An Experimental Study of Selective Exposure

Lauren Feiler, Jacob K. Goeree, and Leeat Yariv*

Current Version: November 2006

Abstract

We provide an array of experimental observations pertaining to agents' endogenous selection of information. We study settings in which agents have to choose one of two actions, the value of which depends on an unknown state of nature. When given a choice between different information sources, more than half of all subjects choose a source that potentially confirms their prior inclinations but is inferior in terms of associated expected payoffs. Introducing externalities between action choices in the form of a majority vote does not alleviate the results, while separating the choice of information structure from the action choice does to some extent. The analysis is potentially important for understanding political and economic phenomena such as people's tendency to read news about their preferred political candidate (rather than political opponents), to learn more about a desired product than about its alternatives, and to track their stocks more frequently when the market is doing well.

JEL classification: C91, D70, D80

Keywords: Selective Exposure, Strategic Voting.

*Division of the Humanities and Social Sciences, California Institute of Technology, Mail Code 228-77, Pasadena, CA 91125, USA and CEPR. 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.

1 Introduction

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 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. She may also decide to buy the model she was initially considering. However, only reading about this model would not give her the chance to choose the best of the available options.

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 paper 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 more valuable signals 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.

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. Since subjects in our study should have no external reason to seek information about one state over another, this allows us to obtain a more pure measure of the extent of the selective exposure bias. We also let subjects make decisions over several periods, so we can examine learning over time. We analyze the effect of majority rule voting on the selective exposure bias. To our knowledge, this institution has not previously been studied in relation to selective exposure.

It is important to understand what may exacerbate or alleviate the selective exposure bias. This bias could lead people to waste money on products that are not best-suited for their purpose or cause them to vote for the less competent of two candidates. From our finding that selective exposure is very prevalent in a simple setting, one might also infer that this bias would be persistent in everyday settings. People who select information that reinforces their preferences may also selectively avoid or enter into certain settings.

Before we fully describe our experiment and its results, we will discuss related work. We will explain the design of the experiment in Section 2. Section 3 discusses the theoretical predictions of behavior in each treatment of the experiment, and section 4 provides hypotheses based on these predictions. Section 5 details the experimental results, and section 6 summarizes the findings.

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.

This paper 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 had to guess whether a selected state was A or B, and they knew that the probability that A would be chosen was always higher than the probability that B would be chosen. They were also told that the probability that A would be chosen was either high or low (for example, 1 or .51), and they could learn which probability would be used by paying a fee. Although this information was not useful for guessing the state, nearly 80% of subjects were willing to pay the fee under certain conditions.

Eliaz and Schotter tested various models with respect to their data, including the disjunction effect (Tversky and Shafir, 1992), which they did not find support for. We will also consider this effect in our analysis. 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. Tverksy and Shafir found that the disjunction effect disappears when decision makers first state which choice they would make in each potential state of the world.

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. Along the lines of our interpretation of the selective exposure bias is work by Karlsson, Loewenstein, and Seppi (2005), who found that people checked the performance of their investments more often when a general market index had risen than when it had declined or remained the same. Other related work examines agents who selectively ignore information in order to justify making self-serving decisions (Dana, Weber, and Kuang, forthcoming; Feiler, 2006), 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).

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. (Although confirmation bias appears in a variety of contexts; see Klayman, 1995, for a review.) The classic tool used to study 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 has also been found in an experiment that allowed subjects to choose to read articles that were either in support of or against their opinion about insurance payments for alternative medicine (Jonas, Schulz-Hardt, Frey, and Thelen, 2001). This study showed that confirmation bias could be self-reinforcing, as subjects who chose articles sequentially exhibited a larger bias than those who chose the articles before reading any. Confirmation bias has also been seen when subjects do not choose the source of information themselves (Dave and Wolfe, 2004).

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.

The empirical literature has generally focused on selective exposure 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 (Huckfeldt and Sprague, 1988). While such studies are beneficial, we find it important to examine selective exposure in a carefully-controlled setting that is free of outside context.

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

[1] For 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.

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 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,

[2]This is true for any payoff levels when the states are equally likely, as long as the bonus for one state is positive and the payoff for correct guesses is higher than the payoff for incorrect ones.

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 both of 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 votes on the 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 could see the number of votes for the red and blue state in addition to the payoff information. In one voting treatment, payoff pairs are selected for each individual subject, so different members of a group may earn different amounts. In the other, the payoff pairs are the same for each member of a group. We will refer to the first treatment as *voting with individual payoffs*, and the second as *voting with shared payoffs*.

Partners

The partners treatment divides the experimental task amongst two players with shared payoffs. One subject decides which indicator to ask for, and his 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.

[3]The experiment was programmed and conducted with the z-Tree software (Fischbacher, 1999). Instructions are available at www.hss.caltech.edu/~lauren/expts/

Table 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	Session 1: 11, Session 2: 12	\$13.05
Voting, individual payoffs	2	Session 1: 15, Session 2: 15	\$14.52
Voting, shared payoffs	1	10	\$15.00
Partners	1	18	\$12.60

Table 1: Summary of Sessions. Average Payoff excludes \$5 show-up fee.

3 Theoretical Predictions

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 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 is for an agent to choose the red (blue) indicator when her a-priori preferences favor blue (red), and to follow her signal.*

First, 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 indicator a subject asked for.

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. Upon receiving a blue signal, she would guess blue. Upon receiving a null signal, the expected value of guessing red is much higher than that for guessing blue, so she would guess red. Her guess will be correct whenever the state is red. When the state is blue, there is a .5 chance that she will receive a null signal and guess the wrong state. Therefore, an agent who picks the blue indicator will guess correctly more often than the 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 more the more valuable source of information.

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 shared 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 he 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, for the shared bias, selective exposure is part of an equilibrium.

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 selective exposure occurs at the lower bonus level.

Behavioral factors may play a larger role in explaining data in the voting treatments than the formal equilibrium predictions do. 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 bother carefully considering their decisions, since there is a dilution of responsibility. A subject's guess only matters when the subject is pivotal.

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.

This treatment is designed to emphasize the importance of the indicator choice. 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.

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

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 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 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 shared 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 shared 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. There will be more selective exposure in the voting treatment with shared payoffs than in the voting treatment with individual payoffs. Within the voting treatment with shared payoffs, there will be more selective exposure at bonus levels of 20 than at bonus levels of 100.

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

Finally, it may be the case that subjects choose information poorly because they do not properly consider their actions further along the decision tree. In the partners treatment, subjects who choose indicators in the second half of the experiment will have already spent time making decisions at the later node. They should be better able to backward-induct and choose indicators than subjects with no experience.

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

5 Results

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.[5]

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, Shared Payoffs	55.5%
Partners	45.6%

Table 2: Proportion of selective exposure choices across treatments.

Table 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 shared payoffs. The amount of selective exposure was significantly greater than 0 in all treatments. ($p<.001$ for all treatments) [6]

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

[5]In the voting treatment with shared payoffs, the equilibrium condition requires subjects to follow their signals, which would mean guessing 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.

[6]In 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 $p = 0.057$) and voting treatment with shared payoffs ($t_{199} = 1.561$, one-tailed $p = 0.060$). The partners treatment was the only one in which subjects performed better than chance ($t_{359} = -1.691$, one-tailed $p = 0.046$).

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 3 shows the percentage of payoff maximizing guesses after a null signal, by subjects 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 payoff amounts.

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 more often when the bonus was high than when it was low. These differences are significant ($\chi^2 = 33.90$, $p<.001$ for the individual choice treatment; $\chi^2 = 55.21$, $p<.001$ in the voting treatment). The proportion of good guesses did not significantly differ between the bonus levels in the partners or voting with shared payoffs treatments. This result means that many subjects

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, Shared Payoffs

Indicator Choice	80, 60 Payoffs	150, 50 Payoffs	Total
Selective Exposure	70.7% (29/41)	75.0% (33/44)	72.9% (62/85)
Payoff Maximizing	81.3% (26/32)	100.0% (32/32)	90.6% (58/64)
Payoff maximizing guesses across indicator choice and bonus level: 80.5%			

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

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. It 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.

5.2 Voting Treatments

As we expected, there was somewhat more selective exposure in the voting treatments than in the individual choice treatment. Across both voting treatments, selective exposure occurred 57.6% of the time, versus 53.7% for individual choice. However, this difference was only marginally significant in a one-tailed t-test when we count each member of a group as an individual (t_{1258}=1.353, one-tailed p=0.088). Comparing the average amount of selective exposure in the voting groups each period to the individual choice data, the results are not significant ($t_{618} = 0.968$).

The proportion of subjects who made the best guess following a null signal also decreased in the voting treatments, dropping to 77.1% from 82.7% in the individual choice treatment (t_{948}=2.046, $p<.05$). Part of this decrease could be attributed to the greater amount of selective exposure in the voting treatments. Looking only at guesses after selective exposure, 73.7% made the best guess in the individual choice treatment, compared to 66.8% in the voting treatments. This difference is significant in a one-tailed test (t_{564}=1.695, one-tailed $p<.05$).

Since selective exposure is justified in the voting treatment with shared 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. However, we cannot reject the hypothesis that the same amount of selective exposure occurred in each treatment (t_{798}=0.702, p= 0.483 at the individual level; t_{158}=0.761, p=0.448 at the group level).

In the treatment with shared 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 a marginally significant effect in the opposite direction, with selective exposure rates of 52.6% for bonuses of 20 and 58.1% for bonuses of 100 (At the group level, t_{38}=1.521, one-tailed p=0.068). Hypothesis 4 is not supported by the data.

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 $p<.05$). This is the only treatment in which selective exposure rates dipped below 50%, but 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 3 shows. This difference is marginally significant (t_{483}=1.630, one-tailed p=0.052).

We expected that subjects with experience in reacting to signals to guess the state 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 amounts 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, one-tailed p=0.183). Experience in choosing an indicator might have helped the subjects who guessed the state in the second half.

Although we did not find support for a disappearance of the disjunction effect, subjects who reacted to signals in the first half of the experiment seemed to base their later indicator choices on their experience. Selective exposure in the second half of the experiment was negatively correlated with the payoffs earned when partners in the first half chose the payoff maximizing indicator, and positively correlated with payoffs earned when partners chose the selective exposure indicator. Interestingly, the sheer amount of selective exposure someone experienced in the first half of the experiment was not correlated with her indicator choices in the second half. Table 4 shows results from regressing the total number of times subjects in the second half chose the selective exposure indicator on three variables from the first half: the total amount of selective exposure seen, payoffs, and the number of best guesses the subject made. This last variable is included because subjects who are better at guessing the state are more likely to understand the strategy in the game, and they are also likely to earn more in the first half regardless of the indicator chosen by their partner. These results should be taken with a grain of salt, as they may be driven by the subjects who always or never picked the selective exposure indicator. Nevertheless, the independent variables all have the expected effect on behavior.

5.4 Learning

Figure 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 exposure on period alone (z=-

Independent Variables	Coefficient	Standard Error	t	p-value
Partners' selective exposure (total)	-0.060	0.091	-0.66	0.548
Selective exposure payoffs	0.006	0.002	3.44	0.026
Non-selective exposure payoffs	-0.007	.002	-3.87	0.018
Best guesses (total)	-1.637	0.311	-5.27	0.006
Constant	21.626	2.049	10.56	0.000
Dependent variable: Selective exposure in second half, N=9, F=244.77, R^2=0.996				

Table 4: Regression results: Number of selective exposure choices in second half of experiment on outcomes in first half.

5.16, p<.001).[7] 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.[8]

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 various interactions. None of these variables were significant. The only variable that consistently had a significant effect was the period. In the treatment with group bonuses, period was not significant, nor were any other variables.

We find that the data seem to be driven by subject types rather than learning over time. Table 5 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 consis-

[7]No other effects were significant. In the individual choice treatment, z=-1.09, p=0.274. In the voting treatment with shared payoffs, z=-0.51, p=0.613. 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.

[8]We saw that subjects who chose indicators in the second half of the partners treatment may have learned from their experience in the first half. Our discussion here focuses on learning while making indicator choices.

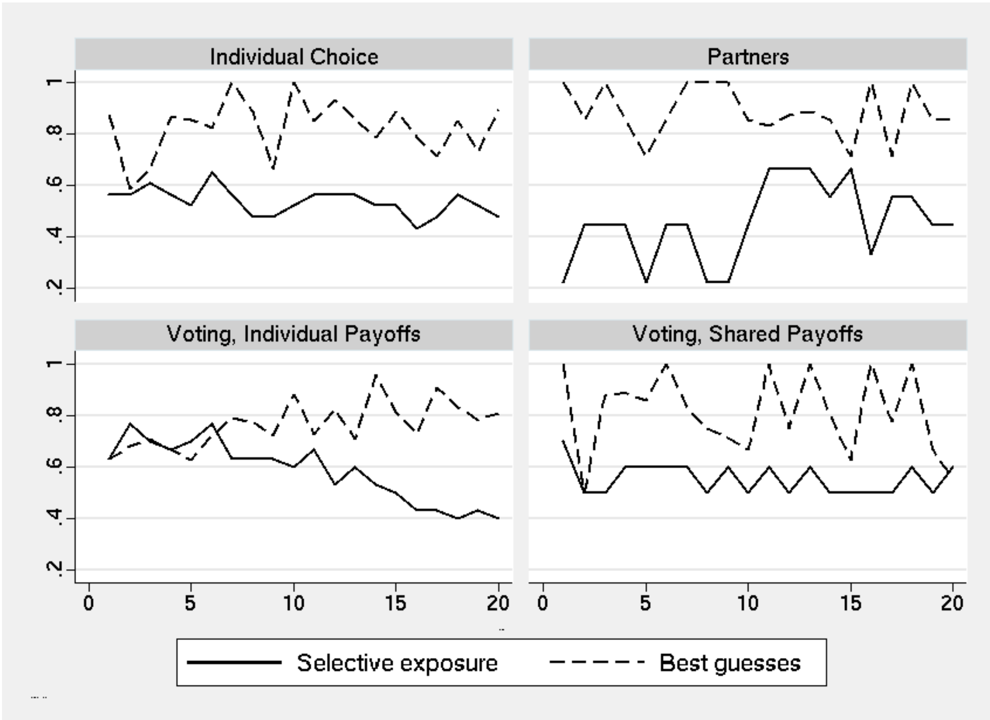

Figure 1: Average amount of selective exposure and payoff maximizing guesses, by period.

tently with only 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 table 5—those who made different indicator choices throughout the experiment.[9] 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 she does not appear to be choosing it as part of a random pattern. Only 3 subjects in the individual choice treatment learned to pick the correct indicator by this definition. In the voting treatment with individual payoffs, 9 subjects seemed to learn over time. For the group payoff treatment, 1 subject quickly learned the strategy. In the partners

[9]It 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.

	Individual Choice	Voting, Individual Payoffs	Voting, Group Payoffs	Partners
Always Payoff Maximizing	7	3	0	6
Payoff Maximizing, 1 or 2 exceptions	0	3	3	0
Always Selective Exposure	8	7	4	3
Selective Exposure, 1 or 2 exceptions	1	1	0	2
Other	7	16	3	7

Table 5: Frequency of subjects' indicator choice patterns.

treatment, of the subjects who chose indicators for the first half of the experiment, 5 subjects did not choose consistently, and only 1 appeared to learn the strategy. Only 2 of the subjects who chose indicators in the second half were not consistent, and 1 of them appeared to learn which indicator to choose. On the whole, 15 out of the 33 subjects (45.5%) in the "Other" category 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.

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, 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, and some happen to start picking the correct signal.

6 Conclusion

This paper finds that selective exposure to information about a preferred state holds fairly consistently across a variety of experimental contexts. The amount of selective exposure seen in four different treatments ranged from around 46% to around 58%. The diffusion of responsibility in majority rule voting settings somewhat increased selective exposure in comparison to the individual choice treatment. 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 voting 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.

Selective exposure may manifest itself in this way outside of the laboratory. Someone might know that she should not buy her initially preferred car if she finds another model that is equivalent in quality but less expensive, but she might not think about the fact that she is most likely to find this model by conducting research on cars other than the one she is drawn

[10] In 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, $z=0.47$, $p=0.640$.

to. We have seen a strong effect of selective exposure in an experiment with limited choices, in which people should have no intrinsic ties to one state or another. We might expect this bias to be even more prevalent when a choice set is large and the initially preferred option is focal.

References

[1] 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.

[2] Caplin, Andrew and John Leahy. (2004) The supply of information by a concerned expert. *The Economic Journal*, 114(July), 487-505.

[3] 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.

[4] Dana, Jason, Daylian M. Cain, and Robyn M. Dawes. (2006) What you dont know wont hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes*, 100(2), 193-201.

[5] Dana, Jason, Roberto A. Weber, and Jason Xi Kuang. (forthcoming) Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Economic Theory.*

[6] Dave, Chetan and Katherine Willey Wolfe. (2004) On confirmation bias and deviations from Bayesian updating. Working paper.

[7] Eliaz, Kfir and Andrew Schotter. (2006) Paying for confidence: An experimental study of preferences over beliefs. Working paper.

[8] Eliaz, Kfir and Ran Spiegler. (2006) Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior*, 56(1), 87-104.

[9] Eyster, Erik. (2002) Rationalizing the past: A taste for consistency. Working paper.

[10] Feiler, Lauren. (2006) Patterns of information avoidance in binary choice dictator games. Working paper.

[11] Fischbacher, Urs. (1999) z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual, Working Paper Nr. 21, Institute for Empirical Research in Economics, University of Zurich.

[12] Gentzkow, Matthew and Jesse M. Shapiro. (2006) Media bias and reputation. *Journal of Political Economy*, 114(2), 280-316.

[13] 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.

[14] 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.

[15] Jones, Martin. (2003a) Positive confirmation in the acquisition of information with an analysis of different subject types. Working paper.

[16] Jones, Martin. (2003b) Positive confirmation in rational and irrational learning. Working paper.

[17] Jones, Martin and Robert Sugden. (2001) "Positive confirmation bias in the acquisition of information" *Theory and Decision*, 50(1), 59-99.

[18] Karlsson, Niklas, George Loewenstein, and Duane Seppi. (2005) The Ostrich Effect: Selective attention to information about investments. Working paper.

[19] Klayman, Joshua. (1995) Varieties of confirmation bias. *The Psychology of Learning and Motivation*, 32, 385-418.

[20] Lazear, Edward P., Ulrike Malmendier, and Roberto A. Weber. (2006) Sorting in experiments with application to social preferences. Working paper.

[21] Mullainathan, Sendhil and Andrei Shleifer. (2005) The market for news. *The American Economic Review*, 95(4), 1031-1053.

[22] Tversky, Amos and Eldar Shafir. (1992) The disjunction effect in choice under uncertainty. *Psychological Science*, 3(5), 305-309.

[23] Wason, Peter C. (1966) Reasoning. In B.M. Foss (Ed.), New Horizons in Psychology. Harmondsworth, U.K.: Penguin.

[24] Yariv, Leeat. (2005) Ill see it when I believe it: A simple model of cognitive consistency. Working paper.

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 6 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.

	Red Indicator		Blue Indicator	
Signal Received	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 6: 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.

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

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

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

Similar logic applies for the payoff pair of (80, 60). Table 7 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.

	Red Indicator		Blue Indicator	
Signal Received	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 7: 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:

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

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

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 5^{th} 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 signals. 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,

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

$$\Leftrightarrow p_{80} = \frac{20}{60} = \frac{5}{12},$$

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 $p \leqslant \frac{5}{12}$, while for an agent who a-priori prefers blue, selective exposure can be part of an equilibrium if $p \geqslant \frac{7}{12}$. 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 Shared 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:

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

$$\Leftrightarrow \frac{(2-\alpha)^2 \alpha^2}{(2-\alpha)^2 \alpha^2 + (1+\alpha)^2 (1-\alpha)^2} = \frac{2}{9}$$

$$\frac{(2-\alpha)^2 \alpha^2}{(2-\alpha)^2 \alpha^2 + (1+\alpha)^2 (1-\alpha)^2} = \frac{2}{9}$$

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 $p_{80} = \frac{4}{11}$, 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:

$$\frac{(2-\alpha)^k \alpha^k}{(2-\alpha)^k \alpha^k + (1+\alpha)^k (1-\alpha)^k} = p$$

$$\Leftrightarrow \frac{1-\alpha^2}{(2-\alpha)\,\alpha} = \left(\frac{1-p}{p}\right)^{1/k}.$$

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 $\frac{1}{2}$.