. *Cognitive, Affective, and Behavioral Neuroscience,* **3**(1) 27-38.

Delgado, M. R., L. E. Nystrom, C. Fissell, D. C. Noll, & J. A. Fiez. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of neurophysiology,* **84** 3072-3077.

Delgado, M. R., A. Schotter, E. Y. Ozbay, & E. A. Phelps. (2008). Understanding overbidding: Using the neural circuitry of reward to design economic auctions. *Science,* **321** 1849-1852.

Fehr, E., & C. F. Camerer. (2007). Social neuroeconomics: The neural circuitry of social preferences. *Trends in Cognitive Sciences,* **11** 419-427.

Frank, M. J., B. Loughry, & R. C. O'Reilly. (2001). Interactions between frontal cortex and basal ganglia in working memroy: A computational model. *Cognitive, Affective, and Behavioral Neuroscience,* **1**(2) 137-160.

Friston, K. J., A. P. Holmes, K. J. Worsley, J.-P. Poline, C. D. Frith, & R. S. J. Frackowiak. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping,* **2**(4) 189-210.

Genovese, C. R., N. A. Lazaar, & T. Nicholos. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage,* **15** 870-878.

Hare, T. A., J. O'Doherty, C. F. Camerer, W. Schultz, & A. Rangel. (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *Journal of Neuroscience,* **28**(22) 5623-5630.

Hess, E. H., & J. M. Polt. (1960). Pupil size as related to interest value of visual stimuli. *Science,* **132**(3423) 349-350.

Hoffmann, B. (1972). *Albert einstein − creator and rebel*. New York: Viking Press.

Kahneman, D. (1973). *Attention and effort.* New Jersey: Prentice-Hall.

Kahneman, D., & W. S. Peavler. (1969). Incentive effects and pupillary changes in association learning. *Journal of Experimenteal Psychology,* **79** 312-318.

Kakade, S., & P. Dayan. (2002). Dopamine: Generalization and bonuses. *Neural Networks,* **15** 549-559.

Kang, M. J., M. Hsu, I. M. Krajbich, G. Loewenstein, S. M. McClure, J. T.-y. Wang, & C. F. Camerer. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science,* **20**(8) 963-973.

King-Casas, B., D. Tomlin, C. Anen, C. F. Camerer, S. R. Quartz, & P. R. Montague. (2005). Getting to know you: Reputation and trust in a two-person economic exchange. *Science,* **308** 78-83.

Knutson, B., C. M. Adams, G. W. Fong, & D. Hommer. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *Journal of Neuroscience,* **21**(159) 1-5.

Knutson, B., A. Westdorp, E. Kaiser, & D. Hommer. (2000). Fmri visualization of brain activity during a monetary incentive delay task. *Neuroimage,* **12** 20-27.

Leon, M. L., & M. N. Shadlen. (1999). Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron,* **24** 415-425.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin,* **116**(1) 75-98.

Maril, A., A. D. Wagner, & D. L. Schacter. (2001). On the tip of the tongue: An event-related fmri study of semantic retrieval failure and cognitive conflict. *Neuron,* **31**(4) 653-660.

McClure, S. M., M. K. York, & P. R. Montague. (2004). The neural substrates of reward processing in humans: The modern role of functional magnetic resonance imaging. *The Neuroscientist,* **10** 260-268.

O'Doherty, J. P. (2004). Reward representations and reward-related learning in the human brain: Insights from neuroimaging. *Current Opinion in Neurobiology,* **14** 769-776.

O'Doherty, J. P., P. Dayan, K. Friston, H. Critchley, & R. J. Dolan. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron,* **38**(2) 329-337.

O'Reilly, R. C., & M. J. Frank. (2006). Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation,* **18** 283-328.

Paller, K. A., & A. D. Wagner. (2002). Observing the transformation of experience into memory. *Trends in Cognitive Sciences,* **6**(2) 93-102.

Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. New York: Oxford University Press.

Rilling, J. K., D. A. Gutman, T. R. Zeh, G. Pagnoni, G. S. Berns, & C. D. Kilts. (2002). A neural basis for social cooperation. *Neuron,* **35** 395-405.

Rogers, R. D., A. M. Owen, H. C. Middleton, E. J. Williams, J. D. Pickard, B. J. Sahakian, & T. W. Robbins. (1999). Choosing between small, likely rewards and large, unlikely rewards activates inferior and orbital prefronal cortex. *Journal of Neuroscience,* **19** 9029-9038.

Seymour, B., N. Daw, P. Dayan, T. Singer, & R. J. Dolan. (2007). Differential encoding of losses and gains in the human striatum. *Journal of Neuroscience,* **27**(18) 4826-4831.

Simon, S. R., M. Meunier, L. Piettre, A. M. Berardi, C. M. Segebarth, & D. Boussaoud. (2002). Spatial attention and memory versus motor preparation: Premotor cortex involvement as revealed by fmri *Journal of Neurophysiology,* **88**(4) 2047-2057.

Wagner, A. D., D. L. Schacter, M. Rotte, W. Koutstaal, A. Maril, A. M. Dale, B. Rosen, & R. L. Buckner. (1998). Building memories: Remembering and forgetting of verbal experiences as predicted by brain activity. *Science,* **281** 1188-1191.

Watanabe, M. (1996). Reward expectancy in primate prefrontal neurons. *Nature,* **382** 629-632.

Wise, R. A. (2004). Dopamine, learning and motivation. *Nature Review Neuroscience,* **5** 483-494.

Chapter 3

Hypothetical and Real Consumer Choices Differentially Activate Common Frontostriatal Brain Circuitry

3.1 Introduction

In the study of choices people and groups make, the canonical choice of most interest is a real choice—a binding commitment to a course of action, like buying a house or casting a vote. However, many of the data collected to understand determinants of real choices are actually hypothetical statements about likely or future choices (they are guesses, plans, but not real choices).

The reliance on hypothetical choice data presumes that hypothetical choices are quantitatively similar to actual choices, and involve the same neural processes that lead to real choices. But saying "We should get married!" is not the same as saying "I do" (which is a legally binding real choice). Indeed, many studies in economics have found that hypothetical valuations of goods reflect a systematic "yes bias" compared to real valuations: i.e., hypothetical intentions to buy are overstated relative to real intentions (Blumenschein et al. 2007; Cummings et al. 1995; Johannesson et al. 1998; List 2001; List & Shogren 1998; Little & Berrens 2004; Murphy et al. 2005). This hypothetical bias has also been demonstrated in statements of intent to vote (Keeter & Samaranayake 2007) and to undertake a plan of action, which subjects later do not take (Ariely & Wertenbroch 2002; O'Donoghue & Rabin 2008; Tanner & Carlson 2009).

These empirical phenomena raise a question: Are different neural computations carried out in the brain for hypothetical and real decisions? This question is important for the cognitive neuroscience of choice. Answering it may provide tools to more accurately predict real choices from hypothetical choice data.

One possibility is that hypothetical and real decisions are differently processed in the brain and recruit different neural systems. Real choices are typically precise, immediate, have higher stakes, and are often more emotionally charged. Hypothetical choices might be rapid and mindless, requiring fewer cognitive resources. This hypothesis predicts no necessary functional overlap in the two types of choices.

Another possibility is that while real and hypothetical decisions share a common neural circuitry, additional activities are dedicated to each type of decision. For example, hypothetical decisions might be evaluated by forecasting a possible real choice, which requires probabilistic reasoning, anticipation of the future and simulation of available choice scenarios, recruiting a wider neural network. Alternatively, a decision in a hypothetical situation could be evaluated in isolation, since the hypothetical choice will not have any consequences. In contrast, a real decision requires thorough consideration of possible consequences, comparisons between alternative options and integration of opportunity costs—the foregone utility from the alternatives, into a valuation process. This hypothesis predicts additional activations special to real and hypothetical choices.

The results of this study clearly support the latter possibility. There is substantial overlap in neural circuitry during hypothetical and real choices, however the magnitude of

activity is greater in real decisions and there are additional activations during real choices.

In our study, 24 human subjects were scanned using functional magnetic resonance imaging (fMRI) while they made hypothetical and real purchase decisions regarding consumer products that are encountered in their daily lives. Subjects reported their willingness to pay (WTP) for 200 consumer products before scanning (figure 3.1 and 3.2). WTP is the maximum amount of money that an individual would be willing to spend to purchase the item (Hare et al. 2008; Plassmann et al. 2007). The median WTP was used to set a constant price for the scanning session; this price is different for each subject. In the scanner, subjects made binary (Yes or No) decisions about purchasing products at the constant price, and also expressed Strong or Weak preference (i.e., strong/weak yes or strong/weak no). The first block of 50 decisions was hypothetical. The second block of 50 decisions was based on different products (different from the first block, but still part of the original 200), and these decisions were unexpectedly real (i.e., an agreement to buy could be enforced by taking money from their initial payment for participation and giving them the purchased good—this will be explained further) (figure 3.3). After the scanning, they were then shown the same products that had been presented in the hypothetical trials (also unexpectedly) and they then made real (binding) purchase decisions for those products.

Previous studies have found that the brain encodes goal values and decision values during everyday economic decision-making scenarios (Hare et al. 2008; Knutson et al. 2007; Plassmann et al. 2007). A goal value is the value of a choice object. A decision value is any quantity that reflects the net benefit of making a decision to reach the goal, and can

typically be linked to observed choice probabilities. The key decision value (DV) in the current study is "consumer surplus," the maximum WTP for an item minus the price of the item, an economic measure of the net benefit of the decision. The specific goals of this study are to investigate differences in the brain regions used during hypothetical and real decision making, and to study differences in how the decision value is encoded in both types of decision making.

3.2 fMRI Study

Experimental Design and Methods

Participants

Male subjects were recruited from Caltech and the neighboring community in the fMRI study, (N=24; mean ± st. dev. age = 20.9 ± 6.1 year; age range = 17–47). Seven additional subjects participated in the experiment. Four were excluded from the analysis because they reported, during a debriefing, misunderstanding the instructions of the second fMRI part (real choice) which compromised internal validity. Three subjects were excluded because modification of decision value to fit choices was not reliable. All subjects were right handed and healthy; had no history of psychiatric diagnoses, or neurological or metabolic illnesses; had normal or corrected-to-normal vision; and were not taking medications that interfere with the performance of fMRI. Subjects were informed about the experiment and gave written consent upon arrival at the laboratory. The institutional review board of Caltech (Pasadena, CA) approved the study.

Stimuli

Two hundred consumer products were selected from various categories appealing to males (e.g., MP3 players, DVDs; for the complete product list, see table A3.1) in order to generate balanced purchase decisions across different individual preferences. The retail price of the products ranged from $20 to $100. Approximately 300 products were initially collected by the experimenters and about 100 of the most unpopular items were screened out by a pilot study, leaving 200 products used in the study. The products were presented to the subjects using high-resolution color pictures (72 dpi). The stimulus presentation and response recording was controlled by E-prime (Psychology Software Tools). Inside the scanner, the visual stimuli were presented using video goggles.

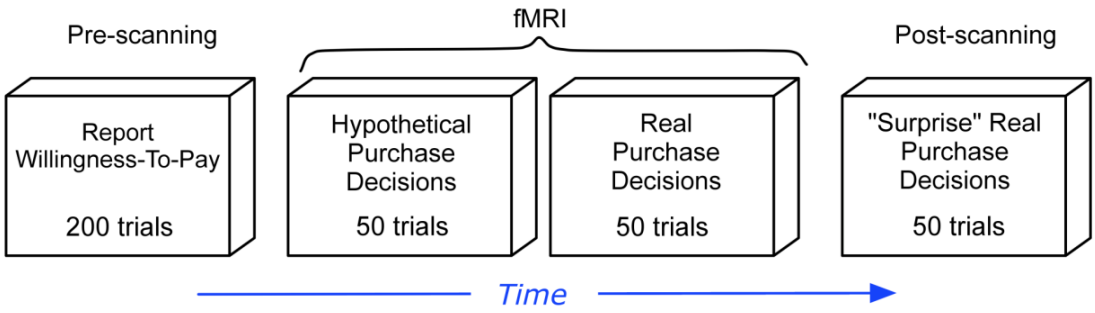

Figure 3.1 Timeline of the entire experiment.

Experimental Procedure and Task

At the start of the experiment, subjects were informed that they would receive a $60 participation fee upon the completion of the experiment. The experiment consisted of three parts—pre-scanning, scanning and post-scanning (figure 3.1). The scanning part itself had 2 decision-making tasks—one hypothetical and the other real. Subjects were only informed that there were 3 decision-making parts, and the detailed instructions for each part were not given until each part began.

In the pre-scanning part, subjects were shown images of 200 consumer products one at a time and asked to state WTP for each item in US dollars. Subjects learned that their reporting of the WTPs was a hypothetical exercise, and thus they were not bound to pay whatever they reported. When evaluating items, they were instructed to note the following points. First, the value of each item should be rated independently from the others, as if it were the only item. Second, the products should be evaluated from the subject's own perspective, not that of someone else. That is, the subjects should not consider an item for the purposes of resale or gifting to someone else. Third, the current ownership of a particular item might affect the rating of the item. For example, if they already own item X, their rating of item X might be high or low depending on whether they want a second one for himself or not. Products were presented in random order. In each trial, subjects were allowed to enter an amount between $0 to $50 using a sliding scale with a $1 increment (figure 3.2).

Figure 3.2 An example screen for the pre-scanning trials.

Each trial started with a 1-second-duration fixation cross and ended as soon as the subject entered their WTP, with the next trial beginning immediately. Upon completion of the

first part, the computer ranked products in a descending order of the subject's reported WTPs and then paired up each two adjacent products ({1^{st}, 2^{nd}}, {3^{rd}, 4^{th}}, …, {199^{th}, 200^{th}}). Then among 100 pairs it selected 50 pairs (the 17 pairs with the highest WTP, the 16 pairs with the medium WTP and the 17 pairs with the lowest WTP) and randomly chose one product of each pair and assigned it to the hypothetical decision-making block and the other to the real decision making block. This pre-selection ensured that the distributions of WTP in both blocks were comparable (see figure 3.4 in **Results**) and there were no repeated appearances of the same product in both blocks. The median WTP of the 200 items was determined as the price for the scanning part (table A3.2). This median price scheme ensured participants made a sufficient number of purchase decisions for statistical power (see figure 3.5 in **results**).

In the scanning part, subjects went through two types of purchase decision-making tasks having identical structure. Both tasks were the same except that the first block of decisions was hypothetical and the second was real, which is a crucial difference for this study. In each block, subjects went through 50 trials, and in each trial they were offered the chance to buy the product shown at a fixed price, which was determined by the median WTP from the first part. Fifty products for each block were pre-selected in the same way as described above; these products were then presented in random order within the block.

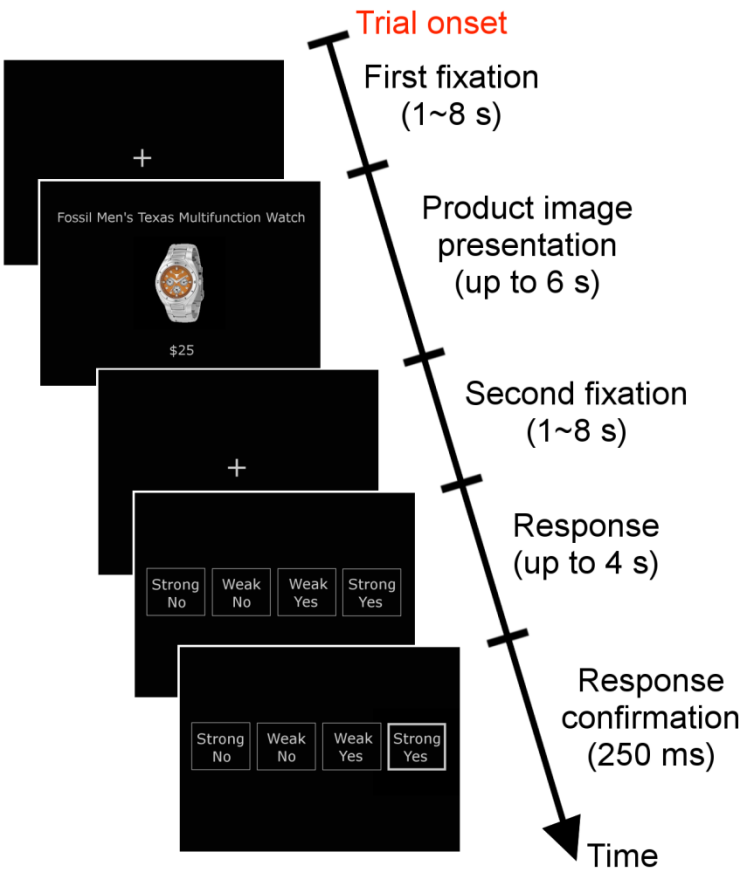

Figure 3.3 Time structure of an individual trial in the scanning part. The structure of the hypothetical and real trials was identical. There was no repetition of a product between the hypothetical and real trials and subjects were asked to specify their decision ('Yes' or 'No') as well as the confidence level ('Strong' or 'Weak') in the decision.

Figure 3.3 describes the time structure of each trial in the scanning part. Following the fixation cross, which lasted for 1 to 8 seconds, subjects were shown a consumer product image with the price below for up to 6 seconds. During the 6 seconds of image presentation they were asked to indicate when they had reached a decision by pressing any button. Either by this first button press or after 6 seconds had passed, the subjects proceeded to the next fixation screen, and then submitted their decision. The subjects were given up to 4 seconds to respond. When submitting a decision, subjects were

instructed to specify their level of confidence in the decision as well, resulting in 4 different response types—strong yes, weak yes, weak no, and strong no. Once a decision was entered, the selected response type flashed to confirm the submission and the subsequent trial began immediately. Subjects were given two button boxes, each of which had two buttons and was held by each hand. Buttons assigned to each response were counterbalanced across blocks and subjects.

In the hypothetical trials, the decisions made were purely hypothetical and did not count. However, in the real trials, subjects learned that one of the 50 trials in this block would be chosen at random and whatever decision they had made in the chosen trial would be implemented whether it be 'buy' or 'not buy.' The subjects were further informed that if the decision was made to 'buy', then the price of the item would be subtracted from their $60 participation fee, and they would receive the item as well as the remaining fee. Since only one trial counted as real, subjects did not have to worry about spreading a budget over the different items and they could treat each trial as if it were the only decision that counted; indeed, they were instructed to act as such. As in the pre-scanning part, subjects were instructed to evaluate each item for their own use and to make every decision independently of the others.

Upon finishing the fMRI part, subjects were asked to complete another task outside the scanner—a 'surprise' real block of decisions. The procedure was the same as in the scanning part except for shortened durations of the fixation screens (each of 500 ms now). The same 50 items presented in the hypothetical trials were shown again. This time, subjects were asked to make a real decision on these items. They were informed that exactly one trial out of the total of 100 real trials – 50 from the real trials and 50 from

this surprise real part—would be randomly selected and implemented based on their decision made in the selected trial. The purpose of this post-scanning part was to keep track of switches in hypothetical decisions for the matched set of items.

fMRI Data Acquisition

fMRI data were acquired on a Siemens 3-T TRIO MRI scanner (Erlangen, Germany). Blood-oxygenation-level-dependent (BOLD) contrast was measured with gradient echo $T2^*$-weighted echo-planar images. To optimize functional sensitivity in the orbitofrontal cortex, we used a tilted acquisition sequence at 30° to the AC-PC line (Deichmann et al. 2003) and an eight-channel phased array coil that yields a 40% signal increase in this area as compared to a standard coil. Each volume consisted of 32 axial slices in 3 mm thickness and 3 × 3 mm in-plane resolution; slices were collected in an interleaved ascending manner. The first 3 volumes in each session were discarded to permit T1 equilibration. Other imaging parameters were as follows: repetition time, TR=2000 ms; echo time, TE=30 ms; field of view, 192×192 mm. A high-resolution T1-weighted structural scan was acquired from each subject to facilitate localization and coregistration of functional data. Functional imaging data were acquired in two separate sessions of ~14 min each.

fMRI Data Preprocessing

fMRI data analysis was performed by using SPM5 (Wellcome Department Imaging Neuroscience, Institute of Neurology, London, UK). Images were corrected for slice acquisition time within each volume, motion corrected with alignment to the first volume, spatially normalized to the standard Montreal Neurological Institute EPI template, and spatially smoothed using a Gaussian kernel with a full width at half maximum of 8 mm.

Intensity normalization and high-pass temporal filtering (using a filter width of 128 s) were also applied to the data. The structural T1 images were coregistered to the mean functional EPI images for each subject and normalized using parameters derived from the EPI images. All regression models included six regressors indexing residual motion and two regressors for session baseline.

fMRI Data Analysis

1. Primary Model

Statistical analysis proceeded in three steps. First, we estimated a general linear model (GLM) with AR(1) and the following regressors that capture the main events in our experiment:

(*H1*) An indicator function denoting product image presentation in the hypothetical trial;
(*H2*) *H1* modulated by mDV;
(*H3*) *H1* modulated by an indicator function denoting decision to purchase in the given trial;
(*H4*) An indicator function denoting a first button press during product image presentation in the hypothetical trials;
(*H5*) An indicator function denoting response phase (from the start of the response screen up to the point a subject submits a decision) in the hypothetical trials;
(*R1*) An indicator function denoting product image presentation in the real trial;
(*R2*) *R1* modulated by mDV;
(*R3*) *R1* modulated by an indicator function denoting decision to purchase in the given trial;
(*R4*) An indicator function denoting a first button press during product image presentation in the real trials;
(*R5*) An indicator function denoting response phase in the real trials;

The regressors *H1*, *H5*, *R1*, and *R5* were modeled using box-car functions with subjects' response time as a duration. The regressors *H4* and *R4* were modeled using stick

functions. We orthogonalized *H4* and *R4* with respect to *H1* and *R1*, respectively. Each of the regressors was convolved with a canonical hemodynamic response function. We also included regressors of no interest.

After estimating a GLM, we then calculated the following 1st-level single-subject contrasts: (1) the real versus the hypothetical trial during image presentation (*R1*-*H1*), (2) image presentation in the hypothetical trials modulated by mDV (*H2*), (3) image presentation in the real trials modulated by mDV (*R2*), and (4) the real versus the hypothetical trials during image presentation modulated by mDV (*R2*-*H2*).

We then calculated 2nd-level group contrasts using a one-sample *t*-test. Anatomical localizations were then performed by overlaying the t maps on a normalized structural image averaged across subjects, and with reference to an anatomical atlas (Duvernoy 1999).

2. Supplementary Model

This model was estimated in order to compare non-switch hypothetical trials with real trials. It is identical to the primary model except that the hypothetical trials were sorted into switch and non-switch trials and modeled separately. The non-switch hypothetical trials are defined as the hypothetical trials whose decision remained the same in the surprise real trials. The switch trials are defined as the hypothetical trials in which a 'Yes' ('No') decision changed to 'No' ('Yes') in the surprise real trials.

The model included the following regressors of interest:

(*NS*-*H1*) An indicator function denoting product image presentation in the non-switch hypothetical trials;
(*NS*-*H2*) *NS*-*H1* modulated by mDV;

(*NS-H3*) *NS-H1* modulated by an indicator function denoting Yes choice in the given trial;

(*NS-H4*) An indicator function denoting a first button press during product image presentation in the non-switch hypothetical trials;

(*NS-H5*) An indicator function denoting response phase in the non-switch hypothetical trials;

(*S-H1*) An indicator function denoting product image presentation in the switch hypothetical trials;

(*S-H2*) *S-H1* modulated by mDV;

(*S-H3*) *S-H1* modulated by an indicator function denoting Yes choice in the given trial;

(*S-H4*) An indicator function denoting a first button press during product image presentation in the switch hypothetical trials;

(*S-H5*) An indicator function denoting response phase in the switch hypothetical trials;

(*R1*) An indicator function denoting product image presentation in the real trials;

(*R2*) *R1* modulated by mDV;

(*R3*) *R1* modulated by an indicator function denoting Yes choice in the given trial;

(*R4*) An indicator function denoting a first button press during product image presentation in the real trials;

(*R5*) An indicator function denoting response phase in the switch real trials;

The regressors *NS-H1*, *NS-H5*, *S-H1*, *S-H5*, *R1* and *R5* were modeled using box-car functions with subjects' response time as a duration. The regressors *NS-H4*, *S-H4* and *R4* were modeled using stick functions. We orthogonalized *NS-H4*, *S-H4* and *R4* with respect to *NS-H1*, *S-H1* and *R1*, respectively. The regressors of interest were convolved with a canonical hemodynamic response function. The model also included motion parameters and three regressors for session baseline as regressors of no interest.

The 1st-level single-subject and 2nd-level group contrasts were created as in the primary model.

3. Psychophysiological interactions (PPIs)

Psychophysiological interaction analysis (Friston et al. 1997) was performed to test the hypothesis that functional connectivity between the ACC and the mOFC would differ between the two types of trials, real and hypothetical. The analyses proceed in three steps.

First, we extracted individual average time series of a region of interest (ROI). The ROI was defined as a 4 mm sphere surrounding individual subject peaks within the functional mask of the ACC shown in figure 3.12 (in orange, left panel and figure 3.16A). Individual subject peaks within the ACC mask were identified based on the areas having the highest Z-values in the Real*mDV – Hyp*mDV contrast. Variance associated with the six motion regressors was removed from the extracted time-series. The seed time course was then deconvolved, using the canonical hemodynamic response, in order to estimate the underlying neuronal activity in the ACC (Gitelman et al. 2003).

Second, we estimated a GLM with the following regressors:

(PPI-Reg1) An interaction between the neural activity in the seed region and a trial-type variable (Real trials coded as 1 and hypothetical trials as -1);
(PPI-Reg2) A trial-type variable;
(PPI-Reg3) The original BOLD eigenvariate (within the 4 mm sphere).

The first two regressors were modeled as box-car functions with a duration that is equal to subjects' response time during product image presentation in the given trial. These regressors were also convolved with a canonical hemodynamic response function. The model also included motion parameters as regressors of no interest. Note that the first regressor identifies areas that exhibit task-related functional connectivity with the ACC;

specifically, it identifies areas in which the correlation in BOLD activity with the ACC increases differentially during real trials (as compared to hypothetical trials).

Third, single subject contrasts for the first regressor were calculated, and then a second level analysis was performed by calculating a one-sample *t*-test on the single-subject contrast coefficients. The results of the second level contrast at $p<0.001$, uncorrected and with a 15-voxel extent threshold are depicted in figure 3.11. As shown in table 3.8, this PPI analysis also identified other regions in which activity correlated differentially with the ACC during the real trials in which subjects made real purchase decisions (as compared to hypothetical trials).

Behavioral Results

Bias in Hypothetical Decisions

Given the procedure for selecting products for the scanning part, the distributions of WTP and decision value for the hypothetical and the real trials are almost identical (figure 3.4, table 3.1). Therefore, if there is no hypothetical bias (overstatement of WTP in hypothetical trials), there should be similar product purchase rates in all three types of trials. However, as found in previous economic studies (Blumenschein et al. 2007; Cummings et al. 1995; Johannesson et al. 1998; List 2001; List & Shogren 1998; Little & Berrens 2004; Murphy et al. 2005), subjects did show a significant pro-purchase bias during hypothetical decisions (figure 3.5). Subjects purchased 53% ± 1.74% (mean ± S.E.M.) of the products they saw in the hypothetical trials, a significantly higher rate than 38% ± 3.10% and 40% ± 3.15% purchases in the real and surprise real trials, respectively (paired t-tests, $t(24)=4.17$, $p< 0.0001$ and $t(24)=3.61$, $p< 0.0008$ versus real and surprise real, respectively).

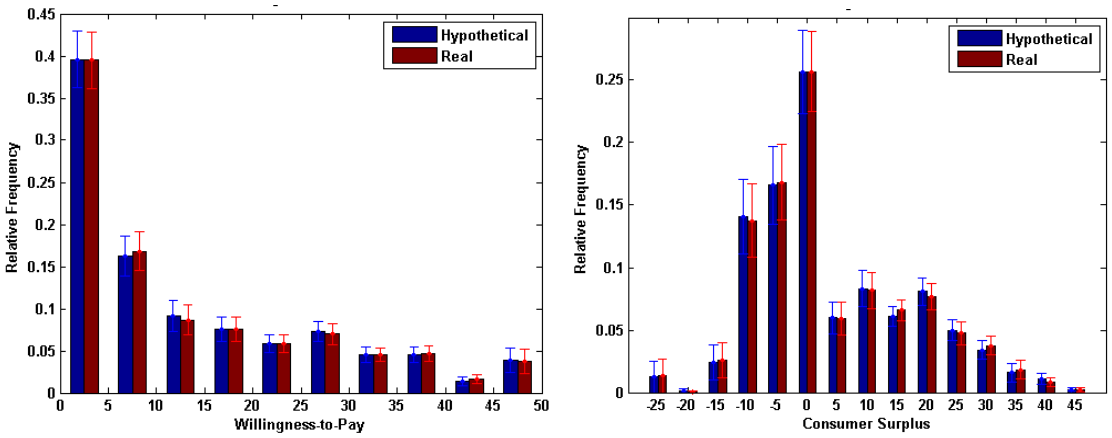

Figure 3.4 Histogram of willingness to pay (left panel) and decision value (right panel) by trial type. The size of a bin is 5.

Table 3.1 Descriptive statistics for willingness to pay, decision value by trial type

WTP	Mean	St. dev.	Min	Max	1st quartile	Median	3rd quartile
Hypothetical	14.33	14.02	0	50	3	10	25
Real	14.34	14.06	0	50	3	10	25

DV	Mean	St. dev.	Min	Max	1st quartile	Median	3rd quartile
Hypothetical	3.91	13.19	-27	48	-6	0	13
Real	3.92	13.23	-27	48	-6	0	13

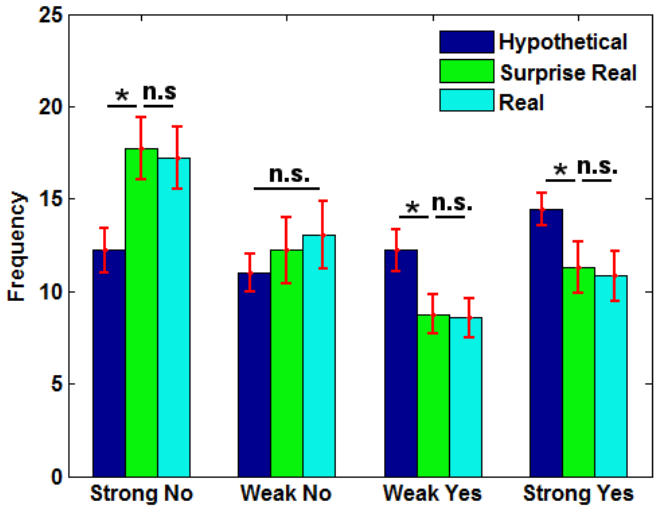

Figure 3.5 Choices show a hypothetical "yes" bias. Choice data averaged across all 24 subjects. Subjects made more 'yes' purchase decisions in hypothetical trials (mean frequency ± S.E.M.). *p<0.005, n.s. = not significant.

For a logistic regression analysis of estimated DV and choice probability, we pooled the data from the hypothetical/real/surprise real trials and incorporated dummy variables for the hypothetical and the surprise real trials. Hence the logistic regression model includes purchase decision as a dependent variable (1 if a decision is 'yes'; 0 otherwise), DV, Hypothetical Dummy (1 if a given trial is hypothetical; 0 otherwise), Surprise Real Dummy (1 if a given trial is surprise real; 0 otherwise) as independent variables with subject random effects. The results of a logistic regression analysis also show hypothetical bias (table 3.2, figure 3.6 top panel). By the definition of WTP, DV, and rational stochastic choice, a person should be more (less) likely to purchase a product when the DV of that product is high (low). At the true DV of zero, a person should be indifferent to buying and not buying, resulting in a purchase probability of 50%. Indeed, in the hypothetical trials purchase probabilities are close to 50% when the measured DV is close to zero (figure 3.6 top panel).

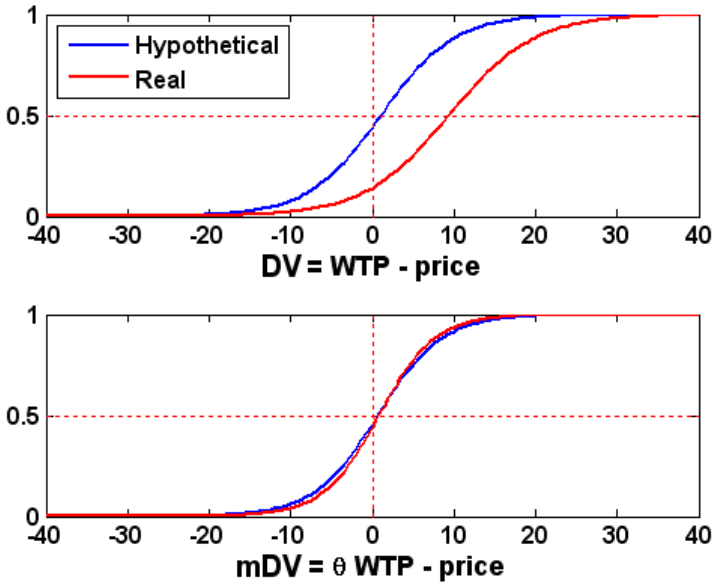

Figure 3.6 Correction of biases in DV. The curves show fitted probability of 'Yes' purchase decisions before correction (top) and after correction (bottom). Y-axis: probability of 'Yes' decision.

Table 3.2 Results of random-effects logistic regression analysis for hypothetical dummy specification[1]

	Coefficient	S.E.	z	*p*-value	95% Confidence Interval	
DV	0.182	0.006	29.14	0.00	[0.169,	0.194]
Hypothetical Dummy	**1.139**	0.120	9.51	**0.00**	[0.904,	1.373]
Surprise Real Dummy	0.066	0.119	0.55	0.58	[-0.167,	0.299]
Constant	-1.352	0.100	-13.49	0.00	[-1.548,	-1.156]

However, in the real trials, the likelihood of purchase is lower than in the hypothetical trials, controlling for measured DV using the initial pre-scanner WTP values, and is much less than 50% for zero DV (figure 3.6 top panel). We hypothesized that when making real decisions, subjects adjust the WTP that was casually stated in the hypothetical pre-scanner phase downward, which lowers DV and makes purchases less frequent.

The initial overstatement of value is probably a result of preference uncertainty— people are not sure precisely what they would pay—combined with a default tendency to think about the value of the product in an ideal world (leading to overestimates of product value) (Tanner & Carlson 2009). This hypothesis is supported by a separate behavioral study comparing hypothetical and real WTP for the same product (See Section 3.5). Those subjects tended to report hypothetical WTPs which were reduced by 60% when real WTP was elicited (figure 3.17, 3.18 and table 3.11).

Correction of Hypothetical Bias

To investigate the brain circuitry that encodes the implied decision values in hypothetical versus real decision making, we corrected for this apparent reduction in hypothetical WTP during real decision making. To do so, a discount factor θ^R is estimated for each

[1] Only the hypothetical dummy variable is significant—subjects' purchase behavior in the surprise real trials was not significantly different from that in the real trials. The trials with missed responses (purchase decisions) were removed, leaving 3590 trials, instead of 3600.

subject, which creates a modified DV (mDV) of $\theta^R \cdot WTP - price$. The value of θ^R was estimated for each subject by imposing the requirement that the fitted probability of purchase at the estimated mDV of zero is 50%. For comparability we also estimated subject-specific discount factors θ^H using the same requirement.[2]

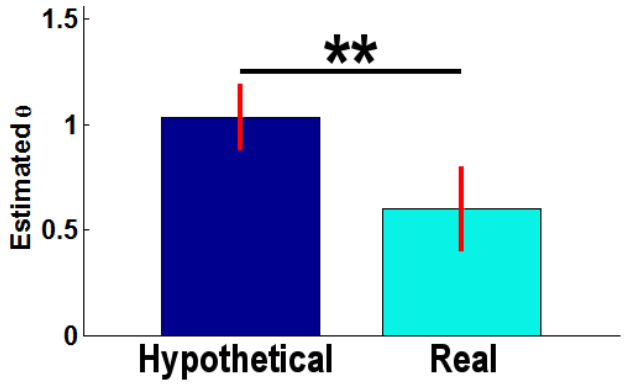

Figure 3.7 Median estimated θ for hypothetical and real trials. Signed rank test, $^{**}p<0.0003$. Error bars indicate standard errors.

Figure 3.6 (bottom panel) shows the smoothed choice probability curves using the modified DV. Figure 3.7 shows the median θ^H and θ^R for hypothetical and real trials, respectively (table 3.3). The median θ^H (1.03) is not significantly different from one, but the median θ^R (0.60) is significantly less than one (signed rank test, $p>0.5$ and $p<0.0015$, respectively). The median difference between θ^H and θ^R is also significantly different from zero (signed rank test, $p<0.0003$, figure 3.7).

This weighting pattern suggests that people act as if they are using the originally stated WTP when making later hypothetical decisions, but are adjusting that WTP when making

[2] Another conceivable way to modify DV is to keep WTP fixed, but to adjust prices by multiplying by a corrective factor θ. However, note that for hypothetical choices the unmodified DV, and unadjusted price, generate a reasonable choice probability fit. And to explain the low rate of purchase in the real condition would require multiplying the price by a factor greater than one. Assuming the perceived value of money is likely to be similar in real and hypothetical cases, the natural approach is therefore to assume that WTPs are adjusted downward ($\theta^R < 1$) in the real choice condition.

real purchase decisions. This adjustment creates two measures of mDV which are comparable in their decision implications across the hypothetical and real conditions.

Table 3.3 Estimated θ for modifying the decision value

Subject ID	θ^H	θ^R
1	1.70	0.49
2	0.74	0.78
3	0.88	0.78
4	3.71	0.20
5	0.42	0.38
6	1.27	1.21
7	0.96	0.68
8	1.10	0.93
9	0.53	0.25
10	1.05	0.57
11	1.09	0.64
12	1.74	4.87
13	1.27	1.75
14	2.76	1.08
15	1.02	0.40
16	0.25	0.16
17	1.21	0.45
18	0.89	0.48
19	0.84	0.35
20	1.12	0.63
21	0.95	0.77
22	0.95	0.39
23	1.06	0.65
24	1.01	0.44
Median	1.04	0.60
Mean	1.19	0.80
Std. Dev.	0.73	0.94

Note: $\theta^H = \theta$ for the hypothetical trials, $\theta^R = \theta$ for the real trials.

Decision times were also substantially faster in real choices than in hypothetical choices when mDV was positive (figure 3.8, $p<.008$); there was no difference in decision times when mDV was negative. The sensitivity of choice probability to modified DV is also greater, though modest, for real choices than for hypothetical ones (table 3.4, signed rank

test, one-sided $p<0.032$) The fact that real choices are faster and more sensitive to value input than hypothetical ones is behavioral evidence that neural processing may differ between real and hypothetical choices.

Table 3.4 Individual differences in sensitivity of choice probability to mDV[3]

Subject	β for Real trials	B for Hypothetical trials
2	0.113	0.164
3	0.185	0.279
4	2.580	1.588
5	0.598	0.398
6	0.101	0.121
7	0.333	0.348
8	0.143	0.077
9	0.573	0.317
10	1.020	0.617
11	0.736	0.828
12	0.082	0.106
13	0.248	0.196
14	0.472	0.233
15	0.383	1.268
16	0.708	0.654
17	1.528	0.457
18	0.311	0.334
19	0.935	0.380
20	0.458	0.565
21	1.088	0.127
22	0.267	0.369
23	0.647	0.468
24	0.390	0.275
Median	0.458	0.348
Mean	0.604	0.442
Std. Dev	0.564	0.367

[3] The following logistic regression model was estimated for hypothetical and real trials, at an individual subject level. Decision = $\alpha + \beta \times \text{mDV} + \varepsilon$, where decision is 1 if a subject chose to buy and 0 otherwise. β coefficient for real trials are larger than those for hypothetical trials (signed rank test, one-sided, *p*-value<0.032, *z*-value=1.860). Subject 1 was not included in the analysis since he used a cut-off rule in hypothetical trials and therefore his β coefficient for these trials could not be reliably estimated.

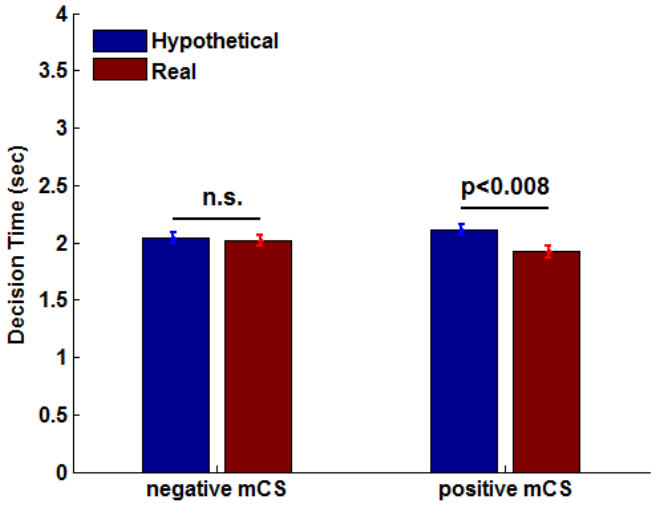

Figure 3.8 Average decision time in hypothetical and real trials by sign of mDV. Error bars indicate standard errors.

fMRI Results

Brain Areas with Higher Activation in Real Decision Making

We first identified brain regions that were generally more active during the real trials than during the hypothetical trials at the time of product image presentation. Stronger activation in this phase for real than for hypothetical choice was shown only in ventromedial prefrontal regions, including the medial orbitofrontal cortex (mOFC) (*x, y, z* = -3, 39, -9; *z* =3.74) and subgenual cingulate (*x, y, z* = 0, 12, -9; *z* =3.64) (table 3.5, figure 3.9). There was no significant activation in the opposite contrast at $p<.001$.

Table 3.5. Areas showing significant activity in the contrast of real versus hypothetical decision making

Region	Laterality	BA	MNI coordinates			Z	Voxels
Anterior Cingulate/ medial OFC	L	11/32	-3	39	-9	3.74	16
Subgenual/ Subcallosal Gyrus	L	25	0	12	-9	3.64	8
Frontal Lobe Sub-Gyral	R		27	30	15	3.49	13
Medial Frontal Gyrus	L	25	-12	36	-12	3.37	5

Height threshold, $T = 3.485$, $p<0.001$, uncorrected; extent threshold ,$k = 5$ voxels. L: left; R: right.

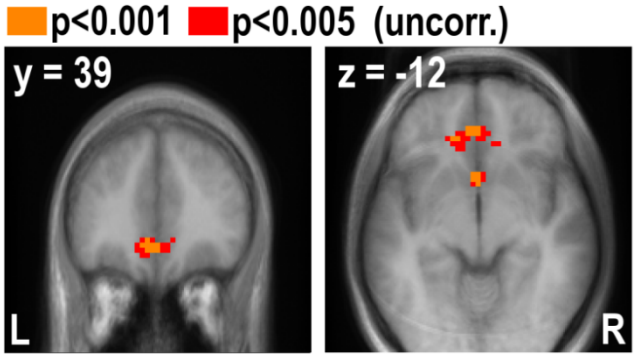

Figure 3.9 Brain areas more engaged in real versus hypothetical decision making.The mOFC is more actively engaged when real decisions are being made, compared to hypothetical decisions (contrast without decision value modulator: Real versus Hypothetical; $k \geq 5$ voxels).

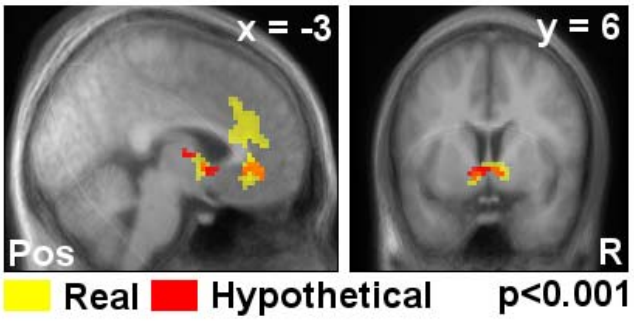

Figure 3.10 Modulation of mDV in real versus hypothetical decisions. (A) The brain encodes modified DV in the shared areas in real and hypothetical trials (overlay of contrasts Real*mDV and Hyp*mDV). The "real" areas are a superset of the "hypothetical" areas. $p<0.001$, uncorrected.

Modulation of Brain Activity by mDV

To investigate which brain areas are responsible for encoding decision value in the hypothetical and the real trials, we identified areas that are correlated with the mDV regressor in a GLM during the product image presentation phase, using the differently adjusted mDV values in the two types of trials. We found common regions, including mOFC and the ventral striatum (VStr) in both types of trials (figure 3.10, table 3.6 and 3.7). However, these areas active during hypothetical valuation are subsets of larger areas

active during real valuation, and the dorsal anterior cingulate cortex (ACC) is correlated with the mDV only in real trials (*x, y, z* = -6, 33, 12; *z* =4.30) (figure 3.10).

Table 3.6 Areas showing significant activity in the contrast of the parametric regressor for real decision making (real*mDV)

Region	Laterality	BA	MNI coordinates			Z	Voxels
Medial Frontal Gyrus/Medial OFC	L	10	-9	36	-6	4.95	223
Anterior Cingulate	R	32	9	33	-9	4.12	
Caudate Head/Putamen/Thalamus	R		12	15	-3	4.49	522
Sublobar	R		9	0	-9	4.30	
Anterior Cingulate/Corpus Callosum	L	32	-6	33	12	4.30	
Inferior Frontal Gyrus	L	13/47	-33	18	-12	4.47	107
Middle Frontal Gyrus	L		-24	48	3	3.55	7
Anterior Cigulate	R	32	12	27	30	3.37	9

Height threshold, $T = 3.485$, $p<0.001$, uncorrected; extent threshold, $k = 5$ voxels. L: left; R: right.

Table 3.7 Areas showing significant activity in the contrast of the parametric regressor for hypothetical decision making (hyp*mDV)

Region	Laterality	BA	MNI coordinates			Z	Voxels
Caudate/ Lentiform Nucleus			-6	3	6	4.12	54
Extranuclear	L		-12	-2	-12	3.75	
Anterior Cingulate	L	32	-3	30	-9	3.98	47
Medial Frontal Gyrus/Medial OFC	L	10/11	-6	39	-6	3.80	
Subcallosal Gyrus	L		-12	21	-15	3.17	
Thalamus	L		-3	-6	3	3.79	14
Medial Frontal Gyrus/Medial OFC	L	10	12	33	-12	3.40	6
Caudate Head	R		6	15	0	3.34	8

Height threshold, $T = 3.485$, $p<0.001$, uncorrected; extent threshold, $k = 5$ voxels. L: left; R: right.

A conjunction analysis (Nichols et al. 2005) between activities modulated by mDV in the real and the hypothetical trials confirmed that the mOFC (*x, y, z* = -6, 39, -6; *z* =3.8) and the VS (*x, y, z* = -3, 3, -6; *z* =3.66) are jointly recruited in both real and hypothetical trials (figure 3.11, table 3.8).

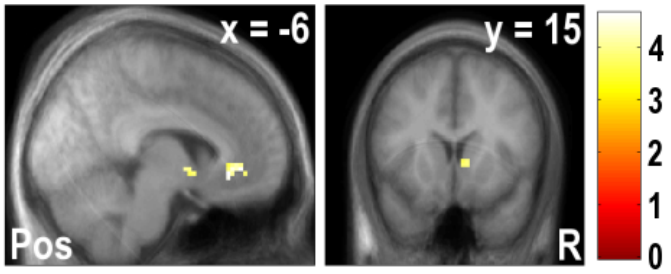

Figure 3.11 Conjunction analysis of Real*mDV and Hyp*mDV contrasts was performed to test the conjunction null hypothesis at a threshold $p<0.001$, uncorrected. The mOFC and VStr show conjointly significant activations modulated by mDV in both real and hypothetical decision making.

Table 3.8 Conjunction analysis: areas showing conjointly significant activations modulated by mDV in both real and hypothetical decision making

Region	Laterality	BA	MNI coordinates			Z	Voxels
Anterior Cingulate	L	32	-6	30	-9	3.86	43
Medial Frontal Gyrus	L	10	-6	39	-6	3.8	
Sub-lobar	L		-3	3	-6	3.66	28
Medial Frontal Gyrus/ Medial OFC	R	11	9	36	-12	3.35	6
Caudate Head	R		9	15	0	3.34	7

Height threshold, $T = 3.485$, $p<0.001$, uncorrected; extent threshold, $k = 5$ voxels. L: left; R: right.

The previous contrasts tell us that *common* areas are correlated with the decision value (mDV) in both the hypothetical and the real trials. We are also interested in *differences* in the sensitivity of those regions to mDV. Regions which are more sensitive to mDV in the real evaluation phase than in the hypothetical evaluation phase include the medial OFC and the ACC (x, y, z = -12, 33, -3; z =4.36), the caudate (x, y, z = 3, 3, 3; z =3.59), the inferior frontal gyrus (IFG) (x, y, z = -27, 9, -18; z =3.8) (figure 3.12, table 3.9). The region of ACC activation observed in the real trials at $p<.001$ is only active in hypothetical trials at $p<.05$. No areas are more sensitive to mDV in hypothetical decision making.

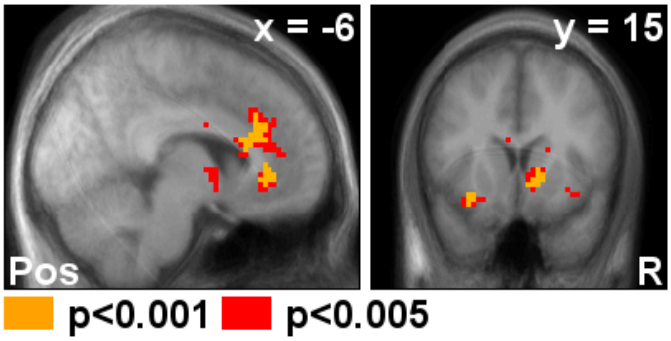

Figure 3.12 Significantly active areas in the [Real*mDV – Hyp*mDV] contrast. The mDV modulates mOFC and VStr activations in hypothetical decision making to a lesser extent. $k \geq 5$ voxels; orange: $p<0.001$, uncorrected; red: $p<0.005$, uncorrected.

Table 3.9 Areas showing activity in the difference of the parametric regressors (real*mDV – hyp*mDV)

Region	Laterality	BA	MNI coordinates			Z	Voxels
Anterior Cingulate/ media OFC	L	11/32	-12	33	-3	4.36	46
Caudate Head	R		12	15	-6	4.05	22
Corpus Callosum	R		12	24	12	3.81	8
Inferior Frontal Gyrus	L	47	-27	9	-18	3.8	39
Caudate Head	R		3	3	3	3.59	11
Middle Frontal Gyrus	L		-30	48	9	3.54	5
Anterior Cingulate	L	32	-6	30	24	3.53	42
Anterior Cingulate	R	32	6	36	24	3.22	8

Height threshold, $T = 3.485$, $p<0.001$, uncorrected; extent threshold, $k = 5$ voxels. L: left; R: right.

To further examine the difference in sensitivity to mDV, we conducted a region-of-interest (ROI) analysis on the common (overlapped) mOFC regions identified in figure 3.10 (mOFC area in orange). Parameter estimates (β's) from Real*mDV and Hyp*mDV contrasts were extracted within the overlapped mOFC region and averaged across all voxels in the region for each subject, and then the individual average β's were averaged across subjects (figure 3.13). While the mOFC activation increases as a function of mDV in both trial types, it shows significantly higher sensitivity to mDV in real trials.

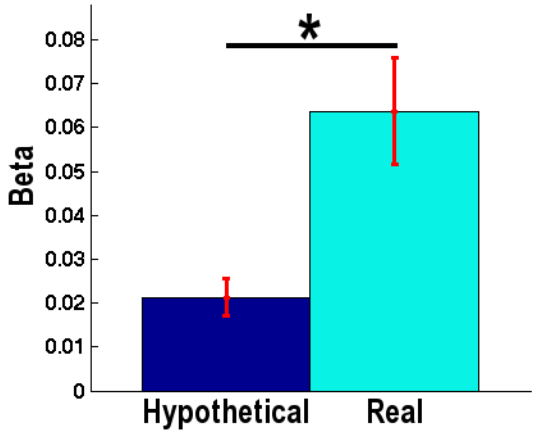

Figure 3.13 Subjects exhibited higher sensitivity to mDV in real trials compared to hypothetical trials (mean β ± SEM). *p<0.003, paired two sample t-test.

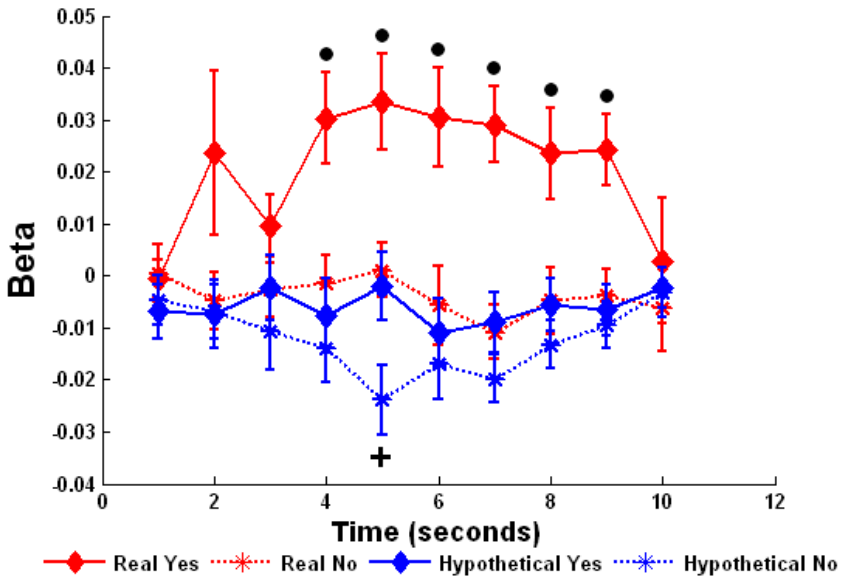

Figure 3.14 Time-course plots of the same overlapped mOFC region in figure 3.13 by response type and trial type. Circles (cross) indicate points in time where there is a significant difference between real and hypothetical yes (no) on paired t-test at $p < 0.001$.

We also extracted trial averaged time-course data from the same mOFC area in figure 3.13. This time-course graph was created by running another GLM using finite impulse response basis functions. First, trials were split according to the Yes/No purchase

decisions, resulting in Hypothetical Yes, Hypothetical No, Real Yes and Real No trials. These four types of trials were then modeled for ten 1 sec intervals after stimulus onset covering a total of 10 sec. This time course clearly shows the interaction of the response type with the trial type (figure 3.14). Although the mOFC shows a similar time course in response to both the hypothetical and the real No decision, it shows significantly higher activity in response to real Yes than hypothetical Yes.

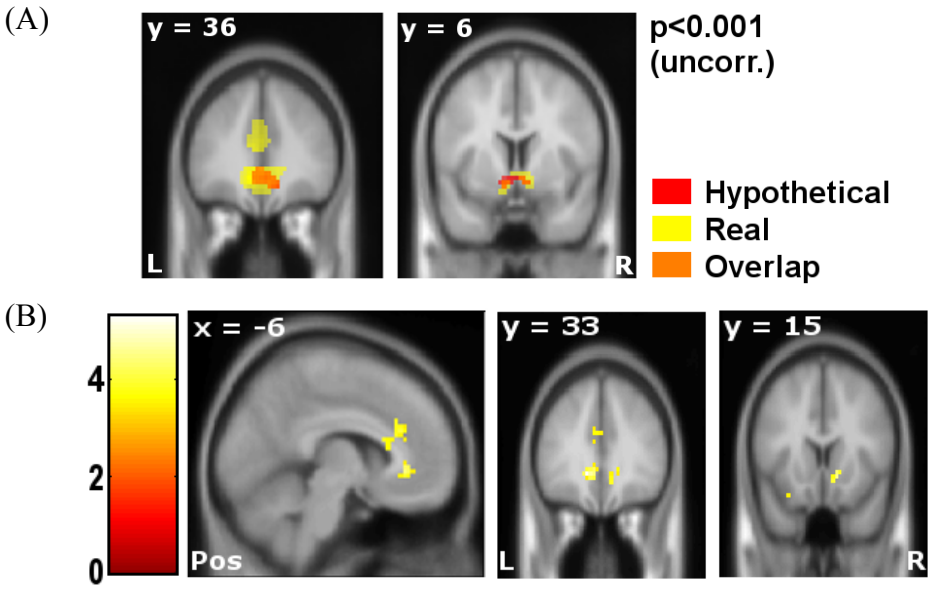

Figure 3.15 Supplementary model. (A) Overlapping regions of the contrast images, Real*mDV and Hyp*mDV. (B) The regions that are significant at $p<0.001$, uncorrected, in the contrast, Real*mDV – Hyp*mDV.

The analyses reported so far (*primary model*) cannot rule out the possibility that less strong activations in hypothetical trials are due to noise added by inconsistent hypothetical decisions that later changed in the surprise real trials. To rule out this possibility, we estimated an additional GLM to compare only non-switch (consistent) hypothetical trials with real trials (*supplementary model* in **experimental design and**

methods). The resulting contrasts are reported in figure 3.15 and we found little difference from the results reported above.

Task-Dependent Functional Connectivity between the ACC and the mOFC

The ACC activation in response to value appeared strongly only in real trials. ACC is often involved in value computation, action selection, and executive function (Botvinick et al. 1999; van Veen et al. 2001; Yeung et al. 2006). This suggests that the special contribution of ACC in real choice might be in adjusting hypothetically stated WTP toward an adjusted value used in real choices (for example, integration of opportunity costs), which could also modulate activity in mOFC. In fact, a psychophysiological interactions (PPI) analysis did show increased functional connectivity between ACC and mOFC in real trials compared to hypothetical trials (figure 3.16, table 3.10). This empirical finding suggests ACC is specially recruited during real choice to adjust or control value signals transmitted to mOFC.

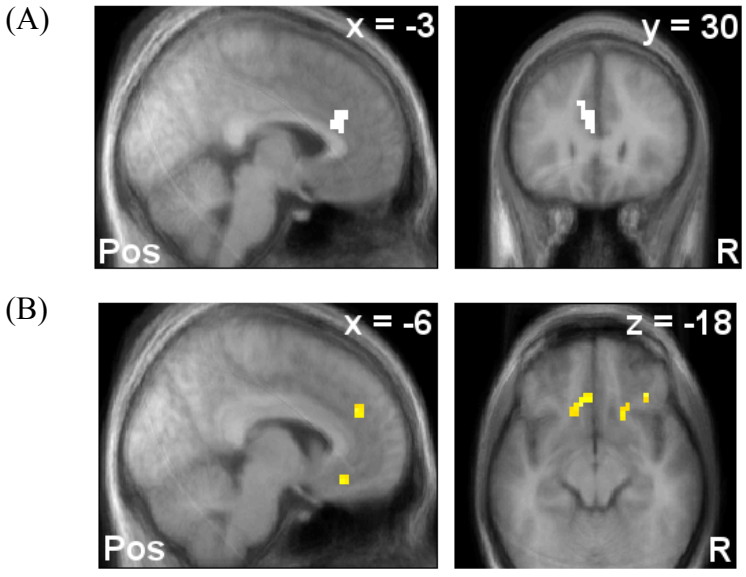

Figure 3.16 PPI analysis. (A) ACC mask used to identify individual seed regions. (B) The functional connectivity between the ACC and the mOFC

becomes enhanced selectively in real trials, not in hypothetical trials. $p<0.001$, uncorrected, extent threshold k=15 voxels.

Region	Laterality	BA	MNI coordinates			Z	Voxels
Middle Frontal Gyrus	L	10	-42	36	21	4.51	77
Inferior Frontal Gyrus	L	46	-39	30	9	3.87	
Medial Frontal Gyrus	R	9	6	39	27	4.10	62
Inferior Frontal Gyrus	L	47	-42	15	-3	4.00	34
Inferior Frontal Gyrus	R		24	24	-3	3.86	24
Superior Frontal Gyrus	L	10	-24	39	30	3.86	44
Caudate Head/ Subcallosal Gyrus	R	25	9	6	-12	3.83	32
Medial Frontal Gyrus/ Medial OFC	L	11/47	-12	21	-15	3.81	20
Inferior Frontal Gyrus	R	11/47	27	24	-21	3.72	18
Caudate Head/Putamen	L		-12	9	-3	3.64	28

Height threshold, T=3.485, p<0.001, uncorrected; extent threshold, k = 15 voxels. L: left; R: right.

Table 3.10 PPI analysis: Regions showing task related functional connectivity with the ACC (dfferentially for real > hypothetical choices).

3.3 Behavioral Experiment for Hypothetical and Real Willingness to Pay

Comparison of hypothetical and real purchase decisions in the fMRI study implies there is a gap between hypothetical WTP and real WTP for the same object—more specifically, hypothetical WTP is overstated compared to real WTP for the same object. We tested this hypothesis with a different group of subjects and report the results in this section.

Experimental Design and Methods

Participants, Stimuli, and Experimental Procedure

Eleven Caltech male students participated in the experiment (mean age = 20.55 ± 3.42 year; age range: 18–29). Subjects were paid $60 for participation. The experiment consisted of two parts of 200 trials each. In the first part, subjects were presented with the same 200 consumer products as in the pre-scanning part of the fMRI study and asked

to report hypothetical willingness to pay, ranging $0~$50, for each item. The second part was identical to the first part except that they were unexpectedly asked to report *real* WTP for the same 200 items. This design allowed us to directly compare hypothetical and real WTP for the same product.

To elicit real WTP in the second part, we employed a Becker-DeGroot-Marschak auction mechanism (Becker et al. 1964; Hare et al. 2008; Plassmann et al. 2007). The mechanism worked as follows: Subjects reported WTP for each of the 200 products. At the end of the experiment, one of the 200 trials from the real WTP part would be randomly selected. The computer generated a random integer between 0~50 (each integer over the interval was equally likely) and set it as a price for the item in the chosen trial. If the subject's reported WTP was greater than the randomly generated price, say $X, then the subject paid $X and got the item. Otherwise, the subject did not get the product and paid nothing. The optimal strategy for a subject under this mechanism is to report exactly what one is willing to pay for the item presented (Becker et al., 1964).[4] Since only one trial counted as real, subjects did not have to be concerned about spreading a $60 budget over the different items; they could treat each trial as if it were the only one that counted. These points were explained and emphasized during the instruction.

Results

Subjects exhibited a tendency to overestimate hypothetical WTP compared to real WTP. Figure 3.17 displays pairs of hypothetical and real WTP for the same product (blue dots).

[4] Since the price for an item is determined by the randomly generated number, subjects cannot affect the price paid and have no incentive to report more or less than the true WTP. Deflating WTP could lead to loss of an opportunity to buy a product at a lower price than they are actually willing to pay. In contrast, subjects could end up paying more than they are willing to if they inflate WTP.

If hypothetical and real WTPs were equivalent, the dots should be placed along the 45° line (dashed line in red). However, most of the pairs are under the 45° line, indicating that subjects reduced initial hypothetical WTP by approximately 60% when later asked to report real (binding) WTP for the identical object.

This pattern is also found at an individual subject level. Figure 3.18 shows individual subjects consistently overestimated WTP when the task was hypothetical, although there is individual variation in the tendency to over-report (table 3.11). These data clearly support the hypothesis that when making real choices, subjects downward-adjusted WTP, which was casually stated WTP in a hypothetical situation.

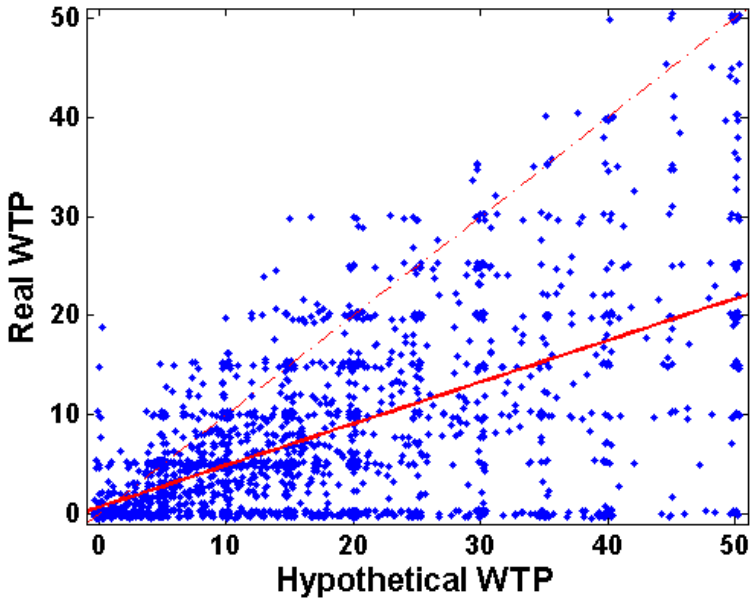

Figure 3.17 Comparison of hypothetical and real WTP for the identical object. In the figure, real and hypothetical WTP are both jittered by adding uniform random numbers over the interval [-0.4, 0.4] to convey data density. The solid line is a fitted regression of real WTP on hypothetical WTP (Real WTP = 0.67 + 0.42(Hyp WTP), R^2 = 0.40).

Table 3.11 Individual differences in hypothetical bias.

Subject	β	α
1	0.698	1.706
2	0.386	0.173
3	0.529	-1.834
4	0.783	-0.065
5	0.660	0.503
6	0.273	-0.320
7	0.496	0.502
8	0.185	-0.441
9	0.688	-0.961
10	0.119	-0.650
11	0.351	-1.103
Mean	0.470	-0.226
Median	0.496	-0.320

Note: The following regression model was estimated at an individual subject level: *Real WTP* = α + β × *Hypothetical WTP* + ε. The estimated parameters are reported in the table below. Subjects' real WTP for a product is substantially reduced compared to the hypothetical WTP, but the degree of reduction varies across subjects

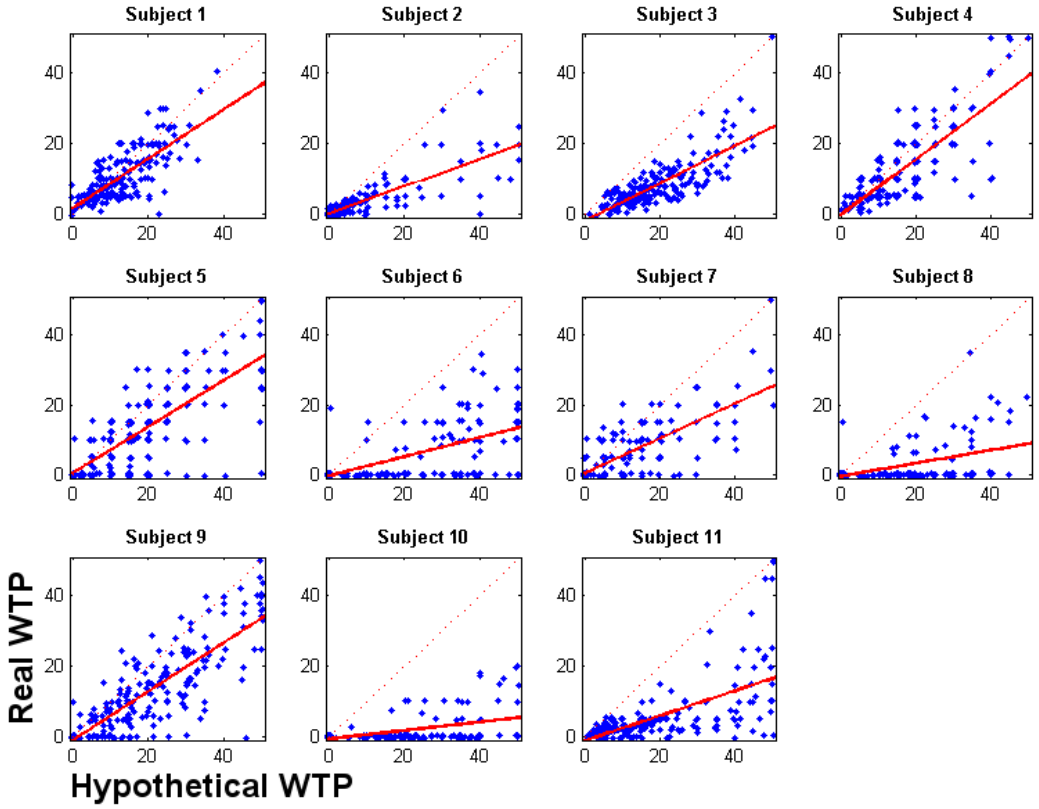

Figure 3.18 Subjects consistently over-reported hypothetical WTP. Blue dots indicate hypothetical and real WTP for the same product (x-axis: hypothetical WTP, y-axis: real WTP). The figures are jittered as in figure 3.17. Solid red lines are fitted regression lines of real WTP on hypothetical WTP. The slopes of the solid lines are reported in table 3.11.

3.4 Behavioral Experiment for Ordering Effects

In the fMRI study, the hypothetical-then-real order was deliberately *not* counterbalanced for the following reasons. First, it is possible that there *might* be an ordering effect in which thinking about real choices first would spill over to affect valuation in hypothetical tasks. Identifying such an effect, however, would require counterbalancing and substantially more subjects. Second, we also estimated the primary GLM in the main text with a trial number included as an additional regressor and then repeated the analysis; there was little difference in results. Any plausible model of ordering effects should show a within-block effect of trials (due to practice effects, for example), so any such effects are controlled for in the GLM as modified above. Third, hypothetical decision followed by real decision is a natural order for forecasting purposes. In most applications, hypothetical decision data are gathered in advance of real decisions—e.g., polls are only useful before elections. Hence, this particular order of events (hypothetical followed by real decision) is of the most interest. Finally, we tested if opposite real-then-hypothetical order would vanish hypothetical bias in a separate behavioral study outside of the scanner and confirmed that the trial type order has little effect on behavior. In the following section, this behavioral experiment for testing ordering effects is briefly described and the results are reported.

Experimental Design and Methods

Participants, Stimuli and Experimental Procedure

Fifteen Caltech students participated in the behavioral experiment (all male; mean ± st. dev. = 20.5 ± 2.97 year; age range = 17–27). Four additional subjects initially participated in the experiment, but upon completion of the first part, their price (median WTP) was $0. Hence, they were not asked to continue.

The experimental procedure, task and stimuli were identical except that the order of hypothetical and real blocks in the second part had changed. In the first part, participants stated willingness to pay for each of 200 consumer product as in the fMRI experiment. In the second part, however, they first went through 50 real purchase decision-making trials and then 50 hypothetical decision-making trials; this reversal in hypothetical and real blocks allowed us to test if hypothetical bias is solely due to "hypothetical-then-real" ordering effects. Lastly, they were unexpectedly showed the products presented in the hypothetical trials and asked to make 'surprise' real decisions on these items. Unlike the fMRI experiment, all parts of the behavioral experiment were conducted outside the scanner and subjects were run in a batch (multiple subjects at a time).

Results

The participants in this behavioral study also showed hypothetical bias—significantly higher purchase percentage in the hypothetical trials than the real/surprise real trials (figure 3.19).

However, the overall purchase percentage in the behavioral study is approximately 10% lower for all types of trials than the fMRI study. This overall decline might be due to

session variability such as: (1) the statistical variance of price was larger for the fMRI subjects (35.21 compared to 12.11 of the behavioral subjects); (2) the behavioral subjects might not have been from the same subject pool as the fMRI subjects—the subjects in the behavioral experiment might have behavioral characteristics that are different from the fMRI subjects (such as a risk attitude—some subjects only participate in behavioral experiments due to a fear of the fMRI technique); (3) unlike the fMRI study, multiple subjects (about 10) were run simultaneously.

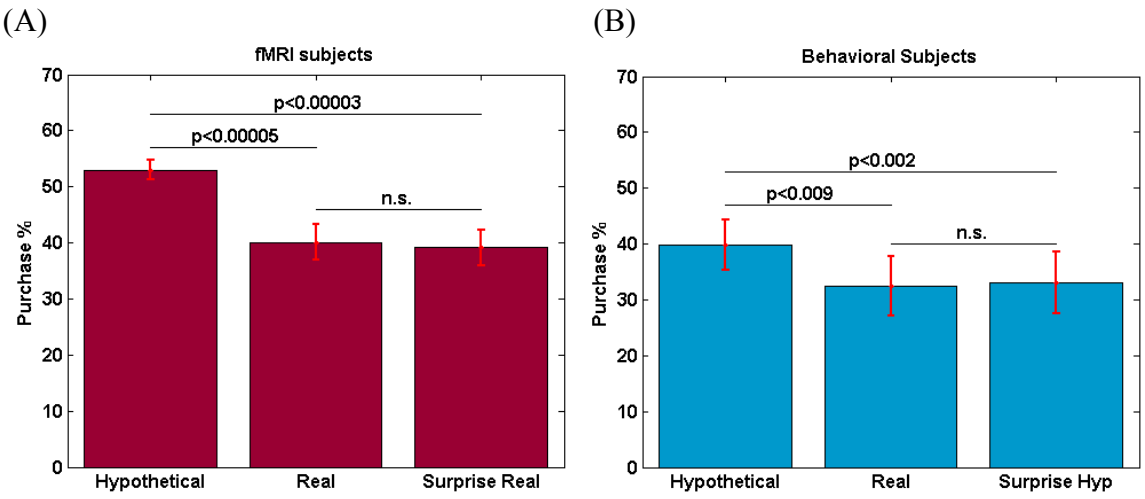

Figure 3.19 Regardless of the order of the trial types, subjects exhibit hypothetical bias (paired two sample t-tests, two-sided). (A) Purchase percentage by the participants in the fMRI study. (B) Purchase percentage by the participants in the behavioral study of ordering effects.

3.5 Discussion and Conclusions

Our study was motivated by the fact that hypothetical statements of intended choice are often different from real choices. In commercial marketing research, for example, self-reported intentions to buy goods are one of the most widely used measures to forecast sales of existing products, to test and plan the launch of new products (Chandon et al.

2004; Infosino 1986; Jamieson & Bass 1989; Silk & Urban 1978; Urban et al. 1983), and to assess the effectiveness of marketing programs (Raghubir & Greenleaf 2006; Schlosser et al. 2006). However, the fact that so many new products test well and then sell poorly suggests that forecasts based on reported intentions are imperfect. In the public policy arena and markets for public goods (e.g., clean air, foreign policy), data on what people actually choose are limited. There are no special markets for these goods, so voters' preferences for them are often coarsely inferred from reactions to political positions. Therefore, public economists and health economists often try to measure actual preferences from hypothetical responses to survey questions about possible actions (Carson et al. 1996; Diamond & Hausman 1994; Mortimer & Segal 2008). In political science, polls provide valuable information about voters' perceptions and preferences (Crespi 1989), but pre-election questions about voting intentions are always hypothetical. The newest research treats survey responses as deliberate products of a psychological and economic process by respondents, in order to understand their likely biases better and improve surveys (McFadden et al. 2005).

Furthermore, even actual choices often have hypothetical future choices embedded in them. For example, a student might make a real choice of a college to attend, planning to major in pre-medical studies. The planned major is an intended or a hypothetical future choice because it can be easily reversed.

The major behavioral result of our study is that subjects did exhibit a tendency to say "Yes" to questions about purchasing consumer goods more often in hypothetical choices than in real choices. This "yes bias" is present when the real choices are different goods and when they are the same goods, and is unaffected by counterbalancing the condition

order (see Section 3.6). The bias can be characterized parametrically by the hypothesis that an initially stated WTP (elicited before the fMRI scanning) is reduced during real choice by a multiplier θ^R, which is typically less than one (a median of .60 in these data). Intuitively, subjects change their minds and make real choices as if they will pay about 40% less than they initially said they might. When the initial WTP is adjusted to fit actual choices, subjects' choice frequencies then appear to be a smooth logistic function of the modified decision value, thati is, $\theta^R \cdot WTP - price$.

FMRI compares differential activity in response to mDV during presentation of the product image—the time at which we assume valuation occurs—in real trials and in hypothetical trials. Three ROIs emerge in this comparison: mOFC, VStr, and ACC. Regions active in hypothetical choice in mOFC and VStr are subsets of the regions active in real choice. ACC is highly significantly active in response to mDV only in real choice.

Medial OFC activity is correlated with decision value in both hypothetical and real decision making. Studies have found that the OFC, especially the medial part, is responsible for encoding the reward value of food, pleasant smells, attractive faces, and abstract rewards such as money or avoiding an aversive outcome (Anderson et al. 2003; Gottfried et al. 2003; Kim et al. 2006; O'Doherty et al. 2001; O'Doherty et al. 2003; Small et al. 2001). Our study uses consumer objects that have many product features and therefore require substantial abstract processing (including integration of the monetary price) which expands the scope of rewards that the mOFC appears to encode.

The medial OFC region identified in this study is similar to areas found to encode economic decision values in previous studies (Hare et al. 2008; Knutson et al. 2007;

Plassmann et al. 2007). Interestingly, activity in mOFC is stronger (more sensitive to decision value) and also more spatially widespread in real choice than in hypothetical choice. It is quite possible that the hypothetical bias might be related to an enhanced value signal in the brain, but the heightened mOFC activity during real choice casts doubt on this possibility.

The VStr activation in response to mDV could also be encoding decision value, as shown in many studies (Hariri et al. 2006; Kable & Glimpcher 2007; McClure et al. 2007; Yacubian et al. 2006). Or it could be encoding a prediction error—the deviation of a new reward from a predicted reward (in our case, the mDV of a good compared to recent mDVs), as suggested by a recent study which separates values and prediction errors (Hare et al. 2008). If a stronger VStr response reflects reaction to prediction error, then learning about values could be more rapid when choices are real, which is a new implication of these results that could be tested in future studies.

An important open question is what neural activity is actually creating the "yes bias." Linking the stated object WTP dollar values to later choices suggests the bias is due to overstatement of object values in the pre-scanning phase. This overstatement may be generated by preference uncertainty or ignoring other alternatives and opportunity costs—that is, evaluating an option in isolation from others in an ideal world (Tanner & Carlson 2009). Our data show that in hypothetical choices, those overstated values seem to be used (since the discounting factor θ^H is around 1), but in real choices they are adjusted downward by the discount factor θ^R. This discount factor θ^R is .60 in the fMRI study and is comparably less than one, .40 in a separate behavioral study.

A candidate region for implementation of this adjustment of stated WTP is the ACC. The ACC is only activated strongly, in response to mDV, in real decision making. Furthermore, functional connectivity between the ACC and the mOFC is stronger in real trials as compared to hypothetical trials. In addition, in real choice, the probabilities of purchase are more sensitive to modified value and response times are faster (for "Yes" decisions); these are possible behavioral manifestations of additional activity in ACC and mOFC. These results are consistent with the hypothesis that ACC is implementing a stronger reduction of casually-expressed WTP, to create an adjusted mDV that guides decisions, when real money is on the line. This possibility is consistent with the established role of ACC in executive function, cognitive control, and Stroop tasks which require effortful overrides of highly learned automatic responses (Botvinick et al. 1999; van Veen et al. 2001; Yeung et al. 2006). While we have not found any cross-subject differences in ACC activity which are related to the degree of WTP adjustment θ^R, studies aimed more closely at discovering such relations are promising.

One potential way to use fMRI and other measures is to classify trial-by-trial hypothetical "Yes" choices into those which are likely to switch to real "No" choices and those which are likely to stick as real "Yes" choices. Economists have tested some simple behavioral adjustments to estimate the overall switch minus stick rate, but there is no ideal technique that works reliably (Cummings & Taylor 1999; Knoepfle et al. 2009; List 2001). Much as in lie detection (Ben-Shakhar & Elaad 2003; Davatzikosa et al. 2005; Gamer et al. 2007), reading hidden intentions (Haynes et al. 2007) and other domains (Haxby et al. 2001; Mitchell et al. 2008), it is possible that neural activation could provide another basis for some kind of low-level "mind-reading" of this sort. Such an application could

conceivably improve high-stakes forecasts for political polling, product design, and personal commitment substantially. Unfortunately, there are not enough switches in our data to permit a powerful classification based on neural activity of switches versus sticks (table A3.3) as in Grosenick et al. (Grosenick et al. 2008). However, a simple classification based on mDV and RTs predicts hypothetical-to-real switches and sticks with about 70% out-of-sample accuracy compared to a normalized 50% baseline—namely, when mDV is low and responses are fast, hypothetical "Yes" decisions are likely to turn into real "No" decisions. (table A3.4). This baseline might be improved by more targeted studies using fMRI and other methods.

Finally, our study has an important methodological implication for scientific practice. In many experiments, especially in psychology and neuroscience, it is common to elicit hypothetical choices or ask hypothetical questions which cannot be actually implemented for practical reasons (e.g., for very high stakes, payments with long delays, creating unusual highly controlled social events, or morally charged consequences (Delgado et al. 2005; Greene et al. 2004; Greene et al. 2001; Hariri et al. 2006; Monterosso et al. 2007; Takahashi et al. 2009)).

Generalizing claims about neural processing based on hypothetical choices to real choice assumes that the neural processes in hypothetical and real choice are highly overlapping. Fortunately, our study shows that this overlap is largely present, at least in the narrow domain of consumer goods purchase. That is, suppose our study had used only hypothetical choices, concluded that mOFC and VStr are encoding decision value of consumer goods, and then guessed that the same regions would be active in response to valuation during real choices of consumer goods. Our results show that this guess would

have been correct (while understating signal strength and spatial extent). However, such a study would have failed to show that ACC responds to decision value, and might have concluded, incorrectly, that ACC plays no role in real choice.

As we have noted, for many phenomena in natural and social sciences, collecting hypothetical choice data is all we can do, even though the goal of collecting those data is to understand and predict real choices. Further studies should therefore continue to explore both similarities and differences in hypothetical and real choices, in different choice domains and with an eye to interesting applications.

3.6 Appendix for Chapter 3

Table A3.1 List of 200 Consumer Products Used in the Study. The retail price of the products ranged from $20 to $100. Approximately 300 products were initially collected by the experimenters and about 100 of the most unpopular items were screened out by a pilot study, leaving 200 products used in the study

Product Name	Product Name
90's Decade Box Gift Basket - Classic 90's Candy	Logitech Quickcam for Notebooks Deluxe
Abercrombie and Fitch Boxers - Set of 2	Lomographic Fisheye Number 2 Camera
Abercrombie and Fitch Canvas And Rubber Flip Flops	M-51 Engineers Field Bag - Military Style
Abercrombie and Fitch Cap New York	Math Formula Tie
Abercrombie and Fitch Kilburn Low Rise Boot Jean	Maxtor−OneTouch 4 Mini 120GB External Hard Drive
Abercrombie and Fitch Short Sleeve Tee Hopkins Trail	Megatech Hydro-Fly 18" Radio Controlled Land/Air/Sea Vehicle
Accudart Classic Bristle Dartboard	Microsoft Natural Ergo Keyboard
Adidas 3-Stripes Dazzle Tearaway Mens' Warm-up Pants	Millafleur Glass Electric Guitar Lamp
AeroBed® Sport Minute Twin airbed	Monopoly Deluxe Edition
Age of Empires III-Windows	Motorized Bumper Boat
Allied 180-pc. Household Tool Set	Motorola 25 Mile, 22 Channel 2-Way Radios (Pair)
Andy Mac Zon Complete Skateboard	Mrs. Fields Soho 12 Cookie Box
Antworks - Space Age Ant Habitat	Nikon Action 8x40 Binocular
Apple 1 GB iPod Shuffle AAC/MP3 Player	Nintendo Game Boy Advance
Arrested Development− The Complete Series (Seasons 1, 2, 3) (2003)	Nintendo Game Boy Micro
Audio X Rocker with built-in subwoofer and speakers	Northface Campus Backpack
Austin Bazaar 38" Black Guitar with Carrying Bag and Accessories	Northface M Amp Hybrid Jacket
Bach Tie	Northface Utility Waist–Sport Hiker
Banner American Desktop Work Station	Number Pad Mouse
Black and Decker 12 cup SmartBrew plus Coffeemaker	Obus Forme Memory Foam Travel Pillow
Black and Decker 15.6 Volt HEPA Cyclonic DustBuster	Olympus 128 MB Digital Voice Recorder with PC Link
Black and Decker 4-Slice Stainless Steel Toaster Oven	Omron Digital Premium Pedometer
Boss Seating Fabric Adjustable Task Chair	Oral B Professional Care Power Toothbrush
Braun Electric Water Kettle	Oregon Scientific ExactSet Projection Clock with Cable-Free Weather Forecaster
Braun Smart Control 3 Shaver	Organic Nut Harvest Basket
Brookstone 12-Language Translater	Oster Egg Cooker
Brookstone 9-in-1 Multi Tool with LED	Pail of Treats−Mrs. Fields 48 bite-sizes cookies and 18 brownie bites

Brookstone Digital Camera Keychain	Panasonic Upper Arm Blood Pressure Monitor
Brookstone Laptop Essentials Kit	Periodic Table Shower Curtain
Brookstone Lighted Lap Desk	Perpetual Calendar
Brookstone MP3 Stereo Speaker	Philips Norelco Rechargeable Razor
Brookstone Rechargeable Booklight	Philips Progressive-scan DVD Player
Brookstone Wireless FM Transmitter	Planet Earth−The Complete BBC Series (DVD)
Brother Personal Labeler	Plantronics Wireless Bluetooth Headset
Bushnell ImageView 7x18 VGA Pocket Digital Camera Binocular	PowerBar® Performance The Original Energy Bar, Peanut Butter (Pack of 24)
Calvin Klein− Men's Bamboo Dress Socks 2 Set of 3-Pack (6 pairs)	Premium Diamond Suited Poker Chip Set
Canon Color Image Scanner	Prodikeys MIDI Keyboard
Casio 2.7" Shock Resistant Portable LCD TV	Puma−Men's Quarter Crew Socks 2 Set of 3-Pack (6 pairs)
Casio Mens G-Shock Classic Watch	Ravensburger Oceanic Wonders 3000 Piece Jigsaw Puzzle
Celestron ExploraScope 80mm Reflector Telescope	Raytek MiniTemp No-Contact Thermometer with Laser Sighting
Classic Pillow by Tempur Pedic	Razor A Kick Scooter
Clif Bar Nutrition Bars, Variety Pack 2.4-Ounce Bars (Pack of 24)	RCA Small Wonder Digital Next Generation Camcorder
Coby 7" Widescreen Digital Photo Frame with MP3 Player	Real Theater Popcorn Kit - 20 Pack
Coca-Cola Personal Fridge	Roll-Up Keyboard
Columbia Sportswear Men's Trail Hiking Boot	Roller Slide−Ab Slider with Computer LCD Read-out
Columbia Sportswear's Men's Steens Mountain Sweater	SanDisk Sansa 2GB Photo MP3 Player
Columbia The Zone 0° Mummy Sleeping Bag	See's Candies−Toffee-Ettes and Almond Royal
Cooler Master NotePal Notebook Cooler	Seinfeld−Seasons 1 and 2 (1993)
Cranium (Board game)	Sharper Image All-In-One Games Set - chess, checkers, backgammon, and poker
Deluxe Turntable Scrabble	Sharper Image Atomic Projection Clock with Indoor Temperature
Desktop Organizer, 5/8" Wood, Corner Radius, Three-Way Use	Sharper Image Bright As Day Power-Port Desk Lamp
Desktop Personal Air Conditioner	Sharper Image Digital 130X USB Microscope Camera
Digital WiFi Detector	Sharper Image Sports Duffle Bag
Dymo LabelWriter 400 Label Printer	Sharper Image Turbo Ear and Nose Hair Trimmer
Fold-Away Wooden Chess Set	Sigma PC 15 Heart Rate Monitor
Fossil Men's Texas Multifunction Watch	Smallville−The Complete First Season (2001)
Fossil Mens Rimless Sunglasses - Rider - Polarized	SolidTek DigiMemo 6"x9" Digital Notepad with Memory
Gaiam Abs Ball Workout Kit	Sony CD/Cassette Portable Boombox
Garmin eTrex Handheld GPS Navigator	Sony DJ Style Monitor Series Headphones
GODIVA 1lb Milk Chocolate Dipped Pretzel Canister	Sony Micro Vault Tiny GB USB Flash Drive with

	Virtual Expander
GODIVA 50 pc. Biscuit Assortment Gift Tin	Sony Quick Battery Charger with 4 AA Batteries
GODIVA Biscuit Collection Gift Tin (50 pc)	Sony Stereo CD Clock Radio with Dual Alarm
GODIVA Chocoiste Solid Milk Chocolate Bars (24 pc)	Sony Water-Resistant Weather Band Shower Radio
GODIVA Spring Milk Chocolate Gift Box (22 pc)	Sony Wireless Headphone System
GODIVA Thank you Ballotin (70pc)	Spinmaster Air Hogs Battling Havoc R/C Helicopters
GPX Karaoke System with Integrated 5-1/2" Black and White Monitor	Star Wars Episode VI - Return of the Jedi (1983, 2004 Versions, 2-Disc)
Green Laser Pointer (5 mW, Class IIIa Laser Product)	Star Wars Trilogy (with Bonus Disc) (1977)
Guitar For Dummies (Lifestyles Paperback)	Starwars R2-D2 mimobot® 2GB USB Drive
Gundam X: G-Falcon Unit Double X Model Kit 1/100 Scale	SuperPen Graphics Tablet
Hamilton Beach Pump Espresso/Cappuccino maker	Swiss Gear 7x7-Foot 3-Person Sport Dome Cheval Tent
Hamilton Beach TrueAir Plug-Mount Odor Eliminator	T-Fal Avante Deluxe 4-Slice Toaster
Harbinger Pro Series Workout Glove	Tangle DNA Speakers
Harry and David Fancy Fruit Buffet - Ready to Serve Dried Fruit Medley	Tanita Duo Scale with Body Fat/Water Monitor
Harry and David Super Moose Munch® Party Drum - Caramel Popcorn, Chocolates, S'Mores etc.	TAO 1.5" Digital Keychain Photo Frame
Heelys Men's Octane	Tempur-Pedic Neck Pillow
Hitch iPod2iPod USB Transfer Device	Texas Instruments TI-83 Plus Graphing Calculator
Hohner Meisterklasse Harmonica, Key of D	Thanko FM/MP3 Watch 1GB
Holmes Medium Room Cool Mist Humidifier	The Adventures of Indiana Jones - The Complete DVD Collection (1981)
HoMedics Therapist Select Quad-Action Percussion Massager	The Complete Idiot's Guide to Weight Training and Body Sculpting (2 Paperbacks)
HoMedics Therapist Select Shiatsu Massaging Cushion	The Lord of the Rings - The Motion Picture Trilogy (Platinum Series Special Extended Edition) (2004)
Honeywell .54-Cubic-Foot Shelf and Floor Anti-Theft Safe with Digital Keypad	The Office−Season One (2005)
Honeywell Platinum Air Purifier	The Simpsons−The Complete Tenth Season (1989)
Hoover Tempo Widepath Bagged Upright Vacuum	Timberland Men's Chest Logo Full Zip Sweatshirt
HP Photosmart Printer	Timberland Men's Lexington Sport Oxford
imation 2GB USB 2.0 Clip Flash Drive with Carabineer and Rubberized Shell	Timberland Stratham Claremont Laptop Messenger Bag
iNeed™ Lumbar Massage Pillow	Timberland Treeline Travel Gear Back Pack Large
James Bond Ultimate Edition−Vol. 1	Timex Men's Chrono Alarm Watch
KEM Playing Cards (Paisley, 100% Cellulose Acetate Plastic)	Tortuga Original Caribbean Rum Cake, 33 Ounce Cake
Kensington DomeHub USB 2.0 (7 ports) Weighted Hub	Tripod Speaker for MP3 Player
Kodak EasyShare 6.2MP Digital Camera with 3X Optical Zoom	Valeo Yoga Kit

Koolatron Kool Fridge	Victorinox Hanging Toiletry Kit by Swiss Army
KORG Digital Metronome and Instrument Tuner	Victorinox Swiss Army Credit Card-Size Multi-Tools with LED Light
LAMPS PLUS Blue and Yellow Silver Base Lava Lamp	Victorinox Swiss Army Explorer Multitool Knife
Latin Percussion Mini Tunable Wood Conga	Victorinox Travel Wallet First Class by Swiss Army
Levi's Relaxed Straight 559 Jeans	Video Watch with OLED Screen
Levitating Desktop Globes	VIOlight Ultraviolet Toothbrush Sanitizers
Linksys Wireless-G Router	VIOlight Ultraviolet Travel Toothbrush Sanitizers
Logitech 2.4 GHz Cordless Presenter with Laser Pointer	Water Powered Multifunction Alarm Clock
Logitech Cordless Laser Mouse	Welcom WearEver Deluxe Aluminum Hi-Back Backpack Chair
Logitech Gaming Keyboard	Zelco Bookmark Dictionary

Table A3.2 Price for Individual Subjects.

Subject ID	Price
1	17
2	11
3	20
4	2
5	10
6	27
7	10
8	14
9	6
10	5
11	5
12	5
13	8
14	18
15	10
16	3
17	3
18	10
19	12
20	9
21	15
22	12
23	8
24	10
Mean	10.42
Median	10
St.Dev	5.93

Table A3.3 The Number of Decision Switches Between Hypothetical and Surprise Real Trials. The trials with missed responses were removed, hence leaving 1197 trials, instead of 1200.

	Yes→Yes	**Yes→No**	**No→Yes**	**No→No**	**total**
Count	**450**	**190**	**29**	**528**	**1197**
%	**38%**	**16%**	**2%**	**44%**	

Yes→Yes (No→No): Hypothetical Yes (No) decision which remained the same in the surprise real trials.
Yes→No (No→Yes): Hypothetical Yes (No) decision which switched to No (Yes) in the surprise real trials

Table A3.4 We performed a linear discriminant analysis for each subject with RT and mDV as independent variables (See (A) in the following page) and decision switch as a dependent variable (switch = 1, stick = 0). The detailed classification procedure was as follows for each subject:

1. Select all hypothetical Yes trials.
2. Among all the hypothetical Yes trials, select a training set (n=20).
    a. If there are more than 10 switch trials (hypothetical Yes → surprise real No: YN) and 10 stick trials (hypothetical yes →surprise real yes: YY), randomly select 10 from each of switch and stay trials. In this case, 50% is a base success rate.
    b. If there are more than 10 YN trials but less than 10 YY trials, select all YY trials and randomly select (20 – # of YY trials) of YN trials.
    c. If there are less than 10 YN trials but more than 10 YY trials, select all YN trials and randomly select (20 – # of YN) of YY trials.
    d. Subjects whose YY or YN are less than 5 are excluded in the classification analysis.
3. Leave one observation out for test.
4. Conduct a linear discriminant analysis (LDA) with the rest of the data (n=19) and predict the one left out (switch [YN] or stick [YY]).
5. Repeat steps 2~3 1000 times.

We also repeated the same classification analysis with RT and confidence (Strong, Weak) as independent variables and decision switch as a dependent variable. The results are reported in (B).

(A) Classification Analysis of Decision Switches Using RT and mDV.

Subject ID	# of hypothetical yes	# of sticks (YY)	# of switches (YN)	Base success rate	% of correct predictions	Sensitivity	Specificity
1	27	19	8	0.60	0.53	0.39	0.63
2	25	19	6	0.70	0.59	0.55	0.61
3	25	18	7	0.65	0.58	0.44	0.65
8	30	18	12	0.50	0.70	0.71	0.69
9	23	12	11	0.50	0.76	0.68	0.83
10	26	20	6	0.70	0.69	0.77	0.66
11	29	22	7	0.65	0.72	0.72	0.72
14	41	24	17	0.50	0.66	0.69	0.63
15	23	10	13	0.50	0.75	0.76	0.73
17	32	26	6	0.70	0.71	0.69	0.72
18	25	18	7	0.65	0.63	0.49	0.70
19	23	8	15	0.60	0.76	0.67	0.90
20	27	21	6	0.70	0.53	0.52	0.54
21	27	18	9	0.55	0.66	0.65	0.67
23	26	15	11	0.50	0.84	0.87	0.82
24	26	12	14	0.50	0.61	0.63	0.60
mean	27.19	17.5	9.69	0.59	0.67	0.64	0.69
St. dev.	4.45	4.98	3.65	0.09	0.09	0.13	0.09
s.e.	1.11	1.24	0.91	0.02	0.02	0.03	0.02

Notes:

(1) The classification success rate is significantly higher than the baseline success rate (paired two sample t-test, one-tail, $p<0.028$).

(2) Sensitivity = % of correctly predicted switches; specificity = % of correctly predicted sticks

(B) The Same Analysis of Decision Switches Using RT and Confidence (Yes/No).

Subject ID	# of hypothetical yes	# of sticks (YY)	# of switches (YN)	base success rate	% of correct predictions	Sensitivity	Specificity
1	27	19	8	0.60	0.54	0.46	0.60
2	25	19	6	0.70	0.35	0.36	0.35
3	25	18	7	0.65	0.73	0.56	0.82
8	30	18	12	0.50	0.65	0.60	0.70
9	23	12	11	0.50	0.82	0.63	1.00
10	26	20	6	0.70	0.67	0.96	0.55
11	29	22	7	0.65	0.76	0.88	0.69
14	41	24	17	0.50	0.55	0.63	0.46
15	23	10	13	0.50	0.66	0.58	0.82
17	32	26	6	0.70	0.75	0.88	0.70
18	25	18	7	0.65	0.88	0.98	0.82
19	23	8	15	0.60	0.87	0.78	1.00
20	27	21	6	0.70	0.71	0.83	0.66
21	27	18	9	0.55	0.71	0.70	0.72
23	26	15	11	0.50	0.80	0.64	0.96
24	26	12	14	0.50	0.53	0.63	0.44
mean	27.19	17.5	9.69	0.59	0.69	0.69	0.70
St. dev.	4.45	4.98	3.65	0.09	0.14	0.18	0.20
s.e.	1.11	1.24	0.91	0.02	0.03	0.04	0.05

Notes:

(1) The classification success rate is significantly higher than the baseline success rate (paired two sample t-test, one-tail, $p<0.021$).

(2) Sensitivity = % of correctly predicted switches; specificity = % of correctly predicted sticks

Bibliography

Anderson, A. K., K. Christoff, I. Stappen, D. Panitz, D. G. Ghahremani, G. Glover, J. D. E. Gabrieli, & N. Sobel. (2003). Dissociated neural representations of intensity and valence in human olfaction. *Nature Neuroscience,* **6**(2) 196-202.

Ariely, D., & K. Wertenbroch. (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science,* **13**(3) 219-224.

Becker, G., M. DeGroot, & J. Marshak. (1964). Measuring utility by a single response sequential method. *Behavioral Science,* **9**(3) 226-232.

Ben-Shakhar, G., & E. Elaad. (2003). The validity of psychophysiological detection of information with the guilty knowledge test: A meta-analytic review. *Journal of Applied Psychology,* **88**(1) 131-151.

Blumenschein, K., B. G. C., M. Johannesson, N. Horn, & P. Freeman. (2007). Eliciting willingness to pay without bias: Evidence from a field experiment. *The Economic Journal,* **118**(525) 114-137.

Botvinick, M., L. E. Nystrom, K. Fissell, C. S. Carter, & J. D. Cohen. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature,* **402** 179-181.

Carson, R. T., N. E. Flores, K. M. Martin, & J. L. Wright. (1996). Contingent valuation and revealed preference methodologies: Comparing the estimates for quasi-public goods. *Land Economics,* **72**(1) 80-99.

Chandon, P., V. G. Morwitz, & W. J. Reinartz. (2004). The short- and long-term effects of measuring intent to repurchase. *Journal of Consumer Research,* **31** 566-572.

Crespi, I. (1989). *Public Opinion, Polls, and Democracy*. Bourder, CO: Westview Press.

Cummings, R. G., G. W. Harrison, & E. Rutstrom. (1995). Homegrown values and hypothetical surveys: Is the dichotomous choice approach incentive-compatible? *American Economic Review,* **85**(1) 260-266.

Cummings, R. G., & L. O. Taylor. (1999). Unbiased value estimates for environmental goods: A cheap talk design for the contingent valuation method. *American Economic Review,* **89**(3) 649-665

Davatzikosa, C., K. Ruparelb, Y. Fana, D. G. Shena, M. Acharyyaa, J. W. Lougheadb, G. R.C., & D. D. Langleben. (2005). Classifying spatial patterns of brain activity with machine learning methods: Application to lie detection. *NeuroImage,* **28**(13) 663-668

Deichmann, R., J. A. Gottfried, C. Hutton, & R. Turner. (2003). Optimized epi for fmri studies of the orbitofrontal cortex. *NeuroImage,* **19** 430-441.

Delgado, M. R., R. H. Frank, & E. A. Phelps. (2005). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nature Neuroscience,* **8**(11) 1611-1618.

Diamond, P. A., & J. A. Hausman. (1994). Contingent valuation: Is some number better than no number? . *Journal of Economic Perspective,* **8**(4) 45-64.

Duvernoy, H. M. (1999). *The Human Brain: Surface, Three-Dimensional Sectional Anatomy with MRI, and Blood Supply.* (2nd ed.): New York: Springer.

Friston, K. J., C. Buechel, G. R. Fink, J. Morris, E. T. Rolls, & R. J. Dolan. (1997). Psychophysiological and modulatory interactions in neuroimaging. *NeuroImage,* **6** 218-299.

Gamer, M., T. Bauermann, P. Stoeter, & G. Vossel. (2007). Covariations among fmri, skin conductance, and behavioral data during processing of concealed information. *Human Brain Mapping,* **28** 1287-1301.

Gitelman, D. R., W. D. Penny, J. Ashburner, & K. J. Friston. (2003). Modeling regional and psychophysiologic interactions in fmri: The importance of hemodynamic deconvolution. *NeuroImage,* **19** 200-207.

Gottfried, J. A., J. P. O'Doherty, & R. J. Dolan. (2003). Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science,* **301**(5636) 1104-1107.

Greene, J. D., L. E. Nystrom, A. D. Engell, & J. M. Darley. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron,* **44**(2) 389-400.

Greene, J. D., R. B. Sommerville, L. E. Nystrom, J. M. Darley, & J. D. Cohen. (2001). An fmri investigation of emotional engagement in moral judgment. *Science,* **293**(5537) 2105-2108.

Grosenick, L., S. Greer, & B. Knutson. (2008). Interpretable classifiers for fmri improve prediction of purchases. *IEEE Transactions on Neural Systems and Rehabilitation Engineering,* **16**(6) 539-538.

Hare, T. A., J. O'Doherty, C. F. Camerer, W. Schultz, & A. Rangel. (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *Journal of Neuroscience,* **28**(22) 5623-5630.

Hariri, A. R., S. M. Brown, D. E. Williamson, J. D. Flory, H. de Wit, & S. B. Manuck. (2006). Preference for immediate over delayed rewards is associated with magnitude of ventral striatal activity. *Journal of Neuroscience,* **26**(51) 13213-13217.

Haxby, J. V., M. I. Gobbini, M. L. Furey, A. Ishai, J. L. Schouten, & P. Pietrini. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex *Science,* **293**(5539) 2425-2430.

Haynes, J.-D., K. Sakai, G. Rees, S. Gilbert, C. Frith, & R. E. Passingham. (2007). Reading hidden intentions in the human brain. *Current Biology,* **17** 323-328.

Infosino, W. J. (1986). Forecasting new product sales from likelihood of purchase ratings. *Marketing Science,* **5**(4) 372-384

Jamieson, L. F., & F. M. Bass. (1989). Adjusting stated intention measures to predict trial purchase of new products: A comparison of models and methods. *Journal of Marketing Research,* **26** 336-345.

Johannesson, M., B. Liljas, & P.-O. Johansson. (1998). An experimental comparison of dichotomous choice contingent valuation questions and real purchase decisions. *Applied Economics,* **30**(5) 643-647.

Kable, J. W., & P. W. Glimpcher. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience,* **10**(12) 1625-1633.

Keeter, S., & N. Samaranayake. (2007). Can you trust what polls say about obama's electoral prospects? [Electronic Version]. *Pew Research Center Report.*

Kim, H., S. Shimojo, & J. P. O'Doherty. (2006). Is avoiding an aversive outcome rewarding? Neural substrates of avoidance learning in the human brain. *PLOS Biology,* **4**(8) 1453-1461.

Knoepfle, D., C. F. Camerer, & A. Rangel. (2009). Are hypothetical valuation biases driven by inattention to opportunity costs? Unpublished manuscript, California Institute of Technology.

Knutson, B., S. Rick, E. Wimmer, D. Prelec, & G. Loewenstein. (2007). Neural predictors of purchases. *Neuron,* **53**(1) 147-156.

List, J. A. (2001). Do explicit warnings eliminate the hypothetical bias in elicitation procedures? Evidence from field auctions for sportscards. *American Economic Review,* **91**(5) 1498-1507.

List, J. A., & J. F. Shogren. (1998). Calibration of the difference between actual and hypothetical valuations in a field experiment. *Journal of Economic Behavior & Organization,* **37**(2) 193-205.

Little, J., & R. Berrens. (2004). Explaining disparities between actual and hypothetical stated values: Further investigation using meta analysis. *Economics Bulletin,* **3**(6) 1-13.

McClure, S. M., K. M. Ericson, D. I. Laibson, G. Loewenstein, & J. D. Cohen. (2007). Time discounting for primary rewards. *Journal of Neuroscience,* **27**(21) 5796-5804.

McFadden, D. L., A. C. Bemmaor, F. G. Caro, J. Dominitz, B.-H. Jun, A. Lewbel, R. L. Matzkin, F. Molinari, N. Schwarz, R. J. Willis, et al. (2005). Statistical analysis of choice experiments and surveys *Marketing Letters,* **16** 183-196.

Mitchell, T. M., S. V. Shinkareva, A. Carlson, K.-M. Chang, V. L. Malave, R. A. Mason, & M. A. Just. (2008). Predicting human brain activity associated with the meanings of nouns. *Science,* **320** 1191-1195.

Monterosso, J. R., G. Ainslie, J. Xu, X. Cordova, C. P. Domier, & E. D. London. (2007). Frontoparietal cortical activity of methamphetamine-dependent and comparison subjects performing a delay discounting task. *Human Brain Mapping,* **28**(6) 383-393.

Mortimer, D., & L. Segal. (2008). Comparing the incomparable? A systematic review of competing techniques for converting descriptive measures of health status into qaly-weights. *Medical Decision Making,* **28**(1) 66-89.

Murphy, J. J., P. G. Allen, T. H. Stevens, & D. Weatherhead. (2005). A meta-analysis of hypothetical bias in stated preference valuation. *Environmental and Resource Economics,* **30**(3) 313-325.

Nichols, T., M. Brett, J. Andersson, T. Wager, & J.-B. Poline. (2005). Valid conjunction inference with the minimum statistic. *NeuroImage,* **25** 653-660.

O'Doherty, J. P., M. L. Kringelbach, E. T. Rolls, J. Hornak, & C. Andrews. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience,* **4**(1) 95-102.

O'Doherty, J. P., J. Winston, H. Critchley, D. Perrett, D. M. Burtb, & R. J. Dolan. (2003). Beauty in a smile: The role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia,* **41**(2) 147-155.

O'Donoghue, T., & M. Rabin. (2008). Procrastination on long-term projects. *Journal of Economic Behavior & Organization,* **66**(2) 161-175.

Plassmann, H., J. P. O'Doherty, & A. Rangel. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *Journal of Neuroscience,* **27**(37) 9984-9988.

Raghubir, P., & E. A. Greenleaf. (2006). Ratios in proportion: What should the shape of the package be? *Journal of Marketing,* **70** 95-107.

Schlosser, A. E., T. B. White, & S. M. Lloyd. (2006). Converting web site visitors into buyers: How web site investment increases consumer trusting beliefs and online purchase intentions. *Journal of Marketing,* **70** 133-148.

Silk, A. J., & G. L. Urban. (1978). Pre-test-market evaluation of new packaged goods: A model and measurement methodology. *Journal of Marketing Research,* **15**(2) 171-191

Small, D. M., R. J. Zatorre, A. Dagher, A. C. Evans, & M. Jones-Gotman. (2001). Changes in brain activity related to eating chocolate: From pleasure to aversion. *Brain,* **124**(9) 1720-1733.

Takahashi, H., M. Kato, M. Matsuura, D. Mobbs, T. Suhara, & Y. Okubo. (2009). When your gain is my pain and your pain is my gain: Neural correlates of envy and schadenfreude. *Science,* **323**(5916) 937-939.

Tanner, R. J., & K. A. Carlson. (2009). Unrealistically optimistic consumers: A selective hypothesis testing account for optimism in predictions of future behavior. *Journal of Consumer Research,* **35**(5) 810-822.

Urban, G. L., G. M. Katz, T. E. Hatch, & A. J. Silk. (1983). The assessor pre-test market evaluation system. *Interfaces,* **13**(6) 38-59.

van Veen, V., J. D. Cohen, M. M. Botvinick, V. A. Stenger, & C. S. Carter. (2001). Anterior cingulate cortex, conflict monitoring, and levels of processing. *NeuroImage,* **14** 1302-1308.

Yacubian, J., J. Glascher, K. Schroeder, T. Sommer, D. F. Braus, & C. Buchel. (2006). Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *Journal of Neuroscience,* **26**(37) 9530 –9537.

Yeung, N., L. E. Nystrom, J. A. Aronson, & J. D. Cohen. (2006). Between-task competition and cognitive control in task switching. *Journal of Neuroscience,* **26**(5) 1429-1438.

Chapter 4

Optimal-Timing Decisions: An Experimental Study of a Multiagent Timing Game

4.1 Introduction

We often face situations where we have to decide when to take an action. Rewards can grow over time, so patience can pay off. But, while waiting, you are risking the possibility that someone else may take the rewards you have long waited to grow, or that the rewards may decrease in value – a stock market selling decision is a relevant example. When rewards from waiting versus taking an immediate action are weighed against each other, what is the optimal timing decision? This is a much studied question in industrial economics and management (Dutta et al. 1995; Fudenberg & Tirole 1985; Lambkin 1988; Robinson & Fornell 1985). The timing of R&D investments, new technology adoption, new product launches, and entry to new markets determines the success and survival of firms in the market. The benefits of being a first mover always trade off against those of being a follower[1] and this issue at its core is a firms' optimal-timing problem.

[1] First movers or early adopters have advantages that arise from initial larger market share; technological leadership (proprietary experience, patents); preemption of assets (geographic space, controlling natural resources, plants, and equipment); customer loyalty or switching costs. In contrast, followers can benefit from the ability to freeride on pioneering firms' investments; resolution of technological and market uncertainty; incumbent inertia (slow to respond to change in market conditions or customer needs) (Lieberman & Montgomery 1988).

A timing decision familiar to academics is when to stop analysis and send a paper out for publication. Spending more time on extra analysis and putting more materials in the paper can improve the quality of the paper and hence increase the chance of publication in a prestigious journal. However, if many other researchers have been working on a similar topic, then you are risking the chance of being 'scooped.' So academics are often torn between the choice of "stop and send it in" and that of "work a little bit more."

Another timely example, perhaps, is suggested by the recent housing market crash. Whoever invested in a house for speculative purposes in the past several years and did not sell before the crash would probably be regretting his or her bad timing decision.

In this experimental study, we consider an optimal timing decision where strategic motives for delaying and taking an immediate action must be weighed against each other. We adopt the clock game first introduced by Brunnermeier and Morgan (2008) (BM) in order to provide a framework for optimal timing decisions in the presence of the trade-off between rewards from waiting versus risk of preemption.

We study this game for the following reasons. First, it has a formal model and a very simple, closed-form optimal strategy (timing decision) that a rational player should use in equilibrium. This sharp prediction about behavior makes it easy to detect any deviation from the optimal strategy or abnormality of behavior.

Further, the game has a rich structure and can be used as a general paradigm to study related psychological processes. The basic decision that a player faces in the game is whether to wait for a reward to grow (risk a preemption) or whether to cash out and sell at a sure gain (but risk the regret of knowing that gains would have been larger from

waiting). This trade-off requires strategic reasoning of what others will do, as well as managing real-time emotions (e.g., anxiety and stress), and revising a strategy based on feedback (learning based on reinforcement and regret). These are all interesting, important components that guide decision making.

The specific goals of this study are three-fold: (1) we build on the work of BM and re-examine the theory's implications for optimal timing decisions, (2) we examine the strategic equivalence of dynamic and static implementations of the clock game, and (3) we test for subject learning/adaptation during the experiment.

In a lab experiment, BM tested the theoretical predictions about timing decisions, but they were not borne out by the data. This might be due to their experimental design whereby the subjects played repeatedly with 11 other human subjects in many rounds. After a while, the game essentially became a repeated game.[2] In the presence of naïve subjects, rational players would have adjusted their behavior according to others' unsophisticated behavior. Or it could be just an artifact of data censoring (further explained later), which could have biased the results. We overcome these issues by using computerized rivals who employ a theory-predicted strategy; and by implementing a strategically equivalent form of the clock game that addresses the potential for data censoring. An alternative explanation for the failure of the theory to describe the data would be that the theory might have a missing element that is crucial to the decision

[2] BM randomly and anonymously re-matched subjects into two different groups every trial, but it was not actually random. The matching was based on the order of the log-in to the experiment server—that is, after each trial (or round), subjects had to log in again and the first 6 people who logged in early were grouped together and the last 6 who logged in late were grouped together. Hence, it is possible that subjects who were always quick to log in were grouped together in all 45 rounds.

process. We investigate this possibility in our own experimental data and by re-analyzing the BM's data.

The clock game can be implemented in two different forms: dynamic and static, and we implemented both. Interestingly, the two versions have an identical equilibrium solution, which is called 'strategic equivalence.' We demonstrate that this strategic equivalence is not found in an experimental setting. Finally, although this has been ignored in BM, we look at the process of learning and see individuals adapt their strategy over time.

The Structure of the Clock Game and Theoretical Predictions

Dynamic Clock Game

Here we describe the "dynamic" clock game in detail, which is the original form of the clock game with unobservable actions as proposed by BM. This version of the game is played in real clock time and hence we call it "dynamic." For ease of explanation of the experiment later, we will frame the game environment as a stock market with a speculative bubble.

In this game, a finite number, I, of players participate in the stock market. At the start of the game, all players are in the market—in a sense that they hold one share of an stock, and their task is to make an selling decision; i.e., when to sell the stock. Each player can make one and only one decision before the game (or market) ends—once they sell their stock, that is a final decision and they cannot buy it back.

The price of the stock starts at 1 and exponentially increases at the rate of g; that is, at the period of t, the price is e^{gt}. At any time period, sellers can ***privately*** sell their stock as long as the market is still open.

At one point t_0, a speculative bubble in the price starts. t_0 is a random variable that follows an exponential distribution[3] with parameter λ, whose probability distribution function is $f(t_0) = \lambda e^{-\lambda t_0}$. One interpretation of t_0 is that it is the period in which the current stock price starts exceeding the stock's true worth. The distribution of t_0 is a common knowledge, but the exact value of t_0 is unknown to sellers. Only after a random delay, do sellers receive incomplete and private information about this and learn that this event has happened. More specifically, seller i privately receives a signal at period t_i indicating that t_0 has happened (or the price bubble has started). t_i is also a random variable and is uniformly distributed over the interval, [t_0, $t_0 + \eta$]. By the period $t_0 + \eta$, everybody has received the signal, and hence η is called the 'window of awareness.' The signal arrival period is private information and sellers cannot observe others' signal arrival time.

Once K sellers have sold, the market crashes. If a seller successfully sells the stock before the market crash, he will receive the amount of money that is equal to the price of the stock at the period of selling—that is, if he sells the stock at period x, he will earn e^{gx}. His selling price is not observed by others. Once the market crashes, the rest of the sellers who did not sell on time only receive the post-crash price, e^{gt_0}, which is the price at the bubble starting period, t_0. Note that a seller can be worse off if he sells too early. Selling before t_0, say at t_x, means practically losing money since if he had waited, he could have earned e^{gt_0}, which is larger than e^{gt_x}.

[3] An important property of the exponential distribution is that it is memoryless. Due to this "memorylessness" property, each period has a constant probability of λ of being t_0.

There are two ways the market ends (crashes): (1) once K ($< I$) sellers have sold, it crashes; (2) it ends at the period of $t_0 + \tau^*$ if fewer than K sellers have sold the item by this period. τ^* is common knowledge and is assumed to be large enough to ensure players will not take this exogenous ending period into consideration of their selling strategy (Brunnermeier & Morgan 2006). If more than one seller sell exactly at the crash period, then a tie-breaking rule applies—one of the tied sellers will be randomly (with an equal probability among them) chosen to sell his item. At the end of the market, the earnings of all sellers are publicly announced without individual identification.

This game structure has very nice properties, which are conducive to experimental tests. First, the dynamic clock game has a unique symmetric Bayesian perfect equilibrium. Second, the equilibrium strategy is characterized by a fixed constant. More specifically, in the unique symmetric equilibrium, each player (risk neutrality assumed) employs the following simple strategy: wait for τ periods after signal arrival and then sell, where

$$\tau = \frac{1}{g}\log\left(\frac{\lambda\Phi(K,\ I,\ \eta(\lambda\text{-}g))}{g-(g-\lambda)\Phi(K,\ I,\ \eta\lambda)}\right) \text{ and}$$

$$\Phi(a,b,x) = \frac{(b-1)!}{(b-a-1)!(a-1)!}\int_0^1 e^{xz}z^{a-1}(1-z)^{b-a-1}dz\ .$$

The formal proof of the existence and the uniqueness of the equilibrium can be found in an older version of BM (Brunnermeier & Morgan 2006).

Static Clock Game

The static version of the clock game is identical to the dynamic clock game with one difference: the game is played in a static environment. The static clock game is to the

dynamic game as the first-price sealed bid auction is to the Dutch auction.[4] All the properties of the dynamic clock game except for the dynamic structure will carry over to the static clock game.

In the static clock game, all sellers are given their signal arrival time, t_i at the beginning of the auction. Each seller is then asked to ***privately*** submit his selling time. Once all selling times are submitted, they are then compared and the market ending period is determined as the *K*-th earliest selling time. The *K* sellers with the earliest selling times can sell at the price of their choosing. The (*I* - *K*) sellers whose selling times are later than the market ending period, or who did not sell in time, receive the post-crash price, e^{gt_0}. t_0 and t_i are determined as in the dynamic game.

Interestingly, in their revised paper (Brunnermeier & Morgan 2008), BM solve the equilibrium using this static structure and reach the identical equilibrium solution. In other words, the dynamic clock game is strategically equivalent[5] to the static clock game: players in both types of the games should employ an equivalent strategy in equilibrium, which is that upon receiving the signal, wait for τ periods and sell (in the static game, submit the selling time which is τ periods after the signal arrival time). We test this theoretical isomorphism[6] between the dynamic and static clock games.[7]

---

[4] In a first-price sealed-bid auction, bidders submit their bid privately (historically, in a sealed envelope) and the bidder making the highest bid (hence, the name first-price) claims the object and pays the amount he has bid (Milgrom & Weber 1982). The Dutch auction (descending, clock) auction is an auction where the auctioneer begins with a very high asking price and then successively lowers the price until some bidder stops and claims the good for that price (Milgrom & Weber 1982)

[5] This means that the sets of strategies between the two types of games are identical and the outcome rules that transform strategies into allocations are identical (Milgrom 1989).

[6] We use the term 'isomorphism' in a loose sense, not as is rigorously defined in game theory, and interchange 'strategic equivalence' with 'isomorphism'.

[7] Strategic equivalence between different types of auctions is often observed in the standard forward auctions. In standard models, the first-price sealed bid and the Dutch auctions are strategically equivalent

4.2 Experimental Design and Methods

Participants and Stimuli

Thirty-eight subjects were recruited from UCLA and its neighborhood area (19 males, 19 females; mean age = 20.9 ± 4.1 year; age range = 17 – 42) and twenty-five subjects were recruited from Caltech (21 males, 4 females; mean age = 21.3 ± 2.2 year; age range = 18 – 27). Seven additional subjects participated in the experiment. They were excluded from the analysis because their response time in many static trials was too short[8] or their submitted selling times in static trials were highly correlated with initial, default values (explained further in the next section), implying that those subjects might have skipped trials by quickly pressing the response key instead of submitting their actual selling time decisions and thus added noise to the data.

Before the experiment, subjects were informed about the experiment and gave informed consent to participate according to a protocol approved by Institutional Review Boards of the University of California, Los Angeles CA (for the UCLA sessions), and of the California Institute of Technology, Pasadena CA (for the Caltech sessions).

In addition to the $5 show-up fee, subjects were paid whatever they earned during the experiment with the conversion rate of 100 Experimental Dollars equal to 0.50 US dollar.

---

and assuming that each bidder knows the value to himself of the object being auctioned, the English and the second-price sealed-bid auctions are also strategically equivalent (Milgrom 1989). Despite the theoretical equivalence, lab and field experiments have found that the strategic equivalence does not hold in practice (Cox et al. 1982; Cox et al. 1983; Kagel et al. 1987; Kagel & Levin forthcoming; Katok & Kwasnica 2006; Lucking-Reiley 1999).

[8] These subjects either showed less than 2 seconds of average response time (RT) in static trials or responded less than a second for more than 30% of all static trials. The average RT across all other subjects was 6.66 seconds (*SD* = 2.91).

Stimulus presentation and the timing of all stimuli and response events were achieved using Matlab (www.mathworks.com) and the Psychtoolbox (www.psychtoolbox.org) on computers running Microsoft Windows.

Experimental Procedure and Task

The experimental procedure followed the basic structure of the clock game experiments by Brunnermeirer and Morgan (Brunnermeier & Morgan 2006, 2008). The experiment was framed as an asset market in which subjects made selling time (price) decisions on the stock they were holding. Each participant played in his own market with computer players programmed to follow a theory-predicted equilibrium strategy. Participants knew that they were playing with five computer players (i.e., $I = 6$), who would receive the same amount of information as they did and who would employ a strategy constant throughout the experiment. The strategy of the computer players remained unknown to the participants.

An experimenter read out instructions (appendix 4.6.2) aloud while participants were provided with copies of the instructions to read. The figures for the instructions (appendix 4.6.3) were displayed on the computer screen in front of them. To ensure that subjects understood the game structure and procedure, a quiz was administered at the end of the instructions. Subjects who failed the quiz (answered two or more questions incorrectly; see appendix 4.6.4) twice were not allowed to proceed and dismissed ($n = 3$).

Each individual completed 100 trading rounds or trials—5 blocks of 10 dynamic clock game trials and 5 blocks of 10 static clock game trials. The dynamic blocks were

interleaved with the static blocks and half of the subjects started with a block of dynamic trials and the other half with a static round block.

In a dynamic trial, subjects were asked to make a real-time decision as to at what period (price) to sell the asset they were holding. At the start of each trading round, everyone (the subject and five computer players) received a share of the same asset with an initial price of $1. This price grew exponentially at the rate of 4% in each period (i.e. $g = 4\%$) as trading time periods passed in real-time. Specifically, each period lasted 250 milliseconds. The computer screen displayed a trajectory of the price graphically as well as the current price numerically (figure 4.1). The current price in each time period was the same for all 6 sellers (the human subject and 5 computer rivals).

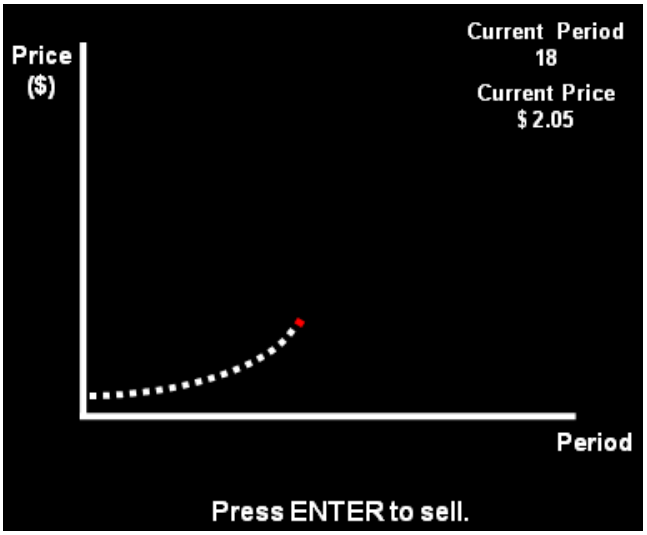

Figure 4.1 A dynamic clock game trial. The price trajectory from the start to the current period is displayed in a price-period graph. The current price and period are numerically displayed in the upper right corner of the screen.

At the start of each round, the computer determined the post-crash price, which was worded, in instructions to the participants, as the maximum "true value" of the asset, or the value when the true value stopped growing (i.e., t_0). In each period, there was a 2.5%

probability that the true value stopped growing (i.e., λ = 2.5%). After the true value stopped growing, a subject received a message (or a signal) indicating that the current price of the asset is above its true value; this message is delayed by an amount of time that is equally likely to be anywhere from 0 to 60 periods (i.e., η= 60), (figure 4.2). Subjects were told that the delay between the time the true value stopped growing (t_0) and the signal arrival (t_i) was randomly chosen separately for each player in each round.

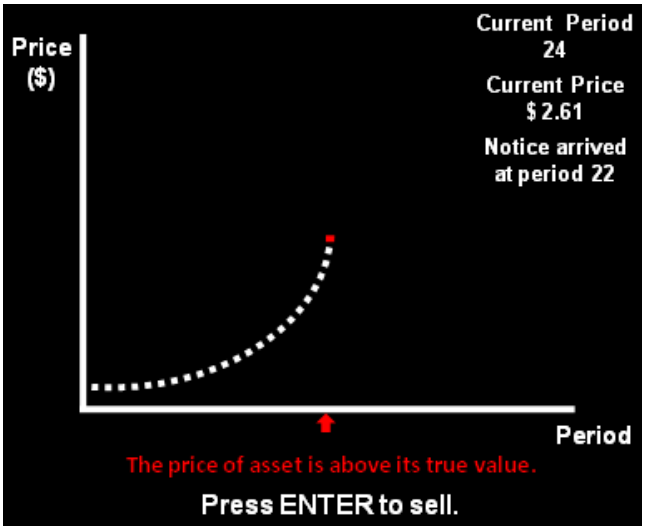

Figure 4.2 Subjects receive a message about the true value with a random delay.

Once three or more sellers in a market sold their asset, the round ended (i.e., K = 3). At this time, subjects received feedback about their current earnings as well as the prices at which other sellers sold their asset (figure 4.3). Subject's earnings in a dynamic round were determined as follows: if the subject successfully sold the asset (i.e., he was among the first three sellers to sell), he received the price of the asset at the time of selling. Otherwise, the subject received an amount equal to the maximum true value of the asset (post-crash price).

Figure 4.3 Once the market ends, subjects receive information on their earnings and others' selling price and the true value (the earnings of the last three subjects to sell).

In summary, the values of the parameters used in this experiment are as follows:

(a) Number of participants: $I = 6$

(b) Number of sellers necessary for market crash: $K = 3$

(c) Price growth rate: $g = 4\%$

(d) The probability that the true value stops growing in each period: $\lambda = 2.5\%$

(e) Exogenous ending parameter: $\tau^* = 300$

(f) Window of awareness: $\eta = 60$.

Under this set of parameter values, the equilibrium delay, τ, is 21 periods. That is, in equilibrium, once a subject receives a signal, he should strategically delay for 21 periods until he sells this asset.

In a static trial, subjects were presented with all of the information that was revealed across time in the dynamic rounds, all at once. More specifically, the signal arrival time t_i

was known to subject *i* at the beginning of a trial, who was then asked at what price (in what period) they would sell if they had received the signal in period t_i. The signal arrival period, t_i, varied between trials and was determined as in dynamic trials. The period in which a message arrived was shown on the computer screen. Subjects were allowed to explore the stock price over time by moving a cursor on the screen (figure 4.4). Once the cursor was placed at the desired selling time, the participants would press the enter button to choose that time and proceed. Once they submitted their decision, the computer compared the selling times of all players (a human subject and 5 computer players) and determined earnings as in a dynamic round—the computer first determined the time of market ending as the earliest period by which 3 players had sold and then the first 3 sellers earned their selling price and the other 3 earned only the maximum true value (post-crash price).

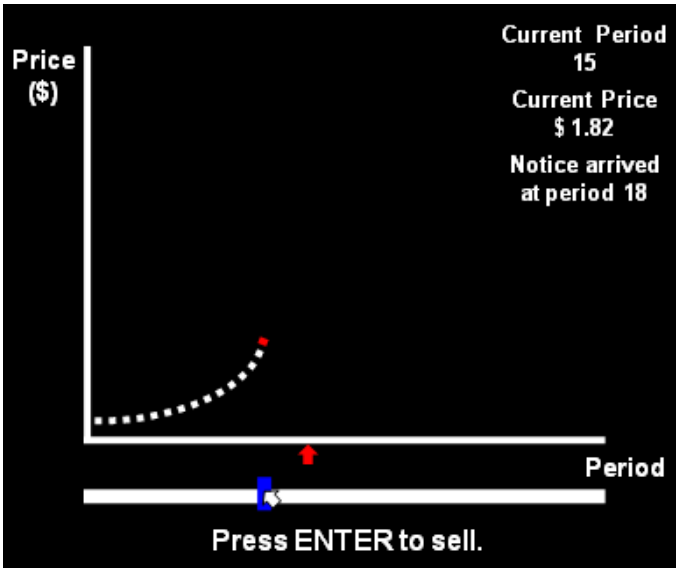

Figure 4.4 A static trading round. Subjects explore the price over time by moving the blue rectangular cursor along the horizontal bar at the bottom of the screen.

The initial position of the selling-time cursor is called an anchor and was randomized in every trial in order to avoid any potential anchoring effects.[9] The location of the initial position was only recorded during the sessions at Caltech and these data were used as a check for subjects' engagement in the task. Three participants at one of the Caltech sessions showed 40%~60% correlation between their selling times and anchor periods (significant at 0.005) in static trials and hence were discarded for further analysis. Their average RT in static trials was also less than 2 seconds.

Design Considerations

The use of computer players has a number of advantages, and this trick has been used in many experimental studies on auctions (Dorsey & Razzolini 2003; Engelbrecht-Wiggans & Katok 2007; Kagel & Levin 2001; Neugebauer & Selten 2006). First, we can control "collusive" strategies. Subjects cannot coordinate to wait longer than they would otherwise, and split the high payoff outside the lab. Second, participants do not face strategic uncertainty in regards to other players' behavior and have no need to worry if common knowledge assumptions are satisfied (that is, whether other players understand the instructions and play accordingly or not). Without computer players, our concern was that the presence of some abnormal or strategically naïve subjects with out-of-equilibrium behavior would "contaminate" others' decisions. In this case, it would be necessary to isolate subjects whose behavior is suboptimal, due to limited strategic thinking, from those who employ suboptimal strategy out of a rational expectation of others' unsophisticated behavior, which is a daunting task. Finally, using computer

[9] In many decision making situations, an individual's final decision or behavior is influenced by an initial value or a reference point. This phenomenon is called 'anchoring' (Tversky & Kahneman 1974).

opponents also prevents information leakage[10] among subjects in the lab since each subject participates in an individual game independently from others.

Our design permits within-subjects comparisons of the treatment effects of the two versions of the game, dynamic and static, which enhances the power of the statistical tests. Further, it rules out the possibility that any measured differences in behavior are due to the idiosyncrasies of a particular group (Cox et al. 1982).

Static trials are valuable as a test bed for the strategic equivalence between the two different implementations of the isomorphic games. Further, an important, additional benefit of static trials is that all selling decisions are observable. In dynamic trials, if the market ends before a seller would have sold, his strategic delay (the length of waiting periods between the signal arrival and the selling) is either not observed or right censored, and this lost information could bias results. Static trials do not suffer from this problem. We take advantage of this fact during the analysis.

4.3 Results

We start with definitions of the terminology used in this section.

Delay, an empirical measure of strategic delay τ, is the dependent variable of the most interest, and is defined as follows: In dynamic trials,

(1) If a subject received the signal and then he sold the asset afterwards, *Delay* is **uncensored** and defined as the length of periods between the signal arrival period and the selling time;

[10] This was indeed a problem in BM since mouse-clicking noise could serve as an auditory cue about selling time to other subjects. To get around this issue, they had subjects "hover" the mouse over a button to sell and this led to many erroneous sellings.

(2) If the subject sold prior to or at the signal arrival period, *Delay* is **left-censored** and coded as zero (left-censored at 0);

(3) If he received the signal but did not sell prior to the market crash, *Delay* is **right-censored** and coded as the amount of periods between the signal arrival period and the market ending time;

(4) If he neither received the signal nor sold, then *Delay* is treated as **missing** and hence discarded.

Although all selling decisions in static trials were indeed observed, for comparability of behavior in dynamic and static trials, *Delay* in static trials is defined in the same manner.

Duration was defined as in BM, the length of periods from t_0 to the market ending.

Early exit is an event in which a subject sold prior to the signal arrival and which led to left-censored *Delay*.

Success rate is the number of the trials in which a subject sold before the market ending, divided by the number of all trials, whether sold or not, in a given condition.

Below we define the independent variables of relevance:

(a) *Signal* indicates the signal arrival time in periods;

(b) *Condition* is a categorical variable indicating a trial type (dynamic or static);

(c) *Group* is a categorical variable indicating the group to which a subject belongs (UCLA or Caltech);

(d) *UCLA* (*Caltech*) is a dummy variable for UCLA (Caltech): 1 if a subject participated in one of the UCLA (Caltech) sessions and 0 otherwise;

(e) *Dyn* (*Static*) is a dummy variable for dynamic (static) trials;

(f) *Experience* is a dummy variable to indicate a given trial belongs to the last half of the experiment.

In the following subsections, we first examine the trials where selling decisions were observed (the trials with left-censored and uncensored *Delay*)—hence, therein *Delay* means left-censored or uncensored *Delay*. When the trials with right-censored *Delay* were included, the qualitative aspects of the data did not change and these additional results are reported in the appendix (Section 4.5.1). Note that right-censored *Delay* could underestimate actual *Delay*, and the left-censored *Delay* could overestimate the actual *Delay* that could have been negative, potentially biasing the results.[11] Therefore, we repeat all the analyses using static trials only, with a redefinition of *Delay* as the length of periods between the submitted selling time and the signal arrival time; $Delay_{uncensored}$ denotes this new definition of *Delay*. Note that $Delay_{uncensored}$ permits negative values when a subject sold prior to the signal arrival. Using $Delay_{uncensored}$ yields qualitatively identical to those obtained using the *Delay* measure.

In theory, all sellers are ex ante identical in a sense that in equilibrium they all employ the same unique symmetric strategy. Hence in analyzing data, we do not distinguish the behavior of 1st, 2nd, and 3rd sellers as in BM.

p-values reported herein are two sided unless noted otherwise.

Descriptive Statistics

Table 4.1 shows the percentage of the trials with left-/right-, uncensored *Delay*.

[11] BM is aware of this censoring problem. To get around the problem, they estimated the delay measure using the Tobit procedure, assuming that the right-censored delay is normally distributed with mean τ. However, this is tautological (using τ to estimate τ) and this specification would yield an estimate in favor of τ.

Table 4.1 Percentage of trials with (un)censored *Delay*

	All	Dynamic	Static
Right censored	26%	25%	28%
Uncensored	48%	49%	47%
Left censored	9%	8%	10%
Missing	16%	18%	15%
Total	100%	100%	100%

Signal Arrival Time

Average signal arrival period was 67.55 (*SD* = 38.05). In dynamic trials, on average, the signal arrived at the period of 68.21 (*SD* = 37.89) and in static trials, 66. 89 (*SD* = 38.19). The difference in the signal arrival time between dynamic and static trials was marginally significant (paired t-test, $t(62) = 1.80$, $p = 0.077$; signed rank test, *z-value* = -2.77, $p < 0.01$), and this slight difference was driven by the UCLA group (figure 4.5). The signal arrival period is the primary independent variable of interest and the difference in signal arrival time could possibly lead to different behavior between the two types of trials. Hence, we look at the data from the UCLA subjects and the Caltech subjects separately, as well as the pooled. However, our main qualitative results remain unchanged regardless.

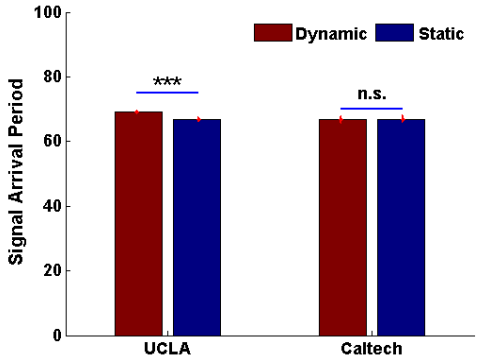

Figure 4.5 Signal arrival period by trial type and group. Error bars indicate standard errors. *** $p < 0.001$, n.s.: not significant, paired t-test.

Success Rate

The success rates between dynamic (M = 0.57, SD = 0.13) and static trials (M = 0.57, SD = 0.14) were not significantly different (paired t-test, $t(62) = -0.17$, $p = 0.87$; signed rank test, *z-value* = -0.003, $p = 0.997$), but they were significantly correlated (figure 4.6, $r(61) = 0.59$, $p < 4.24 \times 10^{-7}$). The success rates of dynamic and static trials were not significantly different between the UCLA and the Caltech subjects (dynamic trials: two-sample t-test, $t(61) = -0.64$, $p = 0.52$; static trials: two-sample t-test, $t(61) = -1.44$, $p = 0.15$)

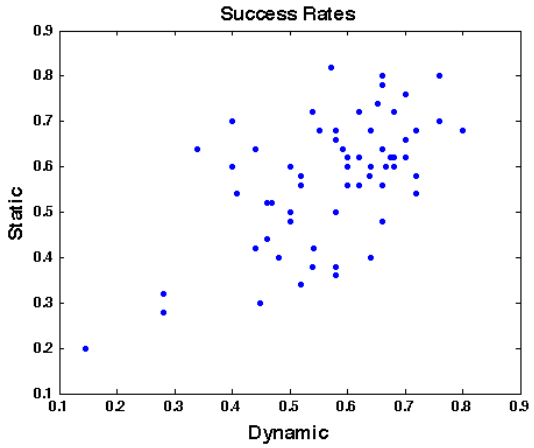

Figure 4.6 Scatter plot of success rates in dynamic and static trials.

Durations

*Duration*s in static trials and in dynamic trials were significantly different (figure 4.7A; paired t-test, $t(62) = -2.68$, $p < 0.0096$; signed rank test, *z-value* = -2.52, $p < 0.012$). They were also shorter than the theory predicts (figure 4.7B). Figure 4.7B shows simulated durations assuming that subjects had played an equilibrium strategy given the same signal arrival periods. The actual duration was slightly shorter than predicted, but the difference was significant in dynamic trials (paired t-test, $t(62) = -6.48$, $p < 0.0001$; signed rank test, *z-value* = -5.15, $p < 0.0001$) and in static trials (paired t-test, $t(62) = -4.63$, $p < 0.0001$;

signed rank test, *z-value* = -4.12, $p < 0.0001$). Mean (median) *Duration*s were not different between the UCLA and the Caltech subjects in dynamic trials (two-sample t-test, $t(61) = 1.53$, $p < 0.13$; rank sum test, *z-value* = -0.75, $p < 0.45$) and in static trials (two-sample t-test, $t(61) = 1.71$, $p < 0.09$; rank sum test, *z-value* = -1.13, $p < 0.26$).

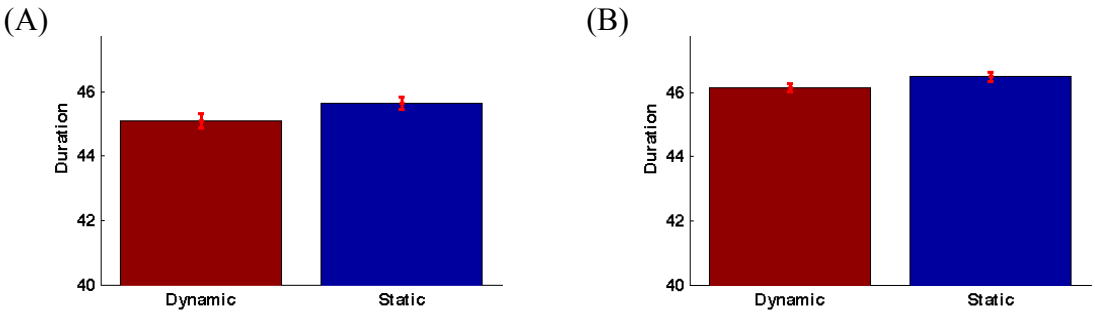

Figure 4.7 Difference in average *Duration* between dynamic and static trials. The y-axis scale is adjusted to display the difference in means. (A) Duration data from the actual subjects. (B) Simulated durations from an equilibrium strategy. Error bars indicate standard errors.

Difference in *Delay* between Dynamic and Static Trials

Mean (median) *Delay* was greater in static trials than in dynamic trials (figure 4.8A; paired t-test, $t(62) = -5.32$, $p < 7.57 \times 10^{-7}$; signed rank test, *z-value* = -4.40, $p < 0.0001$). *Delay*s in both types of trials were significantly lower than the theory predicted delay, τ = 21 periods (dynamic trials: one-sample t-test, $t(62) = -18.86$, $p < 2.29 \times 10^{-27}$; static trials: one-sample t-test, $t(62) = -10.37$, $p < 3.55 \times 10^{-15}$; figures 4.8A and 4.9). This pattern did not change when the trials with right-censored *Delay* were included (figures A4.1 and A4.2).

A two-way analysis of variance (ANOVA) with one within-subject factor, *Condition*, with two levels (dynamic or static) and one between-subject factor, *Group*, with two levels (UCLA or Caltech) was conducted to evaluate the effect of static trials and groups

on *Delay*. Significant main effects were observed for *Condition*, $F(1, 120) = 14.12$, $p < 0.0002$, and for *Group*, $F(1, 120) = 11.66$, $p < 0.0009$, but their interaction did not reach statistical significance ($F(1, 120) = 0.13$, $p < 0.72$). Post hoc t-tests showed that *Delay* for the UCLA subjects was significantly lower than that for the Caltech subjects in dynamic trials (two sample t-test, $t(61) = -3.73$, $p < 4.20 \times 10^{-4}$), but not in static trials (two sample t-test, $t(61) = -1.02$, $p = 0.31$) (figure 4.8B).

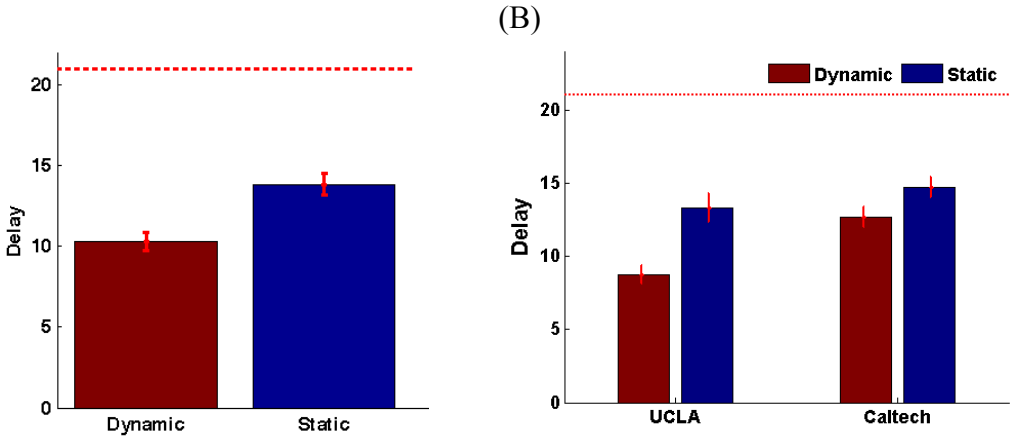

Figure 4.8 Difference in average *Delay* between dynamic and static trials. Dashed lines in red indicate the theory-predicted delay, $\tau = 21$. Error bars indicate standard errors. (A) *Delay* by trial type. (B) *Delay* by group and trial type.

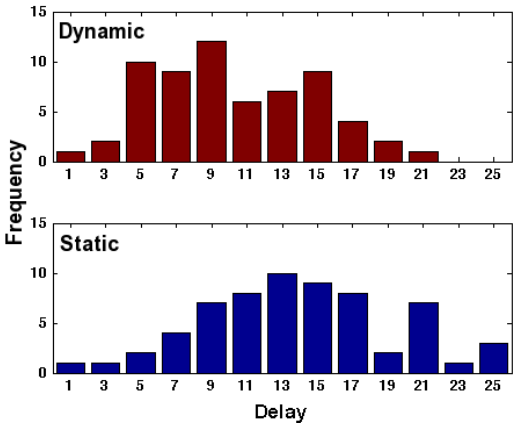

Figure 4.9 Distribution of individual average *Delay* between dynamic and static trials.

Although $Delay_{uncensored}$ ($M = 19.23$, $Med = 16.76$, $SD = 10.35$) was not significantly different from $\tau = 21$ in a parametric test (one sample t-test, $t(62) = -1.36$, one sided $p =$

0.18), it was significantly lower than τ in a non-parametric test (one-sample signed-rank test, *z-value* = -2.27, p = 0.02), which is a more appropriate test to apply given non-normality of the data (figure 4.10). Furthermore, while $Delay_{uncensored}$ for the Caltech group (M = 16.80, *Med* = 15.56, *SD* = 5.62) was significantly lower than τ (one-sample t-test, $t(24) = -3.73$, $p < 0.001$; one-sample signed-rank test, *z-value* = -3.00, $p < 0.003$), in the UCLA subjects it was not significantly different from τ (M = 20.82, *Med* = 18.25, *SD* = 12.34, one-sample t-test, $t(24) = -0.09$, p = 0,93; one-sample signed-rank test, *z-value* = -0.70, p = 0.49).

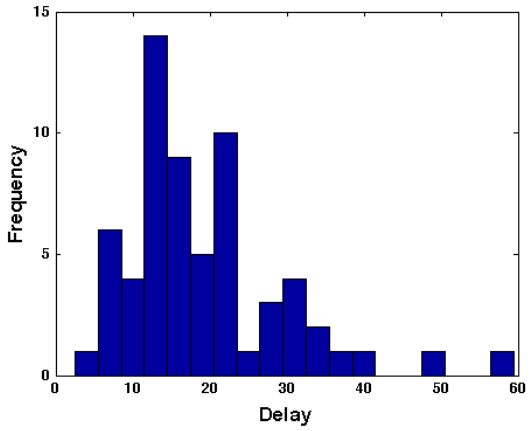

Figure 4.10 Distribution of individual average $Delay_{uncensored}$ in static trials.

Signal Dependence of *Delay*

Although the dynamic and static clock games are strategically equivalent, this predicted equivalence was not borne out by the data. Figure 4.11 shows scatter plots of *Delay* against the signal arrival time in dynamic and static trials. To investigate if there is any systematic difference in *Delay* across time, we further binned the signal arrival periods into 6 bins, by 25 periods starting from 0, except for the last bin—due to rarity of the observations, signal periods over 125 were binned all together—and then compared *Delay* between dynamic and static trials in each of the 6 time bins. *Delay* decreased as a

function of the signal arrival time in both dynamic and static trials. Further, in each time bin mean *Delay* in static trials was longer than in dynamic trials (figure 4.12). This pattern was preserved when the trials with right-censored *Delay* were included (figure A4.3).

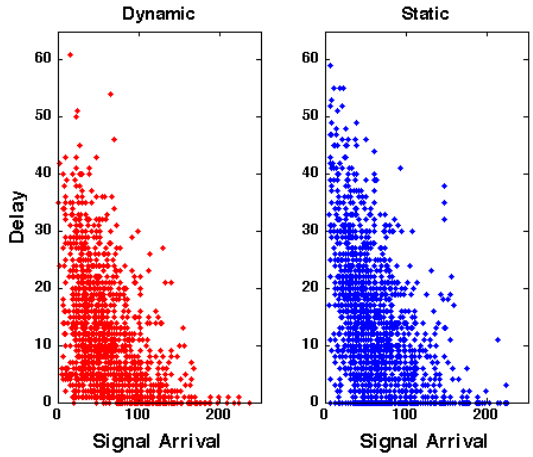

Figure 4.11 Distribution of *Delay* against the signal arrival period by trial type. Note that negative *Delay*s are censored at 0.

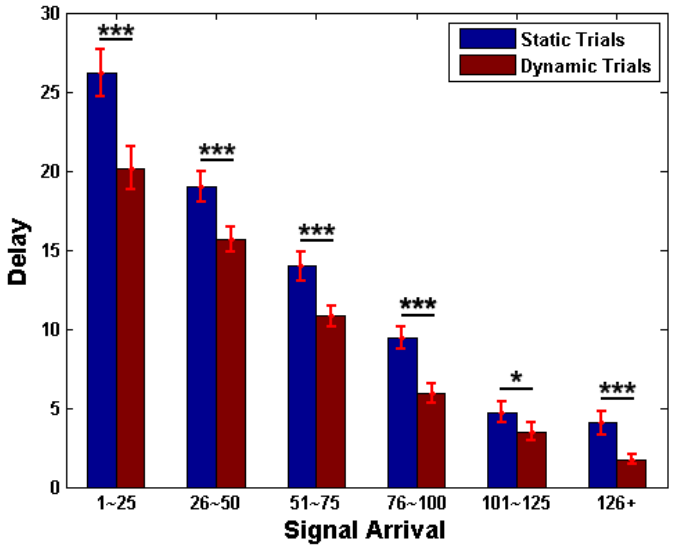

Figure 4.12 *Delay* as a function of the signal arrival period. *Delay* decreases as a signal arrives later. Error bars indicate standard errors. *** $p < 0.001$, * $p < 0.05$ (paired t-test, one sided).

Table 4.2 Results of random-effects Tobit regression analyses for *Delay* censored at 0 (subject random-effects incorporated)

Variable	A	B	C	D	E	F
Constant	30.724†	29.15†	28.892†	28.612†	30.572†	29.418†
	(0.430)	(0.398)	(0.472)	(0.486)	(0.430)	(0.672)
Signal	-0.196†	-0.194†	-0.194†	-0.190†	-0.204†	-0.202†
	(0.005)	(0.004)	(0.004)	(0.007)	(0.005)	(0.009)
*UCLA*Static*			3.231†		3.651†	4.537†
			(0.420)		(0.442)	(0.906)
*Caltech*Dyn*			0.521		-4.902†	-1.768*
			(0.573)		(0.799)	(0.996)
*Caltech*Static*			2.580†			1.461
			(0.576)			(0.971)
*Signal*UCLA*Static*				-0.009		-0.02
				(0.010)		(0.013)
*Signal*Caltech*Dyn*				0.01	0.038†	0.036**
				(0.007)	(0.011)	(0.013)
*Signal*Caltech*Static*				0.016		0.017
				(0.010)		(0.012)
Caltech	-0.25					
	(0.491)					
Static		2.787†		3.816†		
		(0.320)		(0.651)		
Log likelihood	-11527.73	-11498.79	-11488.36	-11491.68	-11491.21	-11477.68

Left-censored at *Delay* = 0	584
Uncensored	3012
Right-censored	0
Included Observations	3596

Standard errors are reported in parentheses. † $p < 0.0001$, ** $p < 0.01$.

We further examined if there is any group specific effect on *Delay* by conducting Tobit regression analysis (Greene 2002) for *Delay* (table 4.2). The structure of the Tobit models in this analysis is expressed as

$$Delay_i = \begin{cases} Delay_i^* & if\ Delay_i^* > 0 \\ 0 & if\ Delay_i^* \leq 0 \end{cases}$$

where $Delay_i^* = x_i\beta + e_i$, $e_i \sim N(0, \sigma)$, x_i is a vector of independent variables.

*Delay** is a latent variable that is observed for values greater than 0 and is censored for values less than or equal to 0.

The signal arrival time remained significant across various model specifications, and *Delay* decreased as the signal arrival was delayed. Both of the UCLA and the Caltech groups waited longer to sell in static trials than in dynamic trials, although on average the length of delay in static trials was longer in the Caltech group. Furthermore, the Caltech subjects exhibited less sensitivity of *Delay* to the signal arrival period in dynamic trials than in static trials. A similar pattern was found when the trials with right-censored *Delay* were included (table A4.1).

$Delay_{uncensored}$ also decreased as a function of the signal arrival period (figure 4.13), although the length of delay was longer due to the inclusion of the long delays that were previously excluded.

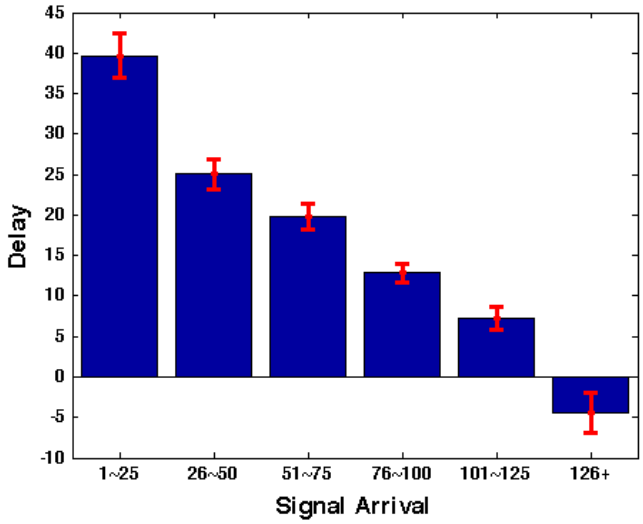

Figure 4.13 $Delay_{uncensored}$ as a function of the signal arrival period.

Furthermore, $Delay_{uncensored}$ showed very different behavior between the two groups than shown by the previous analyses (figure 4.14 and table 4.3). The Caltech subjects showed

less variance in *Delay* over all signal periods compared to the UCLA subjects. Caltech subjects's average length of $Delay_{uncensored}$ was shorter and the decrease in their $Delay_{uncensored}$ measure was less sensitive to the signal delay. In both subject groups, the signal arrival period was a highly significant predictor of $Delay_{uncensored}$, but the effect of the signal arrival period differed between the two groups—it was smaller (in an absolute sense) for the Caltech participants than for the UCLA participants.

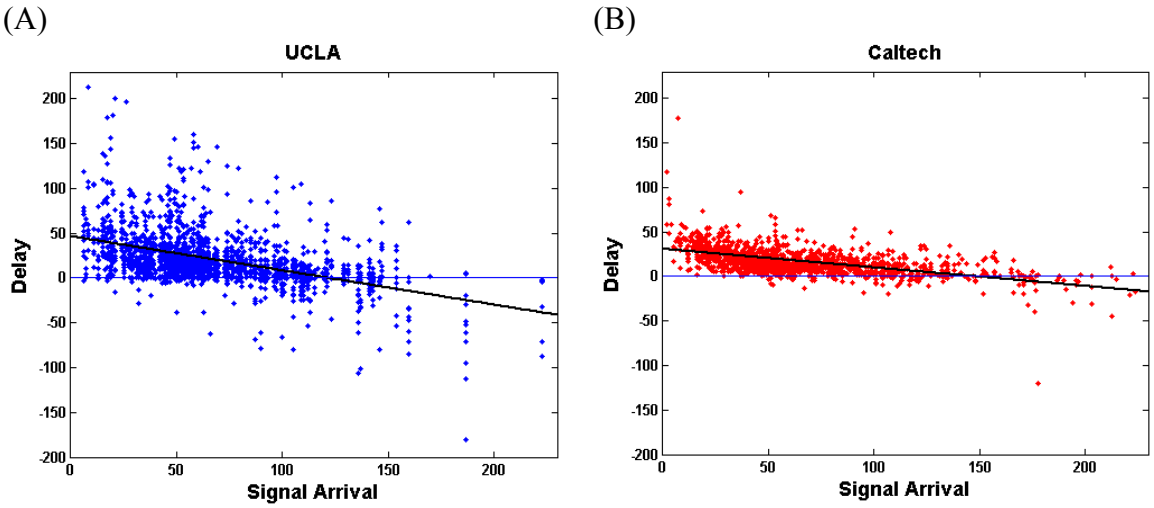

Figure 4.14 Scatter plot of $Delay_{uncensored}$ versus the signal arrival period. Straight lines indicate fitted regression lines from B in table 4.3. (A) UCLA subjects. (B) Caltech subjects.

Table 4.3 Results of random-effects regression analyses of $Delay_{uncensored}$ (subject random-effects incorporated)

Variable	A	B
Constant	40.813†	45.875†
	(1.726)	(1.787)
Signal	-0.299†	-0.374†
	(0.010)	(0.013)
Caltech	-4.019	-15.280†
	(2.547)	(2.798)
*Signal * Caltech*		0.168†
		(0.019)
R^2	0.217	0.234

Included Observations	3150

Standard errors are reported in parentheses. ✝ $p < 0.0001$

Effect of Experience on *Delay*

We examined whether subjects' selling strategy was adaptive over time. As a first attempt, we compared behavior in the first half of the experiment (trial number 1~50) and the last half of the experiment (trial number 51~100), by which point substantial learning and adaptation would have occurred. The dynamic vs. static difference in *Delay* was more pronounced in the first half of the game; *Delay* in the static case was substantially longer in each time bin as compared to the dynamic case. In later trials, the difference was reduced, but still persisted and appeared to stabilize (figure 4.15). Regardless of experience, however, *Delay* decreased as a function of the signal time.

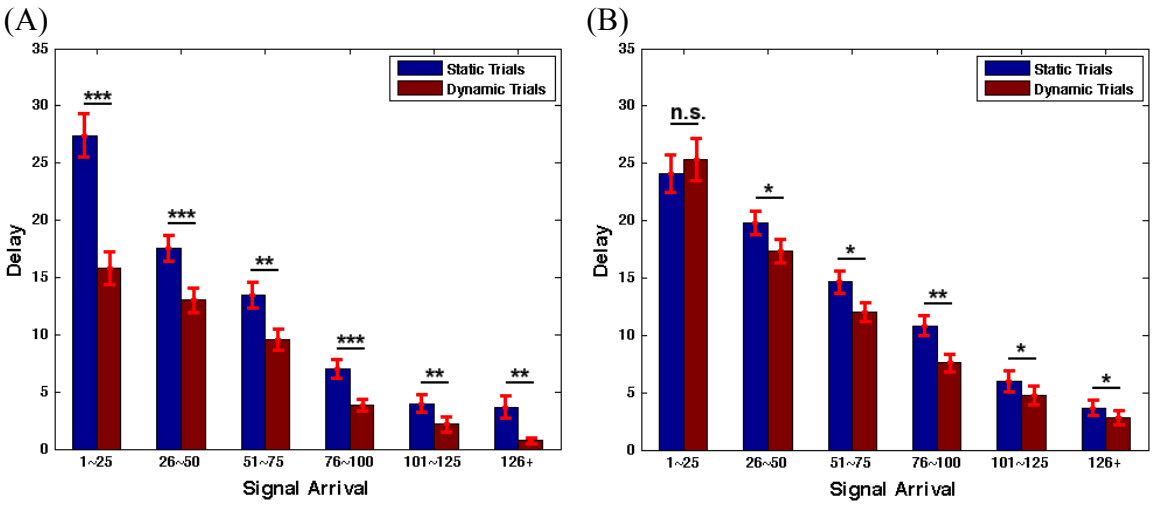

Figure 4.15 *Delay* as a function of the signal arrival period, moderated by experience. (A) First 50 trials. (B) Last 50 trials. *** p <0.001, ** p <0.01, * p <0.05, n.s.: not significant (paired t-test, one sided)

While the length of delay was generally longer, the same pattern was found when the $Delay_{uncensored}$ measure was used (figure 4.16).

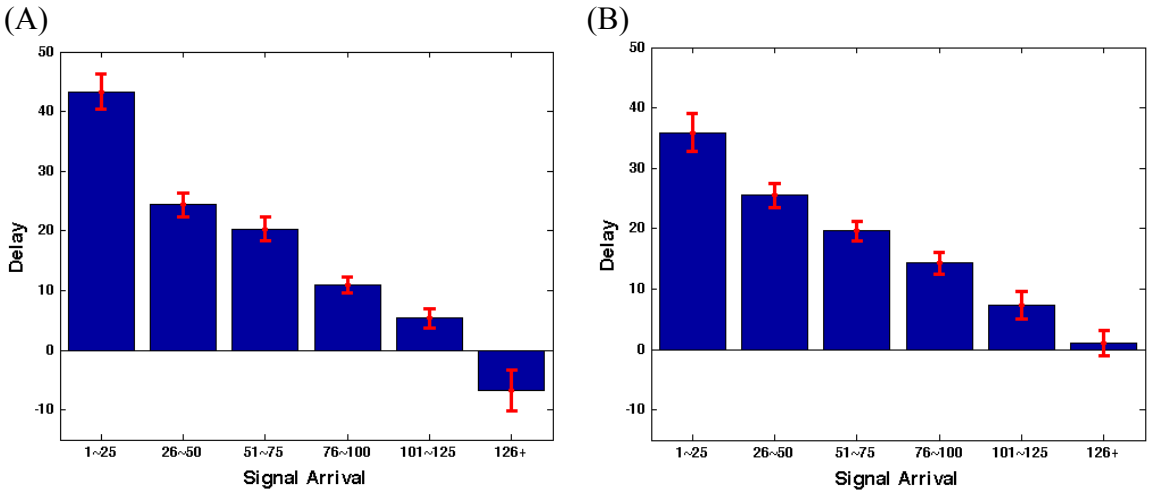

Figure 4.16 $Delay_{uncensored}$ as a function of the signal arrival period, moderated by experience. (A) First 50 trials only. (B) Last 50 trials only.

Further regression analyses of $Delay_{uncensored}$ revealed that significant learning effects were found in both groups (table 4.4). While the signal arrival time remained a significantly predictor of $Delay_{uncensored}$ throughout the experiment, its effect was moderated by experience in both groups—$Delay_{uncensored}$ decreased by a smaller amount in response to one unit of signal delay in later trials than in earlier trials.

Table 4.4 Results of regression analyses of $Delay_{uncensored}$ on the signal arrival time and experience

	UCLA subjects only			Caltech subjects only		
Variable	A	B	C	D	E	F
Constant	46.502†	46.103†	50.864†	29.858†	31.004†	31.270†
	(1.454)	(1.276)	(1.764)	(0.760)	(0.701)	(1.000)
Signal	-0.382†	-0.392†	-0.444†	-0.211†	-0.230†	-0.233†
	(0.017)	(0.017)	(0.022)	(0.009)	(0.011)	(0.014)
Experience	-0.226		-9.897†	2.073**		-0.521
	(1.218)		(2.544)	(0.730)		(1.402)
Signal*Experience		0.030◊	0.146†		0.033†	0.039*
		(0.016)	(0.034)		(0.009)	(0.018)
R-square	0.2168	0.2183	0.2245	0.3077	0.3102	0.3103
Adjusted R-square	0.216	0.2174	0.2232	0.3066	0.3091	0.3086
F-statistic	262.58	264.84	182.92	277.13	280.42	186.86
Prob(F-statistic)	0.000	0.000	0.000	0.000	0.000	0.000
Included Observations	1900	1900	1900	1250	1250	1250

Standard errors are reported in parentheses. † $p < 0.0001$, ** $p < 0.01$, * $p < 0.05$, ◊ $p < 0.1$

We also defined a measure of deviation from the equilibrium to see if subjects' delay strategy was adaptive to the equilibrium over time. This measure, denoted by δ, the standard deviation of *Delay* from τ, was defined as $\delta = \sqrt{(Delay_i - \tau)^2}$. Most of the subjects, on average, exhibited higher δ in the first half of the experiment than in the second half (figure 4.17)

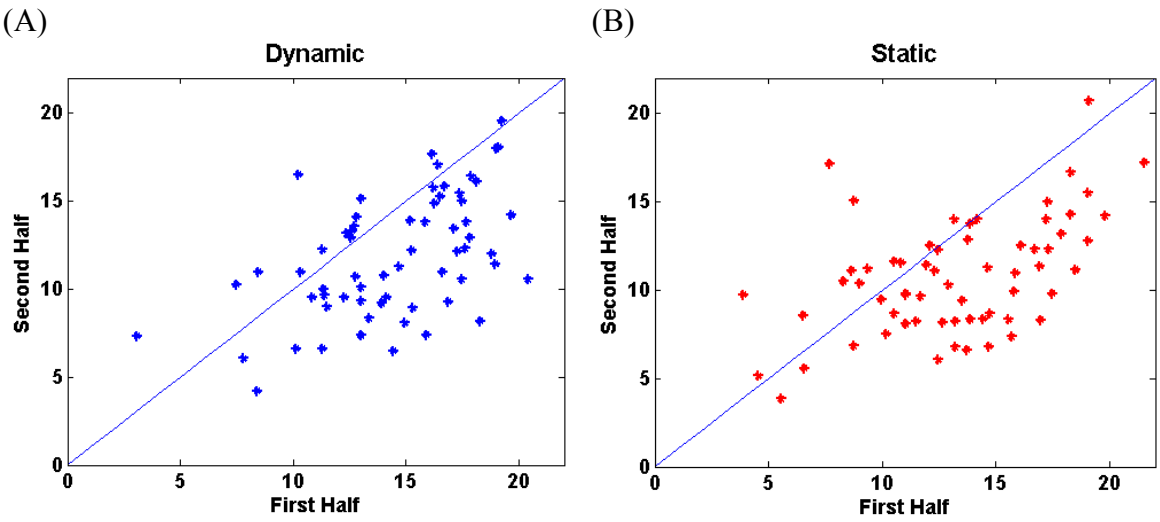

Figure 4.17 Scatter plot of average δ's for the first and the last halves of the experiment. (A) Dynamic trials. (B) Static trials. Blue lines indicate 45˚ line.

A two-way within-subjects ANOVA was conducted to evaluate the effect of static trials and experience on δ. The within-subjects factors were *Condition* (dynamic vs. static) and *Experience* (first half vs. last half). Significant main effects were found for *Condition* ($F(1, 244) = 6.65$, $p < 0.02$), and for *Experience* ($F(1, 244) = 29.75$, $p < 1.20 \times 10^{-7}$). However, there was no interaction between *Condition* and *Experience* ($F(1, 244) = 0.97$, $p < 0.33$; figure 4.18). Posthoc t-tests showed that mean δ's in the first half of the trials were significantly greater than in the last half of the trials in both dynamic (paired t-test, $t(62) = 6.14$, $p < 6.75 \times 10^{-8}$; signed rank test, *z-value* = -5.09, $p < 3.64 \times 10^{-7}$) and static

trials (paired t-test, $t(62) = 5.78$, $p < 1.80 \times 10^{-6}$; signed rank test, *z-value* = -4.60, $p < 4.21 \times 10^{-6}$). δ's for the first and the last half of the trials in the static condition were lower than those in the dynamic condition (first half: paired t-test, $t(62) = 2.84$, $p < 0.01$; signed rank test, *z-value* = -2.76, $p < 0.01$, last half: paired t-test, $t(62) = 2.15$, $p < 0.05$; signed rank test, *z-value* = -2.13, $p < 0.05$).

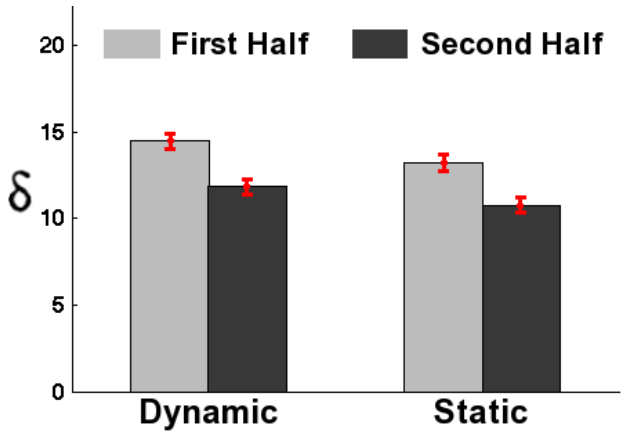

Figure 4.18 Deviation measure δ by trial type (Dynamic vs. Static) and experience (First vs. Last half of the trials).

There was a significant main effect for *Group* (UCLA vs. Caltech) on δ in both dynamic and static trials, but the interaction between *Group* and *Experience* was only found in static trials (figure 4.19). In dynamic trials, both of the UCLA and the Caltech subjects showed diminished deviations over time, but the Caltech subjects deviated significantly less than the UCLA subjects throughout the experiment. However, in static trials, the UCLA subjects' δ was substantially reduced as they gained more experience and it was not significantly different from the Caltech subjects' deviation in the second half of the trials.

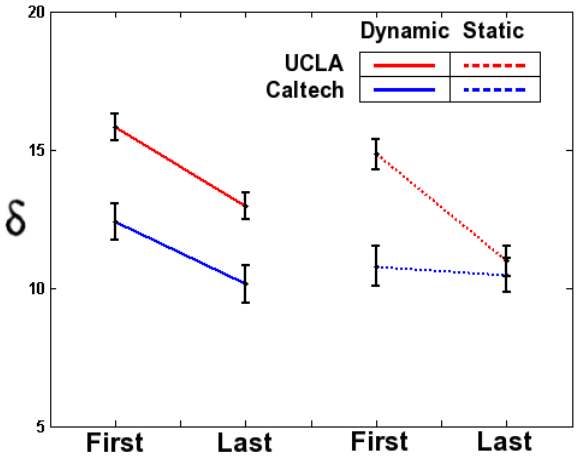

Figure 4.19 Interactions between trial type, experience, and group on deviation measure δ. Red (blue) lines indicate UCLA (Caltech) subjects. Solid (dashed) lines indicate dynamic (static) trials.

Further, we defined $\delta_{uncensored}$ using $Delay_{uncensored}$ in static trials only as follows:

$$\delta_{uncensored} = \sqrt{(Delay_{uncensored,i} - \tau)^2} \ .$$

$\delta_{uncensored}$'s for the first half of the trials in static trials (M = 18.59, SD = 4.31) were lower that for the last half (M = 14.42 SD = 3.80) (paired t-test, $t(62) = 4.74$, $p < 1.31 \times 10^{-5}$; signed rank test, *z-value* = -4.31, $p < 1.66 \times 10^{-5}$). A two-way ANOVA with a within-subjects factor, *Experience*, and a between-subjects factor, *Caltech*, showed that there were significant main effects of *Caltech* and *Experience* on $\delta_{uncensored}$ (*Group*: $F(1, 120) = 7.07$, $p < 0.009$, *Experience*: $F(1, 120) = 5.77$, $p < 0.018$), but no interaction between *Caltech* and *Experience* ($F(1, 120) = 0.39$, $p < 0.535$).

Posthoc regression analyses (table 4.5) revealed that while controlling for the signal arrival time, the trial number was negatively correlated with $\delta_{uncensored}$, implying that the deviation from equilibrium diminished over time as the subjects obtained more experience. The group difference was also observed: the UCLA subjects showed a higher

deviation to start with but it was substantially reduced over time relative to the Caltech subjects.

Table 4.5 Results of regression analyses of $\delta_{uncensored}$ on the signal arrival period and the trial number

Variable	UCLA only	Caltech only
Constant	21.035†	9.008†
	(1.402)	(0.761)
Signal	0.061†	0.061†
	(0.014)	(0.007)
Trial #	-0.212†	-0.055†
	(0.035)	(0.021)
R-square	0.0290	0.0541
Adjusted R-square	0.0280	0.0526
F-statistic	0.0280	35.69
Prob(F-statistic)	0.0001	0.0001
Included Observations	1900	1250

Note: 1) Standard errors are reported in parentheses. † $p < 0.0001$
2) *Trial #* is the trial number and ascends in actual calendar time.
3) The interaction between *Signal* and *Trial #* was not significant and thus not included.

Table 4.6 Random-effects logistic regression of early exit on the signal arrival period, *Experience*, and *Condition* (subject random-effects incorporated)

Variable	Coefficients	Std. Err.	z-stat	p-value
Signal	0.045	0.002	25.59	0.0001
Experience	-1.503	0.127	-11.81	0.0001
Condition	0.508	0.117	4.33	0.0001
Constant	-5.951	0.257	-23.13	0.0001
Log likelihood	-1095.28			
Number of Groups	63			
Included Observations	5249			
Wald Chi^2 test stat. =	684.61	p-value =	0.0001	

Note: The dependent variable was a dummy variable for early exits.
The trials with missing *Delay* were not included.

Although early exits should not occur in theory, we observed early exits in 9.3% of the total trials (584 incidents out of 6278). A logistic regression analysis (table 4.6) revealed that the subjects were more likely to sell before receiving the signal as the signal arrived later. The probability of early exit was reduced in later trials and was higher in static trials than in dynamic trials (figure 4.20).

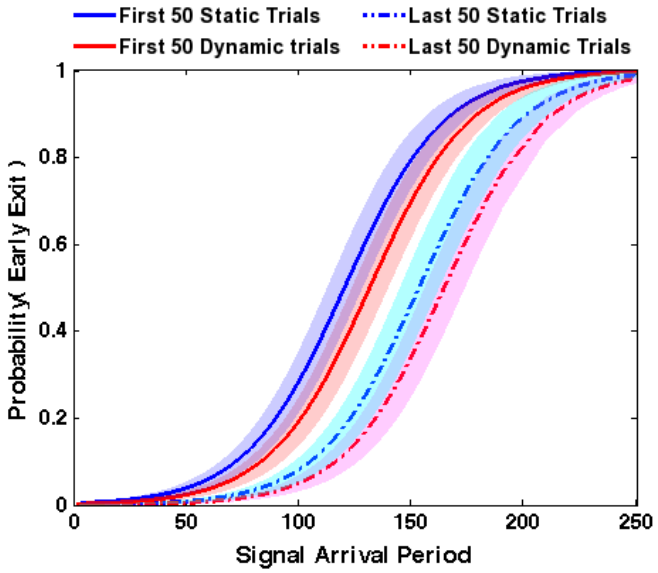

Figure 4.20 Probability of early exit as a function of *Signal*, *Condition* and *Experience*.

Individual Differences in Sensitivity to Signal Arrival vs. Risk Preference

In this subsection, we report preliminary results on the relationship between strategic delay and risk attitude. The estimates for risk aversion parameters for 13 Caltech subjects who participated in our study were obtained from Frydman et al. (Frydman et al. 2009). In Frydman et al. study, individual risk attitude was measured by empirically estimating exponents, ρ^+ and ρ^-, of the utility function for money of the form

$$u(x)=\begin{cases} x^{\rho+} & if \;\; x \geq 0 \\ -\lambda(-x)^{\rho-} & if \;\; x < 0. \end{cases}$$

A smaller ρ^+ (ρ^-) represents more risk aversion (seeking) over gains (losses).

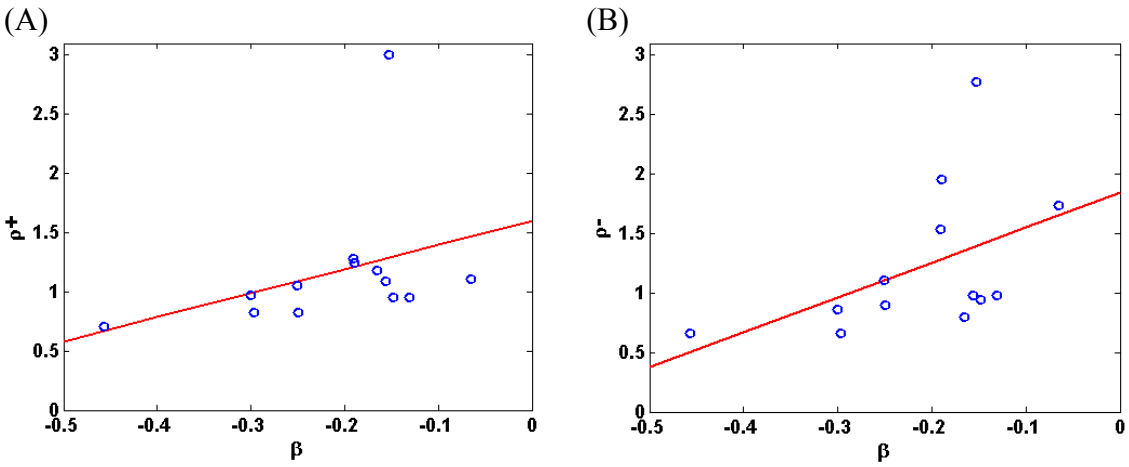

Figure 4.21 Risk aversion measures, ρ^+ and ρ^-, versus β coefficients from the individual regression models (A) The red line is a fitted regression line, $\rho^+ = 1.60 + 2.03\beta$, $R^2 = 0.12$, $F(1, 11) = 1.57$, $p < 0.24$. (B) The red line is a regression line, $\rho^- = 1.85 + 2.94\beta$, $R^2 = 0.23$, $F(1, 11) = 3.24$, $p < 0.10$.

β is a slope coefficient for *Signal* in the individual regression model of $Delay_{uncensored}$ on *Signal* and a constant. β represents sensitivity of subjects' strategic delay to the signal arrival time in static trials. For most of the subjects, β was negative, which implies the subjects waited less before selling as the signal arrived later. A more negative (steeper) β means a greater decrease in strategic delay in response to one period of signal arrival delay. A β value of zero means a subject waited for a fixed length of periods after the signal before selling, regardless of the signal arrival time, as the theory predicted. Due to an insufficient amount of data, β was only modestly correlated with ρ^+ ($r(11) = 0.35$, $p = 0.24$), but without an outlier on the upper right corner, the correlation was significant

($r(10) = 0.64$, $p = 0.025$). β was significantly correlated with ρ^- at a 10% significance level ($r(10) = 0.48$, $p = 0.099$).

4.4 Re-analyzing the BM data

In this section, for the purpose of comparison with our data, we re-analyze and report results from the BM data on the unobservable treatment, all of which were not previously investigated. The unobservable treatment has an identical structure to the dynamic trials in our experiment except for different parameter values. In the original BM experiment's unobservable treatment, there were 6 experimental sessions and 12 human subjects participated in each session. In every trial, there were 2 independent stock markets; 6 subjects participated in one and the rest of them in the other. Each session consisted of 45 trials (called "rounds" in BM). Each experimental period lasted a half second in real time. The following parameter values were used for the experiment:

(a) Number of participants: $I = 6$

(b) Number of sellers necessary for market crash: $K = 3$

(c) Price growth rate: $g = 2\%$

(d) The probability that the true value stops growing in each period: $\lambda = 1\%$

(e) Exogenous ending parameter: $\tau^* = 200$

(f) Window of awareness: $\eta = 90$.

Under this parameter set, the equilibrium strategic delay τ is 23 periods.

Prior to analysis, we discarded observations in which selling occurs within the first 10 periods after the start of the trial, as in BM. For comparability with our results, we did not distinguish sellers according to the order of their signal receiving time. All the definitions

of the variables remain the same. Herein only trials with uncensored and left-censored delay (trials with successful selling) were analyzed.

As shown in our data, mean *Delay* was shorter than the prediction, $\tau = 23$, in all sessions (figure 4.22A) and most of the subjects exhibited mean *Delay* that is significantly shorter than 23 periods (one sample t-test, $t(71) = -18.81$, $p < 2.74 \times 10^{-29}$; signed rank test, *z-value* = -7.33, $p < 2.23 \times 10^{-13}$; figure 4.22B).

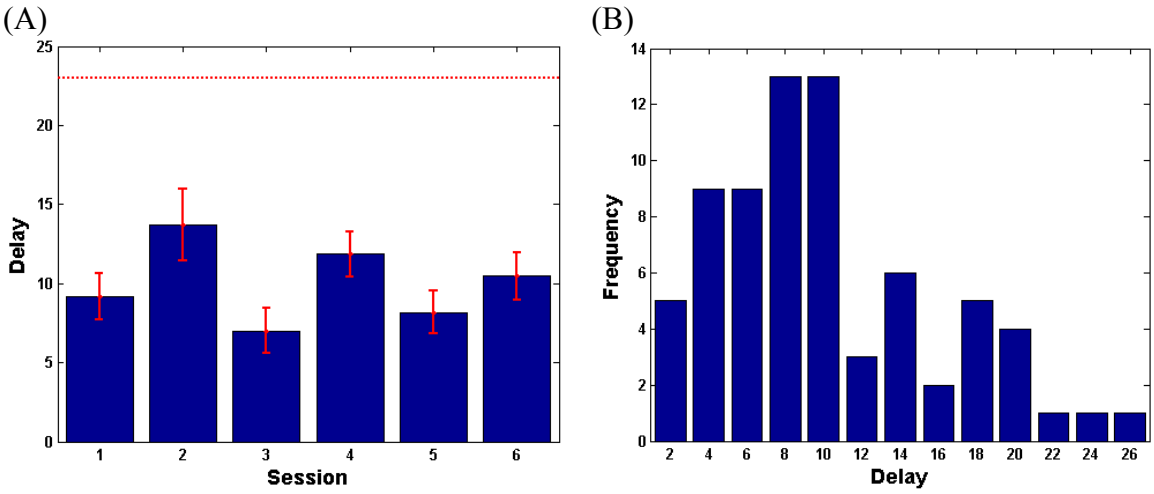

Figure 4.22 Average *Delay*. (A) *Delay* in different sessions. A dashed line in red indicates the theory-predicted delay, $\tau = 23$. Error bars indicate standard errors. (B) Histogram of *Delay*. Bin center: even numbers, starting at 2 and ending at 26; bin size: 2.

Delay was also dependent on the signal arrival time and was negatively correlated with the signal arrival time (figure 4.23, table 4.7). Note that the coefficient for the signal arrival time in the Tobit model (-0.209 in table 4.7) is very close to those in our data ($-0.190 \sim -0.204$; see table 4.2).

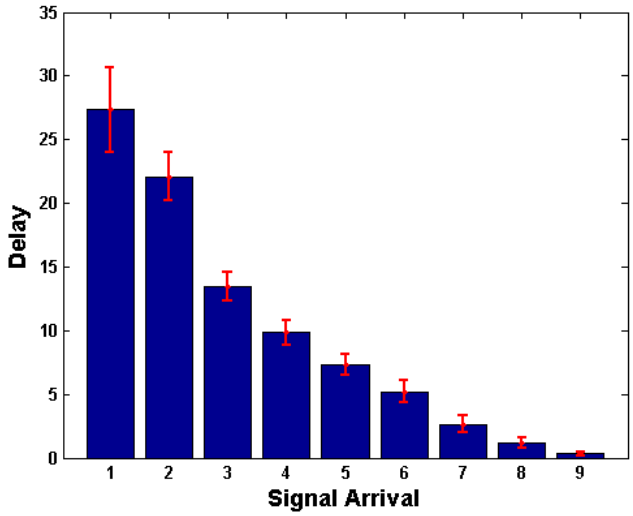

Figure 4.23 *Delay* as a function of the signal arrival period. Error bars indicate standard errors. X-axis: signal arrival period where **1** indicates signal arriving between 0 and 25; **2** between 26 and 50; **3** between 51 and 75; **4** between 76 and 100; **5** between 101 between 125; **6** between 126 and 150; **7** between 151 and 175; **8** between 176 and 200; **9** from 201 onwards.

Table 4.7 Results of random-effects Tobit regression analyses of *Delay* on the signal arrival time (subject random-effects incorporated)

Dependent Variable	*Delay*			
Variable	Coefficients	Std. Err.	z-stat	p-value
Signal	-0.209	0.007	-29.11	0.0001
Constant	36.667	0.794	46.19	0.0001
Log likelihood	-4399.32			
Number of Groups	72			
Wald Chi^2 test statistics = 847.16		p-value =	0.0001	
Left-censored Observations		452		
Uncensored Observations		1055		
Included Observations		1507		

Further, we examined the effect of learning on *Delay*. Regardless of experience, the signal arrival time remained as a significant predictor of *Delay* (figure 4.24, 4.25 and table 4.8). However, *Delay* increased in the last 25 trials compared to the first 20 trials

(figure 4.25 and table 4.8) as observed in our data (see *Delay*s the dynamic trials in figure 4.15).

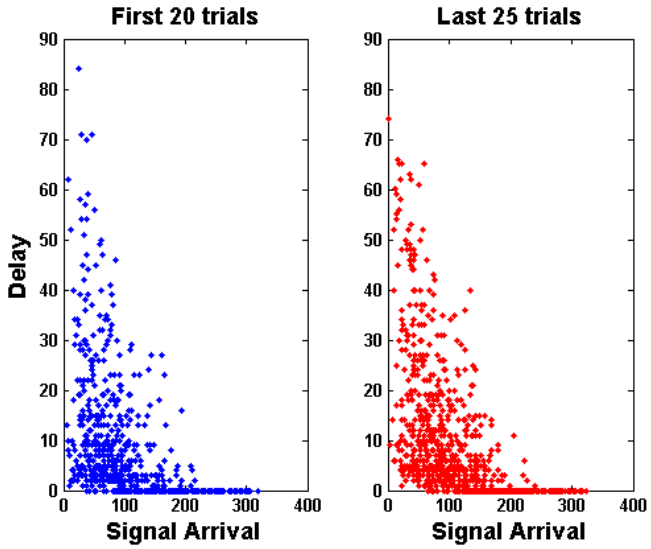

Figure 4.24 Scatter plot of *Delay* and the signal arrival time by experience.

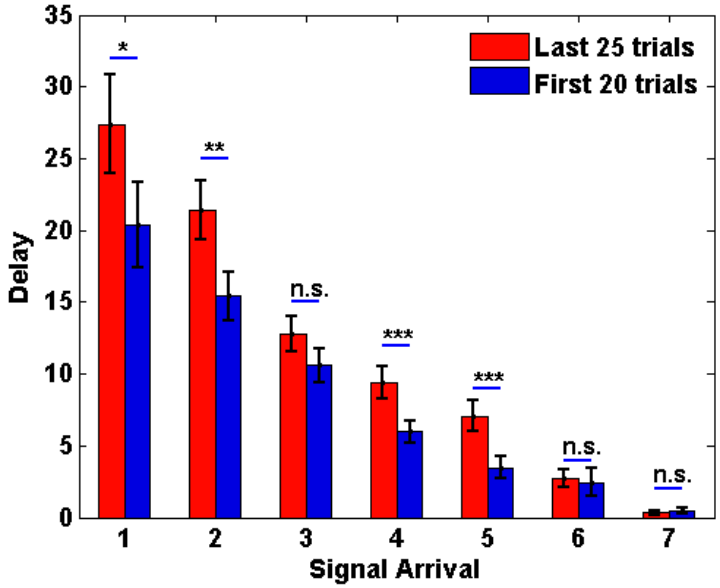

Figure 4.25 *Delay* as a function of the signal arrival period, moderated by experience. *** p <0.01, ** p <0.02, * p <0.07, n.s.: not significant (paired t-test, one-sided). X-axis: signal arrival period where **1** indicates signal arriving between 0 and 30; **2** between 31 and 60; **3** between 61 and 90; **4** between 91 and 120; **5** between 121 and 150; **6** between 151 and 180; **7** from 181 onwards.

Table 4.8 Results of random-effects Tobit regression analyses of *Delay* on the signal arrival time and experience (subject random-effects incorporated)

Variable	A	B
Constant	34.973†	34.342†
	(0.872)	(1.118)
Signal	-0.211†	-.202†
	(0.007)	(0.011)
Experience	2.907†	4.143*
	(0.697)	(1.432)
*Signal * Experience*		-.0139
		(0.014)
Log likelihood	-4388.37	-4388.93
Number of Groups	72	72
Wald Chi^2 stat. (p-value)	867.93 (<0.001)	870.51 (<0.001)
Left-censored Observations	452	452
Uncensored Observations	1055	1055
Included Observations	1507	1507

Standard errors are reported in parentheses. † $p < 0.0001$, * $p < 0.005$.

The deviation measure δ increased as a function of the signal arrival time, but decreased in later trials (table 4.9).

Table 4.9 Results of random-effects regression analyses of δ on the signal arrival period and the trial number (subject random-effects incorporated)

Dependent Variable	δ			
Variable	Coefficients	Std. Err.	z-stat	p-value
Signal	0.038	0.002	15.82	0.0001
Trial #	-0.023	0.006	-3.53	0.0001
Constant	14.235	0.514	27.71	0.0001
R2	0.1427			
Number of Groups	72			
Wald Chi^2 test statistics = 260.86		p-value =	0.0001	
Included Observations		1507		

Note: The interaction between *Signal* and *Trial #* was not significant and thus not included.

In summary, despite differences in design and procedures (e.g., human subjects competing against each other instead of computerized rivals; different parameter values), the BM data showed qualitative results similar to ours.

4.5 Discussion and Conclusions

We have presented four major findings that are not explained by existing theory. First, the subjects' strategic delay decreased as a function of the signal arrival period, despite the theoretical prediction of a fixed delay, independent of the signal arrival time. Second, the theoretical isomorphism was inconsistent with the experimental data: the subjects consistently delayed less in the dynamic clock game than in the static game. Third, we found that the subjects' strategy evolved as they gained more experience. Fourth, albeit preliminary and not conclusive, there appeared to be a relationship between risk preference and strategic delay.

Regardless of group, condition, or experience, the length of strategic delay was negatively correlated with the signal arrival time, contrary to theory. This pattern was persistent after a significant amount of experience, so this rules out the possibility that the effect is solely attributable to subjects' mistakes. This was true when human subjects competed against other human participants (as in the BM data) or when they played against computer players employing a theory-predicted strategy. The use of computer players precludes the possibility that this pattern is attributable to the best response of the sophisticated subjects to the suboptimal strategy of the naïve. However, one could still argue that the signal dependence might have been caused by subjects' false belief about computer players' strategy (the strategy was unknown to the subjects). This possibility

could be tested by conducting an experiment where the subjects know the computer players' strategy. Alternatively, the signal dependence might be due to a misperception about the bubble starting time (or generating process), whose memoryless property might not be intuitive. This could be tested by using a different generating process for bubbles such as a uniform distribution.

The signal dependence might also be attributable to subjects' risk aversion. In this game, the decision to be made in every period is a gamble between reaping the sure payoff by selling immediately or waiting for another period to get a risky but (exponentially) higher reward. One possible heuristics that could have been used is as follows: When the market is still open and a subject has not received his signal yet, he can think of two salient possibilities: the event t_0 has not happened yet (hence, everyone is getting their signal late); or the event t_0 has occurred but he happens to be the unlucky one who has a long delay between the event t_0 and the signal arrival. When he is in a later time period, the probability of the former event decreases and the probability of the latter increases. Although the size of reward exponentially increases as a function of time, it might have not been enough to compensate for the subject's risk-aversion. Recall that the theoretical equilibrium strategy was derived with an assumption of risk neutrality. The possibility that risk aversion may be the cause for decreasing delay against the signal time is modestly evidenced by the preliminary findings on the correlation between the β coefficient and the risk aversion measure ρ. However, more data would be necessary to make a conclusive argument or perhaps, the risk aversion parameter could be estimated from the current data by way of simulations.

A more interesting finding is that the strategic delay is shorter in the dynamic condition than in the static condition, controlling for the signal arrival time and experience.

One possibility is that given time constraints, reasoning and strategic computations in dynamic trials were not as precise as in static trials and led to more erroneous decisions. If this were the case, this error should be diminished as the subject gain more experience. Indeed, in later trials the dynamic and static differences were attenuated, but did not disappear and remained significant in later trials, by which point the subjects would have acquired substantial experience. This needs further investigation in experiments by manipulating clock-times of one experimental period to see if the dynamic vs. static differences disappear as the clock-time length of a period becomes longer, or by manipulating cognitive loads in static trials (giving distractions or extra tasks such as memorizing phone numbers, applying time pressure).

One could argue that the reason for earlier selling times in the dynamic condition is that there is an implicit cost of waiting. If the subjects had wanted to finish the experiment earlier to do something else or just for leisure, every second they waited could have incurred an implicit cost—an opportunity cost or a shadow price of waiting. However, because they were competing with the computer rivals, they could not actually speed up the experiment as they would have wished. Had they sold early in order to terminate the round, they would have still had to wait until the other two computer players sold, not earning as much. Furthermore, had there been better opportunities to earn rewards, they would have not shown up for the experiment in the first place. Also, they could have actually sped up the static trials by keeping on pressing the enter key and submitting the randomly selected anchor values in every static round. However, most of the subjects

indeed took time to explore and submit their own selling times, and this was evidenced by the fact that their selling times were uncorrelated with the default anchor values.[12] In addition, note that the average delay signal for each time bin in the dynamic case actually increased, rather than decreasing in later trials. This seems inconsistent with the shadow price argument since by then, distracted subjects (by other fun outside activities) most likely would have been even more bored of playing the same game over 50 times (which means increased shadow price).

Different kinds of implicit costs from waiting are possible. One such disutility of waiting (or utility of early quitting) might come from timing of uncertainty resolution. People prefer resolution of uncertainty earlier than later (Mossin 1969; Wu 1999) because of its informational value that allows them to psychologically and physically better plan prior to outcome delivery. The shorter delay in dynamic trials might be due to this preference for the earlier resolution of uncertainty.

Early selling might have reduced some other psychological toll on the subjects. Recall that in dynamic trials, a subject makes a flow of decisions between selling now or waiting another period in every period, in a faced-paced environment, under time pressure—each period only lasts a quarter second. Participating in a rapidly progressing, dynamic game

[12] On a side note, this kind of waiting game could be quite engaging. In Cox et al. (Cox et al. 1982), they compared the Dutch auction with the first-price sealed bid auction using lab experiments to test the strategic equivalence between the two mechanisms. They found that the prices in the Dutch auction were lower than those in the first-price auction. One of their explanations of this discrepancy was that subjects got additional utility, the "utility-of-suspense" from playing the Dutch auction "waiting game", and hence they waited longer, decreasing the prices. This was based on many subjects' comments that they enjoyed the Dutch auction experiment more than the other auction formats because of the "suspense of waiting". In personal conversations, some of the subjects in this study also made similar comments.

can be stressful, anxiety provoking, and make one feel "out of control."[13] Decision to wait means another decision in the next period, and this might have come at a psychological cost such as stress or anxiety, resulting in the shorter delay in dynamic trials[14]. This conjecture could be tested by inducing higher stress and anxiety levels (Porcelli & Delgado 2009)—inducing anxiety in the dynamic case only would further reduce the delay, and anxiety in the static case only would decrease the gap between the two conditions. If empirically proven, a model incorporating risk aversion and distaste for anxiety built in the utility function could provide a more accurate description of timing decision making.

Lastly, we have sufficient data to study learning. Although it is not discussed here, it would be interesting to fit different types of learning models, reinforcement-based or regret-based, to see how the strategic delay is adapted over time.

[13] Jap (2003), in field interviews with suppliers participating in reverse auctions, found that "the compressed time frame of open-bid auctions creates a stressful context for the supplier and many suppliers complained that the format prevented them from carefully considering price bids and gave them a sense of being 'out of control.'"

[14] Wood and Schweitzer (Wood & Schweitzer 2009) in their experiments induced different levels of anxiety (high vs. low) and found negotiators experiencing high levels of anxiety make steeper concessions and exit bargaining situations earlier.

4.6 Appendix for Chapter 4

4.6.1 Additional Results

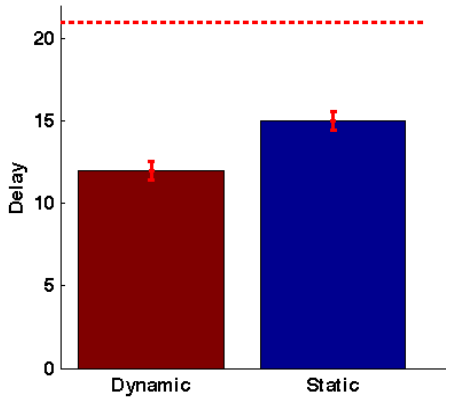

Figure A4.1 Difference in average *Delay* between dynamic and static trials. The trials with right-/left-censored or uncensored *Delay*s were included. A dashed line indicates a theory-predicted delay, $\tau = 21$. Error bars indicate standard errors. Paired t-test, $t(62) = -5.46$, $p < 8.93 \times 10^{-7}$; signed rank test, *z-value* = -4.57, $p < 0.0001$.

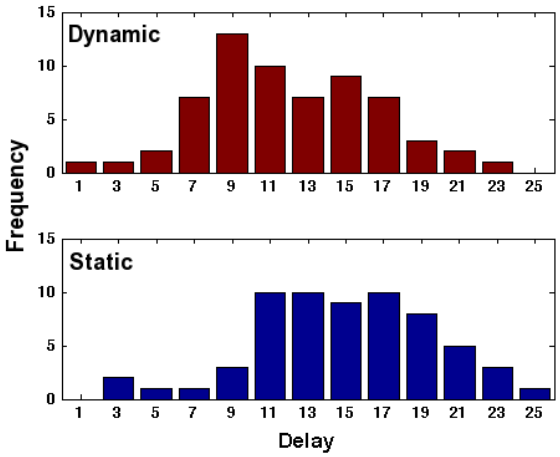

Figure A4.2 Distribution of individual average *Delay* between dynamic and static trials. The trials with right-/left-censored or uncensored *Delay*s were included.

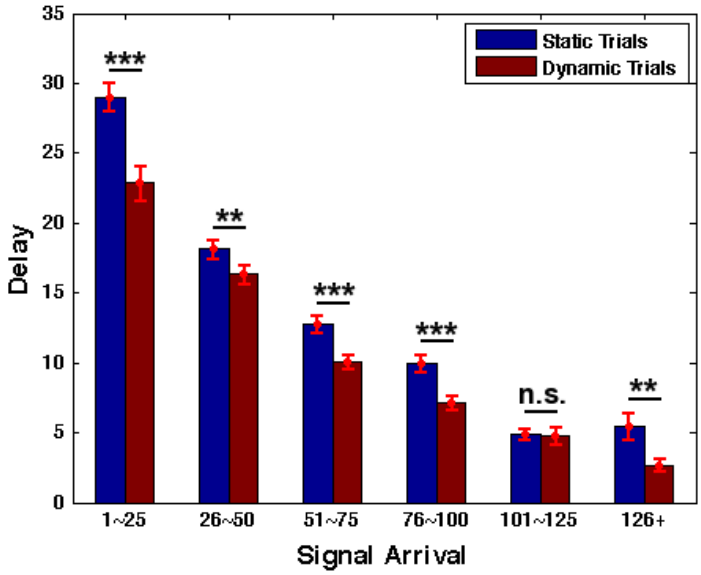

Figure A4.3 Delay as a function of signal arrival period. Delay decreases as a signal arrives late. The trials with right-/left-censored or uncensored *Delay*s were included. Error bars indicate standard errors. ***$p < 0.001$, ** $p < 0.01$, n.s.: not significant (paired t-test, one sided).

Dependent Variable	Delay			
Included Observations	5249			
Variable	Coefficients	Std. Err.	z-stat	p-value
Signal	-0.250	0.006	-42.62	0
Condition	2.729	0.413	6.6	0
Constant	30.836	0.477	64.59	0
Log likelihood	-13721.039			
Pseudo R-square	0.061			
Left-censored at *Delay* = 0		584		
Uncensored		3012		
Right-censored		1653		

Table A4.1 Results of censored (normal) regression analyses for *Delay*. The trials with right-/left-censored or uncensored *Delay*s were included

4.6.2 Instructions

Instructions (A)

Thank you for participating in this experiment on the economics of investment decision making. By participating, you have already earned a $5 show-up fee. If you follow the instructions carefully you may be able to make additional money that will be paid to you at the end of the experiment privately, in cash. The dollar amounts mentioned below (and on your screens) are 'experimental dollars.' We will use a fixed conversion rate to convert experimental dollars to real cash. The conversion rate is 100 Experimental Dollars = 0.50 US Dollar

You are about to make selling decisions 100 times in a row. Each round represents one trading round. You will be paired with 5 other **computer** players, who are also sellers, in every round. Note that you are **NOT** playing with other human participants in the room. To make sure you understand the procedure, please complete the quiz following these instructions. If you cannot answer the quiz questions yourself (with a little guidance from us) you will not be allowed to proceed. Please see the power point slides in front of you for the figures referred to below.

Half of the rounds will be played in a dynamic condition and the other half in a static condition. We will explain these two conditions next.

Dynamic Trading Rounds
In a dynamic trading round, your job is to decide in real time when to sell an asset you are holding. At the start of each trading round, everyone gets one share of the same asset. The price of the asset begins at $ 1 (figure 1) and grows exponentially as trading time periods pass (figure 2).

Each period lasts 250 milliseconds. In every period you are making a real-time decision about whether to sell or whether to wait. When you decide to sell your asset, press the ENTER key. Once you have decided to sell, you have no more decisions to make—you just wait until the trading round ends. That is, you make one and only one selling decision in each round.

The price of the asset increases by 4% in each trading period. You will see the price increasing on your computer screen—graphically in the price-period plot at the center of the screen. The *graph* of prices is shown in the center of the screen (figure 3) and the *numerical* price is shown in the top right corner (figure 4). The current price is the same for all sellers (you and five *computer* players).

It is possible that the price graph will go out of the top of the screen in a very long round (figure 5). This is normal and does not change anything about your decisions. Even if the graph hits the top of the screen, the current price will continue to rise and is displayed on the top right of the screen.

There are two ways a trading round can end:
(1) Once 3 players have all decided to sell; **or**
(2) The trading round reaches 300 periods after the ***true value*** (explained below) of the asset has stopped growing, even though fewer than three players have sold.

These rules mean that among the 6 players (including you and five other *computer* players), **only** 3 players can actually sell their stock before the round ends. If you sell before the round ends, you earn an amount in $ equal to the price at which you sold. If three others sell before you do, and the round ends, then you will earn the maximum ***true value*** of the asset, as explained below.

Here is how the ***true value*** is determined. At the start of each trading round, the computer randomly selects a trading period at which the true value growth will stop. The true value is equal to the asset price until this pre-determined stopping period; that is, the true value grows with the price (at the rate of 4%) and stops growing after the pre-determined stopping period is reached, while the price continues to grow (at the rate of 4%) in excess of the true value. The maximum true value of the asset is equal to the asset price in the stopping period. If the trading round continues beyond this period, the asset price still grows as before, but the true value of the asset stays at its maximum true value.

You will not learn right away when the true value has stopped growing. After the true value stops growing, you will only receive a message (notice), with a delay that is equally likely to be anywhere from 0 to 60 periods, indicating that the current price of the asset is above its true value. So you will never know the exact time period in which the true value stopped growing. A red arrow will indicate the period in which you have received that notice and that period will also be posted on the right-hand side of the screen (figure 6, 7).

The random delay between the time the true value stops growing and the message (notice) arrival is randomly chosen ***separately*** for each player. Therefore, when you receive your message, some players may have already received it, and some might not have received it yet. Also, all the messages are private, so you do not know when other players receive messages and they do not know when you received your message.

In each trading period, there is a 2.5% chance that the true value will stop growing (if it has not stopped already). This percentage is constant in each period, which implies that even if the round has just begun, there is a small (2.5%) chance it will end in the next period, and even if the round has lasted a long time, there is a small (2.5%) chance it will end in the next period.

At the end of each trading round, your earnings and the earnings of other players for this round will all be displayed on your screen (figure 8). Note that in this example, the 3 players who did not sell in time will receive the maximum true value of the asset ($3.20 in the sample screen in figure 8).

In each round, you will be grouped with 5 new computer players. They will receive the same amount of information as you do (though at a different time in regards to the maximum true value message). However, they will be playing with a pre-determined strategy, which will remain consistent throughout the experiment, but unknown to you.

Static Trading Rounds

The static trading rounds are the same as the dynamic rounds with one difference: You are presented with your message, and with all of the other information that is revealed across time in the dynamic rounds, all at once (figure 9). More specifically, you get to step outside of time and see what the stock price would be at any time– under the assumption that the market had not ended by that time. The period in which your message is received is shown with a red arrow; remember that the message is still randomly delayed from the time when the true value stopped growing. You then choose at what price along the curve you would sell (and the five computer players will be doing the same).

You can explore the stock price over time by moving a little blue bar on the screen (figure 10). Place a cursor on any point on the white bar below the x-axis of the price-period plot to move the blue bar. You do **NOT** have to hold a button on the mouse. The period and the price corresponding to the location of the blue bar will be automatically updated in the plot.

When you have determined the period in which (or the price at which) you want to sell your asset, press the ENTER key to submit your decision.

Once you submit your decision, the computer compares the selling times of all 6 players (you and 5 other computer players), and determines earnings as in a dynamic round. First, the computer calculates the time of the market ending as the time by which the first 3

players had sold. If your selling time is before the market ended, you will earn an amount equivalent to the asset price at which you sold. If your selling time is after that end time (i.e., three other players' chose selling times earlier than yours), you earn the maximum true value of the stock as in the dynamic round.

To summarize:

1. You will be playing with 5 computer players (not other people in the room).
2. Only 3 players can sell their asset before the trading round ends.
3. In each trading round, the computer randomly chooses a period and the true value of the asset grows with the price up to this period. Afterwards, the price exceeds the true value.
4. Those 3 players who submit their decision in time will receive earnings equivalent to the asset price in the period in which they sold their asset. Others will receive the maximum true value of the asset.
5. In dynamic trading rounds, sellers make a real-time decision and will receive a message with a random delay after the true value has reached its maximum.
6. In static trading rounds, sellers learn when a message would have arrived (still with a random delay), had they played this round in real time, and you can make a selling decision without time pressure.

For you to become familiar with the software, we will ask you to complete 10 practice rounds (5 dynamic rounds followed by 5 static rounds). After completing the practice, you will go through 5 blocks of rounds with each block consisting of 10 dynamic rounds followed by 10 static rounds.

At the end of the experiment, we will sum up the money you earn in all trading rounds and pay you in cash. The conversion rate is 100 Experimental Dollars = 0.50 US Dollar. Your total earnings will be what you earned in the 100 rounds plus your show-up fee of $5. Are there any questions?

You can read these instructions again yourself until you understand how the experiment works. When you are ready, please complete the quiz and let the experimenter know when you are done.

4.6.3 Figures for the instructions

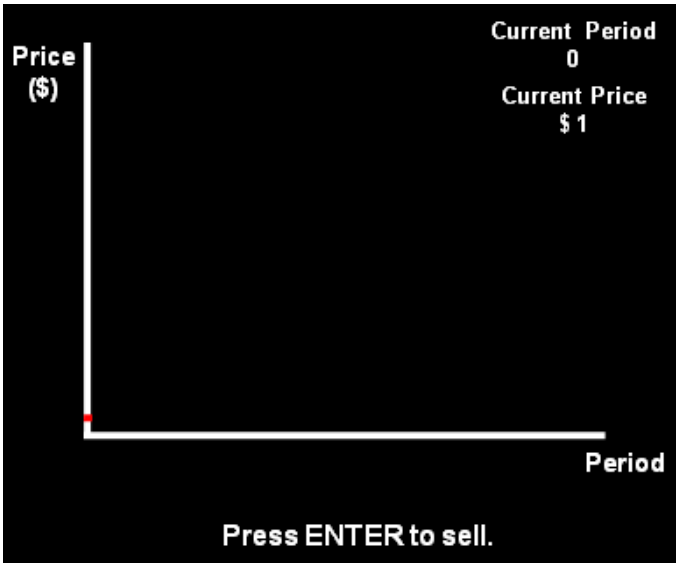

Figure 1 The start of a trading round.

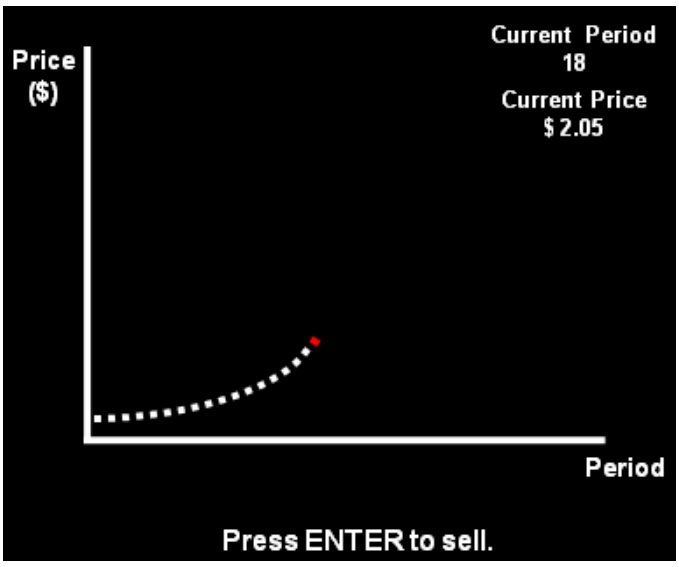

Figure 2 A trading round in action. The price grows exponentially.

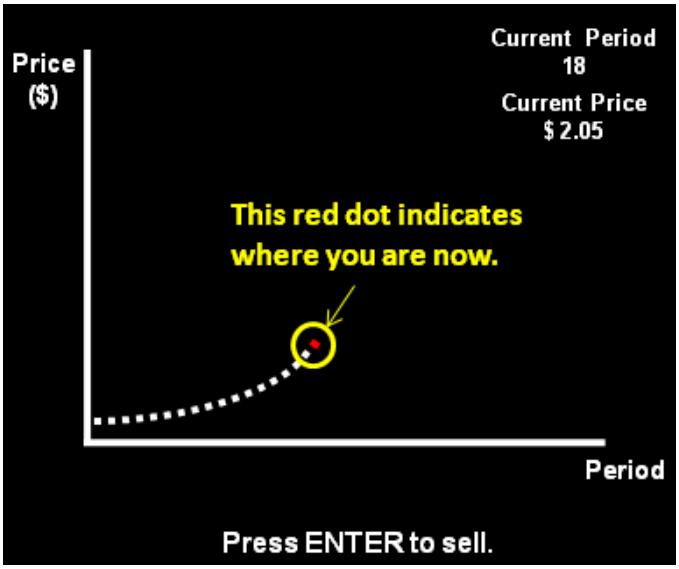

Figure 3

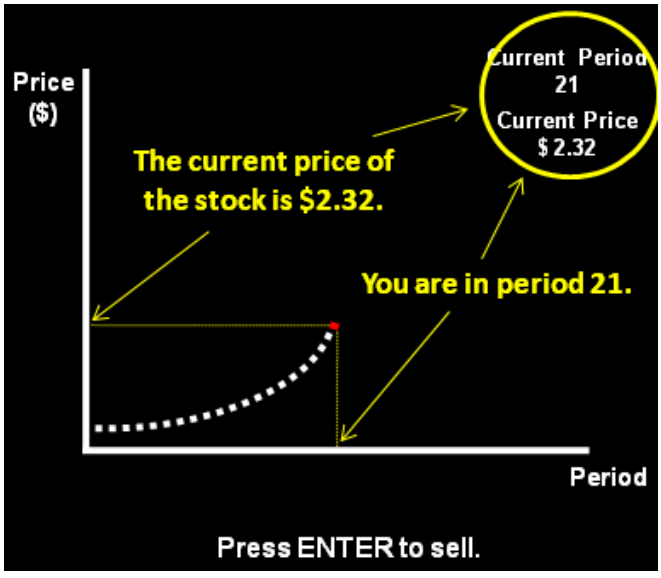

Figure 4

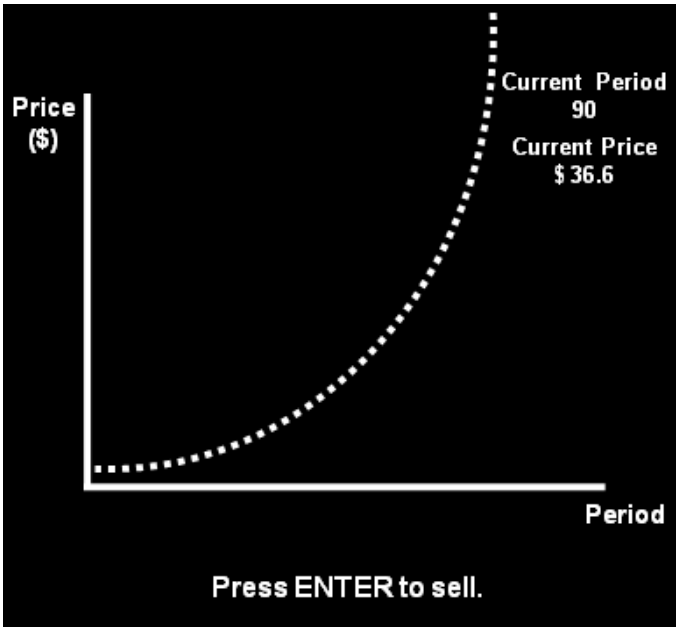

Figure 5 This is not an error.

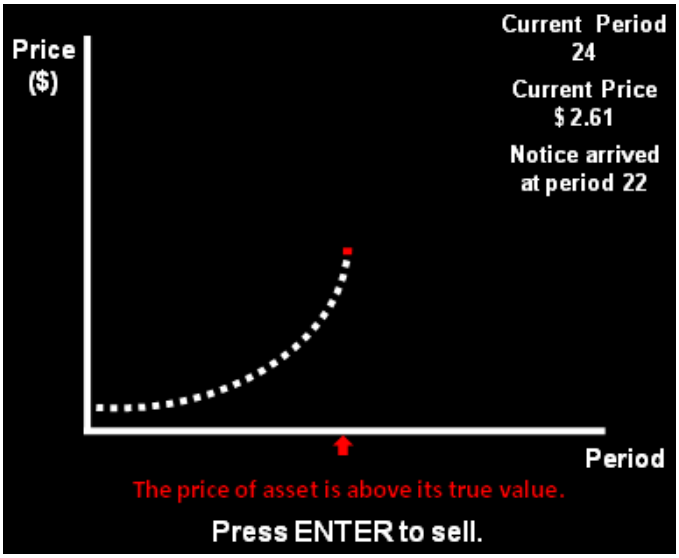

Figure 6 You will be notified that the true value has stopped growing already.

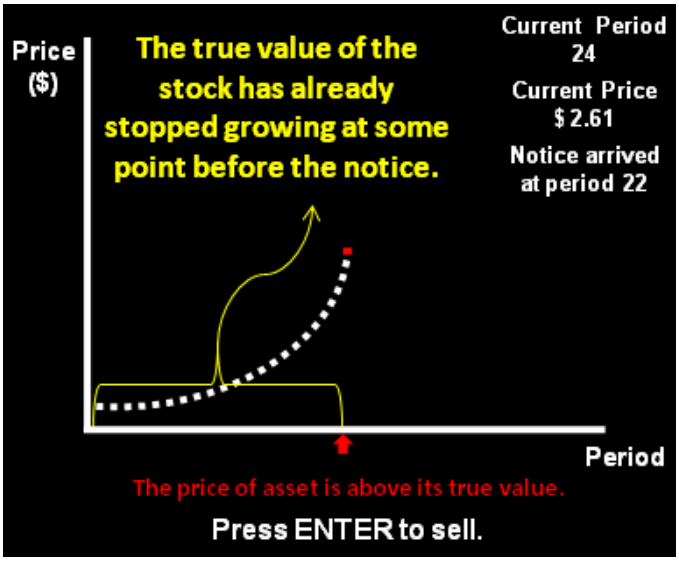

Figure 7 Note that the sentence in yellow will not appear in the actual screen.

Figure 8

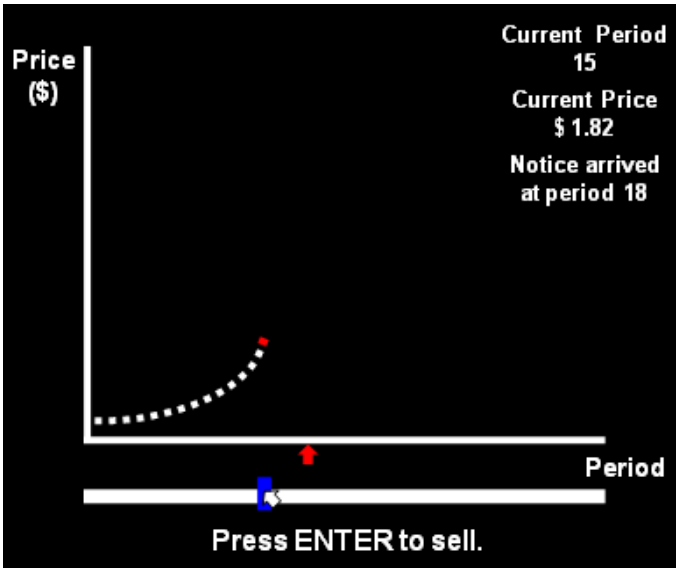

Figure 9 A static trading round.

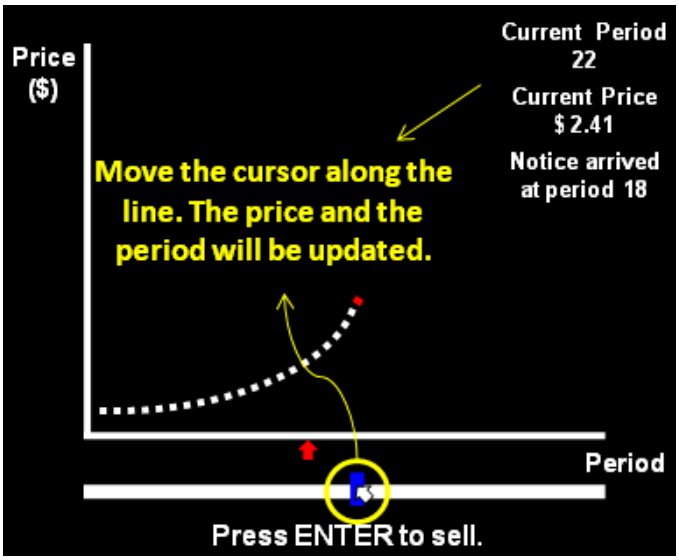

Figure 10 Explore the price in each period with the mouse before submitting your decision.

4.6.4 Quiz

Name:

Quiz

True or False

1. You are playing with other people in the room. (True, False)
2. The price of the asset increases at different rates in dynamic and static rounds. (True, False)
3. The price of the asset never decreases or stops. (True, False)
4. The maximum true value of the asset is the same as the asset price in the period when you receive a notice. (True, False)
5. The true value of the asset always matches the price. (True, False)
6. You earn the maximum true value of the asset when the market ends before you decide to sell. (True, False)
7. In static trading rounds, all players earn an amount equal to the price during the period in which they chose to sell. (True, False)
8. There is a 2.5% chance that the true value stops increasing in every period in static rounds. (True, False)
9. There is a 2.5% chance that the asset price stops increasing in every period in static rounds. (True, False)
10. All the players receive the notice that the true value stops growing, but at different time points. (True, False)
11. Your selling price is always greater than the true value of the asset. (True, False)
12. A trading round ends once three players sold their asset. (True, False)

Bibliography

Brunnermeier, M. K., & J. Morgan. (2006). *Clock games: Theory and experiments*. Unpublished manuscript.

Brunnermeier, M. K., & J. Morgan. (2008). *Clock games: Theory and experiments*. Unpublished manuscript.

Cox, J. C., B. Roberson, & V. L. Smith. (1982). Theory and behavior of single object auctions. *Research in Experimental Economics,* **2** 1-43.

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

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

Dutta, P. K., S. Lach, & A. Rustichini. (1995). Better late than early: Vertical differentiation in the adoption of a new technology. *Journal of Economics and Management Strategy,* **4**(4) 563-589.

Engelbrecht-Wiggans, R., & E. Katok. (2007). Regret in auctions: Theory and evidence. *Economic Theory,* **33** 81-101.

Frydman, C., C. F. Camerer, P. Bossaerts, & A. Rangel. (2009). *MAOA is associated with changes in the computations made in risky choice situations*. Unpublished manuscript.

Fudenberg, D., & J. Tirole. (1985). Preemption and rent equalization in the adoption of new technology. *Review of Economic Studies,* **52** 383-402.

Greene, W. H. (2003). *Econometric analysis* (5 ed.): Prentice Hall.

Jap, S. D. (2003). An exploratory study of the introduction of online reverse auctions *Journal of Marketing,* **67**(3) 96-107

Kagel, J. H., R. M. Harstad, & D. Levin. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica,* **55**(6) 1275-1304.

Kagel, J. H., & D. Levin. (2001). Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic vickrey auctions. *Econometrica,* **69**(2) 413-454.

Kagel, J. H., & D. Levin. (forthcoming). Auctions: A survey of experimental research, 1995-2008. *Handbook of Experimental Economics,* **2**.

Katok, E., & A. M. Kwasnica. (2006). Time is money: The effect of clock speed on seller's revenue in dutch auctions. *Experimental Economics,* **11** 344-357.

Lambkin, M. (1988). Order of entry and performance in new markets. *Strategic Management Journal,* **9** 127-140.

Lieberman, M. B., & D. B. Montgomery. (1988). First-mover advantages. *Strategic Management Journal,* **9** 41-58.

Lucking-Reiley, D. (1999). Using field experiments to test equivalence between auction formats: Magic on the internet. *American Economic Review,* **89**(5) 1063-1080.

Milgrom, P. R. (1989). Auctions and biddings: A primer. *Journal of Economic Perspective,* **3**(3) 3-22.

Milgrom, P. R., & R. J. Weber. (1982). A theory of auctions and competitive bidding. *Econometrica,* **50**(5) 1089-1122.

Mossin, J. (1969). A note on uncertainty and preferences in a temporal context. *American Economic Review,* **59**(1) 172-174.

Neugebauer, T., & R. Selten. (2006). Individual behavior and first-price auctions: The importance of information feedback in computerized experimental markets. *Games and Economic Behavior,* **54** 183-204.

Porcelli, A. J., & M. R. Delgado. (2009). Acute stress modulates risk taking in financial decision making. *Psychological Science,* **20**(3) 278-283.

Robinson, W. T., & C. Fornell. (1985). Sources of market pioneer advantages in consumer goods industries. *Journal of Marketing Research,* **22**(3) 305-317.

Tversky, A., & D. Kahneman. (1974). Judgment under uncertainty: Heuristics and biases. *Science,* **185**(4157) 1124-1131.

Wood, A. M., & M. E. Schweitzer. (2009). *Can nervous nelly negotiate? How anxiety causes negotiators to exit early and make steep concessions*. Unpublished manuscript.

Wu, G. (1999). Anxiety and decision making with delayed resolution of uncertainty. *Theory and Decision,* **46** 159–198.