On Confirmation Bias and Deviations From Bayesian Updating

Chetan Dave*†
Department of Economics
University of Pittsburgh
4S01 W. W. Posvar Hall
Pittsburgh, Pa., 15260
E-mail: chdst16@pitt.edu

Katherine W. Wolfe
Department of Economics
University of Pittsburgh
4S01 W. W. Posvar Hall
Pittsburgh, Pa., 15260
E-mail: kwolfe@pitt.edu

March 21, 2003

Abstract

Psychologists have documented several biases and heuristics that describe deviations from Bayesian updating in individual decision making under uncertainty. One particular notion, Confirmation Bias, predicts that individuals will exhibit systematic errors in updating despite the existence of learning opportunities. This paper formulates an experimental design to evaluate this behavioral heuristic within a non-strategic environment that motivates subjects financially. Subjects report probability estimates of the state of the world at the draw of each signal which may be conservative relative to a Bayesian. Indeed, pilot data indicate that both conservatism and confirmation bias are present in updating behavior.

Keywords: Confirmation Bias, Conservatism and Bayesian Updating

*This is a preliminary draft. We would like to thank David N. DeJong, John Duffy, Jack Ochs, Jean-Francois Richard, Roberto Weber and Lise Vesturlund for helpful comments and guidance. We take responsibility for any remaining errors and welcome all comments.

†Corresponding author.

1 Introduction

Models of decision making under uncertainty are central to theories of strategic and non-strategic economic interactions. In environments with imperfect information economic agents must form judgements about uncertain states of the world. These judgments, or beliefs, are then used to evaluate alternative courses of state-contingent actions. The widespread view among economists is that individuals update their beliefs based on information they receive via the use of Bayes' rule. Agents are assumed to posses some prior beliefs on states of the world, a set of actions that optimize their objectives, a well defined cost of incorrect belief formation and knowledge and skill in the use of Bayes' rule. This presumption about human behavior has been a central component of the contribution of Harsanyi (1967, 1968) who developed the theory of strategic interaction under uncertainty, and Muth (1960, 1961) who introduced to generations of macroeconomists and econometricians the role of rational beliefs in closing expectational models of non-strategic economic behavior.

The rapidly developing field of behavioral economics has identified several reasons why the above presumption on judgement under uncertainty is flawed as an accurate description of the cognitive processes that underlie human decision making. Many of these reasons rely on concepts such as learning and cognitive heuristics that approximate learning, as a better description of the judgment process. Of the many biases and heuristics proposed by the literature little attention has been paid by economists to notions that predict systematic biases in judgment due solely to cognition errors[1]. This paper examines whether one such cognitive bias, Confirmation Bias, inhibits learning.

[1] See Camerer (1995 pp. 608-609).

Confirmation bias is defined as an agents' tendency to misinterpret ambiguous evidence as confirming his current belief. An agent subject to confirmation bias will not hold beliefs that are identical to those held by a Bayesian observer. His belief formation given new information or the opportunity to obtain new information will be biased towards his original belief. This behavior will lead to systematic errors in judgment as agents may never learn to fully update given constant mis-perception of information. This paper provides an experimental analysis in order to investigate the extent to which subjects suffer from confirmation bias in forming probability judgements about the state of the world under economic incentives to update beliefs in a Bayesian manner.

The next section reviews the alternatives to Bayesian updating offered by the behavioral economics literature with the aim of briefly comparing and contrasting confirmation bias with other heuristics. This section also identifies the cognitive processes underlying confirmation bias as suggested by the psychology and behavioral economics literature. Section 3 provides an experimental design to test for confirmation bias and some alternative hypotheses in the laboratory. Section 3 also describes a pilot version of the experimental design that is analyzed in this paper. Section 4 presents the pilot data and a model that tests for confirmation bias, Section 5 concludes.

Descriptive and statistical results from the pilot data indicate that subjects exhibited conservatism and confirmation bias. Indeed, conservatism caused the subjects to report updated probabilities that are lower than that of a Bayesian and the presence of confirmation bias tended to make subjects report probabilities that are higher than that of a Bayesian. Thus in the aggregate it seemed that subjects were Bayesian, however, the micro data estimates demonstrated otherwise.

2 Confirmation Bias and Other Heuristics

2.1 Confirmation Bias

The standard view of judgment under uncertainty and its relation to decision making under uncertainty is as follows. The economic agent is presumed to begin the decision making process with a set of prior subjective beliefs about states of the world that are updated using Bayes' rule as information arrives over time. These updated posterior beliefs are then used as probability judgements that motivate a certain action or set of actions given costs of errors and the optimization objectives.

Confirmation bias is defined as the tendency of agents to update their beliefs in light of new information in a manner more likely to confirm and less likely to disconfirm previously held beliefs relative to a Bayesian observer. The psychology literature offers two suggestions towards identifying underlying cognitive processes which yield confirmation bias.

First, the decision maker is more likely to seek information that can confirm a hypothesis than that which can disconfirm. Wason's (1968) original experiment had subjects engaged in a card-selection task in which two-sided cards could be turned over or not to confirm or disconfirm a rule that the cards followed. The experiment yielded overwhelming evidence in favor of the hypothesis that subjects would turn over only those cards that could confirm the rule and not those cards that could disconfirm. Jones & Sugden (2001) tested for confirmation bias when subjects chose what information to purchase in order to make decisions. They found presence of the bias both when subjects purchased information and when they used it for decision making in a selection task environment. In addition, in their environment the bias persisted even when subjects repeatedly engaged in the selection task.

Second, the agent is more likely to make mistakes in perceiving signals or interpreting evidence so as to support his hypothesis. Lord, Lepper & Ross (1979) and Plous (1991) show that two subjects with opposing beliefs can both interpret ambiguous evidence as supporting their own position. Rabin & Schrag's (1999) model provides a theoretical foundation to this view within a signal-extraction framework. They demonstrate that when an agent with confirmation bias perceives signals he/she is not only under or over confident relative to a Bayesian observer, but can also suffer from wrongness and may not learn despite being given an infinite amount of free information.

These analyses suggest that the cognitive processes underlying confirmation bias are as follows. First agents seek confirmatory evidence suggesting a cost-benefit approach to updating such that there is excessive weight, relative to a Bayesian, placed on the cost of ignoring confirmatory evidence. Second agents mis-perceive evidence to support beliefs even when they do not seek information. It is the latter process that is investigated empirically in this paper, the former has been explored by Jones & Sugden (2001).

These two cognitive processes are related to the main reason for investigating confirmation bias: which of the two following theoretical frameworks for persistent errors in updating, a lack of learning, should be adopted for use in increasing the realism of economic models? First, Jones & Sugden (2001) demonstrate the bias within the context of Bayesian decision making. This suggests a modification to the standard Bayesian decision making task that agents face in any stochastic model. Second, Rabin & Schrag (1999) offer a signal-extraction problem that is solved differently by Bayesian observers and biased agents. This suggests a modification to only the signal extraction problem that is a part of the Bayesian decision task in many economic models. Both of these alternatives are relevant as the psychology literature is clear about the effects of confirmation bias, a lack of learning, which is a strong motivating factor underlying this paper.

2.2 Related Heuristics

It is important to distinguish between confirmation bias and other heuristics[2]. First, a few heuristics and biases are related to errors in Bayesian updating. The conservatism bias describes the situation in which all new information is insufficiently weighted in the updating process. The opposite bias, to overweight all new information, is called overreaction. Anchoring and adjustment contains conservatism in updating, the 'adjustment', in addition to the incorrect priors generated through the choice of the 'anchor'. Confirmation bias differs from conservatism or overreaction in that it expects confirming evidence to be given more weight than disconfirming evidence. The law of small numbers is a heuristic where evidence from small samples is overweighted and evidence from large samples is underweighted; sample size should not matter under confirmation bias. Representativeness is a heuristic that overweights samples that are representative of one particular state of the world. Confirmation bias would imply that evidence exactly representing the disbelieved state of the world would still be underweighted. It is possible for any or all of these biases that cause over or underweighting of new information to be present at the same time. If so, then the weight given to any new piece of evidence will depend on which of the heuristics is operational. The experimental design in this paper allows only for conservatism (or overreaction) as sample sizes are not varied, and provides

[2] The seminal exposition of the psychological issues relevant to judgment under uncertainty is provided by Kahneman, Slovic & Tversky (1982).

treatments and empirical tests to distinguish this heuristic from confirmation bias.

Next, confirmation bias should be distinguished from cognitive search as they have different implications for decision theory. Theories of cognitive search attribute observed irrationalities as agents balancing the costs and benefits of receiving and processing increasing amounts of information: at some point it does not pay to process large amounts of information. The implication of confirmation bias, given the second underlying process noted above, is that actions may simply be based on beliefs that are statistically incorrect.

2.3 Related Literature

Economists have analyzed some of these heuristics in experiments focused on individual decision making and in settings in which subjects interact. The main difference between the psychological and economic experiments being that in the latter subjects are motivated financially whereas in the former there is reliance on intrinsic motivation. An experimental analysis in which several judgment biases were jointly investigated in a group environment was conducted by Camerer (1987). In that analysis the heuristics of representativeness, base rate ignorance, conservatism and overreaction were documented and analyzed in subjects engaged in asset trades. The economic environment was characterized by a double-oral auction and subjects were provided with priors on states of the world as well as sample information that could be used to update priors. An important environmental characteristic of the biases investigated in Camerer (1987) is that subjects observed each others' behavior. Grether (1980) finds that for financially motivated subjects in individual decision making environments Bayesian updating could be rejected in favor of representativeness and/or base rate ignorance. El-Gamal & Grether (1995) formulate a procedure in which the rules of thumb actually used by subjects can be identified. They find that subjects use Bayes' rule, representativeness and conservatism in that order of importance. These studies have primarily elicited responses of the following variety: "Do you think that the [state of the world] is A or B?". This paper is innovative in that it elicits probability judgments. That is, the design in this paper elicits responses of the following variety: "What do you think is the probability of [the state of the world A]?".

In summary, confirmation bias is investigated in an environment where agents infer states of the world through possibly 'ambiguous' evidence via the use of informative signals. The approach is to test confirmation bias as an individual phenomenon as this bias needs to be first investigated in an environment in which subjects have a non-strategic incentive to update in a Bayesian manner. Having characterized the bias in such an environment future analyses can investigate the effects of this bias in strategic environments with increased clarity. The first focus is on the case in which agents do not seek information but are provided costless and clear signals instead. However, confirmation bias may also exist in environments without information seeking if it is based on mis-perception of signals, thus the second focus is on the case in which there are ambiguous signals. Given this dual focus ambiguity is modeled both as differing signal correlations as well as through stimuli that by construction can be mis-perceived. These latter stimuli are of two types, emotive and non-emotive, given that confirmation bias may be more likely with emotive stimuli given the analysis of Lord, Lepper & Ross (1979). Positive empirical results will substantially increase the applicability of the bias in theoretical and applied analyses of dynamic decision making under uncertainty.

3 Experimental Design

3.1 Design

A similar experimental environment to that of Camerer (1987), Grether (1980) and El-Gamal & Grether (1995) is adopted in the design. The experiment allows for two states of the world. Subjects report their initial probability estimate as to which state of the world they think prevails in order to insure that they start the process with uniform priors. Subjects then receive a series of signals about which state of the world prevails. After each signal, they issue new probability estimate reports. In initial treatments signals are unambiguous, so a bias or heuristic should be present only in how subjects update their beliefs about the state of the world. In subsequent treatments it is possible for subjects to physically mis-perceive the signals and to attach emotive content to the stimuli.

3.1.1 Treatments

For the first treatment, states of the world are represented by two bingo cages, one containing seven black balls and three white balls (state A) and the other containing three black balls and seven white balls (state B). The second treatment employs a different correlation between signals and the state of the world by changing the ratios of black to white balls to six black and four white and vice-versa. The two types of balls are clearly contrasted during the instructions at the beginning of the session.

In the third treatment, the white balls and black balls are replaced by grey balls where the grey scales are different but measurable. For the fourth treatment the grey balls are replaced by similar pictures of brand name shoes[3]. This treatment captures the notion that confirmation bias may arise primarily in environments with emotive alternatives.

3.1.2 Automation

The bingo cages are covered so that the contents are not visible on the computer screen. In all treatments the sequence of signals are exactly the same so that the data are comparable across treatments. The state of the world is randomly determined at the start of each of the ten rounds. Subjects are allowed a training session to not only learn how the experiment operates but also to learn how the payoff structure ensures that they receive the highest payoff when they report their estimates truthfully. The draws are videotaped and shown on-screen on a shielded computer. While it is feasible to automate the draw such that subjects see only signals on-screen without the bingo cage theatre, this is not an optimal choice. The main reason for the choice of a physical randomization device is that it is well known that in these types of environments subjects often do not believe that a computer draws randomly; the bingo cage reinforces the belief that the choice of the state of the world and the draws are truly randomly drawn.

The signals are represented by drawing one ball from the selected state of the world bingo cage with replacement, the signals are therefore *i.i.d.* In the third and fourth treatments it is possible that subjects mis-perceive the signal. The subjects record their perception of the signal on their record sheet screen for each of the ten draws per round. Before receiving any signals, the subjects record their probability

[3] The shoes stimuli are motivated by telling the students that the sports director at the university is hypothetically considering two contracts for shoes for athletes.

estimates as to the state of the world on their record sheet. Finally, the history of each subjects reports are displayed in the history screen. Once subjects enter their estimates for the current draw and round they click on 'OK' to receive the next signal.

3.1.3 Payoffs

Subjects are induced to report their probability estimates truthfully by the betting mechanism. They are instructed that at the end of the experiment, one probability estimate will be selected at random to determine the payoff parameters. A random number from 50 to 100 is chosen to be the bonus number. This number determines which of two bets the subject will play for payment: the state of the world bet or the random number bet. If the probability the subject reported for his more likely state of the world is greater than the bonus number, then the subject plays the state of the world bet for that round. If the bingo cage contained the set he reported as more likely, he wins the bonus payment, otherwise he wins nothing. The probability of winning the bonus payment in this bet is the probability that the state of the world was that which he reported as more likely. If the probability the subject reported for his more likely state of the world is less than or equal to the bonus number, then the subject plays the random number bet. A second random number is generated, this time from 1 to 100. If the bonus number exceeds the second random number the subject wins the bonus payment otherwise the subjects wins nothing. The probability of winning the bonus payment in this bet is the bonus number.

The payoff scheme is designed to provide an incentive for subjects to accurately report their probability estimates. If p is the subjects' true estimate of the probability of state A (the more likely state), and d is the randomly selected bonus number, then the two bets are represented as,

$$\begin{aligned} G_1 &= p \cdot [\$20], (1-p) \cdot [\$0] \\ G_2 &= d \cdot [\$20], (1-d) \cdot [\$0] \end{aligned}$$

All subjects, regardless of their risk posture will prefer G_1 to G_2 when $p > d$ and will prefer G_2 to G_1 when $p < d$. They can guarantee their preferred gamble by truthfully reporting p. Figure 1 in Appendix B illustrates this payoff procedure.

3.2 Design Advantages

The design motivates subjects to report $P(A \mid n_a, n_b)$ and $P(B \mid n_a, n_b)$ where n_a and n_b are the number of signals of either state. With the various treatments the design clearly delineates the effect of confirmation bias as an updating problem (more weight on confirming evidence) from that of the psychologically motivated mis-perception issue. The design also captures two important features, the elicitation of beliefs and risk neutrality. The first feature is self evident: elicitation is required in order to compare the reported probabilities relative to those of a Bayesian. The second feature ensures that reported probabilities are not conditioned upon a subjects' desire to maximize payoffs under a concave utility for money such that risk averse subjects would report probabilities reflecting the resulting conservative estimate. The design is similar to that employed by Grether (1980) and El-Gamal & Grether (1995) who only elicited the state of the world the subject thought was more likely. The aim here is to elicit the actual probability estimates subjects use in decision making as they are the main object of analyses.

The objective is to identify confirmation bias over and above conservatism. This identification is necessary as given the experimental environment it is entirely possible that in the treatments without perception errors agents will suffer from conservatism (or overreaction). Given this possibility, introducing perception errors should lead to increased conservatism. The models discussed in the next section allow for identification given the treatments in the design.

3.3 Pilot Experiment

The design discussed above was tested on undergraduates enrolled in an intermediate microeconomics course at the University of Pittsburgh. The payoffs were replaced by a zero show up fee and a bag of chocolates as the bonus prize. Forty students participated in five rounds each of which consisted of ten draws. Figures 2-6 in Appendix B plot the aggregate data for each of rounds. The only treatment administered was the one in which the proportion of black-white balls is 0.7.

The aggregate data indicate that the subjects may well be quasi-Bayesian in their updating behavior. In particular, in rounds 1 and 2 the data suggest the presence of conservatism. In round 3 subjects were much more confused and hesitated to return to the 50-50 mark, this suggests that in the final experiments the sequences be chosen so as to avoid this noisy possibility. Further, despite the fact that round 5 and 1 were similar there seems to be less attention paid by students to the exercise, suggesting that five rounds is near the maximum number of rounds that can be implemented. Finally, despite the average behavior exhibited in these figures, the presence of confirmation bias cannot be clearly seen, for which purpose statistical tests on the micro data need to be carried out.

4 Econometric Model

4.1 Definitions and Data

Let i index subjects (40), j index rounds (5) and k index draws (10). In the pilot experiment data for the following variables were directly collected. The color of the ball drawn, White or Black, denoted as $B_{jk} \in \{W, B\}$ was noted for each draw. Next, each subjects reported probability of the state of the world being A at each draw, denoted as $P_{ijk} \in (0, 100)$[4] was recorded by students. Given these data the following variables were computed. The log posterior odds ratio in each round for each draw, denoted as $\pi_{ijk} = \ln\left(\frac{P_{ijk}}{100-P_{ijk}}\right)$ was computed for each subject. Since the proportion of black to white balls is constant at 0.7, for a Bayesian[5] this ratio is constant at $\ln\left(\frac{0.7}{0.3}\right) = 0.847$. Next, the strength of each subjects beliefs, denoted as $S_{ijk} = |\pi_{ijk-1}|$ was computed. Third, each subjects belief update was computed, denoted as,

$$y_{ijk} = \left\{ \begin{array}{ccc} \Delta\pi_{ijk} & \forall & B_{jk} = B \\ -\Delta\pi_{ijk} & \forall & B_{jk} = W \end{array} \right\}$$

Fourth, a confirming signal dummy was computed, denoted as,

$$x^C_{ijk} = \left\{ \begin{array}{ll} 1 & \begin{array}{l} \text{if} \quad (\pi_{ijk-1} > 0 \cap B_{jk} = B) \\ \text{or } (\pi_{ijk-1} < 0 \cap B_{jk} = W) \end{array} \\ 0 & \text{otherwise} \end{array} \right\}$$

[4]The following truncation was employed on the raw data, which had a range of range of $[0, 100]$. If $P_{ijk} > 99.9$ then $P_{ijk} = 99.9$, if $P_{ijk} < 0.1$ then $P_{ijk} = 0.1$.

[5]See Appendix C for a formal derivation.

In addition a disconfirming signal dummy was computed, denoted as,

$$x^{D}_{ijk} = \left\{ \begin{array}{ll} 1 & \begin{array}{l} \text{if} \quad (\pi_{ijk-1} > 0 \cap B_{jk} = W) \\ \text{or } (\pi_{ijk-1} < 0 \cap B_{jk} = B) \end{array} \\ 0 & \text{otherwise} \end{array} \right\}$$

Finally, two cross variates were computed, confirming strength, denoted as $x^{CS}_{ijk} = x^{C}_{ijk} \times S_{ijk}$ and disconfirming strength, denoted as $x^{DS}_{ijk} = x^{D}_{ijk} \times S_{ijk}$. These last four variables are interpretable given that there are three possible states of beliefs: neutral, confirming and disconfirming, defined as follows.

1. Neutral Beliefs: $P_{ijk-1} = 50 \longrightarrow S_{ijk} = 0$ and regardless of the value of B_{jk} the updating is not subject to confirmation bias.

2. Confirming Beliefs: $P_{ijk-1} > 50 \longrightarrow S_{ijk} > 0 \cap B_{jk} = B$ or $P_{ijk-1} < 50 \longrightarrow S_{ijk} > 0 \cap B_{jk} = W$.

3. Disconfirming Beliefs: $P_{ijk-1} > 50 \longrightarrow S_{ijk} > 0 \cap B_{jk} = W$ or $P_{ijk-1} < 50 \longrightarrow S_{ijk} > 0 \cap B_{jk} = B$.

Therefore the dummy variables (and cross terms) above are to be interpreted relative to the neutral case. Further since neutral beliefs are not of relevance only the confirming and disconfirming dummies are to be included in any regression tests. Finally, this data collection/computation procedure should function well in the other treatments as well, especially as subjects will be requested to note the color of the ball themselves.

4.2 The Model and Estimates

Given the variables above, the presence of confirmation bias in updating can be investigated by estimating and testing the following regression equation,

$$y_{ijk} = \alpha_0 + \alpha_1 x_{ijk}^{C} + \alpha_2 x_{ijk}^{D} + \alpha_3 x_{ijk}^{CS} + \alpha_4 x_{ijk}^{DS} + \varepsilon_{ijk} \tag{1}$$

where the regression error ε_{ijk} is decomposed into individual specific fixed effects, υ_i, and an idiosyncratic error, ξ_{ijk}. Given equation (1), the following are the hypotheses that can be tested,

$$H_0^B \quad : \quad \widehat{\alpha_0} = 0.847 \cap \widehat{\alpha_1} = \widehat{\alpha_2} = \widehat{\alpha_3} = \widehat{\alpha_4} = 0 \tag{2}$$

$$H_0^C \quad : \quad \widehat{\alpha_0} < 0.847 \tag{3}$$

$$H_0^{CB} \quad : \quad \widehat{\alpha_1} > 0 \text{ or } \widehat{\alpha_3} \neq 0 \tag{4}$$

$$H_0^{CB} \quad : \quad \widehat{\alpha_2} < 0 \text{ or } \widehat{\alpha_4} \neq 0 \tag{5}$$

where the hats indicate estimated coefficients. The first hypothesis tests for Bayesian behavior, the second tests whether subjects follow the conservatism heuristic. The third hypothesis tests for confirmation bias and the fourth for disconfirmation bias. The estimates for the above equation using the pilot data are presented in Table 1 of Appendix B.

The coefficient estimates suggest that the following conclusions can be drawn from the pilot data. First, subjects' exhibited use of the conservatism heuristic as $\widehat{\alpha_0} = 0.535 < 0.847$. Second, subjects' exhibited confirmation bias in updating as $\widehat{\alpha_1} > 0$ and $\widehat{\alpha_3} \neq 0$. Third, the impact of confirmation bias diminished as subjects' became more certain of the state of the world. For values of strength greater than zero until approximately 2 (which corresponds to reported probabilities of the more likely state of the world from 51% to 88%), subjects updated more strongly in the confirming case than the neutral or disconfirming cases. Finally, updating behavior given disconfirming evidence is not significantly different than that in the neutral case as $\widehat{\alpha_2} = 0$.

5 Conclusion

There are certain probability judgment biases that can act as obstacles to learning, this paper investigates empirically whether one such bias holds. The experimental design, built on previous analyses of judgement heuristics and biases, was implemented in a pilot experiment with strong results. The pilot data indicated that subjects used the conservatism heuristic and exhibited confirmation biases in updating. In the final experiments, treatments will also be conducted in which subjects can physically mis-perceive stimuli that themselves can be emotive. The pilot data suggest that statistical tests for these remaining treatments will bolster the main finding, that there is reason to believe that confirmation biases exist even under financial incentives to not weight confirming evidence differently from disconfirming evidence.

References

[1] Camerer, C. (1987), "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence", **American Economic Review**, v. 77.

[2] Camerer, C. (1995), "Individual Decision Making", **Handbook of Experimental Economics**.

[3] Edwards, W. (1982), "Conservatism in Human Information Processing", in *Judgment Under Uncertainty: Heuristics and Biases*, (Kahneman, Slovic & Tversky eds.), **Cambridge University Press**.

[4] El-Gamal, M. A. & Grether, D. M. (1995), "Are People Bayesian? Uncovering Behavioral Strategies", **Journal of the American Statistical Association**, v. 90.

[5] Grether, D. M. (1980), "Bayes Rule as a Descriptive Model: The Representativeness Heuristic", **Quarterly Journal of Economics**, v. 95.

[6] Harsanyi, J. (1967, 1968), "Games With Incomplete Information Played by Bayesian Players, Parts 1-3", **Management Science**, v. 14.

[7] Jones, M. & Sugden, R. (2001), "Positive Confirmation Bias In The Acquisition of Information", **Theory and Decision**, v. 50.

[8] Kahneman, D., Slovic, P. & Tversky, A. (1982), *Judgment Under Uncertainty: Heuristics and Biases*, **Cambridge University Press**.

[9] Lord, C., Lepper, M. R. & Ross, L. (1979), "Biased Assimilation and Attitude Polarization: The Effects or Prior Theories on Subsequently Considered Evidence", **Journal of Personality and Social Psychology**, v. 37.

[10] Muth, J. F. (1960), "Optimal Properties of Exponentially Weighted Forecasts", **Journal of the American Statistical Association**, v. 55.

[11] Muth, J. F. (1961), "Rational Expectations and the Theory of Price Movements", **Econometrica**, v. 29.

[12] Plous, S. (1991), "Biases in the Assimilation of Technological Breakdowns: So accidents make us safer?", **Journal of Applied Social Psychology**, v. 21.

[13] Rabin, M. & Schrag, J. L. (1999), "First Impressions Matter: A Model of Confirmatory Bias", **Quarterly Journal of Economics**, v. 114.

[14] Wason, P. C. (1968), "Reasoning About A Rule", **Quarterly Journal of Experimental Psychology**, v. 20.

Appendices

A Experiment Instructions

A.1 General Instructions

A.2 Payments

A.3 Screens

B Figures & Tables

B.1 Figures

Figure 1

Subject forms p

$\downarrow$

Subject reports $\widehat{p}$

$\downarrow$

$d \sim UNI[50, 100]$ drawn

$\downarrow$

$\widehat{p} > d \swarrow \searrow \widehat{p} \leq d$

$G_1 \qquad G_2$

$\Longrightarrow p = \widehat{p}$

Figure 2

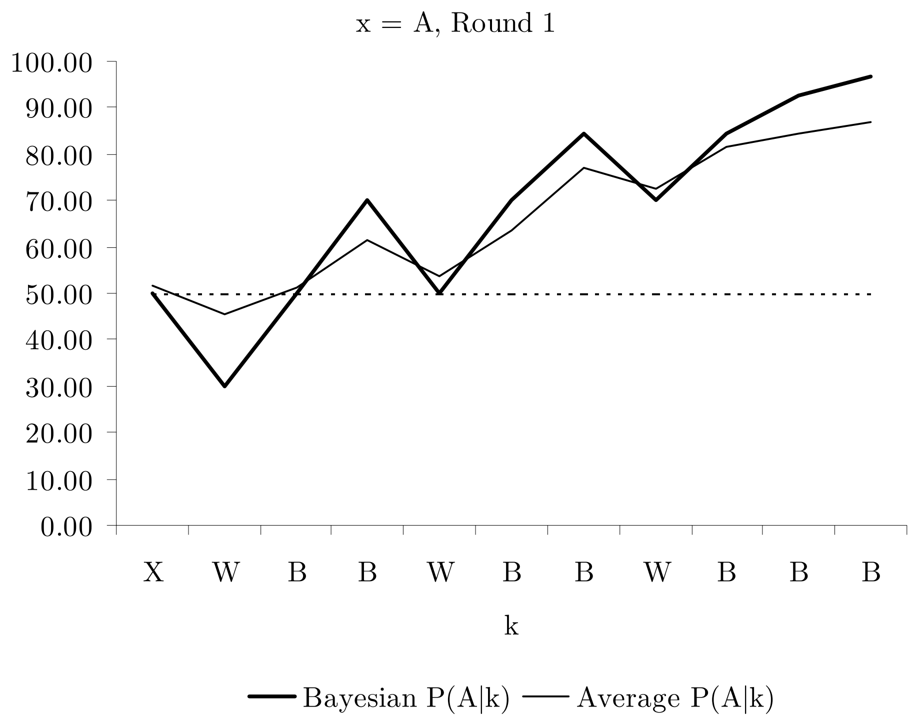
x = A, Round 1
100.00
90.00
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.00
X W B B W B B W B B B
k
Bayesian P(A|k)
Average P(A|k)

Figure 3

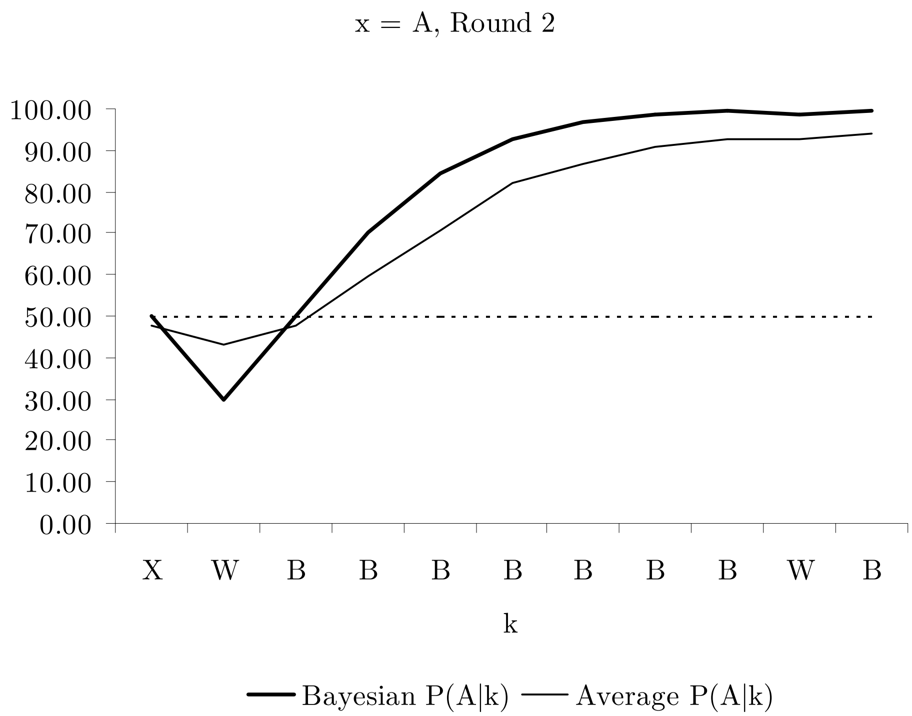
x = A, Round 2
100.00
90.00
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.00
X
W
B
B
B
B
B
B
B
W
B
k
Bayesian P(A|k)
Average P(A|k)

Figure 4

x = A, Round 3

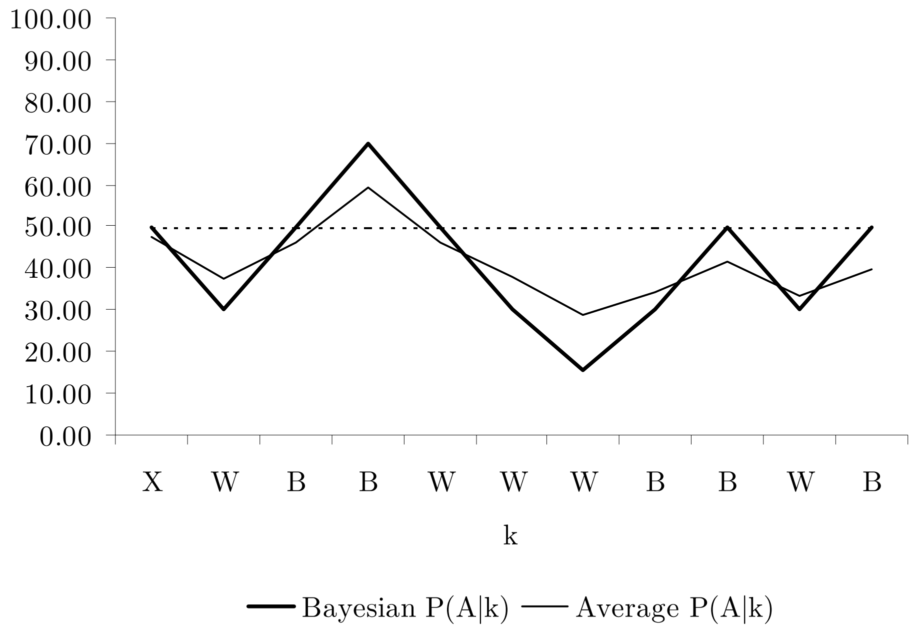

Figure 5

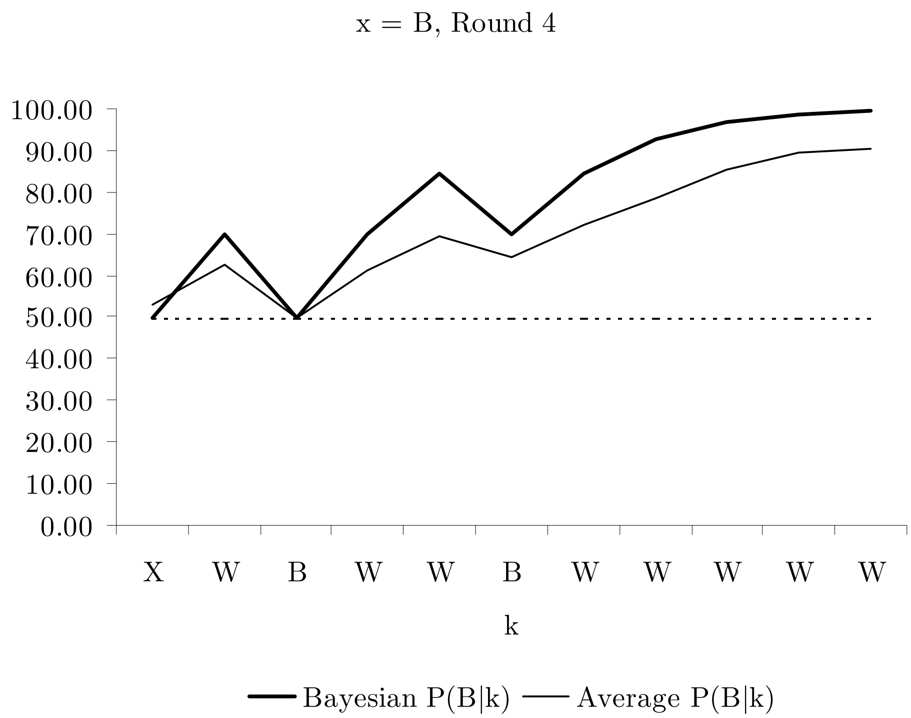
x = B, Round 4
100.00
90.00
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.00
X W B W W B W W W W W
k
Bayesian P(B|k)
Average P(B|k)

Figure 6

x = B, Round 5

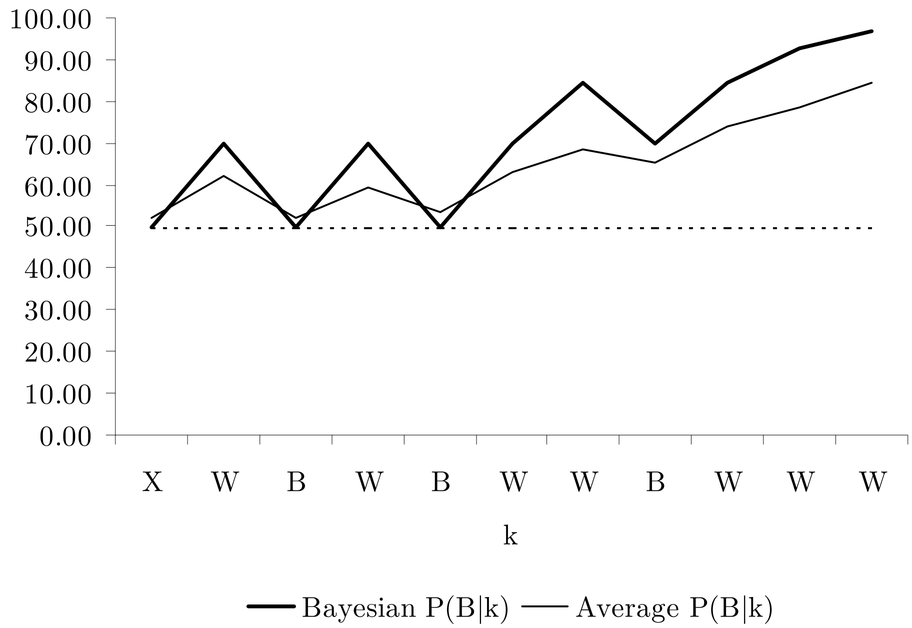

B.2 Tables

Table 1

Coefficient	Estimate	Std. Error	$\lvert t \rvert$
α_0	0.535	0.073	7.32
α_1	0.434	0.095	4.57
α_2	-0.070	0.101	0.69
α_3	-0.134	0.019	7.14
α_4	0.122	0.026	4.69
Within R^2 :	0.0439	F_1 :	22.45
Between R^2 :	0.2696	F_2 :	6.24
Overall R^2 :	0.0238		

- F_1 refers to the F-statistic for the hypothesis that all of the independent variables included have no effect on the dependant variable.
- F_2 refers to the F-statistic for the hypothesis that all of the fixed effects are zero.

C Bayesian Behavior

Within a given round a Bayesian would satisfy the following version of Bayes' Law,

$$\pi_{1k} = L_k \pi_0 \tag{6}$$

where π_{1k} and π_0 are respectively the posterior and prior logs at the k^{th} draw. The likelihood, L_k, takes the following form when there are two possible signals, say 'a' and 'b', that are indicative of two possible states, say 'A' and 'B',

$$L_k = \left(\frac{\theta}{1-\theta} \right)^{Z_k} \tag{7}$$

where θ is the proportion of 'a' signals to 'b' signals and Z_k is the difference between the number of 'a' signals and 'b' signals as of the k^{th} draw. Combining the above two equations and taking logs yields the following.

$$\ln \pi_{1k} - \ln \pi_0 = \beta Z_k, \qquad \beta = \ln \left(\frac{\theta}{1-\theta} \right) \tag{8}$$

First differencing yields,

$$\Delta \ln \pi_{1k} = \beta \Delta Z_k \tag{9}$$

where $\Delta Z_k = \{-1, +1\}$. Taking absolute values of the first differenced equation demonstrates that a Bayesian updates, in log odds terms, at a constant of β, as follows.

$$|\Delta \ln \pi_{1k}| = \beta \tag{10}$$