

Behavioral Biases in Information Acquisition

Thesis by

Lauren Elizabeth Feiler

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy



California Institute of Technology

Pasadena, California

2007

(Defended May 21, 2007)

© 2007

Lauren Elizabeth Feiler

All Rights Reserved

For Tom and his infinite patience.

Acknowledgements

Chapter 4 is joint work with Jacob Goeree and Leeat Yariv. We gratefully acknowledge financial support from the National Science Foundation (SES 0551014), the Alfred P. Sloan Foundation, and the Dutch National Science Foundation. We thank attendees at the 2006 ESA meetings in Atlanta and Tucson for comments.

The experiments in Chapters 2 and 3 were funded by Colin Camerer. I am very grateful for this funding, especially since this work did not involve fMRI data.

I also wish to thank Colin Camerer for his support, encouragement, and advice. I am constantly impressed with his ability to conduct research, quickly write clear and interesting papers, and entertain students both in and outside of class. I am grateful for everything I have learned from him. The members of the Camerer lab group have also given me many useful comments over the years. Jacob Goeree and Leeat Yariv have offered me much valuable guidance and insight, and I have appreciated the opportunity to learn even more while coauthoring with them.

I have also benefitted from working with (and enjoying the enthusiasm of) Charlie Plott and Preston McAfee. I would like to thank J. Morgan Kousser for advice on writing, particularly for Chapter 2 of this thesis, and Antonio Rangel for serving on my committee. I thank P.J. Healy for providing me with both excellent advice and friendship.

I would not be writing this were it not for the encouragement, goodwill, and good humor of my classmates, who always make me grateful that I came to Caltech when I did. Thanks to them, I did not have to spend much time under my desk.

Finally, I thank my family for their love and support, and my husband, Thomas Feiler, for many forms of assistance. Tom created the software used in Chapter 2 and

assisted with programming the experiments in Chapter 3. In addition to letting me put his amazing programming skills to work, Tom has allowed me to drag him across the country not once, but twice now. I cannot express how much I treasure Tom's willingness to join me in life's adventures.

Abstract

This dissertation examines two related biases in information acquisition, information avoidance and selective exposure to information. The first essay focuses on the use of information avoidance to justify self-serving decisions. Participants in an experimental dictator game are given the chance to avoid costless information about a recipient's payoffs. Many dictators choose an allocation that maximizes their own profit without learning whether this allocation will help or hurt the recipient. Even subjects who make equitable choices when non-aligned payoffs are known will avoid information, especially when it is likely that doing so will not hurt the recipient. Through assessing the role of beliefs in information avoidance, this chapter provides an evaluation of several theoretical models of avoidance.

The second essay considers the effects of both information avoidance and information-seeking behavior on charitable donations. Experimental subjects are allowed to avoid or seek out a range of information about a charity before deciding how much to donate to the organization. Donation sizes are positively correlated with the amount of information subjects choose to obtain. When subjects are required to read descriptions of charities, longer descriptions lead to higher donations. This indicates that agents may avoid further information if they have already learned about a person or charity in need, because learning more could obligate them to give more.

The final essay studies selective exposure, the tendency to seek information that could support or validate one's beliefs or preferences but not maximize payoffs. Subjects in a context-free environment have to guess an unknown state of nature, and we induce preferences for one particular state. When given a choice between different information sources, around half of all subjects choose a source that potentially

confirms that the state is the one they prefer, but is inferior in terms of expected payoffs. This finding holds consistently across a variety of contexts. The results of these studies have implications within experimental economics, since experiments tend to impose information on subjects that they might not otherwise gain. They also demonstrate that the ability to selectively acquire or avoid information can have a large impact on economic decisions.

Contents

Acknowledgements	iv
Abstract	vi
1 Introduction	1
2 Patterns of Information Avoidance in Binary Choice Dictator Games	4
2.1 Introduction	4
2.2 Studies of Rationalization	7
2.3 Experimental Procedures	12
2.4 Potential Models of Information Avoidance	15
2.5 Results	25
2.5.1 Choices in the Nonaligned Game	29
2.5.2 Comparison to Dana, Weber, and Kuang	30
2.5.3 Patterns of Avoidance	30
2.6 Conclusion	34
2.7 Appendix: Experimental Instructions	39
3 The Effect of Selective Information Acquisition on Charitable Giving	43
3.1 Introduction	43
3.2 Experimental Design	46
3.3 Hypotheses	49
3.4 Results	51

3.5	Tests of Causality	57
3.5.1	Hypotheses	58
3.5.2	Results	59
3.6	Conclusion	64
3.7	Appendix: List of Organizations	66
4	An Experimental Study of Selective Exposure	70
4.1	Introduction	70
4.1.1	The Selective Exposure Bias	70
4.1.2	Related Literature	72
4.2	Experimental Design	75
4.3	Theoretical Predictions	78
4.3.1	Individual Choice	78
4.3.2	Majority Rule Voting	79
4.3.3	Partners	81
4.4	Hypotheses	82
4.5	Results	83
4.5.1	Overall Performance	83
4.5.2	Voting Treatments	87
4.5.3	Partners	87
4.5.4	Learning	88
4.6	Potential Causes of Selective Exposure	92
4.6.1	Small Payoff Differences	92
4.6.2	Failure to Backward Induct	93
4.7	Conclusion	94
4.8	Appendix 1: Dominant Strategy in the Individual Choice Treatment .	96
4.9	Appendix 2: Equilibria in Majority Rule Voting Treatments	98
	Bibliography	103

List of Figures

2.1	Screenshot from the experiment	13
2.2	Payoff sets (payoff to dictator, payoff to recipient)	14
2.3	Game tree of decisions in Payoff Set 1	19
3.1	Screenshot from stage 2	48
3.2	Average donation, by amount of description read	53
3.3	Time spent per word (milliseconds) on hidden names that have been revealed, by donation and preference for charity	57
4.1	Average amount of selective exposure and payoff maximizing guesses, by period	89
4.2	Frequency of subjects' indicator choice patterns	90
4.3	Subjects who learned to choose the payoff-maximizing indicator	91

List of Tables

2.1	Predictions of revelation patterns	25
2.2	Information decisions and dictator game choices	26
2.3	Probit estimates. Dependent variable is decision to reveal.	28
2.4	Frequency of revelation patterns by game	31
2.5	Total number of cases by revelation pattern and choice at $p = 0$	33
3.1	Summary information	52
3.2	Ordered logit of donation on variables listed and dummy variables for all charities	54
3.3	Summary information, forced reading treatments	59
3.4	Ordered logit of donation on variables listed and charity indicator vari- ables, forced reading treatments	61
3.5	Ordered logit for combined treatment data	62
3.6	First differences: Estimated changes in the probability of donating \$0, \$10, or \$20 when the proportion read increases, by treatment	63
4.1	Summary of Sessions	77
4.2	Proportion of selective exposure choices across treatments	84
4.3	Frequency of payoff-maximizing guesses following null signals, by treat- ment, indicator choice, and payoffs	86
4.4	Expected values of red and blue guesses for payoffs of (150, 50)	96
4.5	Expected values of red and blue guesses for payoffs of (80, 60)	97

Chapter 1

Introduction

This dissertation studies two ways in which people justify economic decisions through selectively acquiring information. In the next two chapters, agents can intentionally choose not to learn something in order to make self-serving decisions; this action will be referred to as information avoidance. One example is failing to pick up a ringing phone in case a charity is calling to ask for money. The last chapter considers the selective exposure bias, which is demonstrated when agents select information sources that can support their preferences or beliefs but do not maximize payoffs. This bias allows agents to justify decisions that they have already made but can prevent them from making better choices. This thesis uses an experimental approach to gain insight on practices that can be difficult to observe empirically.

The next chapter examines the role of beliefs about another person's payoffs in the decision to avoid information in a dictator game, and it tests models that might explain information avoidance. This chapter provides experimental evidence that agents who avoid information before making a self-serving decision are generally ones that would choose the same option under full disclosure. However, some people who choose equitably when they know that the option that is best for them could hurt another person will use the ability to avoid information in order to make more selfish decisions. Such agents avoid information more often when they believe that they are unlikely to hurt another person by doing so.

Chapter 3 expands upon the findings of the preceding chapter to consider more complex situations in which people might choose to obtain varying amounts of infor-

mation. Rather than facing a simple option to gain or avoid information, experimental subjects can choose to read any amount of a description of a charity before deciding how much to donate to the organization. The more subjects read about a charity, the more they donate. Reading more causes larger donations in the experiments presented in this chapter; it does not appear to be the case that people who choose to donate simply have a preference for learning more about the cause they support. The findings of this chapter imply that even people who take the time to give a homeless man money or make a donation to a charity may be preventing themselves from learning more about the man or the cause to avoid having to donate more. This chapter also finds some evidence of guilt among subjects who see the name of a charity they like but do not make a donation, which may be another reason to avoid information.

The final chapter focuses on biased information seeking, not avoidance. As an example of the selective exposure bias studied in this chapter, someone might choose to read favorable articles about a political candidate that she already supports instead of reading about the opponent, so the voter would not learn if the opponent's policies were the ones most closely aligned with her preferences. This bias is studied in a simple laboratory experiment in which preferences for a certain state of the world are induced. Before guessing what the state is, subjects can choose to gain information from a source that can maximize payoffs or a less informative source that could serve to confirm that the state is the one that is preferred. Across a variety of contexts, subjects demonstrate the selective exposure bias by choosing the less informative source in around 50% of all trials.

Together, these chapters demonstrate the impact that a biased selection of information can have on economic decisions. Even in simplified experiments, agents will avoid information to make self-serving decisions or seek out sources of information that support their preferences. The prevalence and persistence of the selective exposure bias in a context-free setting indicates that this bias may often lead to poorly-informed decisions outside of the laboratory. Information avoidance is likely to affect both everyday choices, like a pedestrian's option to give money to a destitute man or cross to the opposite side of the street, and corporate decisions, like a department

head's choice to quickly submit a large budget request or learn what other divisions need first. This research also helps to explain why altruism in economic experiments is often much larger than empirical findings would suggest, and it highlights the importance of considering the effects of imposed information in experimental studies.

Chapter 2

Patterns of Information Avoidance in Binary Choice Dictator Games

2.1 Introduction

There are good reasons we say “ignorance is bliss.” Our ability to avoid information can allow us to save money, evade blame, and engage in activities we might otherwise object to. Opportunities to avoid information arise on a daily basis, and we can use these opportunities to prevent ourselves from falling into difficult situations. We might let our phone keep ringing in case someone is calling to ask for donations. We might cross to the opposite side of the street so we do not have to pass by a homeless person. We might eat the last cookie in the tin without asking a friend if he wants it. These actions do not imply that we are selfish; we might very well donate money to charities, give to the homeless, and let our friend have the last cookie if we were directly asked to do so. But there is some part of us that would rather not do these things, and by avoiding direct choices, we manage to get around the guilt we would feel from flatly refusing to be altruistic.

These examples may seem insignificant, but this type of behavior can have major ramifications. CEOs embroiled in accounting scandals, like Bernie Ebbers of WorldCom, often claim to know nothing about the accounting practices of their own companies to try to avoid sentencing. It is hard to believe that the head of a major company would not have any knowledge about the firm’s key expenditures and profit

sources, but these CEOs might choose to remain in the dark so they can later claim ignorance. The U.S. military has even institutionalized information avoidance in the form of its “Don’t ask, don’t tell” policy towards homosexuals, which encourages members of the armed forces to avoid discussions about sexual orientation. Information avoidance can be a powerful force in justifying decisions.

Many economic experiments indicate that people have much greater tendencies to give to others than rational choice theory would predict. For example, in dictator games, in which one person keeps whatever amount of an endowment he does not give to another player, dictators give away about 20 to 30 percent of the total sum of money. Such behavior has been explained with models ranging from one in which altruists get a “warm glow” from giving (Andreoni, 1989) to ones that specify utility functions that incorporate inequality aversion (e.g., Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). However, the games typically played in the laboratory force subjects to make a stark choice. Outside of the lab, people are able to avoid situations that might involve difficult decisions, such as whether to give money to another person.

Although there is large body of psychological literature on how people justify decisions they are uncomfortable with, there are few models of justification in economics. It is surprising that economists have largely neglected the topic of justification of decisions, since this phenomenon can bridge the gap between predictions of income-maximization and empirical evidence of equitable behavior. People may feel the need to act fair when faced with a direct choice about how much to give someone else, but they will maximize their own income when they are able to justify their decision.

This chapter focuses on information avoidance as one form of justification. I study how the decision to avoid information can lead to a greater number of self-serving actions in a binary choice dictator game. In particular, I focus on how beliefs about the likelihood of helping or hurting someone else affect decisions to avoid information. I conduct an experiment in which a dictator knows her own payoffs but does not know whether they are aligned with the recipient’s payoffs. The dictator is told the probability that payoffs are aligned, and she can choose to see the recipient’s payoffs

at no cost. She might avoid learning the other's payoffs since without knowing for sure that she is hurting the other person, she could choose the action that gives her more money. Ex ante, it is difficult to predict whether a high probability of hurting someone by choosing the income-maximizing option would make a dictator more likely or less likely to avoid information.

For example, if someone thinks he might have HIV, his decision to get tested or not likely depends on his beliefs about the probability of testing positive, but it is unclear how these beliefs affect his choice. He might be more inclined to get tested if he thinks the probability of testing positive is slim, because there is a good chance he will hear the result he wants and he can then get rid of a nagging feeling of fear or guilt when sleeping with other people. On the other hand, he might be more likely to get tested if he thinks there is a high probability that he is HIV positive because he will be more afraid of hurting both himself and others if he does not realize he has the disease.

Ironically, because it lacks the subtleties of normal life, an experiment is the best way to study information avoidance. It would be difficult to study this phenomenon outside of a laboratory setting. We cannot simply ask people if they avoid information to justify making selfish decisions. Even if they were able to answer honestly, it would be difficult to assess the patterns of when they avoid information. Likewise, observing someone's behavior in an uncontrolled setting would not offer much insight, since we could not know whether the person was deliberately avoiding information in a certain situation or simply failed to notice it. The additional control of an experiment can allow us to know when someone is avoiding information that could be easily obtained, and what circumstances cause her to avoid the information. Since the decisions in a laboratory experiment are much more clear-cut than in everyday situations, experimental results can act as a lower bound for the level of information avoidance that we might expect to occur outside of the lab. In their normal interactions, people can select from any number of stories to explain why they did not seek out certain pieces of information or why they were justified in making a self-serving choice; information avoidance is bound to be more prevalent than in the laboratory. However, knowing

why people avoid information in an experiment can offer valuable insight about the patterns of information avoidance in everyday life.

This chapter is a first pass at gaining a better understanding of when people avoid information, in order to guide a theoretical model of information avoidance. There are a few models that show promise for explaining information avoidance, but they provide contradictory predictions. I will assess predictions based on those models to determine which would be the most useful focus of future research. But first, I review the economics and a small portion of the psychology literature on various forms of justification in Section 2.2. Section 2.3 describes the experiment in detail, and Section 2.4 discusses predictions for behavior in the experiment based on economic models. The results of the experiment are discussed in Section 2.5, and Section 2.6 concludes.

2.2 Studies of Rationalization

The problem of personal justification, and the self-deception it entails, has been studied for millennia. Yet the economics literature has only recently begun to explore the desire to justify choices, as well as situations that allow for such justification. This section will review theoretical and empirical research on various forms of justification, and will lead up to the work most closely related to this chapter.

Two studies suggest that information avoidance may stem from a desire to minimize regret. Gilovich and Medvec (1995) find that in the short term, people experience stronger regret for actions taken than for those not taken. People might avoid information because it is less likely to cause regret than learning the information and being forced to act on it. Larrick and Boles (1995) run experiments in which subjects either know they can avoid feedback about the best alternative after reaching an agreement in a negotiation or will receive the feedback whether or not an agreement is reached. They find that subjects who can avoid learning about the alternative are more likely to reach an agreement. Without receiving feedback, the subjects can justify their decision by assuming that the alternative would have resulted in a poorer outcome.

Some theoretical models revolve around the idea that people will alter their actions and beliefs to justify their decisions. Akerlof and Dickens (1982) model workers in dangerous occupations who may not buy beneficial safety equipment because they have chosen to believe that their work is safe. Eyster (2002) posits a model in which agents minimize regret by taking actions that justify past decisions. These agents engage in self-deception, behaving as if they truly prefer the action that makes the past choice optimal, while in reality, they might have taken a different action were it not for the past decision. Yariv (2005) models agents whose utility functions depend on the payoffs associated with their actions and on the consistency of the beliefs that lead to those actions. In this setup, an agent might prefer not to receive an informative signal if it will reduce her confidence in her beliefs.

Models like those of Yariv and Eyster are based on theories that have been well studied in social psychology. One of the most influential is the theory of cognitive dissonance, identified by Festinger in 1957. According to this theory, when people experience a conflict between an action they have taken and their underlying attitude, they will change their attitudes to relieve the tension. For example, in Festinger's original experiment, subjects who have taken part in a boring task are paid to tell other subjects that the task was interesting, and they later rate the task. Subjects who were paid only \$1 to lie about the task rate it as more interesting than the ones who were paid \$20. Festinger's interpretation of the result was that subjects who were only paid one dollar were faced with conflicting (dissonant) thoughts: the task was boring, but they were willing to say it was interesting for a mere dollar. To make these thoughts more consistent, they alter their beliefs about the task itself and judge it to be somewhat interesting. The subjects who were paid \$20, on the other hand, can easily explain the action of lying since they received adequate compensation for it.

Bem (1972) interprets the evidence from this experiment differently. In his self-perception theory, people judge their own emotions and attitudes by evaluating their actions, almost as an outsider would. In the above experiment, a person who "sees" herself telling another person a task was interesting after getting paid only one dollar

would conclude that she must have enjoyed the task. In Bem's theory, a person's actions affect how she views herself, so a need to justify questionable actions may easily arise. Murnighan, Oesch, and Pillutla (2001) find support for Bem's theory in a series of dictator games. Some dictators act to avoid construing themselves as greedy by offering an amount greater than zero but less than half of the pie in the standard game, but choose less fair options when their choices are limited to an equal allocation or one that gives them the bulk of the pie. Bénabou and Tirole (2006) build a model of charitable behavior on the idea that people not only want to signal to others that they are kind, fair people, but they also want to view themselves in a positive light and will take actions to do so.

Our ability to justify decisions may be so well ingrained that it happens instantaneously. Cognitive dissonance and self-perception theory are generally explained as if people think about their past actions and revise their attitudes or beliefs based on those actions. However, a test of attitude change among amnesic and control subjects shows that it does not rely on memory. Lieberman, Ochsner, Gilbert, and Schacter (2001) show that rankings for a series of paintings change after subjects are asked to choose which of two pairs of the paintings they prefer. Control subjects and those with anterograde amnesia switch their rankings in favor of the paintings they had chosen. Since the amnesic subjects could not remember what paintings they had preferred, their attitude change must have occurred immediately upon choosing one pair of paintings over the other. Moreover, the control participants who better remember which paintings they stated a preference for are less likely to change their rankings, indicating that memory could actually disrupt attitude change. These subjects seem to sense that they are more inclined to rank the chosen paintings higher than before, and they try to compensate for their attitude change.

Although there are several theoretical and empirical papers about the justification of decisions in the economics literature, and many more in psychology, most of these papers focus on choices that affect only a person's own welfare. Yet often the choices we feel the need to justify are the ones that were made in our own self-interest, but that potentially hurt other people. Recent empirical work has focused on this use of

justification.

In a laboratory experiment, Cain, Loewenstein, and Moore (2005) find that disclosing conflicts of interest between advisors and the “clients” they give information to causes advisors to bias their advice more than when the conflict is not publicly revealed. When conflicts of interest are revealed, the advisors might feel free to lie, believing that the clients should know to discount what they say. Ambiguity about an outcome for a recipient can also cause less honest advice, as well as less equitable allocations in a dictator game, while fostering perceptions that one’s actions are fair and justifiable (Schweitzer and Hsee, 2002; Haisley and Weber, 2005).

Dana, Cain, and Dawes (2006) and Lazear, Malmendier, and Weber (2006) allow subjects to opt out of a dictator game and receive a sum of money around what they could have kept in the game, while would-be recipients are given nothing and do not know the game could have been played. A substantial proportion of subjects choose to opt out of the game even when this choice is costly, since it allows them to avoid facing an expectant recipient.

Dana, Weber, and Kuang (forthcoming; all future references will be to this paper) similarly examine and find evidence for multiple forms of rationalization. In each of their treatments, subjects have to pick one of two choices in a dictator game that would allocate payoffs to themselves and either one or two other people, and they can always ensure the most equitable allocation. In one of the treatments, two players act as dictators, while a third person receives an amount based on the choice both players make. Although either dictator could make a choice that would ensure each person received the same amount, he could also place the responsibility for choosing the equal allocation on the other dictator. In another treatment, players are told that the computer would cut off the time dictators have to make a choice, picking either the self-serving or equal allocation with equal probability, so a dictator could choose a self-serving allocation without the other player knowing that they actively picked it. A third, information avoidance treatment will be described below. All three treatments give a dictator room to justify choosing the option that maximizes his own payoff, and subjects exploit this “moral wiggle room.” Dictators are less

likely to choose the equitable option in these treatments than in a baseline treatment that only offers a clear choice between an equal allocation and the more selfish one.

The current chapter is most closely related to the “hidden information” treatment, in which two states are equally likely and dictators can reveal the true state by clicking a button. Payoffs for a dictator are the same in both states: \$6 if he chooses “A” and \$5 if he chooses “B.” In one state, payoffs are aligned: the other person would get \$5 if the dictator chooses A and \$1 if he chooses B, so both would benefit from a choice of A. In the other state, payoffs are not aligned: the dictator’s choice of A would give the other person \$1, while B would give \$5. In Dana, Weber, and Kuang’s baseline treatment, subjects have to make a decision in the game with non-aligned payoffs. Three-quarters of them choose the equal allocation (5,5 instead of 6,1). However, in the hidden information treatment, many subjects choose not to reveal the true state. Eighty-six percent of those subjects choose A, giving themselves the higher payoff. This leads to a total of only 38% of all dictators choosing the equal allocation in the non-aligned state.

Dana, Weber, and Kuang offer interesting evidence of information avoidance and other forms of rationalization in the laboratory, but their tests do not allow for any explanation of why the phenomenon occurred. Their results could be explained by a wide variety of models. They use a between-subjects approach, which means that the subjects in the baseline treatment are not the same ones who take part in the hidden information treatment. The baseline treatment might happen to include more people with preferences for equality than the hidden information treatment, especially since the sample sizes were fairly small, so we cannot tell if the option to avoid information is the sole cause of the disparity between the two treatments.

More importantly, a great deal of information is lost through the between-subjects approach. We do not know if the subjects choosing not to reveal would have picked the more equitable choice if they had been placed in the baseline treatment, or if they were people who would pick A no matter what and be indifferent about the choice to reveal. Based on the data, we can guess that some of the subjects who did not reveal the state of the world to themselves would have picked B in the direct baseline choice,

but we cannot know that this is true. We also cannot know how large a proportion of those who avoid information do so because they would feel compelled to be equitable in the state with non-aligned preferences.

To gain more information about the causes of information avoidance, I conduct an experiment in which each subject makes choices both in the equivalent of Dana, Weber, and Kuang’s baseline case and in several games with hidden information. Through this within-subjects approach, we can better understand why people avoid information. I also vary the probability that a subject will see a game with aligned payoffs to gain more insight as to when people avoid information.

2.3 Experimental Procedures

Subjects in this experiment played 20 binary choice dictator games, similar to those of Dana, Weber, and Kuang (henceforth, DWK). Each subject played the role of dictator in every round, choosing allocations that would determine how much the subject and a recipient would be paid. The dictators simultaneously served as the recipients for other subjects. Subjects were paid according to their earnings as a dictator and the amount they received as a recipient, at a rate of 15 francs per dollar.¹ Payments were made in such a way that individuals’ decisions and earnings were anonymous to both the experimenter and other subjects. (See the appendix for full instructions.)

Each subject saw four different sets of payoffs, but the basic structure of the games did not change. Each game was presented in the same format on subjects’ computer screens. Figure 2.1 shows the layout for one round of the experiment. In the instructions to subjects, dictators were referred to as “Player X,” while recipients were called “Player Y.” In every game, a subject would choose either A or B. If she chose A, she would receive her highest possible amount of experimental currency (francs), and if she chose B, she would receive a lower amount. She could see her potential payoffs, but she would not necessarily know the payoff of the recipient.

¹Placing subjects in the role of both dictator and receiver may invoke the Golden Rule or cause subjects to view the experiment as a cooperative game, which would slightly bias the results towards more equitable choices, but this effect is expected to be small.




Figure 2.1: Screenshot from the experiment

Subjects were instructed that the game they were playing would be either “Game 1” or “Game 2.” In game 1, the recipient would receive his highest payoff if the dictator chose A, and a smaller amount if the dictator chose B. In game 2, the recipient would get the opposite payoffs—the lower amount if the dictator chose A, and the higher amount if the dictator chose B. Throughout this chapter, game 1 will be referred to as the “aligned game,” since the dictator’s payoffs are aligned with those of the recipient in this game (A is better for both). Game 2 will be referred to as the “nonaligned game.” For every game, subjects were told the probability (expressed as a percentage) that a given game would have aligned payoffs. A subject could know with certainty which game she was playing by clicking a button labeled “Reveal Game.” Subjects were not charged for clicking this button.

Subjects saw each payoff type for five consecutive rounds. Each of these five rounds had a different probability that the aligned game would occur. The probabilities were 0, 0.2, 0.5, 0.8, and 1, and the order in which they were presented was determined randomly for each subject. In each game, a subject could click on the “Reveal Game” button if she wanted to, and she would have to choose either “A” or “B.” When a new round was displayed, the software would not allow subjects to make any selection

until 30 seconds had elapsed, so a subject could not save time by ignoring the payoff possibilities. Subjects did not receive any feedback after making a decision, and they were informed before the experiment that they would only find out their total payoffs at the end of the session.

	Aligned Game	Nonaligned Game
Payoff Set 1	A 6, 5	A 6, 1
	B 5, 1	B 5, 5
Payoff Set 2	A 6, 5	A 6, -2
	B 5, -2	B 5, 5
Payoff Set 3	A 6, 7	A 6, 3
	B 5, 3	B 5, 7
Payoff Set 4	A 7, 5	A 7, 1
	B 5, 1	B 5, 5

Figure 2.2: Payoff sets (payoff to dictator, payoff to recipient)

The probabilities 0 and 1 were chosen to see whether players would pick A or B in direct choices—ones in which revealing would not matter because the game was known. When $p = 0$, a subject was playing the nonaligned game, and when $p = 1$, a subject was playing the aligned game. The probability .5 was selected both to replicate DWK, and to learn which subjects would reveal at this ambiguous probability. The other probabilities, .2 and .8, were ones at which subjects would be likely to decide that the game was fairly certain to have aligned payoffs or not. The four different payoff structures were chosen to potentially elicit different responses in subjects' direct choices and decisions to reveal which game they were playing. Figure 2.2 shows each type of payoff. Payoff Set 1 allows a comparison to DWK.

When the probability of the aligned game was 0, subjects were essentially playing DWK's baseline game, and when the probability equaled .5, the game was equivalent to their hidden information treatment.² We may consider this payoff structure to be a baseline for comparison with the other payoff sets, which each differ only slightly from this one.

The other sets of payoffs were designed to manipulate the number of subjects who make the more equitable choice in the nonaligned game and examine the ramifications for the decision to reveal. Payoff Set 2 in Figure 2.2 was chosen to make subjects more likely to pick B (the equitable choice) in the nonaligned game. This might also increase the number of subjects choosing to reveal since they could potentially cause more harm to the recipient by blindly choosing A than they would in the first payoff structure. On the other hand, since more subjects were likely to pick B in the unpleasant direct choice, more might want to avoid knowing they were in the state with nonaligned payoffs. In Payoff Sets 3 and 4, fewer subjects were expected to pick B in the nonaligned game. Set 3 has the same difference in total welfare as the baseline set, but if a subject chooses B in the nonaligned game, the other person gets more than she does. Players who lose utility when others earn more than them would be less likely to choose B in this set. In Payoff Set 4, a subject has to give up a bit more than in the baseline to make the equitable decision, so we again might expect more choices of A in the nonaligned game for this set. In one session of the experiment, subjects saw the payoffs sets in the order 1, 2, 3, 4, while in the other two sessions, the payoffs were ordered 1, 3, 2, 4. Changing the order had no effect on results.

2.4 Potential Models of Information Avoidance

Although no model directly predicts what might happen in the experiment explained above, we can examine the predictions implied by several models to later assess which

²Aside from the differences in giving the choices from both treatments to the same subjects, this experiment would not exactly replicate DWK since payoffs here were in terms of experimental currency as opposed to dollars and these choices were made alongside other decisions.

one best fits the results. Several social choice models could be applied to the decisions made in the direct games, but that is not the primary focus here. A good model of information avoidance will predict when and why people will avoid revealing information in the present experiment, and should be applicable to other situations as well. One of the models below (Rabin, 1995) specifically discusses what type of agent seeks full information. For the other models, we must infer who would choose to reveal, based on the general model description. We should not expect any of these models to fare perfectly, since they were not designed to explain this particular experiment and some amount of guessing is necessary to create predictions for the experiment. However, analyzing which models work and which ones fail can help guide future theory and tests of information avoidance.

Inequity Aversion

Since it is well regarded as a useful model of social preferences, the Fehr-Schmidt (1999) model of inequity aversion might predict when people choose to avoid information. The standard Fehr-Schmidt model includes parameters for envy and guilt in a linear utility function, such that an individual's utility is given by his payoff, with terms for envy and guilt subtracted off if payoffs are not equal among players. In a two-person game, the utility of player i , receiving a payoff of x is:

$$U_i(x) = x_i - \alpha_i \max\{x_j - x_i, 0\} - \beta_i \max\{x_i - x_j, 0\},$$

where α is the parameter that measures envy and β is the parameter for guilt. If someone receives a lower payoff than the person she is paired with, her utility decreases because she is envious of the other person. But if she receives more than the other person, she will feel guilty, which also causes utility to drop.

To understand its predictions, we can apply the Fehr-Schmidt model to Payoff Set 1 in Figure 2.2. In the nonaligned choice of 6,1 (choice A) versus 5,5 (choice B), a dictator would choose B if $U(5) = 5 > 6 - 5\beta = U(6)$, or if $\beta > \frac{1}{5}$.³

³Fehr and Schmidt suggest that for standard dictator games, the utility function should be concave in the amount of advantageous inequality (the β term), since linear utility only predicts

The Fehr-Schmidt utility function implies that dictators will always pick A in a direct choice for the aligned game (choosing 6,5 instead of 5,1). Let us assume for now that a dictator will pick B in a direct choice for the nonaligned game (choosing 5,5 instead of 6,1) but will pick A if not revealing the payoffs. We may assume that dictators act to maximize expected utility. If a dictator chooses to reveal the payoffs of the other person, her expected utility is given by:

$$\text{EU}(\text{Reveal}) = p(6 - \beta) + (1 - p)5,$$

where p is the probability that the aligned game is the true game. That is, her expected utility is the probability of seeing the aligned game times the utility she gets from her choice of A in this game plus the probability of seeing the nonaligned game times the utility she gets from her choice of B in this game. If the dictator chooses not to reveal the payoffs of the other person, her expected utility is:

$$\text{EU}(\text{Not Reveal}) = p(6 - \beta) + (1 - p)(6 - 5\beta).$$

Here, she picks A and expects to either feel somewhat guilty or much more guilty about the choice. Such a dictator will choose to reveal if:

$$\begin{aligned} \text{EU}(\text{Reveal}) &> \text{EU}(\text{Not Reveal}) \\ p(6 - \beta) + (1 - p)5 &> p(6 - \beta) + (1 - p)(6 - 5\beta) \\ 5 &> 6 - 5\beta \\ \beta &> \frac{1}{5}. \end{aligned}$$

Note that this holds by the assumption that the dictator picks B in the nonaligned game. If we had assumed that the dictator picked B instead of A when not revealing payoffs, the expected utility of revealing would again be higher than that of not revealing.

 offers of 0 or half the pie. The linear utility function used in this example might not be appropriate, but it is easy to understand and seems to make reasonable predictions for a binary game. Also, as will be discussed, any specification could be refuted by the same results.

revealing. In this setup, anyone who picks B in the direct choice of the nonaligned game should always choose to reveal payoffs regardless of the probability that the game is aligned. Anyone who picks A in the direct choice of the nonaligned game (and when not revealing) will always be indifferent between revealing or not, since the expected utility of revealing is the same as that of not revealing.

As can be seen from the Fehr-Schmidt example, if agents take expectations over the guilt they might feel when they do not reveal payoffs, their choice of whether or not to reveal does not depend on the probability of seeing the aligned game. This result does not arise from the specification of utility functions. Any model that uses expected utility to predict the choice to reveal and assumes that the utility of a certain payoff does not depend on the decision to reveal will imply that the probability does not matter. Expected utility implies that certain agents will always choose to reveal. Under reasonable assumptions, if people who pick B in a direct choice of the nonaligned game do not reveal at some point, then the independence axiom of expected utility theory will be violated. If this occurs, expected utility cannot solely explain information avoidance. Consider Figure 2.3, which depicts the game tree of decisions in one round of the experiment, using Payoff Set 1.⁴ Each payoff allocation is labeled, and the same letter is used to label equal payoffs. We will consider each letter to represent the degenerate lottery that gives the corresponding payoffs with probability 1.

According to the independence axiom, $X \succ Y$ if and only if $pZ + (1 - p)X \succ pZ + (1 - p)Y$, for all $p \in (0, 1)$ and X, Y, Z in the set of simple lotteries. Let us assume agents prefer J to K when they are playing the aligned game. This is a reasonable assumption, since both the dictator and the other player receive higher payoffs in J than in K. Now assume an agent prefers M to L in the direct choice of the nonaligned game (when $p = 0$). By the independence axiom, $M \succ L$ implies $pJ + (1 - p)M \succ pJ + (1 - p)L$. Note that the left-hand side is the lottery faced by someone who reveals the game and picks A in the aligned game and B in the

⁴For a more general discussion of the inconsistencies between expected utility and information aversion, see Wakker (1988) or Safra and Sulganik (1995).




Figure 2.3: Game tree of decisions in Payoff Set 1

nonaligned game. The right-hand side is the lottery for someone who does not reveal and picks A. If an agent who picks A in the aligned game and B in the direct choice of the nonaligned game does not reveal for some value of p , the independence axiom is violated.

Moral Rules and Moral Preferences

Since models that rely on expected utility analysis might easily fail to explain behavior in the experiments, we should consider the predictions of alternative models. Rabin (1995) proposes a unique model in which some people have moral preferences and directly incorporate moral concerns into their utility functions, while others act as if they use a “moral rule.” Those with the moral rule will engage in a potentially harmful activity as long as the probability of causing social harm is less than a certain cutoff point. In this model, agents can either engage in a desirable activity or not, and the activity either causes social harm or not. An agent receives enjoyment $v \in (0, 1)$ if she engages in an activity. She believes the activity causes social harm with probability q . An agent with moral preferences maximizes $U(v, q)$, which equals $v - q$ if she engages in the activity and 0 if she does not. This type of agent will

therefore engage in an activity if the probability of causing harm is less than her value for the activity.

For an agent with a moral rule, on the other hand, $U(v, q)$ equals $v - g(q)$ if she engages in the activity and 0 if she does not. This agent has some $y > 0$ such that $g(q) = 1$ if $q > y$ and $g(q) = 0$ if $q \leq y$. This ensures that this type of agent will not engage in the activity if she believes the probability that it causes social harm is higher than her cutoff, y . However, if q is less than y , the agent will engage in the activity without losing any utility.

Rabin assumes that the cutoff $y = v$, so neither type of agent will engage in an activity if she deems it to cause social harm. That means that for given beliefs, the agents will behave in the same way. However, beliefs are manipulable in the model. An agent first chooses a signal in the form of a probability distribution, f , which is constrained to have an expected value of the agent's initial belief. Then the chosen distribution generates a belief, and the agent chooses an optimal action based on the realized belief. During this belief manipulation, agents with moral preferences take into account the harm they will cause by choosing a signal, while agents with a moral rule do not. Expected utilities are:

$$\begin{aligned} U(f) &= x(f) \cdot v - z(f) - c(f) \text{ for an agent with moral preferences, and} \\ U(f) &= x(f) \cdot v - c(f) \text{ for an agent with a moral rule,} \end{aligned}$$

where $x(f)$ is the probability of engaging in an activity generated by the signal f , $z(f)$ is the probability of causing social harm generated by f , and $c(f)$ is the cost of generating f . This cost is not necessarily monetary; it could be the cognitive cost of changing beliefs.

When $c(f) = 0$, agents with moral preferences will obtain full information and will not cause any social harm. Agents with a moral rule will engage in the activity without seeking any informative signals if their initial beliefs about the probability of causing social harm are below their cutoff point. If their initial beliefs are above their cutoff point, they will seek the signal that maximizes the chance that they will

be able to engage in the activity.

In the current experiment, subjects can receive full information by revealing the game, so those with moral preferences will always reveal. Subjects who use a moral rule will pick A without revealing if their belief that A will cause social harm is less than their cutoff point. If the initial probability of social harm is above such a subject's cutoff point, he might try to manipulate his beliefs to make the probability lower. He might successfully draw from his manipulated belief distribution and be able to choose A without revealing. If this does not work or is too cognitively costly, he will be forced to reveal the game.⁵ Since the probability of causing social harm by picking A (which provides the higher payoff for the dictator) is $1-p$, subjects with a moral rule are more likely to reveal when p is lower. For example, someone with a cutoff point of .6 will pick A without revealing when the probability of playing the aligned game (the game in which A does not cause social harm) is greater than .6, so she would not reveal at $p = .8$. But if $p = .2$ or $p = .5$, the probability of social harm is too high, so the subject would reveal the game in hopes of seeing aligned payoffs.

Note that these calculations would not apply to someone who picks A when she is definitely faced with the nonaligned game ($p = 0$). The utility function specified for agents with moral preferences above assumed that agents place a weight of 1 on causing social harm ($U(v, q) = v - 1q$ if agents engage in the activity), so they will not engage in an activity if the probability of social harm is greater than their value. However, agents need not place so much weight on concerns for social harm. We can assume that agents who pick A when faced with the nonaligned game place very little weight on social harm, so they essentially behave as standard expected utility maximizers, always choosing the option that gives them the highest payoffs. We can expect such subjects to always pick A and be indifferent about revealing or not.

The moral preferences / moral rule model may work well to explain information avoidance, but self-perception theory (Bem, 1972) also appears to be well equipped

⁵Revealing the game is monetarily costless, and should also be cognitively costless since it clarifies the choice of A or B for the player. Belief manipulation, on the other hand, should be cognitively costly, since players are clearly told the probability that the game is aligned. Debriefing statements revealed that subjects believed the given probabilities.

for this task. When avoiding information enables someone to take an action she would not take with full knowledge, she is acting as if she did not know she could easily gain the information. In a sense, she is viewing herself as an outsider might, seeing only the action she took and not the internal decision to avoid learning about the consequences of the action.

Self-signaling

Based on Bem's self-perception theory, Bodner and Prelec (2003) create a self-signaling model, in which agents have preferences over not only outcomes, but also their types, or momentary dispositions, which are unknown. An agent's utility for choosing an outcome from a set consists of two parts, an outcome utility and a diagnostic utility. The outcome utility is simply the value someone gets from the outcome, given her true type. Diagnostic utility is the person's expectation of her utility over types, given her interpretation of the choice of one outcome over other outcomes in the set. Essentially, after making a choice, an agent revises her self-image based on that action and derives utility from this new expectation of her type. For example, someone who wishes to view herself as a patron of the arts might donate to a local museum and gain utility both from the donation itself and from being able to view herself as a supporter of fine art. As Bodner and Prelec discuss, when agents interpret their actions, they might ignore the underlying motives, or they might discount their self-signal, taking into account the fact that they chose an action to maximize both outcome and diagnostic utility.

The self-signaling model does not offer clear predictions for behavior in the current experiment, but we can use the concepts from this model and Bem's self-perception theory (1972) to make a conjecture about what might happen. If viewing herself as an outsider, an agent who does not reveal the true game may interpret her choice as one between two payoffs for herself. Faced with a choice of getting either 6 units or 5, the agent would naturally pick 6, and little, if any, diagnostic utility could be gained from this choice. If this person were to reveal payoffs, however, she could judge herself as being helpful or harmful to another person based on her decision. Let us assume

that subjects want to view themselves as helpful to others, and they will gain positive diagnostic utility when they make the choice that is better for a recipient. They will lose utility when making a choice that is worse for the other person. They will not gain any diagnostic utility when the other person's payoffs are unknown.⁶

To understand what would happen under these assumptions, again consider Payoff Set 1, in which the payoffs for a choice of A are 6,5 in the aligned game and 6,1 in the nonaligned game, and payoffs for a choice of B are 5,1 or 5,5. Let $u(x)$ represent outcome utility, where x is a person's own payoff, and let $v(x \mid \text{game})$ represent diagnostic utility. Diagnostic utility depends on the given game since an agent judges a choice based on the other available choices. Assume agents will choose A if they do not reveal the game (which gives utility $u(6)$) or if the revealed game has aligned payoffs (which gives utility $u(6) + v(6 \mid \text{aligned game})$).

First, let us consider an agent who chooses B when the nonaligned game is revealed, giving her $u(5) + v(5 \mid \text{nonaligned game})$. If $u(5) + v(5 \mid \text{nonaligned game}) > u(6)$, this person will always reveal since the utility she gets from revealing either game is greater than the utility from her choice when she does not reveal. If this person chooses not to reveal at some point, we know that $u(6) > u(5) + v(5 \mid \text{nonaligned game})$. When this is the case, the agent will reveal when the expected utility of doing so is greater than $u(6)$, the utility from not revealing. That is, she will reveal if:

$$p[u(6) + v(6 \mid \text{aligned game})] + (1 - p)[u(5) + v(5 \mid \text{nonaligned game})] > u(6)$$

$$p > \frac{u(6) - [u(5) + v(5 \mid \text{nonaligned game})]}{u(6) + v(6 \mid \text{aligned game}) - [u(5) + v(5 \mid \text{nonaligned game})]}.$$

The larger p is, the more likely this agent will reveal.

Now consider an agent who chooses A when the nonaligned game is revealed. The agent gets some negative diagnostic utility from this choice, which will be denoted as $-v(6 \mid \text{nonaligned game})$ for the sake of convenience. Since the agent will receive

⁶These assumptions are most closely aligned with Bodner and Prelec's (2003) description of face-value interpretations. Agents consider diagnostic utility when making a choice, but then interpret the choice as if it were not at all motivated by a desire for a positive self-signal.

positive diagnostic utility from choosing A in the aligned game, she may wish to reveal for some levels of p . She will reveal when:

$$p[u(6) + v(6 \mid \text{aligned game})] + (1 - p)[u(6) - v(6 \mid \text{nonaligned game})] > u(6)$$

$$p > \frac{v(6 \mid \text{nonaligned game})}{v(6 \mid \text{aligned game}) + v(6 \mid \text{nonaligned game})}.$$

Again, the agent is more likely to reveal as p increases.

This interpretation of the self-signaling model predicts that subjects will reveal more often at higher levels of p . However, if agents do not receive any extra utility from choosing 6,5 over 5,1 in the aligned game, then the predictions for when they should reveal would not depend on p . It is reasonable to expect that agents who consider their motives for picking A over B in the aligned game will not gain any diagnostic utility, since A gives them the higher payoff. In this alternative model based on self-perception, subjects who choose A when they see the nonaligned game would never reveal. Those who choose B would always reveal if $u(5) + v(5 \mid \text{nonaligned game}) > u(6)$, but never reveal if the inequality were reversed.

Table 2.1 provides a summary of each model's predictions for when subjects will reveal, based on their choices at $p = 0$, when they know the game has nonaligned payoffs. Many of the models share similar predictions. For example, if we only knew that all the subjects who picked A at $p = 0$ never revealed the game, we would not be able to rule out any of the models. We must focus on the differences between the models. In particular, it will be necessary to focus on how decisions to reveal depend on p . If subjects either never reveal or always reveal, then the Fehr-Schmidt model or the second model of self-perception are likely to best fit the data. If more subjects reveal when p is larger, then the first model of self-perception will make the best predictions, while if fewer reveal for larger values of p , Rabin's model would fit better.

Model	Pattern of Revealing Implied By:		Comments
	Choice of A in Nonaligned Game	Choice of B in Nonaligned Game	
Fehr-Schmidt inequity aversion	Indifference about revealing	Always reveal	Violation of independence axiom would rule out this model
Rabin's moral preferences / moral rules	Indifference about revealing	Reveal only if p is below cutoff point (May always reveal)	Cutoff point can vary; will look for trend of revealing at lower values of p
Self-signaling (A in aligned game gives diagnostic utility)	Reveal if p is above cutoff (May never reveal)	Reveal if p is above cutoff (May always reveal)	Cutoff point can vary; will look for trend of revealing at higher values of p
Self-signaling (A in aligned game only gives outcome utility)	Never reveal	Either always reveal or never reveal	Cannot know which of the people who picked B should reveal

Table 2.1: Predictions of revelation patterns, based on choice in nonaligned game, $p = 0$

2.5 Results

Armed with predictions from several models, we can now analyze the results of the experiment. This section will first summarize general patterns in the data, then compare these results to those of DWK. Decisions will then be analyzed at an individual level, and finally, we will determine which models make the best predictions.

Fifty-four subjects participated in this experiment. Three sessions were conducted, all of which contained at least 16 subjects. Subjects were undergraduate students at the California Institute of Technology who were recruited from a database of student volunteers.

Table 2.2 shows the proportion of subjects that chose to reveal and the proportion that picked A for each payoff set. A few trends emerge from a cursory glance at the table. One is that with only a few exceptions, the subjects who did not reveal the game picked A, as we would expect. Another pattern is that across payoff sets, fewer subjects revealed when the probability of the aligned game (p) was .8. As a result,

Payoff Set 1	Information Choice	Proportion Choosing A	
		Aligned Game	Nonaligned Game
Direct choice ($p = 1$ or $p = 0$)		52/54 (96.3%)	22/54 (40.7%)
$p = .2$	REVEAL 37/54 (68.5%)	5/6 (83.3%)	7/31 (22.6%)
	NOT REVEAL 17/54 (31.5%)	3/3 (100%)	13/14 (92.9%)
	Total:	8/9 (88.9%)	20/45 (44.4%)
$p = .5$	REVEAL 34/54 (63.0%)	18/18 (100%)	4/16 (25.0%)
	NOT REVEAL 20/54 (37.0%)	10/11 (90.9%)	8/9 (88.9%)
	Total:	28/29 (96.6%)	12/25 (48.0%)
$p = .8$	REVEAL 29/54 (53.7%)	23/23 (100%)	2/6 (33.3%)
	NOT REVEAL 25/54 (46.3%)	17/18 (94.4%)	7/7 (100%)
	Total:	40/41 (97.6%)	9/13 (69.2%)

Payoff Set 2	Information Choice	Proportion Choosing A	
		Aligned Game	Nonaligned Game
Direct choice ($p = 1$ or $p = 0$)		53/54 (98.1%)	20/54 (37.0%)
$p = .2$	REVEAL 38/54 (70.4%)	15/15 (100%)	1/23 (0.04%)
	NOT REVEAL 16/54 (29.6%)	3/4 (75.0%)	11/12 (91.7%)
	Total:	18/19 (94.7%)	12/35 (34.3%)
$p = .5$	REVEAL 38/54 (70.4%)	23/23 (100%)	4/15 (26.7%)
	NOT REVEAL 16/54 (29.6%)	10/10 (100%)	6/6 (100%)
	Total:	33/33 (100%)	10/21 (47.6%)
$p = .8$	REVEAL 36/54 (66.7%)	29/31 (93.5%)	1/5 (20.0%)
	NOT REVEAL 18/54 (33.3%)	12/13 (92.3%)	5/5 (100%)
	Total:	41/44 (93.2%)	6/10 (60.0%)

Table 2.2: Information decisions and dictator game choices

Payoff Set 3	Information Choice	Proportion Choosing A	
		Aligned Game	Nonaligned Game
Direct choice ($p = 1$ or $p = 0$)		51/54 (94.4%)	28/54 (51.9%)
$p = .2$	REVEAL 34/54 (63.0%)	8/9 (88.9%)	7/25 (28.0%)
	NOT REVEAL 20/54 (37.0%)	7/7 (100%)	11/13 (84.6%)
	Total:	15/16 (93.8%)	18/38 (47.4%)
$p = .5$	REVEAL 34/54 (63.0%)	18/20 (90.0%)	4/14 (28.6%)
	NOT REVEAL 20/54 (37.0%)	9/9 (100%)	11/11 (100%)
	Total:	27/29 (93.1%)	15/25 (60.0%)
$p = .8$	REVEAL 30/54 (55.6%)	22/25 (88.0%)	1/5 (20.0%)
	NOT REVEAL 24/54 (44.4%)	22/23 (95.7%)	1/1 (100%)
	Total:	44/48 (91.7%)	2/6 (33.3%)

Payoff Set 4	Information Choice	Proportion Choosing A	
		Aligned Game	Nonaligned Game
Direct choice ($p = 1$ or $p = 0$)		53/54 (98.1%)	24/54 (44.4%)
$p = .2$	REVEAL 37/54 (68.5%)	6/6 (100%)	7/31 (22.6%)
	NOT REVEAL 17/54 (31.5%)	2/2 (100%)	14/15 (93.3%)
	Total:	8/8 (100%)	21/46 (45.7%)
$p = .5$	REVEAL 37/54 (68.5%)	16/16 (100%)	6/21 (28.6%)
	NOT REVEAL 17/54 (31.5%)	8/8 (100%)	9/9 (100%)
	Total:	24/24 (100%)	15/30 (50.0%)
$p = .8$	REVEAL 30/54 (55.6%)	21/21 (100%)	5/9 (55.6%)
	NOT REVEAL 24/54 (44.4%)	18/18 (100%)	6/6 (100%)
	Total:	39/39 (100%)	11/15 (73.3%)

Table 2.2 (continued)

when p was larger, more subjects chose the less equitable choice in the nonaligned game.

Subjects' decisions to reveal were affected by two factors: the chance that the game would be aligned and their choice of A or B when the payoffs were definitely nonaligned.⁷ According to probit estimates, shown in Table 2.3, subjects were less likely to reveal when the probability of playing the aligned game was high. The probability that a subject would reveal decreased by .11 when the probability of the aligned game increased from .5 to .8. The probability of revealing decreased by a total of .13 when the probability of the aligned game increased from .2 to .8. There was no significant difference in the decision to reveal between probabilities of .2 and .5.

Variable	Coefficient	Change in Probability (at .5)
Probability of aligned game = .2	.060 (.121)	.021 (.043)
Probability of aligned game = .8	-.318 * (.115)	-.112 (.043)
Chose B at $p = 0$ (Equitable in nonaligned game)	1.68 * (.246)	.544 (.077)
Payoff Set 2	.210 (.145)	.074 (.046)
Payoff Set 3	.191 (.153)	.067 (.051)
Payoff Set 4	.192 (.163)	.067 (.055)
Order of Payoff Sets = 1, 3, 2, 4	-.157 (.291)	-.055 (.101)
Intercept	-.430 (.264)	
N	648	
Log pseudo-likelihood	-306.449	
Pseudo R^2	.277	

Note: Standard errors corrected for subject effects. All variables are dummy variables.

* indicates $p < .01$

Table 2.3: Probit estimates. Dependent variable is decision to reveal.

An even larger effect came from subjects' decisions in the direct choice for the

⁷The decision to reveal at $p = 0$ and at $p = 1$ is not considered in any of the results presented here, since for those values of p , a subject would know which game she was playing whether or not she revealed.

nonaligned game. Subjects who chose B, the more equitable decision, when the nonaligned game was known ($p = 0$) were much more likely to reveal than subjects who chose A. The predicted probability of revealing increased by .54 in the probit regression when subjects chose B in the direct choice for the nonaligned game.

2.5.1 Choices in the Nonaligned Game

The payoff sets that subjects saw and the order in which they saw them had no significant effect on subjects' decisions to reveal. This is not surprising since varying the payoff sets did not affect subjects' decisions in the direct choice for the nonaligned game as much as expected. The only significant differences in subjects' decisions to choose A in this game (when $p = 0$) were between Payoff Set 3 and Payoff Sets 1 and 2, and between Payoff Set 2 and Payoff Set 4. In Payoff Set 3, picking B would give the other person more money than the dictator would get, and these payoffs caused more people to pick A than in Payoff Sets 1 ($t_{53} = 1.95$, $p < .05$ in a one-sample test of differences) and 2 ($t_{53} = 2.41$, $p < .01$). In Payoff Set 4, subjects would have to give up 2 units to choose the equitable option, while in Payoff Set 2, picking A would cause the other person to lose money. More people chose A in Set 4 than in Set 2 ($t_{53} = 1.66$, $p = .05$).

Information avoidance led to fewer equitable choices in the nonaligned game. Focusing only on decisions in this game, we can examine the difference in each subject's choice at $p = 0$ (when the game was known to have nonaligned payoffs) to his or her choice at other levels of p . A stringent test of this difference is a one-sample t-test that compares the average difference in subjects' decisions at $p = 0$ and another probability to a mean of 0. This test revealed no significant differences between $p = 0$ and $p = .2$. Subjects were more likely to choose B at $p = 0$ than at $p = .5$ in Payoff Set 4 ($t_{29} = 1.95$, $p < .05$), and this difference was marginally significant in Payoff Set 1 ($t_{24} = 1.44$, $p = .08$). When $p = .8$, more subjects switched from picking B to picking A. There was a significant difference between $p = 0$ and $p = .8$ in Payoff Sets 1 ($t_{12} = 2.13$, $p < .05$) and 4 ($t_{14} = 2.65$, $p < .01$), and there was a marginally

significant difference in Payoff Set 2 ($t_9 = 1.50$, $p = .08$).⁸

2.5.2 Comparison to Dana, Weber, and Kuang

Since one of the payoff sets subjects saw (Set 1) was the same as in DWK, we can compare results across experiments. The results are not significantly different, but the effects of information avoidance appear stronger in DWK.⁹ DWK's baseline game is the equivalent of seeing $p = 0$ and knowing payoffs are not aligned in this experiment. In DWK's baseline, 26% of subjects (5/19) picked A, while 41% (22/54) picked A in this experiment. In DWK's hidden information treatment, the aligned and nonaligned games were equally likely, so that treatment corresponds to $p = .5$ in this experiment. For this level of p , a total of 48% (12/25) of subjects in this experiment chose A when the nonaligned game was the actual game, as opposed to 63% (10/16) in DWK. The difference lies in the number of subjects who did not reveal the game. When the game was nonaligned, half of DWK's subjects did not reveal, and all of them chose A, while 37% of the subjects in this experiment did not reveal, and 8 out of 9 chose A.¹⁰ Of the subjects who did reveal, the same proportion chose A across subject pools.

2.5.3 Patterns of Avoidance

Implicit in DWK's between-subject analysis is the idea that subjects in their baseline treatment would have made similar choices to the ones in the hidden information treatment if they had had the chance to avoid revealing payoffs. The within-subjects design of the experiment in this chapter allows us to directly compare the choices people made when they knew the game to the ones they made when they had the option to avoid knowing it. We can therefore determine who generally chooses to

⁸Only six subjects received the nonaligned game at $p = .8$ in Payoff Set 3, and all six made the same choice as at $p = 0$.

⁹Comparing proportions picking A in the nonaligned game at $p = 0$, $\chi^2 = 1.25$, $p = .26$. For those picking A in the nonaligned game at $p = .5$, $\chi^2 = .82$, $p = .36$. For those revealing when the game was nonaligned at $p = .5$, $\chi^2 = .86$, $p = .35$. Comparing proportions revealing at $p = .5$ regardless of the actual game, $\chi^2 = .38$, $p = .54$.

¹⁰It is odd that someone would not reveal, but pick B. Even if all 9 subjects had chosen A, though, the total percentage picking A would rise to only 52%, still lower than in DWK.

reveal—the subjects who make an equitable choice in the direct game with nonaligned payoffs, or the ones who choose the option that gives them the higher payoff but hurts the recipient.

Pattern	Payoff Set 1	Payoff Set 2	Payoff Set 3	Payoff Set 4
Never Reveal	11 (0)	12 (0)	17 (2)	14 (0)
Always Reveal	23 (19)	34 (28)	28 (20)	29 (26)
Reveal at $p = .2$ only	4 (1)	3 (3)	1 (0)	2 (1)
Reveal at $p = .5$ only	2 (1)	2 (2)	2 (0)	2 (1)
Reveal at $p = .8$ only	3 (2)	1 (0)	1 (0)	1 (0)
Reveal at $p = .2$ and $.5$ only	8 (7)	1 (1)	4 (3)	6 (2)
Reveal at $p = .5$ and $.8$ only	1 (1)	1 (0)	0 (0)	0 (0)
Reveal at $p = .2$ and $.8$ only	2 (1)	0 (0)	1 (1)	0 (0)

Table 2.4: Frequency of revelation patterns by game: Total and number (out of total) that had picked B at $p = 0$

As the probit regression indicated, subjects who picked the more equitable choice in the nonaligned game were the ones who were more likely to reveal the game at other levels of p . The subjects who picked A in the nonaligned game were likely to never reveal. Table 2.4 shows the total number of people following each pattern of revelation choices in each payoff set, as well as the number out of that total who had chosen B in the nonaligned game when $p = 0$. The subjects who did not choose to reveal at any level of p were nearly always ones who had chosen A in the direct choice for the nonaligned game. Conversely, the subjects who always revealed were primarily ones who picked B in the direct choice for the nonaligned game.¹¹

¹¹Using results from this experiment, we can infer whether there were differences between DWK's baseline and hidden information subjects. Twenty-eight out of the 32 subjects who picked B at $p = 0$ chose to reveal at $p = .5$ in Payoff Set 1. Applying this proportion to DWK's baseline, an estimated 12 subjects out of the 14 who chose B would reveal at $p = .5$, leading to at least 63% (12/19) of the subjects revealing at $p = .5$. If we also apply the proportion of subjects who picked A

Considering individual revelation patterns across payoff sets, most subjects made consistent decisions in the nonaligned game direct choices. Sixteen out of the 54 subjects (29.6%) chose A in every payoff set at $p = 0$. Of those, 6 never revealed the game, 5 only revealed for one particular payoff set and level of p , 2 only revealed in two instances, 1 always revealed, 1 revealed with one exception, and 2 people had no discernable pattern in their choices to reveal. The subjects who always picked B at $p = 0$ were less dispersed in their revelation patterns. Twenty-two subjects (40.7%) picked B across the four payoff sets. Of those, 12 always revealed, 7 only chose not to reveal in one instance (notably, at $p = .8$ for all 7), and 3 had no clear patterns in their decisions to reveal.

To specifically address the predictions of various models, we must look at particular patterns of behavior. First, we can examine whether the independence axiom of standard expected utility models was violated. As one might expect, this axiom was violated several times. Out of the instances in which subjects chose B at $p = 0$ and A at $p = 1$, the subjects chose not to reveal the game and picked A 36 times, thus violating the independence axiom. It is not the case that a few people violated the axiom multiple times. Sixteen subjects, or 30% of the subject pool, had at least one violation of the independence axiom.

Rabin's (1995) model of moral rules versus moral preferences predicts that some of the subjects who choose B at $p = 0$ will always reveal the game, while others will only reveal when the probability that the game is aligned is below some cutoff point, meaning we should see more choices to reveal for low values of p . As the probit estimates and t-tests indicate, subjects do reveal more at lower levels of p . Table 2.5 shows the total number of cases in which each revelation pattern was used, separated by the decision at $p = 0$. (A "case" is one payoff set seen by one player.) These totals could be found by adding the columns of Table 2.4, but they are presented in a different format here. R indicates revealing at a certain probability and N indicates at $p = 0$ and revealed at $p = .5$ to DWK's baseline, then around 71% would be expected to reveal in the hidden information treatment. Fifty-six percent of the subjects in DWK's hidden information treatment revealed the game, somewhat below what we would expect if they behaved the same way as the baseline subjects.

not revealing at a certain probability. For example, “NNR” means someone revealed at $p = .8$ but not at $p = .2$ or $p = .5$, while “RNR” means someone revealed at $p = .2$ and $p = .8$ but not at $p = .5$.

Choice at $p = 0$	Pattern of Revealing at $p = .2$, $p = .5$, and $p = .8$							
	NNN	NNR	NRR	NRN	RNR	RNN	RRN	RRR
A (Self-serving)	52	4	1	4	1	5	6	21
B (Equitable)	2	2	1	4	2	5	13	93

Table 2.5: Total number of cases by revelation pattern and choice at $p = 0$

We can find the number of cases in which Rabin’s model predicts the revelation pattern shown in the table. The moral rule model predicts that the subjects who pick B at $p = 0$ will always reveal or reveal only if p is below some cutoff. In 93 cases, subjects always reveal, while they avoid revealing at high levels of p in another 20 cases (including 2 for the pattern NNN). This model correctly predicts behavior for the subjects who pick B at $p = 0$ in 113 out of 122 cases (92.6%). Since we assumed that the subjects who pick A at $p = 0$ will always be indifferent about revealing, this expected utility component of the model easily predicts behavior in the 94 cases where subjects picked A at $p = 0$, but this does not tell us much. The model more specifically predicts that those who pick A at $p = 0$ will always pick A. This is true in 73 out of 94 cases (77.7%). Using this more specific prediction, the model correctly predicts a total of 186 cases (86.1%).

We can apply the same analysis to the predictions based on self-perception theories. One interpretation of self-signaling (Bodner and Prelec, 2003) and Bem’s self-perception theory (1972) predicts that subjects will be more likely to reveal at higher values of p . Eight subjects followed this pattern in one payoff set, but all eight followed different patterns in the other payoff sets. This part of the model’s prediction does not hold much support, but we can look at the total number of cases in which the overall prediction held. This model predicts that subjects who choose A at $p =$

0 will reveal if p is above a cutoff (5 cases) or never reveal (52 cases), and those who choose B at $p = 0$ will reveal if p is above a cutoff (5 cases, including 2 for NNN) or always reveal (93 cases). A total of 155 cases (71.8%) are correctly predicted by this model.

The alternate prediction that subjects who pick A in the nonaligned direct choice will never reveal and those who pick B in that choice will either always reveal or never reveal fails, since many subjects choose to reveal only at certain levels of p . Even though this model poorly describes decisions to reveal in the experiment, it correctly predicts behavior in 147 cases (68.1%). This highlights the need to focus on the differences between models.

Both the moral rule model and self-perception models allow for a high number of people never revealing or always revealing. If we exclude the cases in which subjects always or never revealed, the moral rule model makes a correct prediction 18 out of 27 times for subjects who chose B at $p = 0$, compared to only 3 cases in favor of the self-perception model. For subjects who chose A at $p = 0$, the self-perception model can serve to explain only an additional 5 cases. The moral rule model is better at predicting when people choose to reveal when they are affected by the probability of seeing the aligned game.

2.6 Conclusion

This chapter uses a simple experiment to explore when and why people would avoid information about how their choices would affect another person. The results indicate that while many agents who make equitable decisions in a direct choice will always seek information, some will avoid information as long as they know there is a good chance that they will not hurt another person by choosing the option that is best for themselves. In contrast, most of the subjects who make self-interested decisions that hurt another person in a direct choice will decide not to reveal information.

The findings in this experiment lend support to those of Dana, Weber, and Kuang, although those authors found somewhat stronger evidence of information avoidance

than I did. The experiment in this chapter expands upon DWK's hidden information design to provide greater insight on the nature of information avoidance. The within-subjects design allows us to see who chooses to avoid information. Varying the probability of getting a game with aligned payoffs versus one with nonaligned payoffs lets us examine when subjects are more likely to avoid information. This greater knowledge about patterns of information avoidance allows us to assess which models might be best adapted to explain the phenomenon.

Expected utility theory alone is unable to explain information avoidance in this experiment. Perhaps this is not a surprising conclusion, since the idea that less information might be preferred to more is certainly not a tenet of the theory. However, it seems reasonable to consider the possibility that people would take expectations over their potential utility when they choose to avoid information, as if they will find out later that they made a good choice or a bad one. This does not appear to be the way agents think about their choice when avoiding information, though. Nearly one-third of the subjects in this experiment violated the independence axiom that expected utility theory relies on. Any model that uses expected utility to predict the decision to avoid or seek information will not be appropriate. We must consider other models.

Models based on self-perception theory had a good chance of explaining decisions in the experiment, since the self-deception involved in avoiding information is akin to viewing oneself as an outsider would. However, these models also failed to fit the data. One of the models predicted that subjects would either never reveal or always reveal across values of p in a given payoff set, but many subjects based their decisions to reveal on the given probability. The other model predicted that subjects would reveal at higher values of p . This only occurred in eight cases, though.

Rabin's model of moral preferences and moral constraints (1995) best fits the data in this experiment. Several subjects behaved as if they used a moral rule, revealing the game only at low levels of p , when picking A was likely to hurt the other person. Others acted as if they had moral preferences, seeking full information by always revealing the true game. A third group maximized their own payoffs, choosing A

whether or not they knew it would hurt the other person. While some subjects followed other decision rules in certain payoff sets, Rabin's model explains the bulk of the data. This model is likely to be useful for other situations that involve moral choices and the opportunity to avoid information. Part of the reason this model explained behavior so well, though, is that it predicted a variety of outcomes without entirely specifying which subjects should behave in certain ways. A better test of the model might be to define subjects as having moral preferences or using a moral rule based on their behavior in one task, then see if they reveal or avoid information in another task according to their type. Although this model currently leaves some questions unanswered, it correctly predicted when people would choose to reveal information in this chapter. This model offers the most promise for future work on information avoidance.

In some ways, the evidence in this experiment can be reassuring to people seeking donations for charity or those who would like to believe self-deception is a rare occurrence. For the most part, the subjects who avoided information about another person's payoffs were the ones who would not choose equitably when faced with a direct choice. We might assume that information avoidance would not significantly decrease the total amount of donations in a charity drive, for example, since the people who choose not to pick up their phone or answer their doors are most likely the ones who would not make a donation even if they spoke to someone.

However, the results of this experiment indicate that people who do choose equitably in a direct choice tend to follow a pattern in their decisions to avoid information. Subjects were less likely to reveal information when the probability that they would be helping themselves while helping the other person was high. These subjects might have told themselves that they could just go ahead and pick A, since it probably would not hurt the recipient. This type of justification could prove dangerous. Even someone who would take all necessary measures to protect other people if he knew he was HIV positive might not bother to get tested if he thought the probability of testing positive was sufficiently low.

In the experiment in this chapter, there was little room for belief manipulation

since the probability of seeing a particular game was given. In the outside world, though, such exact probabilities are seldom known. In addition, the choices in this experiment were very stark. Anyone who carefully considered the “Reveal Game” button should have figured out that the only benefit to not clicking the button would be the ability to make a self-serving choice without knowing if you were hurting someone. Even though the information avoidance mechanism was obvious, subjects still systematically chose not to reveal information. We should expect a greater ability to choose beliefs and avoid information outside of the laboratory, where beliefs are more flexible and excuses more plentiful.

The experiment in this chapter was a first step toward characterizing patterns of information avoidance. Future experiments could further explain these patterns through slight variations on the current design. Experiments should also focus on distinguishing between subjects with moral constraints versus moral preferences, as mentioned above, to improve upon the Rabin (1995) model. Introducing a nominal cost for revealing might drive down the number of subjects who choose to reveal the game. This cost would also allow us to more definitively predict that the subjects who should have been indifferent about revealing in the current experiment would never reveal. In this experiment, beliefs were not easily manipulable, since a clear probability of playing one of two games was provided to subjects. If only a range of probabilities were provided instead, subjects would have more room to choose the beliefs that would most benefit them, and we could expect more self-serving choices with less revelation.

Although in many ways the results in this chapter could be seen as a lower bound on levels of information avoidance, we might try to lower the number of people who avoid information further by making the avoidance more salient. Subjects might be forced to click a button to either reveal or not reveal payoffs before making their decision in the dictator game. We could also introduce a treatment in which subjects would learn the true game with some probability after they make a choice in the dictator game, to see if the potential for feeling guilt or regret after avoiding information causes more people to reveal. Clearly, there are many experiments that would create

a better understanding of information avoidance. Any of these experiments would provide interesting results, but the most useful ones will be the experiments that can lead to a solid model of information avoidance.

2.7 Appendix: Experimental Instructions

This is an experiment about decision-making. You will be paid \$5 for participating in this experiment, plus an amount that depends on the decisions you and other people make. You will be paid in cash at the end of the experiment in the following way: After we've read these instructions, you will receive a player number. **DO NOT SHOW ANYONE ELSE THIS NUMBER.** When the experiment is done, the experimenter will put each person's earnings in an envelope marked with the appropriate ID number. These envelopes will be given to someone outside of the room and you will have to give this person your number to receive the envelope containing your money. Any information tying your name to the amount you made in the experiment will not be seen by the experimenter. The procedures will ensure that no one, including the experimenter, will know what decisions you made during the experiment.

In this experiment, you will play a series of games, each with another person in the room. You will always make decisions in the role of Player X. Before each game, you will be randomly matched with another person in the room who will be in the role of Player Y. Since you are randomly rematched every round, there is only a small chance that you will be matched with the same person in multiple rounds.

The games you will play will be like the one pictured below. As player X, you will choose either "A" or "B" by clicking on it. Player Y will not make any choice. Both players will receive payments based on the choice of Player X. The numbers in the table are the payments players receive, given in francs. At the end of the experiment, your francs will be converted to dollars, at a rate of ____ francs per dollar. The numbers below are just an example—you will see different amounts in the games you play.

EXAMPLE 1:

Player X's Choices	A	X: 1	Y: 2
	B	X: 3	Y: 4

In this example, if Player X chooses “B,” we look in the bottom square for the earnings. Here, Player X receives 3 francs and Player Y receives 4 francs. Player X’s payments will always be written to the left of Player Y’s payments.

To make sure everyone understands the instructions so far, please answer the following questions:

In this example, if Player X chooses “A” then:

Player X receives _____

Player Y receives _____

If Player X chooses “B” then:

Player X receives _____

Player Y receives _____

You will not necessarily know which game you are playing—it could be one of two games, called “Game 1” and “Game 2.” You will see a screen that looks like the following example.

EXAMPLE 2:

Player X's Choices	A	X: 4	Y: ?
	B	X: 3	Y: ?

Reveal Game

There is a **20%** chance that the game is GAME 1.

GAME 1			
A	X: 4	Y: 3	
B	X: 3	Y: 1	

GAME 2			
A	X: 4	Y: 1	
B	X: 3	Y: 3	

Notice that “GAME 1” and “GAME 2” are the same except that Player Y’s payments are flipped between the two. In both games, Player X gets his highest payment of 4 francs by choosing A. In the game on the left (Game 1), this gives Player Y her highest payment of 3 francs, but in the game on the right (Game 2), it

gives Player Y her lowest payment of 1 franc. If Player X chooses B, he gets a lower payment of 3 francs. In Game 1, this gives Player Y her lowest payment of 1 franc, but in Game 2, it gives Player Y her highest payment of 3 francs.

The payments that you see will be different than in this example, but the same rules will apply. Player X's payments will be the same in both games, and Player Y's will be flipped between the two games. Player X will always get his highest payment by choosing A. In Game 1, this will give Player Y her highest payment, but in Game 2, it will give Player Y her lowest payment. If Player X chooses B, he will get a lower payment. In Game 1, Player Y will get her lowest payment, and in Game 2, she will get her highest payment.

You will not know which of the games you will be playing, but you will see the chance that you will be playing Game 1. In the example above, there is a 20% chance that the game is Game 1. That means that the computer will randomly select a game with a 1 in 5 chance of picking Game 1. You can choose to find out which game you are actually playing, if you want to do so, by clicking on the "Reveal Game" button before you play. If you click this button, you will see Player Y's payoffs and will therefore know if you are playing Game 1 or Game 2.

Although you will always make decisions as Player X, you will also be acting as Player Y for another person. For example, let's say the person with ID number 5, who we'll call person 5, is making a decision as Player X, and he is paired with person 8, who will be in the role of Player Y. Person 5 will make choices that affect how much he and person 8 receive. At the same time, person 7 will act as Player X with person 5 in the role of Player Y. Person 7 will make choices that affect how much she and person 5 receive. Note that your decisions as Player X do not affect what you will get as Player Y. Your earnings will be a total of the amount you earn as Player X and the amount you get as Player Y, which is based on other people's decisions.

After you have made your decision in one round, you will automatically be taken to the next round. You will not see any information on what you earned in that round, and Player Y will not see what decisions you made in that round. You will find out your total payoffs at the end of the experiment.

To make sure everyone understands these instructions, please answer the following questions.

1. Which action gives Player X his or her highest payment? ____
2. In Example 2, if Player X chooses B, then Player Y receives:
 - a) 3 francs
 - b) 1 franc
 - c) either 3 francs or 1 franc
3. In Example 2, what is the chance that the game is:
Game 1? ____
Game 2? ____
4. In Example 2, if the probability of game 1 were 100% and Player X chose A, what would Player Y receive? ____
5. In one round, you make a choice that gives Player X a payment of 2 francs and Player Y a payment of 1 franc. At the same time, you are Player Y for another person who makes a choice that gives Player X a payment of 4 francs and Player Y a payment of 3 francs. What is your total payment for this round? ____

Chapter 3

The Effect of Selective Information Acquisition on Charitable Giving

3.1 Introduction

Research has shown that people will avoid information in order to justify self-serving economic decisions, though most of us know this to be true from personal experience. In past experimental studies (e.g., the previous chapter; Dana, Weber, and Kuang, forthcoming), subjects faced a simple binary choice of obtaining information or not. Decisions to gain information in everyday situations are seldom so clear-cut. This chapter considers whether the amount of information one chooses to learn about a charity affects the amount one donates. I conduct an experiment that allows subjects to obtain varying amounts of information about a non-profit organization before choosing how much money to keep or donate to the cause. The amount of information obtained is positively correlated with the amount donated, and experimental evidence suggests that reading more leads to higher donations. This chapter also finds signs of guilt or self-deception among subjects who do not make donations. These findings indicate that even agents who appear to act generously by making positive donations to charities might still partake in information avoidance to prevent themselves from donating more.

In the main experiment described in this chapter, subjects play multiple rounds of a dictator game with a different charity as the recipient in each period. In one stage

of the experiment, subjects can choose to read a short description about the charity in order to learn about it before deciding how much to donate or keep. Without reading anything, subjects would not know which charity had been chosen. Half of the charities that are selected for a subject are ones she claims to dislike and the other half are ones she claims to like. A subject could therefore justify a decision to donate nothing by ignoring the description and telling herself that the selected charity is probably one she does not like. Subjects may also choose to read just enough to know whether the charity is liked or not, but avoid learning more about it.

Previous studies have found that subjects make substantial contributions to charities in laboratory dictator games (Eckel and Grossman, 1996 and 2000; Benz and Meier, 2006). This chapter does not focus on the size of donations to particular charities or contrast the findings to a standard dictator game. Non-profit organizations are used as recipients primarily because they can be described within a paragraph, but also because one can readily imagine the everyday equivalent of seeking out information about a charity. Reading an entire description of a charity in this experiment can be likened to taking the time to talk to volunteers about the fund they are raising money for. Reading only part of a description is akin to changing channels halfway through a *Christian Children's Fund* commercial.

This chapter is most closely related to studies on the justification of self-serving decisions, and particularly to studies of information avoidance. Giving subjects in a dictator game the opportunity to avoid knowing whether their payoffs were aligned with their recipients' has been found to increase self-serving allocations (Dana, Weber, and Kuang, forthcoming). Subjects avoid information more often in these games when they can reassure themselves that the recipients' payoffs are most likely aligned with their own (see Chapter 2).

Agents may also justify decisions by seeking information sources that can provide support for their beliefs or preferences. They might read about preferred political candidates (Chaffee and Miyo, 1983) or check their stocks only when a market index has risen (Karlsson, Lowenstein, and Seppi, 2005). This confirmatory bias can be so strong that subjects will pay for non-instrumental information in laboratory

experiments (Eliasz and Schotter, 2006; Chapter 4 of this thesis).

Other forms of justification have been studied experimentally. When outcomes for a recipient are ambiguous or conflicts of interest are revealed, experimental subjects give less honest advice because they can justify taking the action that benefits them (Schweitzer and Hsee, 2002; Cain, Loewenstein, and Moore, 2005). Ambiguity about recipients' payoffs in a dictator game leads to less equitable allocations and can even make dictators perceive their actions to be fair (Haisley and Weber, 2005).

Guilt may be a driving factor for generosity in dictator games. In Dana, Cain, and Dawes (2006) and Lazear, Malmendier, and Weber (2006), subjects are allowed to opt out of a dictator game and receive an amount of money similar to what they could have allocated in the game itself. Opting out prevents a potential recipient from knowing the game could have been played, but the recipient earns nothing. Even when opting out provides them with less money than keeping everything in the dictator game, many subjects take this option. These subjects presumably would feel guilty about directly failing to allocate money to a recipient, but they somehow avoid guilt by exiting out of the game.

The current chapter considers a physiological expression of guilt, namely, looking away quickly from charities that one does not donate to. Guilt has been discussed theoretically by economists. For example, Kandel and Lazear (1992) provide an organizational application of guilt, and Becker (1996) discusses the avoidance of beggars on the street as being caused by guilt about not donating. However, direct measurements that relate to guilt are lacking in the literature. (See Elster, 1998, for a survey of the rare appearances of guilt and other emotions within economic research.)

This chapter aims to provide greater insight on forces that might drive information avoidance, such as guilt and the effect of learning about the potential recipient of a donation. The next section presents the primary experiment used to address the questions of this chapter, followed by hypotheses and results. Section 3.5 contains the details of other treatments designed to determine whether people who choose to donate to a charity seek out more information about it, or whether gaining information causes one to donate more. Section 3.6 ties together the findings of the experiments

discussed in this chapter.

3.2 Experimental Design

The primary experiment to assess information seeking and avoidance behavior consisted of three stages, with subjects receiving instructions prior to each stage. The first stage determined which charities subjects would theoretically support. The second stage provided subjects with the opportunity to read a description of a charity before deciding how much money to donate. The third stage was similar to the second, but simply provided subjects with the names of charities, rather than full descriptions.

In the first part of the experiment, subjects completed a questionnaire that asked them to rate 38 different charities in order to determine which causes they supported or disliked. The goal of this stage was to create a list of at least 10 liked and 10 disliked charities for each subject, to be used in the second stage. Since it can be difficult to find charities that people dislike, many of the organizations on the questionnaire had missions that contradicted those of other groups on the list. For example, someone who supported the Brady Center to Prevent Gun Violence should be unlikely to support the NRA Civil Rights Defense Fund. Other groups were listed in hopes that the Los Angeles-based student participants would not deem them worthy of donations, such as the American Tarantula Society and Metropolitan Opera Guild (which supports the Metropolitan Opera in New York City). All of the organizations listed were categorized as tax-deductible by the Internal Revenue Service; this designation was used to indicate a group's legitimacy. A full list of the charities used in this experiment can be found in the appendix.

Each charity on the questionnaire was listed with its name and a brief description of its mission or goals, which was summarized from the group's website. This description sometimes included specific positions the group had taken, such as a pro-life stance. For each charity, subjects answered three questions about how much money they would want this group to receive. These questions were: "Would you want \$50 to be donated to this charity?" (Yes or No), "Which would you prefer: \$50 donated

to this charity or \$10 donated to the charity of your choice?,” and “Which would you prefer: \$20 donated to this charity or \$10 donated to the charity of your choice?”

To give subjects an incentive to provide honest answers, one subject, one charity, and one of the three questions for that charity was selected at random at the end of the experiment. Any donation the chosen subject indicated she would prefer in the chosen question would be sent to the charity directly from experimental funds. This donation would not affect subjects’ payoffs. Subjects were asked to write their favorite charity from the list at the bottom of their questionnaire. This group would receive \$10 if that were the selected option in the randomly chosen question.

If we call the first choice to each question “A” and the second choice “B,” then charities were categorized as *liked* if the response to the three questions was AAA, *disliked* if the response was BBB, *mildly liked* if the response was AAB, and *mildly disliked* if the response was ABB. Subjects were told what each of the response patterns would indicate, and they were not allowed to give other forms of answers, since no other responses would make sense.¹

For the second stage of the experiment, subjects participated in 20 rounds of a dictator game. In each round, a subject allocated \$20 between herself and the charity selected for that round. Subjects could choose to donate \$0, \$10, or \$20, keeping \$20 minus the amount donated. At the end of the experiment, one round from this stage was chosen at random for payment. Subjects were told that half of the charities chosen to be recipients in this stage were ones they liked and the other half were ones they disliked. A subject would not immediately know which charity had been selected in a round, but she could read a brief description of it by using her mouse to drag a cursor over several gray boxes on the computer screen. When the cursor was placed over a box, the text underneath would appear.

This procedure was implemented using MouselabWEB (©2004, Martijn Willemssen and Eric Johnson). This software allows one to track which boxes a subject rolls

¹Some subjects did not initially select enough liked or disliked charities, and they were instructed to change responses for charities they claimed to mildly like to a response of mild dislike (or vice versa), until they reached the minimum of 10 disliked (or liked) charities. I do not have data on these changes from this experiment, but in later experiments, subjects’ initial responses are used to determine whether they like a charity.

her mouse over, as well as how long is spent on each box, and the order in which the boxes are approached. This can reveal how much of a paragraph a subject reads.² Figure 3.1 shows the choice subjects could make in this stage, and the look of the table used to reveal the charity description.




Figure 3.1: Screenshot from stage 2

The charity descriptions used in this stage were short paragraphs of between three and six sentences. They were intentionally written to be vague, with the exception of a key phrase that should clarify the organization's purpose or beliefs. A subject who wanted to figure out whether the organization was one she liked should only have to read up to this phrase, but subjects may choose to learn more by reading beyond it.

The last stage of the experiment was conducted like the second stage. Subjects participated in the same dictator game as before, over the course of 20 rounds. This stage used the same 20 charities as recipients as stage 2 had. The main difference in this stage was that rather than reading a description of an organization, a subject

²One caveat is that we cannot be positive that a subject is reading everything in the box the cursor is on.

could search through the table of gray boxes to find the charity’s name. However, in half of the rounds in this stage—chosen at random—the name of the charity would simply appear “on top” of the gray boxes, so it would immediately be seen.

At the end of the experiment, subjects were paid a \$5 show-up fee, plus the amount they chose to keep in two randomly selected rounds, one from stage 2 and one from stage 3. Any amounts donated in the selected rounds were paid online by the experimenter during the session, while subjects completed a brief survey with questions on demographics and charitable giving. The receipt for the donations was then shown to subjects so they could verify that their desired donations had been made.³

3.3 Hypotheses

These procedures allow for the direct testing of several hypotheses. Unless otherwise noted, the hypotheses refer only to charities that subjects claimed to like.

Hypothesis 1. The average donation in stage 2 will be larger for subjects who read all of a description than for those who read some of a description. In turn, the subjects who read some of a description will donate more than those who read nothing, who are expected to give \$0.

The rationale for this hypothesis can be explained starting at the end of the statement. Subjects who read nothing should not give any money, since they cannot even know if the charity selected in a given round is one they support. Subjects who read some, but not all, of a description are most likely attempting to determine what the charity is. If it is one they like, they may make a donation, but they do not feel compelled to learn more about the charity. Subjects who read everything, on the other hand, are expected to donate more for one of two reasons. It may be the case that a subject decides to donate to a charity she likes, then gets a warm glow

³The ability to quickly make donations and instantly show subjects proof of payment was facilitated by the use of justgive.org, a nonprofit website that allows users to store favorite charities and make donations to multiple organizations at once.

(Andreoni, 1989) by learning more about the organization she is helping. In this case, information seeking is the effect of a decision to donate money. Information seeking might instead cause larger donations, since the more someone learns about a charity, the more she might want to donate. The direction of causation will be discussed in Section 3.5.

The stage 3 counterpart to the previous hypothesis is the more obvious prediction:

Hypothesis 2. Donations in stage 3 will be larger for subjects who find the hidden charity name than for those who do not.

Since previous research has demonstrated that subjects avoid information to make self-serving decisions in dictator games, taking away the ability to justify decisions in this way should increase the amount given to recipients. This leads to the next hypothesis.

Hypothesis 3. Donations in stage 3 will be larger when the charity name is not hidden.

Finally, the next two hypotheses consider whether subjects show signs of guilt or self-deception when choosing not to donate to a charity they like. Even when it is impossible to avoid information, people may quickly turn away from sights that make them uncomfortable. For example, someone who walks past a homeless man without giving him money is likely to look down or away from the man rather than make eye contact with him. This may be done out of guilt, or it may be the case that by quickly looking away, a person can convince herself that she did not see anyone. If subjects in stage 3 of the experiment take less time to look at names of charities they like before donating nothing, that can indicate that they feel guilty about their decision. The following two-part hypothesis is primarily based on conventional wisdom about expressions of guilt, and as such, the results pertaining to it should not be given the same weight as other contributions of this chapter. However, it is a first step at shedding light on some of the emotional underpinnings that may lead to information avoidance.

Hypothesis 4a. Of subjects who see the charity name in stage 3, those who do not donate will spend less time looking at the name than those who make a donation.

4b. For subjects who see the charity name but do not donate in stage 3, less time will be spent looking at names of charities that are liked than at those of charities that are disliked.

3.4 Results

Sixty subjects participated in this experiment, 10 at the California Institute of Technology and 50 at the University of California, Los Angeles. Each session contained 10 subjects. No significant differences were found between subjects at the two schools, so the data are pooled for analysis. Table 3.1 summarizes information and donation decisions in stages 2 and 3 of the experiment, as well as demographic information about the participants.⁴ Donation data here and elsewhere in the experiment is based on the amount subjects indicated they would give, not on the actual amount that was paid in a randomly selected round.

This table provides an overview of the decisions subjects made in the experiment, but before we can explicitly test the hypotheses, I will first define what it means for a subject to read a charity description. The MouselabWEB data provides information on how much time subjects spend on any given box, or table cell, that they reveal. I will consider a subject to have read a cell if she spent more than 50 milliseconds (ms) on that cell. This threshold is considered to be around the minimum amount of time it takes to read one word (Rayner et al, 1981; Rayner and Pollatsek, 1989), so cells that were looked at for less than 50ms were most likely passed over as a subject moved the mouse to see a different cell or choose a donation amount. It is unlikely that a subject would have read everything in a cell if she spent just slightly more than

⁴N should be 600 for variables relating to liked charities, since subjects participated in a total of 1200 rounds in stage 2, half of which involved charities they liked. Due to a technical error, one round had to be excluded from analysis. Two rounds are missing data for stage 3, one of which was for a liked charity.

	N	Mean (std dev)	Median
Stage 2 donation *	599	4.54 (6.68)	0
Proportion read *	599	0.67 (0.37)	0.78
Stage 3 donation *	599	3.77 (6.13)	0
Looked for hidden charity name in stage 3	600	78.33%	-
Age	60	22.13 (3.91)	21
Female	60	51.67%	-
Undergraduate	60	83.33%	-
US citizen	60	75.00%	-

* For liked charities only

Table 3.1: Summary information

50ms on it, but this low cutoff point avoids excluding text that someone processed quickly and should bias the results against the hypotheses, if at all.

We will say that subjects have read an entire description if they have spent more than 50ms on every non-blank cell in the table containing the description. Note that this definition does not imply that subjects read line by line, from left to right. It would be difficult to determine whether someone who reads a paragraph in non-consecutive chunks processes the information differently than someone who reads each word in order. Based on the raw data, though, most subjects seem to look at cells in consecutive order.

Hypothesis 1 is supported by the data. Subjects who read all of the description for a charity donate more than subjects who only read some of it. There is also a significant difference between the amount donated by subjects who partially read the description and those who do not look at it. As expected, subjects who do not read about the charity at all donate nothing, with the exception of one person who donated \$10 in one round.⁵ Figure 3.2 shows the amounts donated for each level of

⁵Comparing donations made by subjects who read nothing to 0, $t_{78} = 1.00$, $p = .32$. Comparing

information subjects sought.




Figure 3.2: Average donation, by amount of description read

When demographic data and information about the charities are incorporated into the analysis, the proportion that subjects read continues to play a significant role in the amount subjects donate. Table 3.2 shows the results of an ordered logistic regression of donations on the amount a subject read, how much a subject claimed to like the charity (restricted to “mildly like” or “like”), demographic characteristics, and indicator variables for each charity. The individual charities are not included in the table; most do not have a significant effect on donation size. As this table shows, the odds of choosing to donate \$10 or \$20 increase as subjects read more, and demographic characteristics also have a significant effect on the size of donations.

Instead of focusing solely on the proportion that subjects read, we can examine differences in donations based on whether subjects reach the key phrase in the charity description. Recall that this phrase should clarify what charity had been chosen as the recipient for a given round. The key phrase, or key, is defined as a range of consecutive cells that state an organization’s mission or disclose information about the group that would distinguish it from other charities. Subjects are considered to reach the key phrase if they read every cell up to one that is included in the key or

subjects who read some of the description to those who read none, $t_{338} = 5.09$, $p < .001$. Comparing those who read some to subjects who read all of the description, $t_{517} = 5.17$, $p < .001$.

Variable	Coefficient
Proportion Read	3.00*** (0.38)
Strongly Like Charity	0.37 (0.23)
Female	1.27*** (0.23)
Age	0.07*** (0.03)
US Citizen	1.54*** (0.30)
Log likelihood = -403.01, N = 599	
*** Indicates $p < .01$	

Table 3.2: Ordered logit of donation on variables listed and dummy variables for all charities

one cell past the end of the key. Subjects are categorized as reading beyond the key if they see every cell up to a point that is greater than one cell past the end of the key.⁶

The key phrase was defined subjectively but the proportion of cells read by subjects who see some of a charity description is positively correlated with the location of the key. As a further test of the defined key, 10 subjects who had not taken part in any related experiments were asked to read charity descriptions and choose the name of the group that was being described. Each subject saw descriptions for all 38 charities, and had to select an answer from a list of the 38 names. Subjects were paid \$0.50 for each correct answer. Half of the descriptions were cut off in the cell before the key phrase started, and the other half were cut off immediately after the key phrase ended. For each charity, 5 subjects saw the shorter description and 5 saw the longer one. If approximately the same number of people correctly guessed a charity when the description was short as when it was long, then we can consider the key phrase to be poorly defined and exclude that charity from the following analysis.

⁶There are five instances in which the key fell at the end of the description, but subjects read all the way to the end of the paragraph. Here, these five observations are counted as reading beyond the key.

Three charities had a key phrase at the end of the description, so separate long and short blurbs could not be created; one of these is excluded from analysis since fewer than 8 subjects guessed it correctly.

Subjects who reach the key phrase donate \$4.53 on average, \$3.60 more than those who read nothing or who see only pieces of a description ($t_{126} = 3.78$, $p < .001$). This result is found using a restrictive definition of the validity of the key phrase. Observations are only included for 17 charities, which were identified by 4 or 5 people in the test who saw the description through the key phrase but at least 3 fewer of the people who only saw the description preceding the key phrase. To compare subjects who reach the key phrase to those who read beyond it, I include an additional 8 charities that were correctly guessed by 2 or 3 testers when the description was short and 4 or 5 testers when it was long.⁷ Donations are significantly larger for those who read beyond the key phrase (averaging \$6.11) than for those who only read up to the point at which they might know what the charity is (average of \$4.14; $t_{333} = 2.065$, $p < .05$).

Based on the support for Hypothesis 1, it is not surprising that Hypothesis 2 also holds. Subjects who find the hidden name of a charity in stage 3 donate significantly more than those who do not find the name ($t_{309} = 5.63$, $p < .001$).

The third hypothesis does not explain the data. There is no significant difference between the amount donated in stage 3 when names are hidden versus when they are automatically revealed by the software ($t_{597} = 0.05$). This result also holds when the only subjects considered are ones who never look for hidden names. This finding might imply that non-donaters in stage 2 do not bother to read an entire charity description because it would have no effect on their donation decision. It might also be the case that they avoid reading out of a feeling of guilt, but their disutility from guilt is less than the disutility they would get from giving up \$10.

To examine whether guilt or self-deception might affect subjects who choose not to donate, we can consider the amount of time subjects spend looking at names that

⁷We would not want to include charities with meaningless keys, but since the focus is on people who read the key phrase and those who read past it, it is less important to weed out charities that could potentially be identified before the key phrase.

could make them uncomfortable in stage 3. We need to separately analyze names that were hidden from names that were always revealed, since subjects gain information differently in each case. For charity names that were hidden, the time spent looking at the name will be defined as the total amount of time that a subject had her mouse on the cell containing the name. Only the subjects that actually saw the hidden name will be considered for this analysis. For charity names that were not hidden, I will define the time spent on the name as the number of milliseconds between when the page loaded and when the subject submitted her allocation decision. This will include all subjects. To account for different lengths of charity names, the measure used in this analysis will be the time spent per word in the name. Outliers below the 5th or beyond the 95th percentile for time spent per word are excluded. These outliers are defined separately for names that are hidden and not hidden.

Of the subjects who reveal a hidden charity name for an organization they like, subjects who choose not to donate spend less time per word than subjects who donate \$10 or \$20. This result is marginally significant ($t_{211} = 1.61$, one-tailed $p = .054$). When the name of a liked charity is automatically revealed on a page, the same effect is found ($t_{256} = 2.29$, $p < .05$). This supports Hypothesis 4a.

Focusing only on subjects who donate nothing, for those who reveal a hidden charity name, less time is spent looking at the names of liked charities than at those of disliked ones ($t_{312} = 1.86$, one-tailed $p < .05$). When names are not hidden, this same effect is found, but at a marginal level of significance ($t_{437} = 1.43$, one-tailed $p = .077$).⁸ Figure 3.3 shows the average amount of time spent looking at charity names that subjects chose to reveal. Although some effects are marginal, it appears that guilt may lead non-donaters to quickly glance away from charities they might consider giving to. When they see the name of charities they dislike, they do not feel

⁸One could argue that the words used in names of liked charities are shorter than the words in disliked charities, so accounting for the number of words in the charity name is not enough. It is in fact true that the liked charities in this experiment contain 1 fewer syllable, on average, than the disliked charities. But evidence suggests that people process entire words when reading, not individual syllables (Crowder, 1982). Likewise, the number of milliseconds spent per character in a word or phrase is nonlinear in the length of the phrase (Trueswell, Tanenhaus, and Garnsey, 1994), so dividing by the number of letters in the names would not necessarily improve the measure.

compelled to turn away because there is no reason to feel badly about not donating.




Figure 3.3: Time spent per word (milliseconds) on hidden names that have been revealed, by donation and preference for charity

3.5 Tests of Causality

We have seen that when subjects read more of a charity’s description, they donate more to that group. However, the experiment presented above cannot test causality. We cannot tell if subjects who have decided to donate to a charity read more to feel good about supporting the cause, or if reading more leads people to donate. To address this question, two other experiments were conducted that required subjects to read a given amount of information about a charity. I will refer to these experiments as *forced reading* treatments, and the experiment described above as an *endogenous choice* treatment, since subjects could decide how much to read. The design of the forced reading treatments nearly matched the endogenous choice procedures, but stage 3 was omitted and these sessions contained variants of stage 2.⁹ In

⁹The questionnaire stage was also changed slightly to make it faster and easier for subjects. Subjects could rate a charity as one they liked, mildly liked, mildly disliked, or disliked. When this choice was made, the three questions used in the primary experiment were automatically filled in so subjects could see them. This change resulted in a more even distribution of preferences across the four categories than in the previous experiment. Subjects also judged charities as liked or mildly

these experiments, subjects were asked to read the entire paragraph about a charity and had to answer a question about what they had read immediately following each donation choice. Subjects were paid \$5 for answering at least 80% of the questions correctly. The descriptions were randomly chosen to be short or long, and subjects were told this, so they should not have viewed the length of the description as an indicator of the charity's quality. In one of these treatments, which will be referred to as the *long option* treatment, a subject who saw the short description could click a button to read the longer description. Subjects were told that they were not obligated to select this option, and they could answer questions about what they read without seeing the longer description.

The long descriptions in these experiments were the same as the descriptions used in the endogenous choice treatment. The short descriptions were created by considering the median last cell read by subjects who read some, but not all, of the description for a given charity in the endogenous choice experiment. Short descriptions were cut off at the end of the sentence that contained this cell. The median cell fell within the last sentence for two charities; for these, the final sentence was simply cut off to make the short description.

3.5.1 Hypotheses

If it is the case that reading more leads someone to donate more, then the following hypothesis would be supported.

Hypothesis 5. Donations will be larger for long descriptions than for shorter ones.

We must be careful, though, because Hypothesis 5 could also be supported for another reason. Subjects may not give to charities with short descriptions because they feel they do not have enough information to determine that the charity is one

 liked more often than in the previous experiment, with 58% of all ratings favorable, compared to 54%. Since the number of charities seen in stage 2 is limited, these differences are not expected to impact results.

they like. The long option treatment allows us to test whether ambiguity aversion plays a role if Hypothesis 5 is supported by the data. If ambiguity aversion leads to lower donations when descriptions are short, subjects should ask to see the longer descriptions when they are given the option, and then donate more. We might also consider whether simply knowing that the option to see a longer description exists has an effect on behavior. If the following two-part hypothesis is supported, then ambiguity aversion is most likely a cause of any differences in donation size for long and short descriptions.

Hypothesis 6a. When faced with a short description, subjects will choose to see the long description in the long option treatment and donate more upon reading the longer paragraph.

6b. Subjects who are given a short description will donate more when they have the option to read a longer description.

3.5.2 Results

	N	Mean (std dev)	Median
Donation *	526	4.41 (6.32)	0
Donation, short descriptions *	243	3.86 (5.95)	0
Donation, long descriptions *	283	4.88 (6.60)	0
Age	52	20.10 (1.46)	20
Female	52	69.23%	-
Undergraduate	52	98.08%	-
US citizen	52	73.08%	-

* For liked charities only

Table 3.3: Summary information, forced reading treatments

Fifty-two UCLA students participated in the forced reading treatments. Four

sessions were conducted, two of which were for the long option treatment. Table 3.3 provides a summary of demographic characteristics and donations to liked charities. Subjects were properly incentivized to read the charity descriptions, and answered 96% of the reading questions correctly. As a second check on whether subjects read the descriptions fully, we can use the timing data from MouselabWEB and define reading each cell as before. Only 19 descriptions were not fully read; one subject accounts for 11 of these observations. The descriptions that were not read completely are excluded from the following analysis.

We cannot reject Hypothesis 5, but the differences in donations after subjects read long or short descriptions are not very large. When the data from both treatments are combined, subjects who read a short description donate \$3.90 on average, while those who read a longer description (including subjects who selected to read it in the long option treatment) donate \$4.82 ($t_{517} = 1.65$, one-tailed $p < .05$). Similar results are found when the rounds with short descriptions that offer a long option are excluded.

Table 3.4 shows the log odds from an ordered logistic regression similar to the one described in support of Hypothesis 1. This regression includes demographic characteristics, the amount subjects claim to like a charity, and the individual charities, but an indicator variable for the length of the description now replaces the proportion of the description a subject read. Longer charity descriptions increase the odds of a larger donation, though this effect is only marginally significant.

Hypothesis 5 appears to be supported by the data, if only weakly, but we should determine whether ambiguity aversion drives the results. In the long option treatment, subjects only ask to see the long description in about 12% of the rounds in which they receive a short description. Subjects choose to see the long description equally often for liked and disliked charities.¹⁰

Although the percentage of times that subjects request the long option is significantly different than 0, it appears to be selected primarily out of curiosity. In the

¹⁰The paucity of decisions to read the long description serves to reinforce the general idea of this and other papers that subjects will avoid costless information about recipients in dictator games.

Variable	Coefficient
Saw Long Description	0.34* (0.20)
Strongly Like Charity	0.32 (0.21)
Female	0.19 (0.22)
Age	-0.05 (0.07)
US Citizen	0.52** (0.24)
Log likelihood = -405.33, N = 519	
* Indicates $p < .10$, ** Indicates $p < .05$	

Table 3.4: Ordered logit of donation on variables listed and charity indicator variables, forced reading treatments

15 rounds in which subjects chose to see long descriptions for charities they liked, the subjects who saw long paragraphs donated nearly the same average amount as subjects who saw only the short descriptions of the same charities ($t_{80} = 0.18$). Hypothesis 6a is not supported, and neither is 6b. Subjects donate the same amount to liked charities with short descriptions when they have an option to see a longer paragraph as they do without that option ($t_{254} = 0.58$).

It appears that reading more about a charity causes subjects to make larger donations to that group. Ambiguity aversion cannot explain why donations to charities with short descriptions are smaller than donations to charities with longer descriptions. Donation sizes do not differ much between short and long descriptions, but this difference seems larger when one considers that the long descriptions only provided, at most, a few sentences of additional information.

This analysis shows that for subjects who have to read at least something about a charity, subjects who read more donate more. From a policy perspective, though, we should ask whether forcing people to learn about a charity would lead to higher donations. To address this question, we can consider an ordered logistic regression that uses data from the endogenous choice treatment and the forced reading treatments.

To make comparisons across treatments, I define the proportion of a description that subjects read as follows. For the endogenous choice data, the proportion subjects read is defined as before. For the sessions in which subjects are given a long description, the proportion read is set to 1.¹¹ For subjects who saw only a short description, the proportion read is defined as the number of cells in the short description for a given charity, divided by the number of cells in the long description for that same organization.

Table 3.5 contains the coefficients from an ordered regression of donations on the proportion subjects read, whether subjects had to read the description given to them, and the interaction of these two variables (i.e., the proportion read when forced), as well as on the amount subjects liked a charity and demographic data. Indicator variables for charities were also included in the regression but are not presented in the table.

Variable	Coefficient
Proportion Read	2.91*** (0.35)
Forced Reading	1.68*** (0.49)
Proportion Read * Forced Reading	-2.28*** (0.56)
Strongly Like Charity	0.35** (0.15)
Female	0.75*** (0.15)
Age	0.04 (0.02)
US Citizen	0.98*** (0.18)
Log likelihood = -836.10, N = 1118	
** Indicates $p < .05$, *** Indicates $p < .01$	

Table 3.5: Ordered logit for combined treatment data. Dependent variable is donation.

¹¹The proportion read is set to missing for rounds in which a subject who should have read everything failed to do so.

One meaningful way to discuss the results of this regression is to compare first differences of the effects of forcing subjects to read.¹² For this purpose, all variables in the logit are set to their means, with the exception of some binary variables that are set to the point at which more donations are made. Comparisons will be made for female U.S. citizens who strongly like the charity. Holding the proportion read at its mean, the average effect of forcing subjects to read descriptions is negligible; there are no significant changes in the probability of donating \$10 or \$20.

<i>Increase in Proportion Read:</i>	Endogenous Choice		Forced Reading	
	<i>0 to Mean</i>	<i>Mean to 1</i>	<i>0 to Mean</i>	<i>Mean to 1</i>
Change in P(Donating 0)	-.39 [-.45, -.34]	-.18 [-.22, -.14]	-.11 [-.25, .04]	-.04 [-.10, .02]
Change in P(Donating 10)	.29 [.24, .34]	.07 [.04, .11]	.07 [-.02, .17]	.02 [-.01, .05]
Change in P(Donating 20)	.10 [.08, .13]	.11 [.08, .14]	.04 [-.02, .08]	.02 [-.01, .05]

Note: Brackets contain 95% confidence interval for estimated probability.

Table 3.6: First differences: Estimated changes in the probability of donating \$0, \$10, or \$20 when the proportion read increases, by treatment

Table 3.6 summarizes the changes in the probability of donating \$0, \$10, or \$20 that result from increasing the proportion read from 0 to the mean proportion read across experiments (.738) and from the mean to 1, in the endogenous choice treatment and in treatments with forced reading. For the forced reading treatments, subjects never would avoid reading altogether, so the figures for changing the proportion read from 0 to the mean are purely based on simulations from the estimated model. As this table shows, the effect of an increase in the amount a subject read is larger in the endogenous choice treatment than in the forced reading treatment. Based on the simulations from the logit model, the effect of an increase in the proportion read for those who were forced to read is small, but nearly significant. We saw before that subjects donated more when they read longer descriptions in the forced reading

¹²These first differences are calculated using Clarify software (King, Tomz, and Wittenberg, 2000; Tomz, Wittenberg, and King, 2001).

treatment, but the difference in donations was small. The last column of Table 3.6 provides further evidence that supplying subjects with additional information may have positive, but not large, effects.

3.6 Conclusion

This chapter adds to the literature on information avoidance by demonstrating that not only does an agent's choice to gain information about a recipient lead to greater donations, but the amount of information that is obtained also matters. The more subjects read about a charity they like, the more they choose to donate to the cause. This effect is strongest for subjects who choose to read descriptions of charities, but it also holds for subjects who are instructed to read the descriptions.

Given that reading more about a charity can cause people to make larger donations, we have good reason to avoid information about people or charities that could use our support. An added benefit of such avoidance might be the ability to prevent ourselves from feeling guilty about not donating money. This chapter provides evidence that people might feel discomfort or guilt when they are confronted with the name of a charity they are not donating money to. Future research should examine this finding more closely with tools that are better suited to examining physiological responses, such as an eye tracker that can measure pupil dilation and subtle movements of the eye. Future studies should also address the question of whether actively avoiding information can prevent guilt.

For the unfortunate people who are unable to avoid learning about a charity or needy person altogether, this chapter offers some hope. Reading only a few facts about a charity leads to lower donations than reading a more complete description. One might not be able to completely avoid donating some amount of money, but limiting one's intake of information about a given cause could at least prevent larger outlays.

As for those on the other side of the equation, who are hoping to raise money for a cause, the experiments described above provide some insight on the benefits of

forcing people to learn about a charity. Simply providing the name of a charity is not likely to coerce a routine avoider to donate money, as evidenced by the lack of support for Hypothesis 3. The forced reading experiments show that giving people additional details about a cause may help to increase donations. However, on the whole, donations in these experiments were no larger than in ones in which subjects could actively seek out information. Although imposing information upon people might garner some donations from those who would have avoided all knowledge of a charity to justify not giving, it is not likely to be more effective than simply making information easy to obtain for those who are willing to seek it.

3.7 Appendix: List of Organizations

The following list contains the non-profit organizations that could be selected as recipients in the experiment, as well as the short description that was listed with the names in the questionnaire.

Abstinence Clearinghouse Promote the teaching and practice of sexual abstinence before marriage by providing a central location for programs, speakers, and materials

American Border Patrol Detects, locates, and reports illegal immigration as it occurs, and makes information available to the public

American Nonsmokers' Rights Foundation Promotes smoking prevention and education about smoking, passive smoke, and the tobacco industry

American Red Cross Helps people prepare for and respond to emergencies, provides disaster relief

Americans United for Separation of Church and State Defends separation of church and state in the United States; is against the teaching of intelligent design in schools

American Tarantula Society Furthers education about tarantulas and other arachnids

Becket Fund Protects the free expression of all religious traditions; believes public schools should not be religion-free and employers should accommodate religious observances of employees

Boys and Girls Clubs of America Inspires and enables young people to realize their full potential as productive, responsible, and caring citizens

Brady Center to Prevent Gun Violence Aims to reform the gun industry through litigation, enact regulations to reduce gun violence, and educate the public to prevent gun violence

Campaign to End the Death Penalty Grassroots organization dedicated to abolishing capital punishment in the United States

Coalition for Humane Immigrant Rights of Los Angeles (CHIRLA)

Advance human and civil rights of immigrants and refugees in Los Angeles, promote harmonious multi-ethnic relations, empower immigrants and allies

The Center for Bioethics and Human Dignity Helps individuals and organizations address bioethical challenges, incorporating Christian values; finds euthanasia and embryonic stem cell research to be unethical

Criminal Justice Legal Foundation Advocates reduced rights for accused and convicted criminals, supports cases upholding the death penalty

D.A.R.E. (Drug Abuse Resistance Education) America Gives kids life skills to avoid involvement with drugs, gangs, and violence

Death With Dignity National Center Supports and promotes a state law allowing physician-assisted death for terminally ill patients

Federation for American Immigration Reform (FAIR) Works to end illegal immigration to the United States and set legal immigration at the lowest feasible levels

Focus on the Family Spread the Gospel of Jesus Christ by preserving traditional values and the institution of the family; positions include being pro-life and against gay marriage

French Institute Alliance Francaise Promotes and enhances the knowledge and appreciation of French culture, located in New York City

Greenpeace Fund, Inc. Research and education branch of an independent, campaigning organization that uses non-violent, creative confrontation to expose global environmental problems and force solutions

Habitat for Humanity International Builds simple, decent, affordable housing in partnership with people in need

Intelligent Design Evolution Awareness Center Promotes intelligent design theory and fosters discussion and a better understanding of intelligent design and the creation-evolution issue

Locks of Love Provides hair prosthetics for children with long-term medical hair loss, free of charge or on a sliding scale

Marijuana Policy Project Foundation Works to legalize medical marijuana and tax and regulate marijuana for general adult use

Marine Toys for Tots Foundation Delivers new toys to needy children at Christmas to motivate them to grow into responsible, productive, patriotic citizens and leaders

Metropolitan Opera Guild, Inc. Supports the Metropolitan Opera in New York City, and encourages the appreciation of opera

National Organization for Women Foundation Devoted to furthering women's rights through education and litigation, aims to protect women's reproductive health options and promote equality

National Right To Life Committee Educational Trust Fund Provides educational programs opposing abortion, infanticide, and euthanasia

NRA Civil Rights Defense Fund Established by the National Rifle Association, provides legal and financial assistance to people defending their right to keep and bear arms

Pandas International Dedicated to panda research and preservation

Parents, Families and Friends of Lesbians and Gays Promotes the health and well-being of gay, lesbian, bisexual and transgendered persons, their family and friends

People for the Ethical Treatment of Animals (PETA) Establishes and protects rights of all animals, believes animals are not ours to eat, wear, experiment on, or use for entertainment

Planned Parenthood Supports reproductive freedom and provides reproductive health care and family planning services

The Princess Project Provides free prom dresses and accessories to San Francisco Bay area girls who could not otherwise afford them

Save the Children Helps children in poverty and children in crisis around the world and in the US

The Shark Research Institute Dedicated to shark conservation, education, and research

Special Olympics, Inc. Provides year-round sports training and athletic competition to people with intellectual disabilities around the world

Stem Cell Research Foundation Provides public education and funds research using embryonic and adult stem cells in an effort to help find treatments and cures for a wide range of diseases

USO (United Service Organizations) Provides morale, welfare, and recreation-type activities to uniformed military personnel

Chapter 4

An Experimental Study of Selective Exposure

4.1 Introduction

4.1.1 The Selective Exposure Bias

Both casual observation and the empirical literature on biases reveal a number of situations in which decision makers choose sources of information that can confirm their preferences but may not lead to the best choice. Voters tend to read more about their preferred political candidate than about an opponent. Similarly, a car buyer may keep learning about the model she is initially drawn to, instead of researching other vehicles. A consumer debating whether to splurge on a purchase may seek out the opinion of an extravagant friend, rather than a thrifty one. In these examples, the decision makers prevent themselves from discovering that another alternative may be a better choice. For instance, if the car buyer reads about other cars, she might find one that has more features she is looking for at the same price as her initially preferred model.

The above examples demonstrate selective exposure to information. This bias can be used to describe differing phenomena. We will use the term *selective exposure* to refer to the choice of an information source that could potentially confirm that the option one prefers is the best alternative when another information source could maximize expected payoffs. The purpose of this chapter is to examine the effects of

institutions on the prevalence of the selective exposure bias. We consider an array of experimental contexts and take full advantage of experimental control by studying selective exposure in a setting that is free of outside influences.

We conduct an experiment that allows us to measure the extent of the selective exposure bias in a fairly simple task. A state of the world is chosen at random each round, and subjects must guess whether the state is red or blue. We induce a preference for one of the states by offering a bonus if the subject correctly guesses that state. Before making her guess, a subject can choose one of two sources of information, which provides a signal about the state. One source has the potential to confirm that the state is the one that is ex-ante preferred, but the other source provides signals that allow for payoff maximization. A subject exhibiting the selective exposure bias would choose the first source.

We embed this simple task in different settings. Our baseline is a treatment in which subjects choose an information source and guess the state individually. We introduce an externality in the form of majority rule voting, and examine the effect of separating the choice of an information source from the guess about the state. Across all contexts, we find a large and persistent amount of selective exposure. More than half of all choices reflect the selective exposure bias. Although many subjects choose their information source poorly, we find that subjects are fairly good at recognizing the most likely state after receiving a signal. In our setup, though, a choice of the most likely state upon receipt of a signal will result in lower payoffs for the subjects who selectively exposed. Our data show a failure to backward induct and choose the payoff maximizing information source.

This study contributes to the current literature by quantifying the extent of the selective exposure bias in a simple experiment, examining learning over time, and considering the impact of different institutions on the bias. Our approach to the study of selective exposure is novel. We induce preferences, whereas much of the experimental work on selective exposure examines the bias with respect to beliefs or preferences subjects already have when they walk into the lab. This allows us to obtain a more pure measure of the extent of the selective exposure bias. Although

selective exposure is often discussed in terms of political choices, to our knowledge, this study is the first to directly examine the impact of majority rule voting on the prevalence of the bias.

It is important to understand what may exacerbate or alleviate the selective exposure bias. From our finding that selective exposure is very prevalent in a simple experiment, one might infer that this bias would persist in everyday settings. This study also has implications for experimental research. Experiments typically impose information on subjects, but the results might not generalize to a larger world in which people would never have gained that information.

Before we fully describe our experiment and its results, we will discuss related work. We will explain the design of the experiment in Section 4.2. Section 4.3 discusses the theoretical predictions of behavior in each treatment of the experiment, and Section 4.4 provides hypotheses based on these predictions. Section 4.5 details the experimental results, and Section 4.6 examines potential causes of these robust results. Finally, Section 4.7 summarizes our findings.

4.1.2 Related Literature

Studies of selective exposure and related biases generally fall into one of the following categories: theoretical models, empirical work, and experimental analysis, which focuses on either the choice of information or the interpretation of information. We will discuss relevant work in each category, with an emphasis on experimental findings.

Several theoretical models show that biased choices of information can occur when agents are allowed to pick their beliefs or when pre-existing beliefs are included in a utility function. (See, for example, Yariv, 2005; Eyster, 2002; Eliaz and Spiegler, 2005; Caplin and Leahy, 2004.) Selective exposure leads to media bias in a model by Mullainathan and Shleifer (2005), in which readers have preferences for biased newspapers. In contrast, in a belief-based model (Gentzkow and Shapiro, 2006), media bias stems from Bayesian agents, who judge a source to be of higher quality if it agrees with their prior beliefs. Interestingly, the former model predicts that

increased competition will result in greater media bias, but the latter predicts the opposite effect.

Empirical studies of media bias demonstrate that the media can persuade people to vote for a party they did not initially support. For example, DellaVigna and Kaplan (forthcoming) show that the introduction of the conservative Fox News channel shifted votes to the Republican party in Senate and Presidential races. A free subscription to a liberal-leaning newspaper led to more democratic votes in a field study by Gerber, Karlan, and Bergan (2007). Such studies are not focused on selective exposure, but they imply that this bias can have significant effects, since people who do not selectively expose could change their opinions.

The empirical literature that focuses directly on selective exposure is also generally framed in the context of political campaigns, examining choices of information about a preferred candidate (for example, Chaffee and Miyo, 1983), or of friends who share one's political views (e.g., Huckfeldt and Sprague, 1988). In the realm of finance, Karlsson, Loewenstein, and Seppi (2005) find that people check the performance of their investments more often when a general market index has risen than when it has declined or remained the same.

Many experimental studies have examined confirmation bias, which is generally defined as the tendency to search for and interpret information that will confirm one's beliefs, rather than disconfirm them. (However, confirmation bias appears in a variety of contexts; see Klayman, 1995, for a review.) An oft-used tool for the study of confirmation bias is Wason's selection task (1966). In this task, a subject sees four two-sided cards and must choose which to turn over to test a rule of the form "If p, then q." The sides facing up serve as p, not-p, q, and not-q. The correct cards for testing the rule are p and not-q, but many subjects choose the p and q cards, since they could potentially confirm that the rule holds for those cards.

Subjects have been found to pay for uninformative cards in tasks similar to the Wason task (Jones and Sugden, 2001; Jones, 2003a and b). Setting the task in a market environment does not improve performance much (Budescu and Maciejovsky, 2005). Confirmation bias is also found in an experiment that allows subjects to choose

to read articles that are either in support of or against their opinion about insurance payments for alternative medicine (Jonas, Schulz-Hardt, Frey, and Thelen, 2001). This study shows that confirmation bias can be self-reinforcing, as subjects who choose articles sequentially exhibit a larger bias than those who choose the articles before reading any. Confirmation bias has also been identified when subjects do not choose the source of information themselves (Dave and Wolfe, 2004).

Our chapter is closely related to an experimental finding of Eliaz and Schotter (2006), who find that people are willing to pay for non-instrumental information to increase their confidence in a decision. Subjects in the baseline treatment of their experiment have to guess whether a selected state is A or B, knowing that the probability A will be chosen is always higher than the probability B will be chosen. They are told that the probability that A will be chosen is either high or low (for example, 1 or .51), and they can learn which probability would be used by paying a fee. Although this information is not useful for guessing the state, nearly 80% of subjects are willing to pay the fee under certain conditions.

Eliaz and Schotter test various models with respect to their data, including the disjunction effect (Tversky and Shafir, 1992). The disjunction effect describes a pattern in which a decision maker will choose the same option in two (or more) different states of the world, yet will pick a different option or pay to learn the state when it is unknown. Tversky and Shafir find that the disjunction effect disappears when decision makers first state which choice they would make in each potential state of the world. Eliaz and Schotter do not replicate this finding in their experiment, but we will consider the disjunction effect in our analysis.

Like Eliaz and Schotter, we find that subjects pay (in terms of expected value) for non-instrumental information. However, our focus is on preferences, not beliefs. Related work examines agents who selectively ignore information in order to justify making self-serving decisions (Dana, Weber, and Kuang, forthcoming; Chapters 2 and 3 of this thesis), or select out of a playing a dictator game, even at a cost, to avoid facing a potential recipient (Dana, Cain, and Dawes, 2006; Lazear, Malmendier, and Weber, 2006).

4.2 Experimental Design

In each round of our experiment, subjects make a binary choice decision about what they think the state of the world is. Each state—red and blue—has a 50 percent chance of being chosen in a given round. Before guessing the state, subjects can ask for information in the form of a red or blue indicator.¹ If a subject asks for a red indicator, and the state is red, she is equally likely to receive a null signal and one that tells her the state is red. If she asks for a red indicator, but the state is blue, she will always receive a null signal. The blue indicator works analogously. In this setup, a null signal indicates that there is a 1/3 chance that the actual state matches the color of the indicator, and a 2/3 chance that it is the opposite state.

We induce preferences for one state over another by providing a bonus if a subject correctly guesses a particular state. In some rounds, guessing red correctly is worth more than guessing blue correctly, and vice versa. Specifically, in each round, one of four payoff pairs is displayed for the subject, indicating the number of points the subject would receive for a correct answer. These pairs are: (150, 50), (50, 150), (80, 60), and (60, 80), where the first element denotes the payoff for correctly guessing red, and the second denotes the payoff for correctly guessing blue. Each payoff pair is equally likely to be chosen in a period and independent of the choice in other rounds. An incorrect guess is always worth 10 points.

For an individual trying to guess the state, the indicator for the state with lower potential payoffs always has a higher expected value than the indicator for the state with higher potential payoffs.² That is, if a subject could earn more by correctly guessing red than correctly guessing blue, she should ask for the blue indicator. To understand why this is the case, assume that the bonus for correctly guessing red is large, so the subject naturally prefers to guess red. By asking for a blue indicator, she may learn that blue is the correct state and be able to guess that. If she gets a

¹For the sake of differentiating terms, we will use the word “indicator” to represent a source of information, and “signal” for the actual information received from the indicator.

²This is true for any payoff levels when the states are equally likely, as long as the payoff for one state is strictly higher than that for the other and the payoff for correct guesses is higher than the payoff for incorrect ones.

null signal, she will know that red is the more likely state and it will make sense for her to guess her preferred choice. The red indicator, in contrast, is uninformative. Since the expected value of guessing red is higher than switching to blue upon receipt of a null signal, the subject will pick red no matter what signal is received. Similar logic applies for low bonuses, which we will discuss in the next section.

We examine the choice of indicators and states in four settings, using an individual choice treatment, two majority rule voting treatments, and a partners treatment. These treatments are described below. In all treatments, subjects played 20 periods. Points were worth 1 cent each. To make the instructions easily understandable, we asked subjects to imagine that they were game show contestants, trying to guess which door a prize was hidden behind. Subjects chose an indicator by asking if the prize was behind the red door or asking if it was behind the blue door. A null signal was conveyed to the subjects as “no answer” while a red or blue signal was given as a “yes” response. At the end of each period, subjects received information about the actual state and their payoffs from that period and the rest of the experiment. A history panel displayed past payoff information throughout the experiment, along with past indicator choices, signals, and state choices.³

Individual Choice

The individual choice treatment is essentially what we describe above. Each subject is randomly assigned potential payoffs each round. She chooses a red or blue indicator, then receives a signal based on the indicator she had requested and guesses the state. Each subject’s guess about the state directly affects her earnings, according to the payoffs realized in a given period.

Majority Rule Voting

In the voting treatments, subjects’ earnings depend on a group decision. Subjects are divided into groups of five, which remain constant throughout a session. After asking for an indicator and receiving a signal, each subject submits a guess of the

³The experiment was programmed and conducted with z-Tree software (Fischbacher, 2007). Instructions are available at www.hss.caltech.edu/~lauren/expts/

correct state. The state with the majority of votes constitutes the group decision. Subjects are paid based on whether the group decision is correct. At the end of each period, subjects can see the number of votes for the red and blue state in addition to the payoff information. We conducted two types of voting sessions, one with independent preferences and one with common preferences. In the first treatment, which we will refer to as *voting with individual payoffs*, payoff pairs are selected for each individual subject, so different members of a group may earn different amounts. In the treatment that we will refer to as *voting with common payoffs*, the payoff pairs are the same for each member of a group.

Partners

The partners treatment is designed to emphasize the importance of the indicator choice. It divides the experimental task amongst two players with shared payoffs. One subject decides which indicator to ask for, and her partner sees the signal choice and receives the signal before guessing the state. Each player in the pair is paid based on this guess. Partners are matched at random each round, with one constraint. Half of the subjects act as the first player for the first 10 rounds and the second player for the other 10 rounds, while the roles are reversed for the other subjects.

Table 4.1 summarizes the number of subjects and average payoff in each treatment. All sessions were conducted at Caltech, using only subjects who had not previously participated in the experiment.

Treatment	Number of Sessions	Number of Subjects	Average Payoff
Individual Choice	2	23 (11 and 12)	\$13.05
Voting, Individual Payoffs	2	30 (15 and 15)	\$14.52
Voting, Common Payoffs	2	20 (10 and 10)	\$15.13
Partners	1	18	\$12.60

Table 4.1: Summary of Sessions. Average Payoff excludes \$5 show-up fee.

4.3 Theoretical Predictions

4.3.1 Individual Choice

In this section and the next, we will provide the theoretical prediction of behavior, along with the intuition behind it. The appendix contains a full analysis of the dominant strategy in the individual choice treatment and the equilibria in weakly undominated strategies in the voting treatments.

For use in the predictions, we say that an agent *follows her signal* if, when asking for the red (blue) indicator, she guesses red (blue) when a red (blue) signal is received and blue (red) when a null signal is received.

Proposition 1. *In the individual choice treatment, the dominant strategy for an agent is to choose the red (blue) indicator when her a priori preferences favor blue (red), and to follow her signal.*

Recall that a subject has a .5 chance of receiving a red or blue signal (and therefore, knowing the state) if her indicator choice matches the color of the true state. If her indicator choice does not match the state, she will receive a null signal with probability 1. Therefore, a null signal means that there is a 1/3 chance that the realized state matches the color of the indicator and a 2/3 chance that the true state is the opposite color of the chosen indicator.

Assume an agent prefers the red state, and payoffs are (150, 50). Consider an agent who chooses the red indicator. Upon receiving a red signal, she should clearly guess the red state. If she receives a null signal, the red state is less likely than blue, but the expected value of guessing red is higher than the expected value of guessing blue. So she will always guess red. This guess will be correct whenever the state is red, and incorrect whenever the state is blue.

Now consider an agent who chooses the blue indicator. If she receives a blue signal, she would guess blue. If she receives a null signal, the expected value of guessing red is much higher than that for guessing blue, so she would guess red. When the state is blue, there is a .5 chance that she will receive a null signal and guess the wrong state.

Her guess will be correct whenever the state is red, and correct with a .5 chance when the state is blue. Therefore, an agent who picks the blue indicator will guess correctly more often than one who chooses the red indicator, and will earn higher payoffs, in expectation.

When payoffs are (80, 60), an agent who chooses the red indicator will again guess red if she receives a red signal. If she receives a null signal, she has a $2/3$ chance of getting 60 points by guessing blue and a $1/3$ chance of getting 80 points by guessing red (ignoring payoffs for incorrect guesses), so she will guess blue. Now she will be correct whenever the state is blue, and correct with a .5 chance when the state is red.

An agent who chooses the blue indicator will guess blue if she receives a blue signal. She will guess red if she receives a null signal, since that gives her a $2/3$ chance of getting 80 points and a $1/3$ chance of getting 60 points. She will be correct whenever the state is red, and correct with a .5 chance when the state is blue. Notice that although the agent who chooses the red indicator correctly guesses the state as often as the one who chooses the blue indicator, her correct guesses are usually for the lower-paying state. The blue indicator is again the more valuable source of information.

4.3.2 Majority Rule Voting

The voting treatments introduce an externality in individuals' decisions, since each subject's choice of information and guess about the correct state affects the other members of her group. The predictions for these treatments are therefore equilibrium predictions of behavior. We find that the treatment with common payoffs supports selective exposure, but the equilibrium strategy for the treatment with individual payoffs does not.

In the treatment with potential payoffs chosen separately for each individual:

Proposition 2. *The unique symmetric equilibrium in weakly undominated strategies entails each agent choosing the red (blue) indicator when her a priori preferences favor blue (red), regardless of the size of the bonus, and following her signal.*

The basic intuition behind this proposition is as follows: If four out of the five agents in a group follow this strategy, the state is equally likely to be red or blue if the fifth agent is pivotal, due to the symmetry of the setup. The fifth agent should then follow this strategy, just like she should in the individual choice treatment. Choosing an indicator and guessing a particular state regardless of the signal is weakly dominated, as is mixing after receiving a null signal. We find that there is no equilibrium in which some agents choose the indicator matching their preferred state.

When potential payoffs are the same for each member of a group, it is again the case that in equilibrium, an agent will follow her signal. However, when payoffs are shared, selective exposure is part of an equilibrium.

In the treatment with common payoffs:

Proposition 3. *A symmetric equilibrium in weakly undominated strategies entails mixing between the red and blue indicators, and following the signal. The probability of choosing the red indicator depends on the size of the bonus.*

Assume all agents have a bias for the red state. In equilibrium, it cannot be the case that all agents choose the blue indicator and follow the signal. Being pivotal in this case would mean that two agents had received blue signals, so blue must be the correct state. Then choosing an indicator and following the signal would not be a best response. Similarly, all agents choosing the red indicator and following the signal is not an equilibrium. It must be the case that agents mix between choosing the red and blue indicators.

By finding the conditional probability of the red state that makes an agent indifferent between picking the red and blue indicators, we can find the probability that each agent places on picking the red indicator, i.e., the probability of selective exposure. For the payoff pair of (150, 50), in equilibrium, each agent should pick the red indicator (hence, selectively expose) with probability 0.286. For the payoff pair of (80, 60), this probability increases to 0.437.⁴ We cannot expect subjects to use these exact probabilities in our setup, but we can determine whether comparatively more

⁴This probability approaches 0.5 as the number of subjects goes to infinity.

selective exposure occurs at the lower bonus level.

There are several behavioral factors that may also play in role in the voting treatments. Collective choice might cause subjects to think more about the proper actions to take, since a mistake could result in a low payoff not only for an individual, but also for her group. Alternatively, subjects might not take the time to carefully consider their decisions, since responsibility is diffused in this setting. A subject's guess only matters when she is pivotal, which reduces the incentive to think through the experimental task. Subjects may free ride in hopes that other group members will collect the proper information to make the best choice.⁵

4.3.3 Partners

In the partners treatment, the two members of a team share payoffs, so the dominant strategy is the same as in the individual choice treatment. When the potential payoff for the red state is larger than for blue, first players should ask for the blue indicator. Second players should then guess the red state if they receive a null signal, and the blue state if they get a blue signal.

Since the only decision first players make is which type of indicator to send to their partners, we expect that they will consider the consequences of their choice more than in the individual choice treatment. This should lead to less selective exposure.

We also expect subjects who spent the first half of the experiment reacting to signals and guessing the state to better understand the importance of the indicator choice. In order to make their guesses, these subjects should have thought about the signal they received as well as the type of indicator their partner requested. They should be able to carry over their experience in processing the signals to choosing an indicator in the second half of the experiment. A somewhat loose interpretation of Tversky and Shafir's disjunction effect (1992) would also lead to this prediction. Tversky and Shafir find that the disjunction effect does not occur when subjects think through a decision tree. In the partners treatment, we are essentially placing half of

⁵The dilution of responsibility can be a powerful effect, even leading to the failure to report an emergency. (See, for example, Darley and Latané, 1968.)

the subjects at a later node in the decision tree for the first part of the experiment. When the roles switch, these subjects should be able to make better choices at the earlier decision node.

4.4 Hypotheses

Based on the theoretical predictions and expectations for behavior described above, we can create a series of hypotheses. If agents facing the experimental task maximize payoffs, then:

Hypothesis 1. (Rational choice of information) No selective exposure will occur, except in the majority rule voting treatment with common payoffs.

Hypothesis 2. (Rational use of information) Upon receiving a signal, subjects will guess the state with the highest expected value.

Alternatively, subjects may exhibit the selective exposure bias, choosing indicators that will not maximize payoffs. We expect subjects to be more prone to the selective exposure bias if their guess about the state is less likely to influence payoffs, and less prone to the bias if they focus solely on the decision of information in the task. In the voting treatment with common potential payoffs, subjects should exhibit selective exposure, as part of an optimal mixing strategy. We therefore propose the following hypotheses:

Hypothesis 3. (Dilution of responsibility) There will be more selective exposure in the voting treatments than in the individual choice treatment.

Hypothesis 4. (Strategic Voting) There will be more selective exposure in the voting treatment with common payoffs than in the voting treatment with individual payoffs. Within the voting treatment with common payoffs, there will be more selective exposure at bonus levels of 20 than at bonus levels of 100.

Hypothesis 5. (Conjunction) There will be more selective exposure in the individual choice treatment than in the partners treatment.

Our final hypothesis can be interpreted as a test of the disjunction effect, or simply learning, in the partners treatment.

Hypothesis 6. (Learning) There will be more selective exposure in the first half of the partners treatment than in the second half.

4.5 Results

4.5.1 Overall Performance

Before presenting the results, it will be useful to offer a few definitions. We will say that a subject chose the *selective exposure indicator* if she chose the indicator that matches the state with higher potential payoffs. We will refer to the other indicator as payoff maximizing. For example, if correctly guessing the red state pays more than correctly guessing blue, the selective exposure indicator is the red indicator, and the payoff maximizing indicator is blue. The amount of selective exposure in a treatment will be defined as the proportion of choices of the selective exposure indicator out of the total number of choices.

In considering the choices subjects make after receiving a signal, we will define a *best guess*, or *payoff maximizing guess*, as a guess for the state with the highest expected payoff, given the information requested and received. Assume again that the potential payoff for the red state is higher than for blue. When subjects received a red or blue signal, they could be certain of the correct state. The best guess in these cases was simply that state. When subjects received a null signal, the best guess depended on which indicator they had requested. Subjects who had asked for the blue indicator should always guess the red state. For subjects who asked for the red indicator, making the best guess was a bit more complicated. When the potential payoffs were (80, 60), subjects should switch to guessing the blue state after receiving

a null signal. However, for potential payoffs of (150, 50), guessing red would maximize expected value.⁶

We can now discuss the levels of selective exposure and best guesses across treatments. We will first provide a broad overview, then we will discuss individual treatments and learning over time.

Treatment	Selective Exposure
Individual Choice	53.7%
Voting, Individual Payoffs	58.3%
Voting, Common Payoffs	46.5%
Partners	45.6%

Table 4.2: Proportion of selective exposure choices across treatments

Table 4.2 shows the total amount of selective exposure in each treatment. We can clearly reject the hypothesis that selective exposure would only occur in the voting treatment with common payoffs. The amount of selective exposure was significantly greater than 0 in all treatments.⁷

Selective exposure could be costly. In the individual choice treatment, for example, subjects who always chose the payoff maximizing indicator earned \$14.63 on average, whereas those who never did earned only \$11.91—nearly 25% less. Part of the reason selective exposure reduced payoffs was that it was associated with poor guesses.

In all cases in which subjects received a red or blue signal, they guessed the correct state—a good indication that they had understood the instructions. The more interesting cases occurred when a subject received a null signal. Table 4.3 shows the percentage of payoff maximizing guesses after a null signal, by subjects

⁶In the voting treatment with common payoffs, the equilibrium condition requires subjects to *follow their signals*, which would mean that someone who had asked for the red indicator should guess the blue state after a null signal, even for payoffs of (150, 50). To prevent confusion, and because the data do not indicate that the equilibrium condition is being adhered to, we will define best guesses for this treatment in the same way as for the others.

⁷In fact, subjects chose the selective exposure indicator more than random chance would predict in the voting treatment with individual payoffs ($t_{599} = 4.137$, $p < .001$). This same result was marginally significant for the individual choice treatment ($t_{459} = 1.588$, one-tailed $p = 0.057$). The partners treatment was the only one in which subjects performed significantly better than chance ($t_{359} = -1.691$, one-tailed $p = 0.046$), but in the voting treatment with common payoffs, the mean was below .5 with marginal significance ($t_{399} = -1.402$, one-tailed $p = 0.081$).

who exhibited the selective exposure bias and those who did not. The first thing to notice in this table is that for the most part, subjects were able to make the payoff-maximizing guess after receiving a null signal, so we find support for Hypothesis 2. In all treatments, subjects guessed the payoff-maximizing state more often than chance would predict, even if they had chosen the indicator poorly. (For all treatments, $p < .001$ in t-tests comparing best guesses after a null signal to 0.5. The same result holds when looking only at subjects who chose the selective exposure indicator.)

The second point to notice in the table is that across all treatments, those who had chosen the payoff maximizing indicator were more likely to make the best guess ($p < .01$ for χ^2 tests in all treatments). It is not surprising that those who made better indicator choices also made better guesses. The subjects who knew to ask for the blue indicator when red was the preferred state were likely to have considered the expected value of each indicator choice, and would know what strategy to follow after receiving a signal. Also, this strategy was simpler than the one for people who asked for the red indicator. Those who asked for the blue indicator simply had to guess the red state every time they got a null signal. But for those who chose the red indicator, the proper guess to make after a null signal depended on the size of the payoffs.

Focusing on the choices made after selective exposure, we find that in the individual choice treatment and voting treatment with individual payoffs, poor guesses were made significantly more often when the bonus was high than when it was low ($\chi^2 = 33.90$, $p < .001$ for the individual choice treatment; $\chi^2 = 55.21$, $p < .001$ in the voting treatment). The proportion of best guesses did not significantly differ between the bonus levels in the partners or voting with common payoffs treatments. Many subjects were switching their guess to the state they had not asked about even at the high bonus level, when simply guessing their preferred color would have been payoff-maximizing. This indicates that subjects realized that the most likely state following a null signal was the opposite color of the indicator they had requested. Yet they were unable to take this reasoning back into the first stage and choose the payoff-maximizing indicator.

Individual Choice			
Indicator Choice	80, 60 Payoffs	150, 50 Payoffs	Total
Selective Exposure	88.9% (104/117)	51.9% (42/81)	73.7% (146/198)
Payoff Maximizing	91.8% (56/61)	97.6% (81/83)	95.1% (137/144)
Payoff maximizing guesses across indicator choice and bonus level: 82.7%			
Voting, Individual Payoffs			
Indicator Choice	80, 60 Payoffs	150, 50 Payoffs	Total
Selective Exposure	85.1% (126/148)	43.0% (58/135)	65.0% (184/283)
Payoff Maximizing	91.7% (77/84)	95.7% (88/92)	93.8% (165/176)
Payoff maximizing guesses across indicator choice and bonus level: 76.0%			
Voting, Common Payoffs			
Indicator Choice	80, 60 Payoffs	150, 50 Payoffs	Total
Selective Exposure	79.7% (51/64)	69.4% (50/72)	74.3% (101/136)
Payoff Maximizing	91.0% (71/78)	97.6% (82/84)	94.4% (153/162)
Payoff maximizing guesses across indicator choice and bonus level: 89.0%			
Partners			
Indicator Choice	80, 60 Payoffs	150, 50 Payoffs	Total
Selective Exposure	76.5% (52/68)	87.5% (56/64)	81.8% (108/132)
Payoff Maximizing	89.2% (66/74)	100.0% (76/76)	94.7% (142/150)
Payoff maximizing guesses across indicator choice and bonus level: 88.7%			

Table 4.3: Frequency of payoff-maximizing guesses following null signals, by treatment, indicator choice, and payoffs

4.5.2 Voting Treatments

Contrary to our hypothesis, there was no difference in the amount of selective exposure in the voting and individual choice treatments. Across both voting treatments, selective exposure occurred 53.6% of the time, versus 53.7% for individual choice ($t_{1458} = 0.034$).

The proportion of subjects who made the best guess following a null signal was also no different in the voting treatments than in the individual choice treatment ($t_{1097} = 1.200$). On average, subjects in the voting treatments made payoff-maximizing guesses 79.7% of the time they received null signals, compared to 82.7% in the individual choice treatment.

Since selective exposure is justified in the voting treatment with common payoffs amongst group members, subjects in that treatment should have exhibited more selective exposure than subjects in the treatment with individually-selected payoffs. We find the opposite effect. ($t_{998} = 3.697$, $p < 0.01$ at the individual level; $t_{198} = 4.092$, $p < 0.01$ at the group level).

In the treatment with common payoffs, theory predicts that more selective exposure should occur when there is a small difference between potential payoffs (bonus size of 20) than when there is a larger one (bonus size of 100). We find no significant difference in selective exposure rates for the two bonus sizes ($t_{398} = 0.872$ at the individual level; $t_{78} = 1.260$ at the group level). The rate is 44.2% for bonuses of 20 and a slightly larger 48.6% for bonuses of 100. At the bonus level of 20, the rate is actually quite close to the equilibrium prediction of 43.7% ($t_{37} = 0.224$ at the group level), but subjects are clearly not following equilibrium strategies. At the higher bonus level, the rate is well above the predicted 28.6% ($t_{41} = 7.791$, $p < .01$). Hypothesis 4 is not supported by the data.

4.5.3 Partners

Dividing the choice of signal and guess about the state in the partners treatment had the expected effect of reducing selective exposure. Recall that the selective exposure

rate in the individual choice treatment was 53.7%; in the partners treatment, this dropped to 45.6% ($t_{638} = 1.855$, one-tailed $p < .05$). Although selective exposure rates dipped below 50%, the bias was still quite prevalent.

The divided decision in this treatment allows us to consider the reactions to signals separately from the choice of indicator. We have seen that subjects often do not properly react to signals in the individual choice treatment, and we can examine whether guesses about the state improve when subjects are detached from the initial indicator choice. The subjects who were only responsible for guessing the state in the partners treatment made better guesses than the subjects who chose indicators and guessed the state in the individual choice treatment, as Table 4.3 shows. This difference is marginally significant ($t_{483} = 1.630$, one-tailed $p = 0.052$).

We expected that subjects who had based their actions upon received signals would be better able to choose indicators than their inexperienced counterparts. Surprisingly, though, subjects who made indicator choices in the second half of the experiment made more biased decisions than the subjects who chose indicators in the first half. This difference was quite large, with a selective exposure rate of 35.6% in the first half and 55.6% in the second half ($t_{178} = 2.734$, $p < .001$). If selective exposure in our experiment could be explained by something akin to the disjunction effect, then we do not find any support for a disappearance of the effect. Despite the higher amount of selective exposure in the second half of the experiment, the percentage of best guesses following a null signal did not decline much, dropping from 91.2% to 86.3% ($t_{139} = 0.908$). Experience in choosing an indicator might have helped subjects guess the state in the second half, even after their partners chose a poor signal source.

4.5.4 Learning

Figure 4.1 shows the average amount of selective exposure and best guesses over time in each treatment. As the figure shows, selective exposure only decreased significantly over time in the voting treatment with individual payoffs. In this treatment, selective exposure decreased by 1.84% per period, based on a probit regression of selective




Figure 4.1: Average amount of selective exposure and payoff maximizing guesses, by period

exposure on period alone ($z = -5.16$, $p < .001$).⁸ Similarly, guesses following a null signal improved over time in the voting treatment with individual payoffs ($z = 2.91$, $p < .01$ in a probit of best guesses on period), but not in the other treatments.

We might think that subjects in the individual-payoff voting treatment learned from the actions of their group members over time. However, there is no evidence of learning based on previous group decisions in this treatment. We regressed selective exposure on a number of factors that might be expected to help an individual learn over time. These included one-period lags for whether the group correctly guessed the state, the number of votes for the correct state, whether the individual was pivotal (that is, whether the rest of the group submitted two votes for red and two votes for blue), whether the individual was correct, sums of previous results, and

⁸No other effects were significant. In the individual choice treatment, $z = -1.09$, $p = 0.274$. In the voting treatment with common payoffs, $z = -0.35$, $p = 0.728$. In the partners treatment, $z = -0.15$, $p = 0.878$ for the first half, and $z = -1.48$, $p = 0.140$ in the second half.

various interactions. None of these variables were significant. The only variable that consistently had a significant effect was the period.




Figure 4.2: Frequency of subjects' indicator choice patterns. Stripes indicate that subject followed pattern with one or two exceptions.

We find that the data seem to be driven by subject types rather than learning over time. Figure 4.2 shows the number of subjects who always (or almost always) chose the payoff maximizing indicator or the selective exposure indicator. Nearly half of the subjects consistently chose the correct or incorrect indicator throughout the experiment, and 59% chose consistently if we include those with one or two exceptions. Always picking the selective exposure indicator was the modal response in all treatments but the partners treatment.

To look for signs of learning in the experiment, we will focus only on the subjects in the “Other” category of Figure 4.2—those who made different indicator choices throughout the experiment. We will consider a subject to have learned to pick the payoff maximizing indicator by the end of the experiment if she picks the correct indicator in at least the last two periods, and her written response to a debriefing

question does not indicate that she chose indicators at random.⁹ Figure 4.3 shows the proportion of subjects in the “Other” category who learned to choose the payoff maximizing indicator. On the whole, 17 out of 37 subjects (45.9%) showed signs of learning. Most seemed to suddenly recognize the optimal strategy for indicator choices, instead of testing that strategy along with the selective exposure one for several rounds.




Figure 4.3: Subjects who learned to choose the payoff-maximizing indicator

As discussed before, factors that we expected to matter in the voting treatment did not have any effect on learning. For the individual choice and partner data, we thought subjects might respond to the difference in the expected payoff between the (150, 50) and (80, 60) payoff levels. Asking for the correct indicator is more valuable when the bonus for guessing a certain state correctly is higher, so selective exposure might be reduced when payoffs were 150 and 50. For the individual choice treatment,

⁹It might be the case that the subjects who chose the selective exposure indicator with two exceptions learned the optimal strategy two rounds before the end of the experiment. We checked the data to make sure that we were not excluding any subjects with this pattern of behavior from the current analysis.

we ran a fixed effects model of selective exposure on the expected payoff difference, period, and the individuals who chose the correct indicator between 3 and 17 times. This revealed a significant effect of the period ($z = -2.38$, $p < .05$), but not the payoff difference ($z = 0.50$, $p = 0.617$). In the partners treatment, we ran the regression on subjects who chose the correct indicator between 3 and 7 times, since each subject only acted as a first-mover in 10 periods. In this treatment, neither the period nor the payoff difference were significant ($z = -1.23$, $p = 0.220$ for period; $z = 0.54$, $p = 0.590$ for payoff difference). In each treatment, similar effects were found when we excluded only the subjects who always chose the indicator correctly or incorrectly.¹⁰ We view this failure to respond to incentives as further evidence that subjects are divided into types.

4.6 Potential Causes of Selective Exposure

Since the prevalence of selective exposure was robust across the institutions we tested, we explored two potential explanations for the behavior in our experiments. We tested whether the high levels of selective exposure might have been exacerbated by small differences in expected payoffs for each indicator choice. We also tested whether the failure to pick a correct indicator was caused by a failure to backward induct.

4.6.1 Small Payoff Differences

We conducted a high-stakes experiment to see whether selective exposure would disappear if decisions were worth more. This experiment was conducted like the individual choice sessions, with the only exception being the method of payment. At the end of the experiment, five rounds were chosen at random for each subject, and subjects were paid \$0.20 for each point earned in those rounds. Seven subjects took part in this session. Average earnings were \$77.57.

¹⁰In the individual choice treatment: for period, $z = -3.37$, $p < .01$; for payoff difference, $z = -1.67$, $p = 0.095$. In the partners treatment: for period, $z = -0.83$, $p = 0.408$; for payoff difference, $z = 0.47$, $p = 0.640$.

In this session, 35.7% of indicator choices exhibited selective exposure. Although selective exposure occurs less frequently than random chance would predict ($t_{139} = 3.515$, $p < .001$), it is still significantly above 0 ($t_{139} = 8.788$, $p < .001$).

As in the other treatments, the rate of selective exposure was largely determined by subject types. One subject always chose the selective exposure indicator, but two subjects never did. A fourth subject quickly learned to pick the payoff-maximizing indicator, and another seemed to learn toward the end of the session. Two other subjects could not be described as making consistent decisions or learning over time.

4.6.2 Failure to Backward Induct

We have seen that subjects are fairly good at figuring out the most likely state after receiving a signal, whether or not they asked for the payoff-maximizing indicator. Many seem unable to take their future actions into account before picking an indicator. For example, when a subject faces a large bonus, choosing the indicator that matches her preferred state should not affect her actions and is less useful than picking the opposite indicator. Yet when bonuses are large in the individual choice treatment, the selective exposure indicator is chosen 49% of the time.

To test whether the selective exposure bias in our treatments was caused solely by a failure to backward induct, we conducted a simplified version of the individual choice treatment. In this treatment, subjects always had a preference for the red state, with payoffs for red and blue selected to be (150, 50) or (80, 60). The blue indicator was the same as in the standard treatments, but the red indicator was always uninformative. If a subject asked for the red indicator, she would receive a red signal with .5 likelihood and a null signal with .5 likelihood, regardless of the state.

Unlike our previous treatments, subjects with a preference for the red state would now gain nothing from asking for the red indicator, since it cannot even serve to confirm that the state of the world is the one that is preferred. If we found that many subjects in this treatment were still selecting the red indicator, then we might

suspect that subjects in our other treatments were unable to fully consider the effect of the indicator choice, or they might simply be matching the color of their selected indicator to their preferred state.

We found that subjects were generally able to recognize the uninformative indicator, choosing it in only 11.4% of all rounds. Subjects who requested the correct indicator made the payoff maximizing choice 93.5% of the time they received a null signal, comparable to the results from the other treatments. Of the eleven subjects in this session, six always asked for the correct indicator, and two other subjects did so with one and three exceptions. One of the remaining three subjects learned over time, but the other two chose indicators at random or based on flawed reasoning involving the payoffs.

4.7 Conclusion

We find that selective exposure to information about a preferred state holds fairly consistently across a variety of experimental contexts. The amount of selective exposure found in our standard treatments only ranged from around 46% to around 58%. While the opportunity to earn more money seemed to improve subjects' ability to avoid this bias, it did not remove the bias altogether. Based on the diffusion of responsibility and equilibrium predictions in the majority rule voting settings, we expected more selective exposure than in the individual choice treatment. However, we found no difference between these two contexts. Allowing subjects to focus only on choosing an indicator reduced the amount of selective exposure in the partners treatment, but this reduction was driven by the subjects who chose indicators in the first half of the experiment. Those who chose indicators after gaining experience in reacting to signals were more prone to the selective exposure bias.

This result was surprising since we would expect subjects who thought through the second stage of the task to be able to perform better in the first stage. However, it appears that across treatments, most subjects understood which state would be most likely upon receiving a null signal. In the individual choice treatment and vot-

ing treatment with individual payoffs, the bulk of the mistakes in guessing the state occurred when subjects who had chosen the selective exposure indicator guessed the most likely—but not payoff maximizing—state. For the most part, subjects understood the proper action after receiving a signal, yet they were not able to make the logical jump to determine the best source of a signal. This failure to backward induct appears to be caused by the ability to gain confirmation of a preferred state in our experiments. In a similar task that did not offer a confirmatory signal, subjects nearly always recognized an uninformative indicator.

Sizable amounts of selective exposure persisted in our simple experiments, across a range of treatments and even when stakes were high. Over one-quarter of the participants in our main treatments always chose the selective exposure indicator, despite the fact that they had nothing to do but think about the experimental task for twenty consecutive periods. Some people seemed more prone to the bias than others, and we might expect to find different types of agents outside of the laboratory. If selective exposure were to manifest itself in the way our results indicate, then agents might make good use of information they have received, but fail to first choose the best source of information. This could lead to a failure to maximize value when voting for a political candidate or making a purchase, and the prevalence of this bias may mean that certain experiments would not generalize. The selective exposure bias could also have more dangerous effects; for example, a doctor might only order tests that could confirm her diagnosis of a patient, and a patient might only see a doctor that will tell her she seems to be fine.

4.8 Appendix 1: Dominant Strategy in the Individual Choice Treatment

We say that an agent *follows her signal* if, when choosing the red (blue) indicator, she guesses red (blue) when a red (blue) signal is received and blue (red) when a null signal is received.

Proposition 1. *In the individual choice treatment, the dominant strategy is for an agent to choose the red (blue) indicator when her a priori preferences favor blue (red), and to follow her signal.*

Proof of Proposition 1: Recall that a null signal means that there is a $2/3$ chance that the true state is the opposite color of the indicator a subject asked for. Assume that a subject has a preference for red; that is, correctly guessing the red state pays more than correctly guessing the blue state.

Table 4.4 shows the expected value (EV) of guessing the red or blue state when correctly guessing red pays 150 points and correctly guessing blue pays 50. The expected values are based on the indicator a subject has chosen and the signal she has received. Note that when a subject has requested the red indicator, the expected value of guessing red is higher than that for guessing blue, regardless of the signal. The subject would guess red for any signal. In contrast, when a subject requests the blue indicator, she would guess blue if a blue signal is received, and red otherwise.

Signal Received	Red Indicator		Blue Indicator	
	EV(Red Guess)	EV(Blue Guess)	EV(Red Guess)	EV(Blue Guess)
Red	150	10	-	-
Blue	-	-	10	50
Null	$\frac{1}{3}150 + \frac{2}{3}10$	$\frac{2}{3}50 + \frac{1}{3}10$	$\frac{2}{3}150 + \frac{1}{3}10$	$\frac{1}{3}50 + \frac{2}{3}10$

Table 4.4: Expected values of red and blue guesses for payoffs of (150, 50)

We can compare the expected value of each indicator based on these strategies, using the .5 probability that the red state is drawn.

$$\text{EV}(\text{red indicator}) = \frac{1}{2}150 + \frac{1}{2}10 = 80$$

$$\text{EV}(\text{blue indicator}) = \frac{1}{2}150 + \frac{1}{2}\left(\frac{1}{2}50 + \frac{1}{2}10\right) = 90$$

Therefore, the blue indicator is more valuable than the red one.

Similar logic applies for the payoff pair of (80, 60). Table 4.5 compares the expected value of red and blue guesses for these payoffs. If a subject requests the blue indicator, she should follow the same strategy as when the payoffs are (150, 50). However, with the smaller bonus for guessing red correctly, a subject who requested the red indicator should guess blue after receiving a null signal.

Signal Received	Red Indicator		Blue Indicator	
	EV(Red Guess)	EV(Blue Guess)	EV(Red Guess)	EV(Blue Guess)
Red	80	10	-	-
Blue	-	-	10	60
Null	$\frac{1}{3}80 + \frac{2}{3}10$	$\frac{2}{3}60 + \frac{1}{3}10$	$\frac{2}{3}80 + \frac{1}{3}10$	$\frac{1}{3}60 + \frac{2}{3}10$

Table 4.5: Expected values of red and blue guesses for payoffs of (80, 60)

We again find that the blue indicator is more valuable than the red one:

$$\text{EV}(\text{red indicator}) = \frac{1}{2}\left(\frac{1}{2}80 + \frac{1}{2}10\right) + \frac{1}{2}60 = 210/4 = 52.5$$

$$\text{EV}(\text{blue indicator}) = \frac{1}{2}80 + \frac{1}{2}\left(\frac{1}{2}60 + \frac{1}{2}10\right) = 57.5$$

4.9 Appendix 2: Equilibria in Majority Rule Voting Treatments

We consider symmetric equilibria in weakly dominated actions.

We first consider the voting treatment in which each individual is assigned payoffs independently of the other members of her group.

Voting with Individual Payoffs

Observation 1 *Each agent choosing the red (blue) indicator when her a priori preferences favor blue (red), regardless of the size of the bonus, and following her signal, constitutes an equilibrium.*

Proof of Observation 1: From the symmetry of the setup, if 4 of the 5 agents behave as in the observation, then

$$\Pr(\text{Red} \mid 2 \text{ red votes}, 2 \text{ blue votes}) = \frac{1}{2}.$$

In particular, the 5th agent behaving according to the suggested profile is a best response according to the same considerations of the individual treatment.

Observation 2 *Choosing an indicator and guessing a particular state regardless of the revealed signal is weakly dominated.*

Proof of Observation 2: Choosing the red indicator and guessing blue regardless of the signal is dominated by choosing the red indicator and following the signal. Choosing the red indicator and guessing red regardless of the signal is dominated by choosing the blue indicator and following the signal. The analogous arguments follow for the corresponding strategies when the chosen indicator is blue.

Observation 3 *In any symmetric equilibrium in weakly undominated strategies agents do not mix after observing the realization of their signal.*

Proof of Observation 3: From observation 2 above, each agent has a positive probability of being pivotal. In particular, when completely informed, the agent must follow that information in equilibrium. Thus, mixing can potentially occur only upon the observation of a null signal. Note that for any prior p that the state is Red,

$$\Pr(\text{Red} \mid \text{Red indicator, null signal}) = \frac{\frac{1}{2}p}{\frac{1}{2}p + (1-p)} = \frac{p}{2-p},$$

while

$$\Pr(\text{Red} \mid \text{Blue indicator, null signal}) = \frac{p}{\frac{1}{2}(1-p) + p} = \frac{2p}{1+p}.$$

For any $p < 1$, no matter what the probability of red is conditional on being pivotal, as long as it is lower than unity, it is strictly better to choose the blue indicator and follow the signal than choose the red indicator and choose always red, or randomize upon a null signal (which, in equilibrium, should generate the same expected value as always choosing red). A similar argument follows for a mix following the choice of a blue indicator.

Observation 4 *In equilibrium, it cannot be the case that both agents with an a priori preference for red, and agents with an a priori preference for blue put positive weight on the indicator matching their a priori preferred color (i.e., selectively exposing).*

Proof of Observation 4: Denote by p_b the prior of the appropriate guess being red when the payoff for red is $b \in \{50, 60, 80, 150\}$ that makes an agent indifferent between either indicator and following the signal. So,

$$p_{80} \left[\frac{1}{2}80 + \frac{1}{2}10 \right] + (1 - p_{80}) 60 = p_{80}80 + (1 - p_{80}) \left[\frac{1}{2}60 + \frac{1}{2}10 \right]$$

$$\Leftrightarrow p_{80} = \frac{25}{60} = \frac{5}{12},$$

while

$$p_{150} \left[\frac{1}{2}150 + \frac{1}{2}10 \right] + (1 - p_{150}) 50 = p_{150}150 + (1 - p_{150}) \left[\frac{1}{2}50 + \frac{1}{2}10 \right]$$

$$\Leftrightarrow p_{150} = \frac{20}{90} = \frac{2}{9}.$$

Analogously, $p_{60} = \frac{7}{12}$ and $p_{50} = \frac{7}{9}$. Thus, if p is the equilibrium posterior of red being the actual state conditional on pivotality, then for any agent who a priori prefers red, selective exposure can be part of an equilibrium only if $p \leq \frac{5}{12}$, while for an agent who a priori prefers blue, selective exposure can be part of an equilibrium if $p \geq \frac{7}{12}$. The claim then follows.

Observation 5 *There is no equilibrium in which some agents who a priori prefer red selectively expose.*

Proof of Observation 5: Assume that agents with an a priori preference for red selectively expose with some probability α . Then,

$$\Pr(\text{red vote} \mid \text{state is red}) = \frac{1}{4} \sum_{b=50,60,80,150} \Pr(\text{red vote} \mid \text{state is red}, b)$$

$$= \left(\frac{1}{2} + \frac{1}{2}\alpha \right) \frac{1}{2} + \left(\frac{1}{2} - \frac{1}{2}\alpha \right) = \frac{3 - \alpha}{4}$$

and similarly

$$\Pr(\text{blue vote} \mid \text{state is red}) = \frac{1 + \alpha}{4},$$

$$\Pr(\text{blue vote} \mid \text{state is blue}) = \left(\frac{1}{2} + \frac{1}{2}\alpha \right) + \left(\frac{1}{2} - \frac{1}{2}\alpha \right) \frac{1}{2} = \frac{3 + \alpha}{4},$$

$$\Pr(\text{red vote} \mid \text{state is blue}) = \frac{1 - \alpha}{4}.$$

Thus,

$$\Pr(\text{state is red} \mid \text{pivotal}) = \frac{\left(\frac{3-\alpha}{4}\right)^2 \left(\frac{1+\alpha}{4}\right)^2}{\left(\frac{3-\alpha}{4}\right)^2 \left(\frac{1+\alpha}{4}\right)^2 + \left(\frac{3+\alpha}{4}\right)^2 \left(\frac{1-\alpha}{4}\right)^2}$$

$$= \frac{(3 - \alpha)^2 (1 + \alpha)^2}{(3 - \alpha)^2 (1 + \alpha)^2 + (3 + \alpha)^2 (1 - \alpha)^2}.$$

Note that $\left(\frac{3-\alpha}{3+\alpha}\right)\left(\frac{1+\alpha}{1-\alpha}\right) > 1$, and so $\Pr(\text{state is red}|\text{pivotal}) > \frac{1}{2}$. From the calculations derived for Observation 4, this is in contradiction to agents best responding.

In particular, the above observations suggest the following corollary:

Corollary 1 *The unique symmetric equilibrium in weakly undominated strategies entails each agent choosing the red (blue) indicator when her a priori preferences favor blue (red), regardless of the size of the bonus, and following her signal.*

Voting with Common Payoffs

Assume that all agents have a strong bias for Red (so they get 150 if red is chosen and correct, and 50 if blue is chosen and correct).

As in the symmetric case, in equilibrium, if an agent chooses an indicator with positive probability, she must follow the generated signal.

Observation 6 *All agents choosing the blue indicator and following the signal does not constitute an equilibrium. Similarly, all agents choosing the red indicator and following the signal does not constitute an equilibrium.*

Proof of Observation 6: In either scenario, pivotality reveals the correct guess (blue for the former, red for the latter) and so choosing an indicator and following the signal is not a best response.

Therefore, a symmetric equilibrium in weakly undominated strategies entails a mix between the blue and red indicators. Assume that α is the probability each agent places on the red indicator.

As before, the conditional probability of Red that makes an agent indifferent between the two indicators is given by p_{150} and calculated according to:

$$p_{150} \left[\frac{1}{2}150 + \frac{1}{2}10 \right] + (1 - p_{150}) 50 = p_{150}150 + (1 - p_{150}) \left[\frac{1}{2}50 + \frac{1}{2}10 \right]$$

$$\Leftrightarrow p_{150} = \frac{20}{90} = \frac{2}{9}.$$

Thus, the indifference condition becomes:

$$\begin{aligned} \Pr(\text{Red} \mid \text{pivotal}) &= \frac{\left(\frac{\alpha}{2} + 1 - \alpha\right)^2 \left(\frac{\alpha}{2}\right)^2}{\left(\frac{\alpha}{2} + 1 - \alpha\right)^2 \left(\frac{\alpha}{2}\right)^2 + \left(\alpha + \frac{1-\alpha}{2}\right)^2 \left(\frac{1-\alpha}{2}\right)^2} = \frac{2}{9} \\ &\Leftrightarrow \frac{(2 - \alpha)^2 \alpha^2}{(2 - \alpha)^2 \alpha^2 + (1 + \alpha)^2 (1 - \alpha)^2} = \frac{2}{9} \end{aligned}$$

and the solution is $\alpha = 0.286$.

To contrast, suppose we considered a group of 5 agents with a weak red bias, then the (indicator) indifference probability would be $p_{80} = \frac{5}{12}$, and the indifference condition would translate into:

$$\frac{(2 - \alpha)^2 \alpha^2}{(2 - \alpha)^2 \alpha^2 + (1 + \alpha)^2 (1 - \alpha)^2} = \frac{5}{12},$$

the solution of which is $\alpha = 0.437$.

In general, suppose that there are $n = 2k + 1$ agents, and that bonuses are such that the (indicator) indifference probability is $p < \frac{1}{2}$. The equilibrium indifference condition is then:

$$\begin{aligned} \frac{(2 - \alpha)^k \alpha^k}{(2 - \alpha)^k \alpha^k + (1 + \alpha)^k (1 - \alpha)^k} &= p \\ \Leftrightarrow \frac{1 - \alpha^2}{(2 - \alpha) \alpha} &= \left(\frac{1 - p}{p}\right)^{1/k}. \end{aligned}$$

Note that $\left(\frac{1-p}{p}\right)^{1/k} \searrow 1$, and that $\frac{1-\alpha^2}{(2-\alpha)\alpha}$ is decreasing in α and $\frac{1-\alpha^2}{(2-\alpha)\alpha} = 1$ when $\alpha = \frac{1}{2}$.

In particular, as n increases, α approaches $\frac{1}{2}$.

Bibliography

- [1] Akerlof, George A., and William T. Dickens. (1982) The economic consequences of cognitive dissonance. *The American Economic Review*, 72(3), 307–319.
- [2] Andreoni, James. (1989) Giving with impure altruism: Applications to charity and Ricardian equivalence. *The Journal of Political Economy*, 97(6), 1447–1458.
- [3] Becker, Gary S. (1996) *Accounting for Tastes*. Cambridge, MA: Harvard University Press.
- [4] Bem, Daryl J. (1972) Self-perception theory. In *Advances in Experimental Social Psychology (Vol. 6)*, Leonard Berkowitz, Ed. New York: Academic Press.
- [5] Bénabou, Roland, and Jean Tirole. (2006) Incentives and prosocial behavior. *The American Economic Review*, 96(5), 1652–1678.
- [6] Benz, Matthias, and Stephan Meier. (2006) Do people behave in experiments as in the field? Evidence from donations. *Working Paper 248*. University of Zurich, Institute for Empirical Research in Economics.
- [7] Bodner, Ronit, and Drazen Prelec. (2003) Self-signaling and diagnostic utility in everyday decision making. In *The Psychology of Economic Decisions. Volume 1: Rationality and Well-Being*, Isabelle Brocas and Juan D. Carrillo, Eds. New York: Oxford University Press, 105–126.
- [8] Bolton, Gary E., and Axel Ockenfels. (2000) ERC: A theory of equity, reciprocity, and competition. *The American Economic Review*, 90(1), 166–193.

- [9] Budescu, David V., and Boris Maciejovsky. (2005) The effect of payoff feedback and information pooling on reasoning errors: Evidence from experimental markets. *Management Science*, 51(12), 1829–1843.
- [10] Cain, Daylian M., George Loewenstein, and Don A. Moore. (2005) The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *Journal of Legal Studies*, 34, 1–25.
- [11] Caplin, Andrew, and John Leahy. (2004) The supply of information by a concerned expert. *The Economic Journal*, 114(497), 487–505.
- [12] Chaffee, Steven H., and Yuko Miyo. (1983) Selective exposure and the reinforcement hypothesis: An intergenerational panel study of the 1980 Presidential campaign. *Communication Research*, 10(1), 3–36.
- [13] Crowder, Robert G. (1982) *The Psychology of Reading: An Introduction*. New York: Oxford University Press.
- [14] Dana, Jason, Daylian M. Cain, and Robyn M. Dawes. (2006) What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes*, 100(2), 193–201.
- [15] Dana, Jason, Roberto A. Weber, and Jason Xi Kuang. (forthcoming) Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Economic Theory*.
- [16] Darley, John M., and Bibb Latané. (1968) Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology*, 8(4), 377–383.
- [17] Dave, Chetan and Katherine Willey Wolfe. (2004) On confirmation bias and deviations from Bayesian updating. *Working paper*.
- [18] DellaVigna, Stefano, and Ethan Kaplan. (forthcoming) The Fox News effect: Media bias and voting. *Quarterly Journal of Economics*.

- [19] Eckel, Catherine C., and Philip J. Grossman. (1996) Altruism in anonymous dictator games. *Games and Economic Behavior*, 16, 181–191.
- [20] Eckel, Catherine C., and Philip J. Grossman. (2000) Volunteers and pseudo-volunteers: The effect of recruitment method in dictator experiments. *Experimental Economics*, 3, 107–120.
- [21] Eliaz, Kfir, and Andrew Schotter. (2006) Paying for confidence: An experimental study of preferences over beliefs. *Working paper*.
- [22] Eliaz, Kfir, and Ran Spiegler. (2006) Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior*, 56(1), 87–104.
- [23] Elster, Jon. (1998) Emotions and economic theory. *Journal of Economic Literature*, 36(1), 47–74.
- [24] Eyster, Erik. (2002) Rationalizing the past: A taste for consistency. *Working paper*.
- [25] Fehr, Ernst, and Klaus M. Schmidt. (1999) A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817–868.
- [26] Festinger, Leon. (1957) *A Theory of Cognitive Dissonance*. Evanston, IL: Row, Peterson.
- [27] Fischbacher, Urs. (2007) z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, 10(2), 171–178.
- [28] Gentzkow, Matthew, and Jesse M. Shapiro. (2006) Media bias and reputation. *Journal of Political Economy*, 114(2), 280–316.
- [29] Gerber, Alan, Dean Karlan, and Daniel Bergan. (2007) Does the media matter?: A field experiment measuring the effect of newspapers on voting behavior and political opinions. *Working paper*.

- [30] Gilovich, Thomas, and Victoria Husted Medvec. (1995) The experience of regret: What, when, and why. *Psychological Review*, 102(2), 379–395.
- [31] Haisley, Emily C., and Roberto A. Weber. (2005) Favorable interpretations of ambiguity and unstable preferences for fairness. *Working paper*.
- [32] Huckfeldt, Robert, and John Sprague. (1988) Choice, social structure, and political information: The information coercion of minorities. *American Journal of Political Science*, 32(2), 467–482.
- [33] Jonas, Eva, Stefan Schulz-Hardt, Dieter Frey, and Norman Thelen. (2001) Confirmation bias in sequential information search after preliminary decisions: An expansion of dissonance theoretical research on “selective exposure to information.” *Journal of Personality and Social Psychology*, 80(4), 557–571.
- [34] Jones, Martin. (2003a) Positive confirmation in the acquisition of information with an analysis of different subject types. *Working paper 158*. University of Dundee Discussion Papers in Economics.
- [35] Jones, Martin. (2003b) Positive confirmation in rational and irrational learning. *Working paper 150*. University of Dundee Discussion Papers in Economics.
- [36] Jones, Martin, and Robert Sugden. (2001) Positive confirmation bias in the acquisition of information. *Theory and Decision*, 50(1), 59–99.
- [37] Kandel, Eugene, and Edward P. Lazear. (1992) Peer pressure and partnerships. *The Journal of Political Economy*, 100(4), 801–817.
- [38] Karlsson, Niklas, George Loewenstein, and Duane Seppi. (2005) The ‘Ostrich Effect’: Selective attention to information about investments. *Working paper*.
- [39] King, Gary, Michael Tomz, and Jason Wittenberg. (2000) Making the most of statistical analyses: Improving interpretation and presentation. *American Journal of Political Science*, 44(2), 347–361.

- [40] Klayman, Joshua. (1995) Varieties of confirmation bias. *The Psychology of Learning and Motivation*, 32, 385–418.
- [41] Larrick, Richard P., and Terry L. Boles. (1995) Avoiding regret in decisions with feedback: A negotiation example. *Organizational Behavior and Human Decision Processes*, 63(1), 87–97.
- [42] Lazear, Edward P., Ulrike Malmendier, and Roberto A. Weber. (2006) Sorting in experiments with application to social preferences. *Working paper*.
- [43] Lieberman, Matthew D., Kevin N. Ochsner, Daniel T. Gilbert, and Daniel L. Schacter. (2001) Do amnesics exhibit cognitive dissonance reduction?: The role of explicit memory and attention in attitude change. *Psychological Science*, 12(2), 135–140.
- [44] Mullainathan, Sendhil, and Andrei Shleifer. (2005) The market for news. *The American Economic Review*, 95(4), 1031–1053.
- [45] Murnighan, J. Keith, John M. Oesch, and Madan Pillutla. (2001) Player types and self-impression management in dictatorship games: Two experiments. *Games and Economic Behavior*, 37, 388–414.
- [46] Rabin, Matthew. (1995) Moral preferences, moral constraints, and self-serving biases. *Economics working paper 95-241*. University of California at Berkeley.
- [47] Rayner, Keith, Albrecht Werner Inhoff, Robert E. Morrison, Maria L. Slowiaczek, and James H. Bertera. (1981) Masking of foveal and parafoveal vision during eye fixations in reading. *Journal of Experimental Psychology: Human Perception and Performance*, 7(1), 167–179.
- [48] Rayner, Keith, and Alexander Pollatsek. (1989) *The Psychology of Reading*. Englewood Cliffs, NJ: Prentice Hall.

- [49] Safra, Zvi, and Eyal Sulganik. (1995) On the nonexistence of Blackwell's Theorem-type results with general preference relations. *The Journal of Risk and Uncertainty*, 10(3), 187–201.
- [50] Schweitzer, Maurice E., and Christopher K. Hsee. (2002) Stretching the truth: Elastic justification and motivated communication of uncertain information. *The Journal of Risk and Uncertainty*, 25(2), 185–201.
- [51] Tomz, Michael, Jason Wittenberg, and Gary King. (2001) CLARIFY: Software for Interpreting and Presenting Statistical Results. Version 2.0. Cambridge, MA: Harvard University, June 1. Retrieved from <http://gking.harvard.edu>.
- [52] Trueswell, John C., Michael K. Tanenhaus, and Susan M. Garnsey. (1994) Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language*, 33, 285–318.
- [53] Tversky, Amos and Eldar Shafir. (1992) The disjunction effect in choice under uncertainty. *Psychological Science*, 3(5), 305–309.
- [54] Wakker, Peter P. (1988) Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making*, 1(1), 169–175.
- [55] Wason, Peter C. (1966) Reasoning. In *New Horizons in Psychology*, B.M. Foss, Ed. Harmondsworth, U.K.: Penguin.
- [56] Willemsen, Martijn, and Eric Johnson. (2004) MouselabWEB: Monitoring information acquisition processes on the web. Retrieved June 19, 2006 from <http://www.mouselabweb.org>.
- [57] Yariv, Leeat. (2005) I'll see it when I believe it: A simple model of cognitive consistency. *Working paper*.