Good News and Bad News are Still News: Experimental Evidence on Belief Updating

Alexander Coutts[*†]
Nova School of Business and Economics

April 3, 2018

Abstract

Bayesian updating remains the benchmark for dynamic modeling under uncertainty within economics. Recent theory and evidence suggest individuals may process information asymmetrically when it relates to personal characteristics or future life outcomes, with good news receiving more weight than bad news. I examine information processing across a broad set of contexts: 1) ego relevant, 2) financially relevant, and 3) non value relevant. In the first two cases, information about outcomes is valenced, containing either good or bad news. In the third case, information is value neutral. In contrast to a number of previous studies I do not find differences in belief updating across valenced and value neutral settings. Updating across all contexts is asymmetric and conservative: the former is influenced by sequences of signals received, a new variation of confirmation bias, while the latter is driven by non-updates. Despite this, posteriors are well approximated by those calculated using Bayes' rule. Most importantly these patterns are present across all contexts, cautioning against the interpretation of asymmetric updating or other deviations from Bayes' rule as being motivated by psychological biases.

JEL classification: C91, D83, D84.

Keywords: beliefs, Bayes' rule, asymmetric belief updating, conservatism, overconfidence

[*]Nova School of Business and Economics, Faculdade de Economia da Universidade Nova de Lisboa, Campus de Campolide, 1099-032 Lisbon, Portugal; alexander.coutts@novasbe.pt

[†]Acknowledgements: This research has been generously supported by a grant from the Russell Sage Foundation. I am grateful to the editor and two anonymous referees for a host of constructive comments. I am heavily indebted to my advisor David Cesarini for numerous discussions and comments. I am grateful for helpful comments from Hunt Allcott, Kai Barron, Thomas Buser, Colin Camerer, Andrew Demers, David Eil, Guillaume Fréchette, Nicole Hildebrandt, Elliot Lipnowski, David Low, Amnon Maltz, Markus Möbius, Joseph Mullins, Giorgia Romagnoli, Tanya Rosenblat, Andrew Schotter, Emilia Soldani, Tobias Salz, Séverine Toussaert, Joël van der Weele, Christopher Woolnough, Sevgi Yuksel, as well as seminar participants at NYU, the 2015 European ESA Meeeting, the 2016 ASFEE meeting, Games 2016, and THEEM 2017. All errors are my own.

1 Introduction

The ability to process new information in forming and updating beliefs is critical for a wide range of important life decisions. Students receiving grades adjust beliefs about their ability to succeed in different majors before declaring, entrepreneurs may be awarded or denied funding for their projects and must update beliefs about the viability of these projects, smokers who are informed of new health statistics on the dangers of smoking must update about these risks in deciding whether to quit.

In modeling such situations, it is typically assumed that individuals use Bayes' rule to update their beliefs. Individuals who receive partially informative signals about states of the world are assumed to incorporate this information in an unbiased, calculated way. While Bayesian updating is the current paradigm theoretically, it is also accepted that it has a strong normative basis.

Given the importance of updating beliefs for decision making in economic contexts, experimental evidence on updating has been studied for some time (e.g., Kahneman and Tversky (1973); Grether (1980); Grether (1992); Camerer (1987); Camerer (1995); Holt and Smith (2009)). These studies greatly contributed to our understanding of how individuals update their beliefs, and highlighted the existence of cognitive biases, as updating deviated from Bayes' rule in systematic ways. Even so, the nature of the updating tasks in these studies differed considerably from the real updating decisions that motivated them. Unlike updating decisions that are economically relevant to individuals, updating in lab experiments typically involved events such as drawing balls from urns, where subjects hold no personal or financial stake in the outcome, beyond an incentive payment for accuracy. Henceforth, I refer to such events as value neutral; information about them is just news, neither good nor bad.

In contrast, value relevant or valenced events are those in which an individual strictly prefers one outcome to another, and news about the outcome can be categorized as good or bad. This distinction may be critically important, as there is now a small but growing body of theory and empirical evidence suggesting that there exist further psychological biases in how information is processed in value relevant contexts, depending on whether it is perceived as good or bad news (e.g., Eil and Rao (2011); Sharot et al (2011); Ertac (2011); Mobius et al (2014); Kuhnen (2014)). Drawing in part on this evidence, Sharot et al (2012) claim: "Humans form beliefs asymmetrically; we tend to discount bad news but embrace good news."

In this paper I examine whether updating differs across value relevant and neutral contexts. My primary focus is on understanding whether there exist additional psychological biases which lead to asymmetric updating when news is good or bad, beyond the cognitive biases which have been previously found for value neutral contexts. I examine binary events that are either 1) ego relevant, 2) financially relevant, or 3) value neutral. These consist of two uncertain events that are objective in nature, involving the rolling of virtual dice; one subjective, involving estimation of historical temperatures; and one that pertains to ego, involving relative performance on a math and verbal skills quiz. Financial relevance is introduced randomly at the subject-event level, with the endowment of additional financial prizes of $80 in the outcomes of interest. As this experiment utilizes a financially incentivized belief elicitation procedure, this introduces a different type of financial stake. To minimize confusion, I refer to this type of financial stake as an accuracy payment.

Information comes in the form of partially informative binary signals regarding the outcomes of the events. I elicit beliefs utilizing the incentive compatible elicitation procedure of Grether (1992), Holt and Smith (2009), and Karni (2009). The primary analysis focuses on between subject variation in updating patterns and follows Mobius et al (2014) in estimating an empirical model of belief updating that nests Bayesian updating as a special case, but allows for differential response to affirmative versus negative signals. The elicitation procedure improves on previous work that utilizes other elicitation procedures, such as the quadratic scoring rule (QSR), that are not invariant to subjects' risk preferences.[1]

The results show that, common to previous studies, updating behavior deviates from the strict mechanics of Bayesian updating. Updating is conservative, with many non-updates, and asymmetric, with negative signals receiving more weight than affirmative signals. I find evidence that observed asymmetry is affected by the sequence of signals received, with more negative sequences showing more negative asymmetry, a new type of confirmation bias. Yet critically, these deviations do not differ across value relevant and value neutral contexts, i.e. regardless of whether signals contain good or bad news, or are simply conveying neutral information. While I am able to reject the Bayesian benchmark with statistical precision, posteriors are well approximated by those calculated using Bayes' rule. These results are consistent with Holt and Smith (2009) who found important deviations from Bayes' rule, yet also that average posteriors appear to approximate their Bayesian counterparts well, particularly for intermediate priors.

Overall the analysis indicates the importance of observing a broad set of counterfactual belief updates, as results of this paper demonstrate how narrowly comparing two events can lead to conclusions that don't hold up to broader comparisons with other events. Specifically, updating patterns appear more asymmetric when updating about one's own performance on the ego relevant quiz rather than another's performance, yet similar asymmetry is present when subjects update about objective dice events. These results thus caution against attributing biased updating patterns to contexts where such bias is psychologically plausible, as updating patterns are similar across settings where such bias is clearly implausible. The remainder of the paper is as follows. The following section discusses recent theoretical and empirical work investigating belief updating. Next, I outline the experimental design, followed by a description of the results, and concluding with a brief discussion.

[1] Antoniou et al (2015) discuss issues that may arise with inference of updating behaviors when elicitation procedures are not robust to risk preferences.

2 Related Literature

This paper is related to a sizeable literature on studying how individuals process information and whether this is well approximated by Bayes' rule. The majority of these studies investigate updating about value neutral events, where subjects have no personal or financial stake in the outcome, excepting a payment for accurate belief reports. These studies noted a number of cognitive biases corresponding to deviations from Bayes' rule.

Kahneman and Tversky (1972), Kahneman and Tversky (1973), and Grether (1980) found evidence in support of a "representativeness" heuristic, whereby individuals neglect prior or base rate information when facing samples that mimic proportions or qualities of their parent population.[2] More broadly, base rate bias/neglect is often used in reference to a general tendency to over respond to information, relative to a Bayesian. Earlier, Edwards (1968) had observed a seemingly opposing bias, conservatism, the tendency for individuals to under respond to information, ending up with posteriors closer to their priors than if they had followed Bayes' rule.

This evidence demonstrated that individuals do not appear to follow the exact mechanics of Bayes' rule. Yet across contexts, Bayesian approximation does relatively well, and fewer deviations are observed for more experienced subjects, see Camerer (1987) and Camerer (1995). More recently, Holt and Smith (2009) find that belief updating is consistent with Bayes' rule at the aggregate level, however they find systematic deviations which are more pronounced for extreme values of the prior, using a similar analysis to this paper.

Beyond these cognitive biases, which are invariant to the valence of news or signals, neuroscientists, psychologists, and economists have posited the existence of additional psychological biases, i.e. updating biases that are present only when information is valenced. There are a number of psychologically plausible motivations for why updating may differ when the context is financially or personally relevant. The proposed theories share a common consequence: an asymmetry that did not feature in the earlier discussion of cognitive biases, specifically, an over-weighting of good news relative to bad news.

The first such motivation is that asymmetric updating may enable individuals to nurture biased beliefs about their abilities or about future outcomes. For example, Landier (2000), Yariv (2005), Mayraz (2014), and Mobius et al (2014) present models where individuals gain utility from holding positive beliefs, and process information in a biased manner (over-weighting good news relative to bad news) in order to nurture such positive beliefs. Second, biased information processing is similarly rational if optimistic beliefs improve health outcomes, e.g. Scheier and Carver (1987). Finally, biased beliefs could also be nurtured for strategic purposes, as in Benabou and Tirole (2002) regarding self-confidence.

A number of studies have begun to investigate these biases. In an unincentivized study, Sharot et al (2011) examined how individuals updated their beliefs about future life events such as being diagnosed with Alzheimer's disease or being robbed. They found

[2] Representativeness bias was also found and studied by Grether (1992) and Holt and Smith (2009).

that individuals updated more in response to good news relative to bad news.[3] Another unincentivized study by Wiswall and Zafar (2015) finds some evidence that college students revise beliefs more when they receive information that average future earnings are greater than expected, relative to receiving information that earnings are less than expected.

As it is typically not possible to financially incentivize the elicitation of future life events, economists have turned to study belief updating about value relevant events in the laboratory. Crucially, these studies have differed from early work, not only in context, but also in analysis. As there is little theoretical motivation for observing asymmetric updating in value neutral contexts, previous work did not examine the relative weight placed on affirmative versus negative signals. With value relevant contexts, this changed. This has important implications for the comparability of results, as the wealth of analysis from earlier work cannot be brought to bear on more recent studies. This feature of the existing literature presents a clear rationale for the examination of robust counterfactual settings, an advantage of the current paper.

Recent and ongoing work on belief updating in value relevant contexts often focuses on a particular context of interest, along with one counterfactual. Most relevant is Mobius et al (2014), who pair an experiment with a theoretical model of optimally biased Bayesian updating in the context of ego relevant events. In the experiment they examine how individuals update beliefs about scoring in the top 50% on an intelligence test, using the same elicitation procedure as this paper. They find evidence that individuals update asymmetrically, over-weighting good signals relative to bad, and conservatively, updating too little in response to either type of signal. To provide evidence that these biases are not present outside of ego relevant contexts they compare the results to a follow-up where a subset of the same subjects complete the updating task for a robot, with the result that both conservatism and asymmetry are reduced.

Regarding financial relevant events, Barron (2016) investigates updating beliefs when individuals have a financial stake in the outcome of drawing balls from two urns. His experiment complements this paper, focusing on one type of event while exogenously varying different values priors, as opposed to the across event variation that is the focus here. Barron (2016) does not find evidence of asymmetric over-weighting of good news. In contrast to the current paper financial stakes are smaller (10 GBP rather than $80), and the focus precludes examining different types of events, e.g. ego relevant.

Other related papers on ego relevant tasks are Buser et al (2018), Eil and Rao (2011), Ertac (2011). Buser et al (2018) is most similar, however their focus is on heterogeneity in deviations from Bayes' rule, as such they do not have an ego irrelevant control. Eil and Rao (2011) and Ertac (2011) differ in their information structure; the former finds evidence of positive asymmetric updating for relative intelligence and beauty, while the latter finds the opposite asymmetry, that bad news is weighted more than good news. While results

[3] Of note is that there is recent work which critiques some of the evidence for asymmetric updating in these types of studies within psychology and neuroscience, see Shah et al (2016).

of Ertac (2011) are similar to those of this paper, the conclusions differ starkly.[4]

A common challenge inherent to this literature involves the construction of an appropriate counterfactual updating task. Clearly, evidence of biased updating in value relevant contexts could relate to cognitive biases identified by earlier work, rather than psychological biases. Individuals may update differently for different base rates or priors, or for different distributions of received signals. Even the objective versus subjective nature of the event may affect updating behavior.[5] To overcome this challenge, the current paper contains a more comprehensive set of updating decisions which form a natural set of counterfactuals that can be used to evaluate behavior. Observing this broader set of decisions may alter the interpretations of deviations from Bayes' rule in important ways, as the results of this paper will show.

3 Experimental Design

The experiment was conducted at New York University, at the Center for Experimental and Social Science (CESS). Recruitment was done via the CESS online system, where undergraduate students are notified by email when an experiment is scheduled, and have the opportunity to sign up. A total of 326 subjects participated, in 32 different sessions for an average of 10 subjects per session.[6] The average subject payment was $24.96 for approximately 75 minutes including a $10 showup fee. The experimental data is also studied in Coutts (2015), there with the aim of distinguishing models of belief bias. That paper focuses exclusively on prior formation, and does not examine updating behavior.

Individuals in the experiment faced four different binary events and a sequence of four incentive compatible belief elicitations for each event. First, their prior beliefs about the probability of the event were elicited. Next they received a binary signal, regarding whether the event had occurred. This signal was true with two-thirds probability, and false with one-third probability. After receiving this signal their beliefs were again elicited, and the same process was repeated two more times.

One concern was that the elicitation procedure or the sequence of signals might be confusing to some subjects. These features of the experiment were presented making use of intuitive explanations that aid subject understanding, following Grether (1980) who recognized their importance for experimental credibility. A large component of the ex-

[4]Grossman and Owens (2012) examine absolute performance, but do not observe the same biases in information processing found in studies on relative performance. Clark and Friesen (2009) find little evidence of overconfidence in a related study on task experience rather than feedback.

[5]Some evidence that the objective versus subjective nature may affect updating is presented in this paper. Among the studies cited above, the majority have one counterfactual updating task, however only Barron (2016) (financially relevant) and Ertac (2011) (ego relevant) examine a counterfactual updating task that is robust to the above concerns, and in the case of Ertac (2011) only for a small subsample.

[6]Instructions are available in Appendix H. Because of a technical failure, one session resulted in data for only one event. 318 subjects participated in all four events.

periment consisted of intuitive explanations and examples, as well as practice with all of the experiment's components before the actual experiment began. A questionnaire was administered to every participant after completing the experiment. This, as well as verbal feedback, suggested that subjects had a good understanding of the various components of the experiment.

3.1 Belief Elicitation

To elicit beliefs I use the method of Karni (2009), utilized also by Grether (1992), Holt and Smith (2009), and Mobius et al (2014), which I henceforth call the lottery method.[7] Incentive compatibility follows from a dominance argument, that individuals strictly prefer a higher probability of earning the same monetary prize. In order to make comparisons of lotteries, the method additionally requires that individuals exhibit probabilistic sophistication, see Machina and Schmeidler (1992). It does not require that individuals maximize expected utility, nor does it require assumptions on risk preferences.[8] The method involves the possibility of earning a lump sum payment a, and works as follows.

Subjects are asked to report a probability $\tilde{\pi}$ that makes them indifferent between participating in a lottery that pays a with probability $\tilde{\pi}$ and 0 otherwise, or participating in a lottery that pays a whenever the event of interest occurs. After indicating $\tilde{\pi}$, the computer draws a random number r distributed uniformly from 0 to 100. If $r \geq \tilde{\pi}$, a subject participates in the lottery that pays a with probability r. If $r < \tilde{\pi}$ the subject faces the lottery that pays a when the event in fact occurs ("event lottery"). When r takes on discrete values, this mechanism is equivalent to a 101 item choice list that requires a choice between the "event lottery" and an objective lottery which pays a with percentages in the set of integers from 0 to 100, with one choice selected at random. In the experiment a is either low ($3), medium ($10), or high ($20), randomized at the session level.

In order to facilitate subjects understanding the lottery method, the experiment made use of an intuitive graphical interface.[9] Subjects were introduced to a gumball machine, that had 100 black or green gumballs. This represented a lottery with the probability of success equal to the number of green gumballs out of 100. One gumball would be drawn from the machine, at random. Subjects were told that the computer had a list of 101 gumball machines, each with a different proportion of green gumballs, and that one of

[7]See Karni (2009) for a more detailed description of the lottery method, though the method has been described in a number of earlier papers, see Schlag et al (2015). The mechanism has also been referred to as the "crossover method", "matching probabilities", and "reservation probabilities".

[8]While this method is a variant of the Becker-Degroot-Marschak (BDM) mechanism, note that it is not subject to the critique of Karni et al (1987), as the method identifies parallel probabilities that lead to indifference between two lotteries; see Healy (2016).

[9]Subject understanding of methods of belief elicitation has been a concern for many experimental economists. Schlag et al (2015) provides a nice overview of the literature discussing the effects of complexity of elicitation procedures on responses. Schotter and Trevino (2014) is a good overview more generally on eliciting beliefs in laboratory settings.

these would be randomly selected, i.e. the discrete uniform distribution.

They were then asked to indicate on a graphical slider, exactly what point they would prefer to base their potential earnings on the "gumball lottery" instead of the "event lottery", which paid off if the event of interest occurred. Figure 1 provides screenshots of the gumball machine, as well as the slider subjects had to move. Subjects were given significant practice with the slider, with non-paid practice events, before the primary experiment began.

Figure 1: Screenshot from the Experiment: Slider

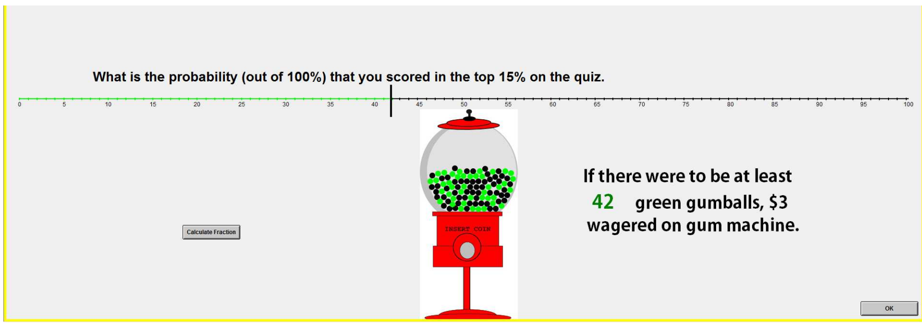

The gumball machine was used to provide an intuitive representation of the lottery method. Subjects would indicate on the slider the minimum threshold, i.e. the number of green gumballs there had to be in the machine before they would prefer to wager the accuracy payment (here $3) on the gumball machine rather than the event, "you scored in the top 15% on the quiz". The proportion of green gumballs in the picture adjusted as the slider was moved.

3.2 Events

The four events, presented in random order, are a key source of variation in the experiment.[10] Of these four events, two involve rolling dice, and their probabilities can be objectively calculated.[11] The outcome of these events was determined by chance, and individuals could not affect these outcomes.

The other two events are subjective, based upon tasks that individuals had completed prior to the beginning of the experiment. Most relevant to previous studies of asymmetric updating is an event where individuals had to estimate the probability they scored in the top 15% on a 5 minute, ego relevant quiz. Percentiles in the quiz were generated by comparing scores to 40 individuals who took the quiz in previous piloting, which was

[10] One of the events (easy dice) was fixed as the final event. The other three events were randomly ordered at the session level. Updating behavior does not differ by order.

[11] It is difficult to find a rigorous definition of what makes a probabilistic process "objective". I follow the definition of Gilboa and Schmeidler (2001).

known to students. The quiz involved multiple choice questions on math and English skills, similar to standardized college entry tests in the USA.[12] In order to generate a sensible control group, every individual in the experiment had a 30% chance of being selected to estimate the performance of a randomly selected anonymous partner in the room, rather than their own performance. 95 out of 318 subjects were randomly selected for this control, independent of any other treatments in the experiment.

The fourth event was whether the individual correctly answered a question about what the weather (mean temperature) was on a randomly selected day in New York City in the previous calendar year. This question is not objective in the sense of the dice questions, but it also does not appear to involve skill or ability.[13] Figure 2 summarizes the four events that all individuals faced.

One important feature of these events are the relatively low probabilities, due to budget constraints. Empirically the events occurred on average 15% of the time, with the implication that there will be significantly more negative than affirmative signals. In the data, 63% of the signals are indeed negative, which, as the results sections will discuss, may affect updating behavior. Importantly, while the probability of events were low, priors were significantly higher.[14] The average prior is 36%, and over one-third of priors are greater than or equal to 50%, which facilitates comparisons with other studies.

3.3 Stakes

Another key source of variation in the experiment involved varying the financial stakes. Within each event, individuals had a 50% chance of receiving an additional \$80 if that event occurred. This financial stake was made salient as subjects physically drew a token from a bag, that was labelled either \$80 or \$0. If they drew a \$0 token, they knew that they would have no financial stake in the event, they could only potentially earn an accuracy payment for their belief report. If they drew an \$80 token, they knew that they could potentially earn \$80, if the event in fact occurred, and they were selected for payment for that specific outcome.

In this latter case, at the end of the experiment *only* one of the financial stake *or* the accuracy payment is paid, chosen at random. This design feature ensured independence between the financial stake and accuracy payment, to maintain incentive compatibility by eliminating hedging opportunities in a manner similar to Blanco et al (2010). Incentive compatibility is preserved assuming a state-wise monotonicity assumption, which is

[12]This quiz was taken before individuals made any choices, and before they had any knowledge of the four events. Subjects were told truthfully that performing better on the test would lead to higher expected payments.

[13]Individuals were correct if the true temperature was within plus or minus 5 degrees Fahrenheit of their estimate. In fact, there is no correlation between beliefs about getting this question correct and actually getting the question correct.

[14]See Coutts (2015) for a discussion of this phenomenon.

Figure 2: Description of Events

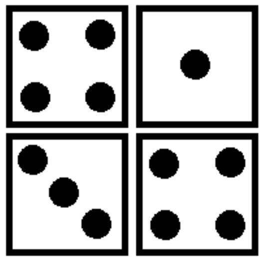

(a) Hard Dice: The computer rolls four dice. Event occurs when exactly two out of those four dice was a specified number (e.g. 4). The probability of this is $\binom{4}{2}\left(\frac{1}{6}\right)^2\left(\frac{5}{6}\right)^2 = \frac{150}{1296}$ or approximately 11.57%.

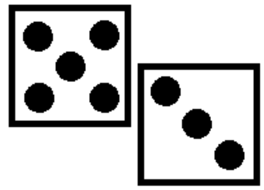

(b) Easy Dice: The computer rolls two dice. Event occurs when two different specified numbers were the only numbers to come up (e.g. 5-3, 3-5, 3-3, 5-5). The probability of this is $\frac{4}{36}$ or approximately 11.11%.

(c) Weather: Event occurs if the individual correctly estimated the average temperature on a specified random day in NYC in the previous year (2013), +/- 5 deg F. In the sample, 25.77% of subjects were in the correct range.

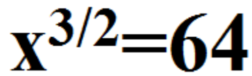

(d) Quiz: Event occurs if the individual scored in the top 15% on an ego relevant multiple choice quiz (self). For a subset of participants the event pertained to a random partner's performance instead of their own (other). Percentiles were generated in comparison to 40 pilot quiz-takers.

required whenever one pays for one randomly selected decision, see Azrieli et al (2018). Further details of this procedure and conditions for incentive compatibility can be found in Appendix A.

3.4 Signals (News)

News comes in the form of noisy, binary signals, after the first elicitation (prior). Signals were explained with the aid of pictures of "gremlins", in a procedure related to that of Mobius et al (2014). In this experiment, individuals were told that there were three gremlins that all knew whether the event had occurred. Two of the three gremlins always told the truth, while one gremlin always lied. The subjects were then told one gremlin had been randomly selected, and that gremlin either provided them an affirmative signal (the event had occurred), or a negative signal (the event had not occurred).

In this way signals were true with probability $\frac{2}{3}$, which differs slightly from the signal strength of $\frac{3}{4}$ in Mobius et al (2014). After receiving the signal, posterior beliefs were elicited again using the lottery method. Subjects were given three independent signals (knowing the structure in advance) and were informed that the gremlins were drawn with replacement in order to maintain a constant probability of $\frac{2}{3}$ that the signal was true.

In total subjects had their beliefs elicited four times for each of the four events: one prior elicitation, and three posterior, for a total of 16 elicitations. One of the elicitation rounds was randomly selected at the end for payment in accordance with the procedure discussed in Appendix A.

Figure 3 depicts screenshots from the experiment that showcase the use of gremlins as graphical aids.[15] In addition, individuals also had practice with receiving signals and updating beliefs with non-paid practice events, before the paid experiment began.

Figure 3: Screenshots from the Experiment: Signals

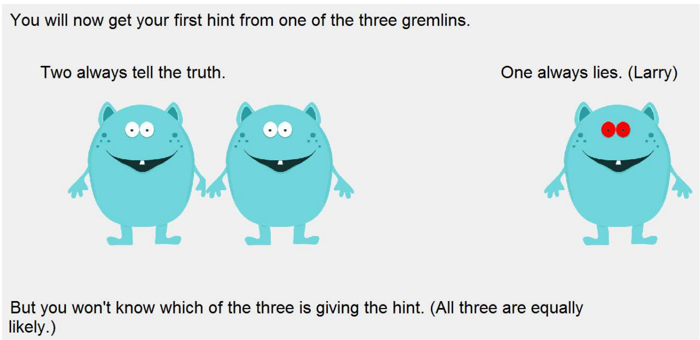

(a) Screenshot introducing signals.

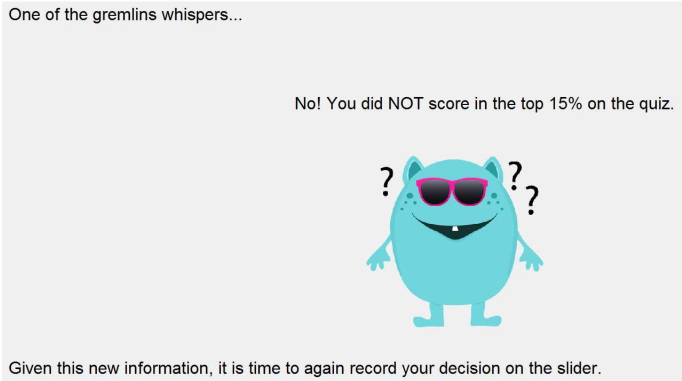

(b) Screenshot of a received (negative) signal.

[15]Images are from www.mycutegraphics.com, by author Laura Strickland.

3.5 Classifying Good/Bad News (Value Relevant) from Just News (Value Neutral)

To summarize the treatments, accuracy payments of $a \in \{3, 10, 20\}$ were randomized at the session level, while financial stakes of $0 or $80 were randomized at the subject-event level. Additionally, for the quiz event, 30% of subjects were randomly allocated to a control treatment where they were asked to update about another randomly selected individual's performance.

Thus, for an event that a given subject was allocated an $80 financial stake, information about whether the event occurred will contain good or bad news. Information which indicates such an event is more likely, corresponds to an increase in the expected probability of earning the $80, and vice-versa.

For those events that a given subject was allocated a $0 stake, whether information/news is good or bad depends on the event itself. To the extent that individuals gain utility from believing they have high ability, the quiz event (estimating own performance) involves a personal, non-pecuniary stake. Thus for this quiz event, binary signals about performance will contain either good news (they are in the top 15%), or bad news (they are not in the top 15%).

These contexts are in contrast with value neutral events where there are no personal stakes: the two dice events and the weather event, *and* those in which subjects held a $0 (rather than $80) financial stake in the outcome. Here signals contain information about outcomes, but these outcomes are irrelevant to individual well being. In other words, news is simply news.[16]

Of 1280 events, 634 involved a financial stake of $80, which means signals regarding these events are categorized as value relevant. The remaining 646 events had no financial stake.[17] However, of these 646 events, 115 were the quiz (self) event which is a personal stake. Thus altogether, I classify 749 events as potentially relating to good/bad news (value relevant), and 531 events as just news (value neutral).

4 Results

4.1 Overview of Pooled Results

As a first pass at examining updating behavior I plot reported subject beliefs and compare these to posterior beliefs that would have resulted if subjects updated using Bayes' rule.

[16] As noted before, accuracy payments are a form of financial stake. However, *ex-ante*, an individual has no reason to care whether an event occurs or not, as long as their belief report is accurate. It may be argued that the weather event involves a personal stake, if individuals derive utility from correctly estimating historical temperatures. I do not find this likely, but regardless, none of the results in this paper hinge on the inclusion of the weather event.

[17] The reason they are not equal is that some sessions had odd numbers of subjects, and the physical drawing of $0 or $80 always involved an equal proportion of both stakes.

Following the discussion in the previous section, I split the sample into events where subjects have a financial or personal stake in the outcome, and events where subjects have no such stake.

The former include all events where subjects stand to earn \$80 if the event occurs, as well as the event that involved whether they scored in the top 15% on the quiz. For these events, information in the form of noisy signals is valenced. The latter events do not involve individual ability and subjects have no additional financial stake in the outcome of the event. Thus, signals provide new information about these outcomes, information which is neither good nor bad. Belief reports are financially incentivized using the lottery method.

Figure 4 plots the average reported belief for the prior, as well as the belief after receiving each signal. Included in the same figure is the path that beliefs would take if individuals were perfect Bayesians, given the subject's first reported belief.

Figure 4: Evolution of Beliefs

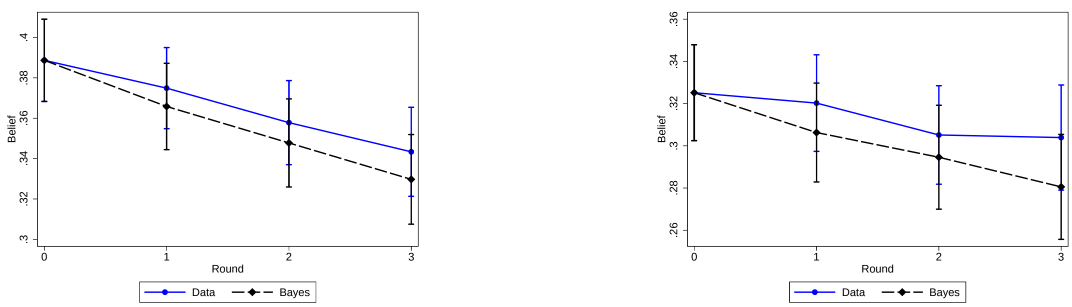

(a) Financial/Personal Stake: Good/Bad News

(b) No Stake: Just News

The path of beliefs starting from the prior (round 0) and after each sequential signal (rounds 1 through 3). **(a)**: Quiz (self) event and all instances where subjects could earn \$80 if the event occurred. **(b)**: All other instances. Average individual responses are the blue solid line, the Bayesian benchmark is marked as the black dashed line. Bayesian benchmark takes prior beliefs, and uses Bayes' rule to calculate how beliefs would evolve given the signals that subjects actually received.. Error bands represent 95% confidence intervals. **(a)** $N = 749$ **(b)** $N = 531$ observations per round.

From Figure 4 one can see that there are slight deviations from Bayes' rule for both subsamples, as individual's update slightly more conservatively. However the difference between reported posteriors and Bayesian posteriors is not significant at any conventional level. There are slight differences in prior belief formation across the two groups, part of which is accounted for by the presence of the quiz (self) event. In fact, priors are biased upwards for all events, as detailed in Coutts (2015).

Beyond this, no substantive differences are apparent in patterns of updating across the two subsamples. Thus from an initial look at the data, updating does not appear to differ when news is good or bad, compared to when news is simply news.[18] The correlation between empirical posterior beliefs and posterior beliefs calculated using Bayes' rule is 0.89, higher than that found in Mobius et al (2014). Pooling across the three updating rounds, the average posterior is 33.8%, while the average Bayesian posterior, calculated using subject priors, is 32.5%. These are remarkably similar, though it is noteworthy that the difference is significantly different from zero at the 1% level using a Wilcoxon rank-sum test. Interestingly, despite the high correlation between actual posteriors and Bayesian posteriors, the aggregate data includes a large number of non-updates. 41% of reported posteriors are identical to reported priors and only 9% of subjects update in every round.

To summarize thus far, updating appears to be well approximated by Bayes' rule, despite a large number of non-updates. However, looking only at aggregated posterior beliefs may potentially obscure how individuals react to signals, and may be affected by the fact that negative signals are more prevalent. I next use a flexible empirical framework to examine how individuals respond to both affirmative or negative signals, which, depending on the event and stake conditions, may be interpreted as good or bad news. This permits a more rigorous investigation into whether individuals update asymmetrically, when the outcomes of events are either ego relevant or financially relevant, as has been found in previous literature.

4.2 Information Processing: Framework

I now provide a more detailed analysis, following Mobius et al (2014) by using a flexible model of updating that retains the structure of Bayes' rule, but allows for the possibility that individuals place different weight on the prior, affirmative signals, or negative signals. The model is a variant of that used originally by Grether (1980), and more recently by Holt and Smith (2009). Bayes' rule can be written in the following form, considering binary signals, $s_t = k \in \{0, 1\}$, and letting $\hat{\mu}_t$ be the belief at time t:

$$\frac{\hat{\mu}_t}{1-\hat{\mu}_t} = \frac{\hat{\mu}_{t-1}}{1-\hat{\mu}_{t-1}} \cdot LR_k \tag{1}$$

where LR_k is the likelihood ratio of observing signal $s_t = k \in \{0, 1\}$. In this experiment, $LR_1 = 2$ and $LR_0 = \frac{1}{2}$, given the signal strength of $\frac{2}{3}$. Taking natural logarithms of both sides and using an indicator function, $I\{s_t = k\}$, for the type of signal observed,

$$logit(\hat{\mu}_t) = logit(\hat{\mu}_{t-1}) + I\{s_t = 1\} \ln (LR_1) + I\{s_t = 0\} \ln (LR_0) . \tag{2}$$

The empirical model nests this Bayesian benchmark as follows,

$$logit(\hat{\mu}_{it}) = \delta logit(\hat{\mu}_{i,t-1}) + \beta_1 I(s_{it} = 1) \ln (LR_1) + \beta_0 I(s_{it} = 0) \ln (LR_0) + \epsilon_{it}. \tag{3}$$

[18] In Appendix G I present these figures for all events, all financial stake, and all accuracy payment conditions separately. From those figures, one can see that there are some slight differences in updating behavior across events, but financial stakes or accuracy payments do not appear to alter updating behavior.

δ captures the weight placed on the log prior odds ratio. β_0 and β_1 capture responsiveness to either negative or affirmative signals respectively. In the context of the experiment, $s_{it} = 1$ corresponds to a signal that YES the event had occurred, while $s_{it} = 0$ corresponds to a signal that NO the event had not occurred. Since $I(s_{it} = 0) + I(s_{it} = 1) = 1$ there is no constant term. ϵ_{it} captures non-systematic errors.

This framework allows for the testing of the primary hypotheses of this paper. Bayes' rule is a special case of this model when $\delta = \beta_0 = \beta_1 = 1$. Additionally, as described in Mobius et al (2014), Bayes' rule satisfies three additional properties: invariance, sufficiency, and stability. When $\delta = 1$, the updating process is said to satisfy invariance, i.e. the change in logit beliefs depends only on past signals. Sufficiency requires that after controlling for prior beliefs, lagged information does not significantly predict posterior beliefs. Finally, stability requires that the structure of updating is stable across rounds.

Given past evidence, it seems unlikely subjects will satisfy the strict requirements of Bayesian updating. As a secondary hypothesis, I significantly weaken the requirement of Bayesian updating to a flexible model that may involve a number of different cognitive biases. The only restriction I impose, is that these cognitive biases do not differ on average across value relevant and value neutral contexts. Thus subjects may be conservative, or suffer from representativeness bias, but if these biases do not lead to differential updating patterns across valenced and non valenced contexts, then they have not entered into the domain of psychological bias.

The key tests of whether there are psychological biases in updating involve whether updating differs across valenced and neutral contexts. Using superscripts V and N respectively on the framework parameters for these contexts, the key hypotheses are presented in Figure 5. The central hypothesis of interest in this paper, is whether there are differences in asymmetric updating, presented on the final line of Figure 5. While this asymmetric updating hypothesis was posited before the experiment was conducted, it was not pre-registered.

Figure 5: Hypotheses of the Empirical Framework

Bayes' Rule

$$\delta = 1; \qquad \beta_0 = 1; \qquad \beta_1 = 1.$$

Cognitive Biases

$$\delta^V = \delta^N; \qquad \beta_1^V = \beta_1^N; \qquad \beta_0^V = \beta_0^N.$$

Psychological Biases

$$\delta^V \neq \delta^N; \quad \text{or} \quad \beta_1^V \neq \beta_1^N; \quad \text{or} \quad \beta_0^V \neq \beta_0^N.$$

Positive Asymmetry:	Negative Asymmetry:
$\beta_1^V - \beta_0^V > \beta_1^N - \beta_0^N$	$\beta_1^V - \beta_0^V < \beta_1^N - \beta_0^N$

Importantly, the relevant benchmark for evidence of psychological bias is not Bayes' rule, but actual observed updating behavior in value neutral contexts. This significantly raises the bar for detecting deviations, as prima facie evidence suggests that individuals already suffer from cognitive biases in information processing.

As a first line of investigation, in Appendix C I examine the three properties of invariance, sufficiency, and stability for Bayes' rule. Overall, the pooled data do *not* support these properties, though the magnitude of deviations is relatively small.[19] I now turn to the main empirical framework of Equation 3.

Table 1 presents the aggregate data, as well as the two subsamples introduced earlier, pooling across all updating rounds. Note that significance is indicated as different from the Bayesian benchmark prediction of one, *not* zero. Standard errors are clustered at the individual level. In the primary sample I do not include posterior beliefs that were updated in the opposite direction that Bayes' rule predicts, which amounts to dropping 4.8% of observations. I include all other subjects, including those who never update their beliefs.[20]

Table 1 provides the finer details of updating behavior that Figure 4 is unable to capture. What is first apparent is that updating behavior deviates from the strict Bayesian prediction that all coefficients are equal to 1. There is substantial conservatism in response to both affirmative and negative signals, as indicated by coefficients less than one for β_1 and β_0. Of note is that the degree of conservatism is less than that found in studies by Mobius et al (2014) and Buser et al (2018).

The coefficient on δ is significantly lower than 1. The significance of this is that subjects are updating as if the priors they held were closer to one-half, i.e. a probability weighting of prior beliefs towards one-half.[21] There is further a strong asymmetric bias that is present across Table 1, with negative signals receiving more weight than affirmative, significant at

[19] One valid concern regarding the OLS analysis is in using priors as a dependent variable. Since priors are lagged posteriors, this creates a potential issue if there is substantial heterogeneity in response to signals, which could lead to upwardly biased estimates of δ, see Mobius et al (2014). Instrumenting with higher order lagged beliefs or lagged Bayesian beliefs is possible, however such techniques do not alter the results reported. Recovering unbiased estimates of δ is also not central to the results of this paper.

[20] 4.8% is less than the approximately 10% in Mobius et al (2014) and Buser et al (2018). As Bayesian posteriors will never be at the boundary for intermediate priors, the framework is agnostic for beliefs of 0 or 1. Following Mobius et al (2014) and Buser et al (2018), these observations are dropped, amounting to 6% of the sample. In Appendix Table F1 I examine implications of these sampling restrictions, following Grether (1992) and Holt and Smith (2009) by replacing boundary observations by 0.01 or 0.99 respectively.

[21] To see this, note that the generalization in Equation 3 implies the following relationship:

$$\frac{\hat{\mu}_t}{1-\hat{\mu}_t} = \left(\frac{\hat{\mu}_{t-1}}{1-\hat{\mu}_{t-1}}\right)^{\delta} \cdot (LR_k)^{\beta_k}.$$

When $\delta < 1$, the effect is to bias or weight the log prior odds ratio $\left(\frac{\hat{\mu}_{t-1}}{1-\hat{\mu}_{t-1}}\right)$ towards 1, i.e. subjects update as if priors $\hat{\mu}_{t-1}$ were closer to one-half. If δ were greater than 1, this effect would lead to updating as if priors greater than one-half were closer to 1, and updating as if priors less than one-half were closer to 0.

Table 1: Updating Beliefs for All Events

Dependent Variable: Logit Posterior Belief			
Regressor	(1) Good/Bad News	(2) Just News	(3) All
δ^V	0.918^{***}		
	(0.012)		
β_1^V	0.594^{***}		
	(0.040)		
β_0^V	0.782^{***}		
	(0.043)		
δ^N		0.907^{***}	
		(0.014)	
β_1^N		0.576^{***}	
		(0.051)	
β_0^N		0.812^{***}	
		(0.051)	
δ			0.914^{***}
			(0.009)
β_1			0.588^{***}
			(0.034)
β_0			0.793^{***}
			(0.038)
P-Value ($\delta = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_0 = 1$)	0.0000	0.0003	0.0000
Diff ($\beta_1 - \beta_0$)	-0.189	-0.236	-0.205
P-Value ($\beta_1 = \beta_0$)	0.0002	0.0003	0.0000
R^2	0.84	0.82	0.84
Observations	1950	1410	3360
P-Value [Chow-test] for $\delta^V = \delta^N$			0.5326
P-Value [Chow-test] for $\beta_1^V = \beta_1^N$			0.7700
P-Value [Chow-test] for $\beta_0^V = \beta_0^N$			0.5852
P-Value [Chow-test] for $(\beta_1^V - \beta_0^V) - (\beta_1^N - \beta_0^N)$			0.5489

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

the 1% level.

However, as with the earlier patterns, there do not appear to be any differences between beliefs about events where subjects have a personal or financial stake versus those where subjects have no stake, and Table 1 shows that I cannot reject equality of any of the three coefficients using Chow tests. Additionally, I cannot reject the hypothesis that $\beta_1^V - \beta_0^V = \beta_1^N - \beta_0^N$. Thus I do not observe any evidence of psychological bias. Whether news is a source of good/bad information or simply neutral information, does not appear to alter how beliefs are updated.

At first glance the results in Table 1 seem at odds with the previous section: posteriors are well approximated by Bayes' rule, yet the framework rejects the Bayesian benchmark and finds significant asymmetry. Similarly contradictory patterns can be seen in Appendix Figure D1, which plots average updating behavior by type of signal received. Asymmetry is not observed for the aggregate data, and is only slightly visible when observing moderate priors.

There are two important factors in reconciling these patterns with the earlier results. First, the response to negative signals is only slightly below the Bayesian prediction, and negative signals represent the majority of signals received. Second is that the framework does not assume unitary weighting of the log odds ratio of prior beliefs. In particular, the findings that $\delta < 1$ and $\beta_0 > \beta_1$ effectively operate in opposite directions in the aggregate data. The reason is that $\delta < 1$ implies that individuals update as if priors were closer to one-half than they actually are. Because approximately two-thirds of the data involve priors less than one-half, this has the overall effect of biasing posteriors upwards, and hence reducing the appearance of this asymmetry in the raw data. The finding that $\delta < 1$ was also found by Holt and Smith (2009), there identified as a cognitive bias in general information processing.

Finally, similar to the previous section, there are no significant differences in updating behavior across the financial stake conditions, including the varying payments for accuracy, presented in Appendix Table B1. Regarding the accuracy payments, this is relevant to studies of the effects of stakes on behavior in lab experiments, especially regarding belief elicitation. It suggests that paying subjects more for accurate beliefs may have little effect on belief updating.[22] In the next section I use the same analysis to investigate wether there are deviations at the event level.

4.3 Information Processing By Event

This section examines whether deviations from Bayesian updating are driven by specific events. Previous evidence has found that individuals update asymmetrically when provided information on their performance on a test, over-weighting good news relative to bad as in Mobius et al (2014) and Eil and Rao (2011), or the opposite asymmetry as in Ertac (2011). I thus focus attention on the quiz (self) event: whether an individual believes they scored in the top 15% on an ego relevant quiz.[23]

[22] Combined with Coutts (2015), which showed that increasing accuracy payments can lead to more biased prior formation, the implications are that ideal incentive payments may be relatively low.

[23] The signal structure and elicitation procedure is comparable with Mobius et al (2014) and Buser et al (2018), who examined beliefs of subjects about scoring in the top 50%, rather than 15%.

Note that in studying updating behavior in one context, it is critical to have an appropriate counterfactual comparison. Deviations in a given context do provide evidence against Bayes' rule, but their attribution to a particular bias or context is only valid if one can rule out that these deviations occur in other contexts. Use of an adequate control group is standard for extrapolating that deviations do not occur across other contexts. This requires that the control group is suitably defined, and has sufficient statistical power to rule out deviations of interest. As I will discuss further, defining a suitable control group for updating on the quiz is not necessarily straightforward, and may explain why few studies in this literature are able to satisfy this requirement.

As an initial control group for the quiz event in this experiment, 30% of subjects did not update about their own performance, but instead updated about the performance of a randomly selected anonymous individual in the lab. This is an intuitive control, yet one issue is that observed prior beliefs about one's own performance tend to be greater than those about another's performance. This is problematic, as in Appendix Figure E1, I present some evidence that updating differs for different values of the prior.

The existence of evidence across contexts of this experiment provide additional comparison groups. There is sufficient variation across the different events and financial treatments such that differences in updating that only appear in one context would be strong evidence that such patterns are indeed context specific. On the other hand, similar differences in updating across contexts would suggest that deviations from Bayes' rule may reflect more general cognitive biases.

Table 2: Updating Beliefs Within Events

Dependent Variable: Logit Posterior Belief					
Regressor	(1) Easy Dice	(2) Hard Dice	(3) Weather	(4) Quiz (S)	(5) Quiz (O)
δ	0.894***	0.872***	0.928***	0.952**	0.912**
	(0.028)	(0.022)	(0.024)	(0.022)	(0.035)
β_1	0.476***	0.404***	0.684***	0.590***	0.709**
	(0.090)	(0.062)	(0.061)	(0.054)	(0.118)
β_0	0.821*	0.886	0.818***	0.834***	0.732***
	(0.099)	(0.080)	(0.047)	(0.060)	(0.099)
P-Value ($\delta = 1$)	0.0002	0.0000	0.0033	0.0259	0.0130
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0000	0.0000	0.0157
P-Value ($\beta_0 = 1$)	0.0719	0.1538	0.0001	0.0063	0.0082
Diff ($\beta_1 - \beta_0$)	−0.345	−0.482	−0.133	−0.244	−0.023
P-Value ($\beta_1 = \beta_0$)	0.0398	0.0000	0.0663	0.0005	0.8942
R^2	0.73	0.77	0.75	0.84	0.79
Observations	836	841	871	565	247

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

Table 2 presents the results of all rounds of updating corresponding to each of the four different events, splitting the quiz event into self or other performance. Examining updating across domains, there are some differences in updating behavior, yet these differences do not appear to fit a consistent pattern.

The most suggestive result is found comparing Columns 4 and 5, which respectively examine updating for own versus other performance on the quiz. Regarding own performance, there is asymmetric under-weighting of good news relative to bad news, significant at the 1% level. Yet for the comparison group, other performance, there is no significant asymmetry in updating.

This result is the opposite asymmetry of that found in Mobius et al (2014), but is consistent with evidence from Ertac (2011). Yet, this result is greatly undermined when comparing estimates across other events. Even larger asymmetries are found for the two dice events, while a slightly smaller asymmetry can be seen for the weather event. Since these events involve outcomes that subjects have no personal stake in (and there is no difference for financial stakes), it is apparent that the differences in updating between own and other performance may be driven by factors that affect information processing more broadly, i.e. general cognitive biases, rather than a specific psychological bias. Importantly, these conclusions do not hinge on particular values of the prior, as even moderate values of the prior are associated with the same asymmetry.[24]

Overall, I am unable to reject the hypothesis that the response to signals for the quiz (self) event is the same as all other events pooled, whether aggregated or partitioned into good/bad news versus just news subsamples. The evidence suggests that the observed deviations from Bayes' rule are not generated by differential responses to good or bad news about performance, as they are present across other events that don't involve performance nor do they involve good or bad news. This suggests caution in interpreting differences in updating patterns for a specific event as evidence of ego relevant psychological bias, without examining for the presence of cognitive bias in other neutral contexts.

4.4 Investigating Signal Structure

The previous section found negative asymmetry across most contexts, independent of whether signals were valenced. This result fits in a literature that has found evidence of positive asymmetry regarding good news versus bad news, as in Mobius et al (2014) and Eil and Rao (2011), no asymmetry as in Buser et al (2018) and Barron (2016), and finally negative asymmetry found by Ertac (2011).[25] As discussed, asymmetry in the framework interacts in important ways with the weight on the prior log odds ratio, δ. Thus, the framework itself, used also in Mobius et al (2014), Buser et al (2018), and Barron (2016), can generate different interpretations of the data. Based on the results, I now consider *ex-post* a further explanation that may account for these mixed results: signal structure.

One difference in the present paper compared with related studies is that signals are skewed in the negative direction, due to events with average probabilities less than one-half. Table E5 examines how subjects update in the last round given the sequence of signals they faced. Subjects could receive three affirmative signals, two affirmative and one negative, one affirmative and two negative, or three negative signals. One issue is that for the Quiz (self) event, the distribution of signals faced may depend on ability as higher scoring individuals are likely to receive more affirmative signals. If anything, excluding this event results in even more pronounced patterns than presented here, as Appendix Table E5 shows.

Contrary to the Bayesian prediction, there are clear differences in updating given different sequences of signals. Examining Columns 1 to 4 in Table E5, there is a pattern of under-weighting signals that are received less often. This also means that when subjects receive two affirmative and one negative signal, the asymmetry is reversed, as affirmative signals receive more weight than negative signals. This is important, because given the unlikely nature of many of the events in this experiment, the distribution of signals is more

[24] Appendix Table E2 presents these results for moderate values of the prior.

[25] Of mention is that Buser et al (2018) find evidence of positive asymmetry when considering updating mistakes in the "wrong" direction.

heavily weighted towards negative signals.

A striking pattern is found comparing Columns 5 and 6, which compares individuals who received exactly the same sequence of signals in the first two rounds (one affirmative, and one negative), but only differed in the order these signals were received.[26] These show that there is more negative asymmetry when the first signal is negative, rather than affirmative. This is surprising in light of both Bayes' rule, as well as considering other known cognitive biases, neither of which can explain this pattern. Appendix Table E6 presents additional robustness checks for this finding, showing that individuals are asymmetric in the positive direction after initially receiving an affirmative signal, and in the negative direction after initially receiving a negative signal.

An implication of these findings is that the observed negative asymmetry may in part be accounted for by the bias towards negative signals. While not conclusive, as the role of differently sized priors and representativeness bias likely play a role, these patterns hint at a different type of cognitive bias, undetected by previous work which traditionally has not been focused on finding asymmetry.[27] It resembles confirmation bias, see for example Rabin and Schrag (1999), but relates to confirmation on signals rather than on priors. This has potentially interesting implications, such as history dependence, where early sequences of signals may exert undue influence on posteriors. An interesting implication is that when individuals have opportunities to exit (e.g. a career or major) they may exit too quickly when facing negative signals early on, or too late when facing an earlier sequence of affirmative signals.

Moreover, this additional cognitive bias could help account for observed differences in updating in recent work, operating through differences in signal structure. While signals in Mobius et al (2014) and Eil and Rao (2011) are balanced between affirmative and negative on average, in Ertac (2011) they were less likely to be affirmative.[28] More generally, across these papers there is also variation in the size of the prior as well as differences across events themselves, suggesting that opposing findings may not be unfounded.

---

[26]Appendix Table E4 shows that updated beliefs after two rounds, i.e. the priors in the regression analysis, do not significantly differ through mean or distribution tests.

[27]Priors are correlated with the types of signals received. Similarly, representativeness bias could account for some of these findings, as when subjects are in the final round and have received either 1 affirmative, 2 negative or 2 affirmative, 1 negative, this matches the signal strength of two-thirds. Thus, one may expect an asymmetric response. Appendix Table E7 shows that the asymmetry in the data remains even if one considers only the first two updating rounds, where the representativeness heuristic cannot be employed. Finally, if subjects update using absolute updating heuristics, e.g. updating by a fixed number of percentage points, this could potentially generate data which look on average asymmetric given average priors less than one-half. This type of strategy can be ruled out by examining whether the asymmetry is reversed for priors greater than one-half. In Appendix Table E3, it can be seen that this is not the case. I thank an anonymous referee for pointing out this possible explanation.

[28]In her paper, top (bottom) was less likely than not top (bottom). Nonetheless, because a signal of top (bottom) would completely reveal the state, it is not clear that the same asymmetry should persist when updating among the remaining possible states.

Table 3: Updating Beliefs in Final Round By Distribution of Signals Received

	Dependent Variable: Logit Posterior Belief					
Regressor	(1) 0 '+' Signals	(2) 1 '+' Signal	(3) 2 '+' Signals	(4) 3 '+' Signals	(5) 1st '−'; 2nd '+'	(6) 1st '+'; 2nd '−'
δ	0.898***	0.881***	0.918**	0.982	0.864***	0.886***
	(0.031)	(0.025)	(0.033)	(0.068)	(0.034)	(0.033)
β_1		0.305***	0.863*	1.214	0.788**	0.828
		(0.099)	(0.079)	(0.152)	(0.105)	(0.127)
β_0	1.126	0.967	0.557***		1.105	0.834*
	(0.104)	(0.078)	(0.106)		(0.108)	(0.095)
P-Value ($\delta = 1$)	0.0011	0.0000	0.0156	0.7958	0.0001	0.0006
P-Value ($\beta_1 = 1$)		0.0000	0.0828	0.1634	0.0453	0.1776
P-Value ($\beta_0 = 1$)	0.2279	0.6791	0.0000		0.3323	0.0830
Diff ($\beta_1 - \beta_0$)		−0.662	0.306		−0.317	−0.006
P-Value ($\beta_1 = \beta_0$)		0.0000	0.0193		0.0439	0.9685
R^2	0.77	0.79	0.78	0.78	0.81	0.81
Observations	253	454	270	68	266	249

Analysis uses OLS regression. Columns (1)-(4): K '+ Signals' refers to K affirmative signals, out of a possible maximum of 3. Columns (5)-(6): Compares individuals who received exactly 1 affirmative and 1 negative signal, only differing in the order these signals were received. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

4.5 Conservatism, Ability, and Gender

In light of the results, I first briefly discuss the substantial conservatism, and next examine the relationship of conservatism and asymmetry with ability and gender. Appendix Table F2 examines only actively revised beliefs, and finds that conservatism is driven entirely by the 41% of non-updates.

These non-updates are not driven by a small subset of conservative individuals: only 9% of subjects update in all 12 rounds, and the median subject updates in 7 of 12 rounds.[29] Table 4 presents some reduced form estimates of the factors that correlate with active updates. An important factor appears to be the size of the prior, as subjects are less likely to update for more extreme values of the prior. Looking across events, subjects update less for the two dice events than for the other more subjective events, though this becomes

[29] Conservatism is correlated across events, within individuals, as found in Buser et al (2018). 30% of the variation in non-updates can be explained by individual fixed effects.

insignificant with individual fixed effects. Additionally, subjects update more frequently in later updating rounds, the probability of an update is approximately 5% greater after receiving the second signal, and 14% greater after the third. Financial payments in the experiment appear to have no effect.

Table 4: Correlates of Active Updating Decision

Dependent Variable: Active Update				
Regressor	(1)	(2)	(3)	(4)
Prior	1.596***	1.629***	1.576***	1.641***
	(0.136)	(0.124)	(0.154)	(0.138)
Prior^2	−1.678***	−1.595***	−1.695***	−1.628***
	(0.145)	(0.133)	(0.154)	(0.140)
Event = Hard Dice			−0.012	−0.016
			(0.022)	(0.023)
Event = Weather			0.068**	0.011
			(0.030)	(0.028)
Event = Quiz (self)			0.053*	0.033
			(0.030)	(0.028)
Event = Quiz (other)			0.100***	0.010
			(0.037)	(0.034)
Round 2			0.045***	0.047***
			(0.016)	(0.017)
Round 3			0.137***	0.138***
			(0.019)	(0.019)
$a = 10$			0.024	
			(0.038)	
$a = 20$			−0.011	
			(0.038)	
Stake = 80			−0.006	−0.023
			(0.019)	(0.019)
Male			0.168**	
			(0.066)	
Percentile Score			0.199***	
			(0.070)	
Male × Percentile Score			−0.236**	
			(0.113)	
Econ Major			−0.014	
			(0.041)	
Constant	0.347***		0.153***	
	(0.025)		(0.053)	
Individual Fixed Effects	NO	YES	NO	YES
R^2	0.06	0.37	0.10	0.38
Observations	3654	3654	3483	3654

Analysis uses OLS regression, dependent variable is binary for whether subject updated prior. Difference is significant from zero at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. Omitted event is Easy Dice. Excludes belief revisions in the opposite direction predicted by Bayes' rule.

Examining Column 3 of Table 4, I also examine the relationship between active updating and the percentile rank on the quiz. Subjects who rank one standard deviation higher on the quiz are 5.8 percentage points more likely to update, significant at the 1% level. Interestingly, the positive association between ability and updating propensity is entirely driven by women. Men who rank poorly on the quiz update more frequently than equal ability women, but scoring higher on the quiz increases the propensity to update for women, the opposite of men.

In Table 5 I examine how gender and ability affect updating behavior more generally. I do this by allowing heterogeneous response to signals by gender and percentile rank.[30] Column 1 examines only interactions with gender, the second examines ability, while the third interacts the two. Similar to Mobius et al (2014) and Ertac (2011) I find that women update more conservatively than men, though this is no longer significant when interacted with ability. I do not find any difference in asymmetry between men and women.[31]

There is also some weak evidence that ability is related to both conservatism and asymmetry, which appears in Column 2, but is no longer significant in Column 3. Subjects with higher ability appear to put more weight on affirmative signals, but not on negative signals. The implication is that those at the top of the ability distribution would no longer exhibit significant asymmetry, and would also be less conservative. This is in contrast to the results of Mobius et al (2014) who found that neither conservatism nor asymmetry were significantly correlated with cognitive ability.

Importantly, the pattern of greater conservatism exhibited by women seen in Column 1 is present across both valenced and neutral contexts. Thus while previous work from Ertac (2011) and Mobius et al (2014) theorized that female conservatism may be related to self-confidence, such conservatism is equally present in ego-irrelevant contexts.[32] Moreover, when comparing the deviations from Bayes' rule of final beliefs, women are only one-tenth of one percentage point further from Bayes' rule than men, a difference that is not statistically significant.

[30]One could also interact gender and ability with the weight on the log prior odds ratio, δ. I do not report these estimates, but interactions with δ are not significant.

[31]One interesting result is that unlike evidence from Mobius et al (2014) and Buser et al (2018) (see Barber and Odean (2001) for an earlier discussion), women are not less confident about their performance than men on the quiz.

[32]Note that if gender differences in updating persist across contexts, the same implications remain: high ability women could end up less confident. An important direction for future research is to understand the source of differences in information processing by gender.

Table 5: Updating Beliefs by Ability and Gender

Dependent Variable: Logit Posterior Belief			
Regressor	(1) Gender	(2) Ability	(3) Gender × Ability
δ	0.910***	0.914***	0.910***
	(0.010)	(0.009)	(0.010)
β_1	0.520***	0.433***	0.394***
	(0.047)	(0.076)	(0.103)
β_0	0.740***	0.776***	0.728***
	(0.048)	(0.071)	(0.097)
$\beta_1\times$ Male	0.141*		0.135
	(0.077)		(0.159)
$\beta_0\times$ Male	0.160**		0.198
	(0.065)		(0.142)
$\beta_1\times$ Percentile		0.315**	0.276
		(0.133)	(0.190)
$\beta_0\times$ Percentile		0.037	0.028
		(0.120)	(0.169)
$\beta_1\times$ Male × Percentile			−0.031
			(0.270)
$\beta_0\times$ Male × Percentile			−0.078
			(0.250)
R^2	0.84	0.84	0.84
Observations	3199	3360	3199

Analysis uses OLS regression. Difference is *significant from 0* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

5 Concluding Discussion

In this experiment I set out to examine whether differences exist in how people process information across varied contexts, focusing especially on response to valenced versus value neutral information. Recent evidence from economics and neuroscience suggests the possibility of additional psychological biases in updating when information contains good or bad news, beyond any cognitive biases that have been observed when information is value neutral.

Unlike some of the findings and claims of previous literature, I do not find evidence of asymmetric over-weighting of good versus bad news. In fact I find evidence of the opposite asymmetry, as in Ertac (2011), yet such patterns are similarly present comparing information processing across valenced or value neutral contexts. Related, while there are differences in updating by gender, as found in to previous studies, these patterns in updating are present across contexts. Thus updating appears similar, whether it involves news regarding events that directly affect subjects' wellbeing, or whether it involves only information about neutral events. This result suggests that asymmetric updating found in previous studies is not a universal property of updating in ego relevant contexts. Moreover, there is evidence that deviations from Bayes' rule may reflect more general cognitive biases, rather than psychological biases in processing good or bad news.

While these results differ from some of the previous literature, they do suggest a possible unifying feature. The distribution of signals received appears to affect how individuals update beliefs, a type of confirmation bias which previously went undocumented. Thus, differences in the direction of asymmetry, e.g. between, Mobius et al (2014) and Eil and Rao (2011), and Ertac (2011), might be explained by differences in signal content. As interest in asymmetry is recent, differences in updating across studies is not unfounded, and the paucity of work on asymmetry in value neutral contexts has hindered our ability to make relevant comparisons.[33]

Overall, while I find that many predictions of the Bayesian model are rejected in the pooled data, these discrepancies are small in magnitude. The average posterior belief is less than one percentage point away from the posterior predicted by Bayes' rule. After three rounds of updating, subjects are less than two percentage points away from the posterior that would have resulted had they used Bayes' rule all the way through, and the final posteriors of men and women are statistically indistinguishable.

Overall, this paper presents counter evidence to recent studies suggesting that psychologically plausible biases may arise in value relevant contexts, through the asymmetric processing of good news relative to bad news. Such evidence is necessary to discipline existing and future theoretical work on updating behavior, in order to better understand how individuals process information.

References

Ambuehl S, Li S (2018) Belief Updating and the Demand for Information. Games and Economic Behavior 109:21–39

Antoniou C, Harrison GW, Lau MI, Read D (2015) Subjective Bayesian beliefs. Journal of Risk and Uncertainty 50(1):35–54

Armantier O, Treich N (2013) Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. European Economic Review 62:17–40

Azrieli Y, Chambers CP, Healy PJ (2018) Incentives in Experiments: A Theoretical Analysis. Journal of Political Economy

[33] There are some differences in implementation between this experiment and that of previous work. In Mobius et al (2014) subjects participated in the experiment online, the relevant event was being in the top 50%, and signal strength differed slightly. On this last point, Ambuehl and Li (2018) suggest that subjects do not react strongly to signal precision.

Barber BM, Odean T (2001) Boys will be boys: Gender, overconfidence, and common stock investment. Quarterly Journal of Economics 116(1):261–292

Barron K (2016) Belief updating: Does the 'good-news, bad-news' asymmetry extend to purely financial domains? WZB Discussion Paper (October)

Benabou R, Tirole J (2002) Self-Confidence and Personal Motivation. The Quarterly Journal of Economics 117(3):871–915

Blanco M, Engelmann D, Koch AK, Normann HT (2010) Belief elicitation in experiments: is there a hedging problem? Experimental Economics 13(4):412–438

Buser T, Gerhards L, Weele JVD (2018) Responsiveness to Feedback as a Personal Trait. Journal of Risk and Uncertainty

Camerer CF (1987) Do Biases in Probability Judgment Matter in Markets? Experimental Evidence. The American Economic Review 77(5):981–997

Camerer CF (1995) Individual Decision Making. In: Kagel JH, Roth AE (eds) The Handbook of Experimental Economics, Princeton University Press, Princeton, NJ, pp 587–703

Clark J, Friesen L (2009) Overconfidence in forecasts of own performance: An experimental study. Economic Journal 119(534):229–251

Coutts A (2015) Testing Models of Belief Bias: An Experiment. mimeo , `RePEc:pra:mprapa:67507`

Edwards W (1968) Conservatism in human information processing. In: Kleinmuntz B (ed) Formal Representation of Human Judgment, John Wiley & Sons Inc, p 51

Eil D, Rao JM (2011) The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself. American Economic Journal: Microeconomics 3(2):114–138

Ertac S (2011) Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback. Journal of Economic Behavior & Organization 80(3):532–545

Gilboa I, Schmeidler D (2001) A Theory of Case-Based Decisions. Cambridge University Press, Cambridge

Grether DM (1980) Bayes Rule as a Descriptive Model: The Representativeness Heuristic. The Quarterly Journal of Economics 95(3):537

Grether DM (1992) Testing bayes rule and the representativeness heuristic: Some experimental evidence. Journal of Economic Behavior & Organization 17(1):31–57

Grossman Z, Owens D (2012) An unlucky feeling: Overconfidence and noisy feedback. Journal of Economic Behavior & Organization 84(2):510–524

Healy PJ (2016) Explaining the BDM - Or Any Random Binary Choice Elicitation Mechanism - To Subjects. mimeo

Holt C, Smith AM (2009) An update on Bayesian updating. Journal of Economic Behavior & Organization 69(2):125–134

Kahneman D, Tversky A (1972) Subjective probability: A judgment of representativeness. Cognitive Psychology 3(3):430–454

Kahneman D, Tversky A (1973) On the psychology of prediction. Psychological Review 80(4):237–251

Karni E (2009) A Mechanism for Eliciting Probabilities. Econometrica 77(2):603–606

Karni E, Safra Z (1995) The impossibility of experimental elicitation of subjective probabilities. Theory and Decision 38(3):313–320

Karni E, Safra Z, Society TE (1987) Preference Reversal" and the Observability of Preferences by Experimental Methods. Econometrica 55(3):pp. 675–685

Kuhnen CM (2014) Asymmetric Learning from Financial Information. The Journal of Finance LXX(5):2029–2062

Landier A (2000) Wishful thinking a model of optimal reality denial. mimeo pp 1–25

Machina MJ, Schmeidler D (1992) A More Robust Definition of Subjective Probability. Econometrica 60(4):745

Mayraz G (2014) Priors and Desires—A Model of Optimism, Pessimism, and Cognitive Dissonance. mimeo pp 1–47

Mobius M, Niederle M, Rosenblat T, Niehaus P (2014) Managing Self-Confidence : Theory and Experimental Evidence. mimeo

Offerman T, Sonnemans J, Van De Kuilen G, Wakker PP (2009) A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes. Review of Economic Studies 76(4):1461–1489

Rabin M, Schrag JL (1999) First Impressions Matter: A Model of Confirmatory Bias. The Quarterly Journal of Economics 114(1):37–82

Scheier ME, Carver CS (1987) Dispositional Optimism and Physical Well-Being: The Influence of Generalized Outcome Expectancies on Health. Journal of Personality 55(2):169–210

Schlag KH, Tremewan J, van der Weele JJ (2015) A penny for your thoughts: a survey of methods for eliciting beliefs. Experimental Economics 18(3):457–490

Schotter A, Trevino I (2014) Belief Elicitation in the Lab. Annual Review of Economics 6:forthcoming

Shah P, Harris AJL, Bird G, Catmur C, Hahn U (2016) A pessimistic view of optimistic belief updating. Cognitive Psychology 90:71–127

Sharot T, Korn C, Dolan R (2011) How unrealistic optimism is maintained in the face of reality. Nature Neuroscience 14(11):1475–1479

Sharot T, Kanai R, Marston D, Korn CW, Rees G, Dolan RJ (2012) Selectively altering belief formation in the human brain. Proceedings of the National Academy of Sciences 109(42):17,058–17,062

Wiswall M, Zafar B (2015) How Do College Students Respond to Public Information about Earnings? Journal of Human Capital 9(2):117–169

Yariv L (2005) I'll See It When I Believe It - A Simple Model of Cognitive Consistency. mimeo

Appendix

A Hedge Proof Design Details

Payoffs are determined in the following way, also described in Figure A1, and earlier utilized by Blanco et al (2010).[34] In order to ensure that individuals have no incentive to hedge their probability reports, the world is partitioned into two disjoint states, the accuracy state and the prize state.[35]

With probability 0.5 the individual is paid solely according to her reported belief $\tilde{\pi}$ about whether event E occurred using the incentive compatible lottery method to elicit beliefs with an accuracy payment of $a > 0$ (accuracy state).

In the other state occurring with probability 0.5, the individual receives a guaranteed payment $\bar{a} \geq a$[36] and receives an additional \$80 if E occurs, but receives nothing extra if E does not occur (prize state). Her report of $\tilde{\pi}$ is no longer relevant in this prize state.

To be clear, two types of hedging are of concern in this experiment. The first is is hedging within the accuracy state, which is solved through use of the lottery method. The second is hedging across accuracy and prize states, which is solved through partitioning. In isolation, the lottery method is incentive compatible under the relatively weak assumption of probabilistic sophistication. However, the experiment design introduces further elements of randomization through the partitioning of the accuracy and prize states, *and* through randomly selection one decision for payment.

As Blanco et al (2010) note, partitioning the world into an accuracy and a prize state is akin to the standard procedure of introducing a new lottery and randomly selecting one lottery for payment.[37] Thus incentive compatibility in the broader experiment design holds for the class of preferences where payment is made by random selection of one task (or lottery). This is true when one assumes a statewise monotonicity condition, see Azrieli et al (2018). This condition is equivalent to saying that subjects never choose dominated gambles, independent of other states.[38]

Two further issues on incentive compatibility deserve some mention here. First, with

[34] It was also independently suggested to me by Christopher Woolnough, who I credit for the design in this paper.

[35] Hedging will be present whenever utility is not linear, for example with a concave utility function and a positive stake in an event an individual would prefer to report a lower than truthful $\tilde{\pi}$, since this will smooth consumption over the different states of the world. Karni and Safra (1995) show that without this partition, no elicitation procedure exists that induces truthful reporting, a fact that is sometimes overlooked in the experimental literature; see Armantier and Treich (2013).

[36] The payment of $\bar{a}$ is to ensure that the prize state is preferred to the accuracy state, required for an earlier theoretical extension; it is not necessary for any of the analysis.

[37] In this case the other lottery is degenerate, as the individual does not make an active choice, but simply has the opportunity to receive a payment.

[38] The assumption of monotonicity is not completely innocuous. Assuming further that subjects reduce compound lotteries, it implies that subject's preferences must conform with expected utility, as detailed in Azrieli et al (2018).

financial incentives it is possible to disentangle accuracy and prize payments. However, if subjects gain utility from beliefs about their ability, evidently the experiment is not able to create an analogous partition. Thus there may be distortions in elicited beliefs about performance on the quiz - note however that most of the results in this paper do not hinge on the inclusion of the quiz (Self) event.

Second, there is a potential concern which arises from paying only for the accuracy *or* the prize state, but never both. The implication is that a subject in the experiment knows with certainty that whenever her belief report is relevant, she will not have an opportunity to win the prize. Or vice-versa, whenever she has a chance to win the prize, her belief report is not relevant.[39]

In this case, the procedure would correctly capture the subject's belief about the event occurring in the event the prize is irrelevant, but the counterfactual belief would not be observed. If such belief patterns are occurring then it remains possible that subjects may hold biased beliefs, but the experiment is not designed to capture them. Under the assumption of monotonicity above, this does not cause an issue as subjects are assumed to form consistent beliefs about an event, which do not depend on the state.[40]

[39]I thank an anonymous referee for bringing this concern to my attention.

[40]Barron (2016) elicits beliefs about events with financial stakes without separating prize and accuracy payments (but addressing the hedging problem retroactively, using the "truth serum" of Offerman et al (2009)). He does not find evidence of differential updating patterns with financial stakes, which suggests that this may not be a concern.

Figure A1: Illustration of Hedge Proof Design

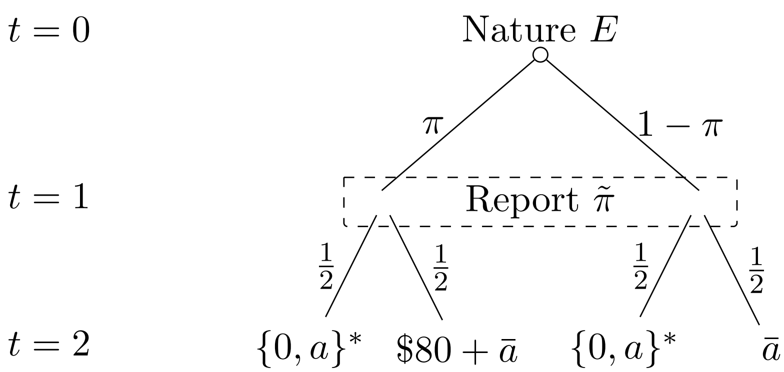

*In the accuracy state the payoff is either 0 or a, depending on the reported belief $\tilde{\pi}$ and whether E occurred, according to the lottery method.

Nature determines outcome of binary event E. Individual submits report $\tilde{\pi}$ without knowing outcome of E, and payoff is determined according to the lottery method elicitation procedure.

B Updating Framework: By Event/Stake/Accuracy Payment

Here I replicate the primary analysis found in Table 1, looking at each of the financial stake and accuracy payment conditions separately. As can be seen in Table B1, there is no clear pattern that emerges within either the accuracy payment or within the financial stake conditions respectively. A formal statistical test confirms that I cannot reject equality between the \$0 and \$80 financial stake conditions, nor between the \$3, \$10, and \$20 accuracy payment conditions. This analysis suggests that different payments for accuracy do not alter updating behavior. Similarly, holding a large financial stake in an event does not alter updating behavior relative to holding no stake.

Table B1: Updating Beliefs for All Events: By Accuracy Payment and Financial Stake

	Dependent Variable: Logit Posterior Belief					
Regressor	(1) Stake = 0	(2) Stake = 80	(3) Acc = 3	(4) Acc = 10	(5) Acc = 20	(6) Total
δ	0.910***	0.918***	0.920***	0.922***	0.898***	0.914***
	(0.012)	(0.014)	(0.017)	(0.014)	(0.016)	(0.009)
β_1	0.587***	0.588***	0.560***	0.662***	0.540***	0.588***
	(0.045)	(0.043)	(0.054)	(0.063)	(0.059)	(0.034)
β_0	0.807***	0.780***	0.774***	0.749***	0.861*	0.793***
	(0.047)	(0.047)	(0.066)	(0.060)	(0.074)	(0.038)
P-Value ($\delta = 1$)	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
P-Value ($\beta_0 = 1$)	0.0001	0.0000	0.0008	0.0001	0.0616	0.0000
Diff ($\beta_1 - \beta_0$)	−0.220	−0.192	−0.214	−0.086	−0.321	−0.205
P-Value ($\beta_1 = \beta_0$)	0.0001	0.0011	0.0042	0.2112	0.0000	0.0000
R^2	0.83	0.84	0.84	0.84	0.82	0.84
Observations	1704	1656	1128	1143	1089	3360

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

In Table B2 I replicate the primary analysis found in Table 2, but excluding any observations where an individual had a financial stake of $80. There do not appear to be any consistent differences in this subsample. In the final column of Table B2 I use the same sampling procedure as Mobius et al (2014), in order to provide a more comparable estimation to their study for the Quiz event. One can see that the sampling procedure does not significantly alter the pattern of observed results.

Table B2: Updating Beliefs Within Events: No Financial Stake Only

Dependent Variable: Logit Posterior Belief						
Regressor	(1) Easy Dice	(2) Hard Dice	(3) Weather	(4) Quiz (S)	(5) Quiz (O)	(6) Quiz (M. et al)
δ	0.839***	0.897***	0.909***	0.924**	0.894*	0.918**
	(0.049)	(0.028)	(0.027)	(0.030)	(0.055)	(0.035)
β_1	0.317***	0.430***	0.683***	0.616***	0.816	0.714***
	(0.146)	(0.092)	(0.089)	(0.078)	(0.200)	(0.090)
β_0	1.073	0.815*	0.783***	0.799**	0.778	0.917
	(0.171)	(0.109)	(0.067)	(0.087)	(0.176)	(0.099)
P-Value ($\delta = 1$)	0.0013	0.0004	0.0009	0.0138	0.0603	0.0226
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0005	0.0000	0.3623	0.0021
P-Value ($\beta_0 = 1$)	0.6711	0.0906	0.0014	0.0224	0.2144	0.4049
Diff ($\beta_1 - \beta_0$)	−0.755	−0.384	−0.100	−0.184	0.038	−0.204
P-Value ($\beta_1 = \beta_0$)	0.0095	0.0136	0.3791	0.0472	0.9040	0.0521
R^2	0.66	0.77	0.73	0.83	0.79	0.84
Observations	435	421	447	294	107	225

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant. Includes only updated beliefs about events where individuals did not hold any additional financial stake in the outcome.

C Additional Tests of Bayes' Rule: Invariance, Sufficiency, and Stability

In this section I investigate three additional properties that are satisfied when updated beliefs follow Bayes' rule. First, the structure of Bayes' rule implies a sufficiency condition, that priors are sufficient statistics for all the information contained in past signals. In other words, after controlling for prior beliefs, lagged information does not significantly predict posterior beliefs. To examine whether updating behavior can be shown to satisfy the sufficiency condition I follow Mobius et al (2014) and include lagged signals as independent variables. Table C1 shows the regressions that include these lagged signals, using only actively revised beliefs. There is some evidence that overall, the updating process may not satisfy the sufficiency condition, as the first signal received has a significant effect on belief updating in round 3.[41]

The next property Bayes' rule satisfies is stability: that updating remains stable across time. Looking across the three updating rounds in Table C2, there appear to be differences. Overall, I can reject equality across rounds 1 to 3 for δ, β_1, β_0 at conventional levels. This

[41] While Mobius et al (2014) do not reject sufficiency, it is worth noting that the ratio of the values of coefficients on lagged signals to current signals is of the same magnitude.

Table C1: Examining Sufficiency

Dependent Variable: Logit Posterior Belief		
Regressor	(1) Round 2	(2) Round 3
δ	0.890***	0.880***
	(0.027)	(0.023)
β_1	1.030***	1.247***
	(0.065)	(0.074)
β_0	1.287***	1.347***
	(0.064)	(0.066)
β_{t-1}	0.052	0.048
	(0.045)	(0.042)
β_{t-2}		0.164***
		(0.042)
R^2	0.82	0.82
Observations	640	670

Analysis uses OLS regression. Difference is *significant from zero* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant. The sample is restricted to include only subjects who actively revised their beliefs in the direction predicted by Bayes' rule. β_{t-k} refers to the k^{th} lagged signal.

provides some evidence that updating is not stable across rounds. Finally, the invariance property is said to hold when $\delta = 1$, that is the change in logit beliefs depends only on past signals. $\delta = 1$ is rejected in the data at the 1% level. However, despite these three conditions not being met in the data, it is important to note that the magnitude of these deviations is reasonably small, in the sense that the resulting posteriors are very close to their Bayesian counterparts.[42]

D Updating by Sequence of Signals Observed

Figure D1 presents an aggregate view of asymmetry, by plotting average posteriors in response to different sequences of observed signals, for both the aggregate data, and for moderate priors between 0.4 and 0.6. One can note that the asymmetry in the framework, observed in Table 1, is not visibly present in the aggregate data. The reason is that the weight on the log odds ratio of prior beliefs is not unity. $\delta < 1$ manifests itself as over-

[42]As in Mobius et al (2014) a concern is that β_1 and β_0 are functions of prior beliefs, but that effects cancel out to give a coefficient of δ closer to 1. To examine if this is a potential issue I check whether there are significant interaction effects between receiving affirmative signals, and the prior. These interactions are never significant at any reasonable significance level, indicating that this is not a problem for the data.

Table C2: Examining Stability

Dependent Variable: Logit Posterior Belief				
Regressor	Round 1	Round 2	Round 3	All Rounds
δ	0.884***	0.926***	0.935***	0.914***
	(0.014)	(0.017)	(0.016)	(0.009)
β_1	0.468***	0.537***	0.800***	0.588***
	(0.047)	(0.046)	(0.062)	(0.034)
β_0	0.687***	0.788***	0.914	0.793***
	(0.045)	(0.053)	(0.055)	(0.038)
P-Value ($\delta = 1$)	0.0000	0.0000	0.0001	0.0000
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0015	0.0000
P-Value ($\beta_0 = 1$)	0.0000	0.0001	0.1205	0.0000
Diff ($\beta_1 - \beta_0$)	-0.219	-0.250	-0.114	-0.205
P-Value ($\beta_1 = \beta_0$)	0.0002	0.0003	0.1592	0.0000
R^2	0.85	0.84	0.83	0.84
Observations	1180	1135	1045	3360
P-Value [Chow-test] for δ (Rounds 1-3)				0.0260
P-Value [Chow-test] for β_1 (Rounds 1-3)				0.0000
P-Value [Chow-test] for β_0 (Rounds 1-3)				0.0005

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

weighting of probabilities for priors < 0.5, the majority of the data of this experiment. This masks the asymmetry in the framework, since it results in an upward shift of posterior beliefs, independent of the types of signals received.[43]

[43] There is some visible asymmetry when restricting priors to be moderate, between 0.4 and 0.6. This is intuitive, as the distortionary weighting of δ is weakest around 0.5, and hence posteriors are more closely matching the response to signals observed in the framework.

Figure D1: Updating in Response to Observed Signals

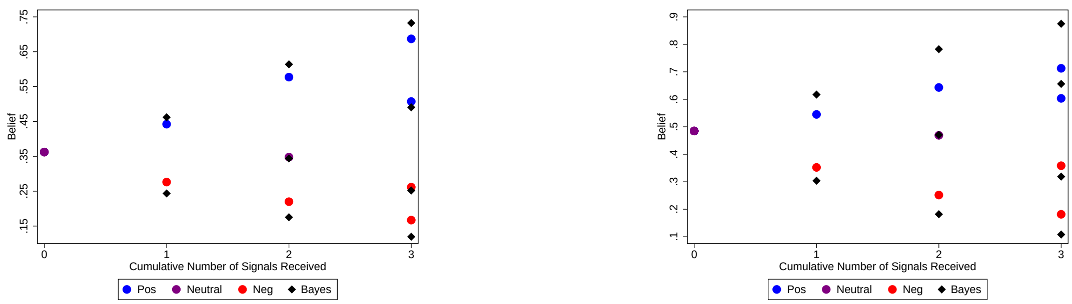

(a) All Data

(b) 40% < Prior < 60%

Average belief update following a sequence of cumulative signals (numbered on the horizontal axis), distinguishing cumulative signals in the positive direction (blue) from negative (red), as well as neutral (purple). For example, when the number of cumulative signals is 2, the possibilities are that a subject received 2 positive signals, 1 positive 1 negative, or 2 negative signals. The Bayesian benchmark is indicated by a black diamond.

E Robustness Checks

E.1 Different Values of the Prior

Of interest is to what extent the results in the paper could be explained by the fact that priors are on average lower than one-half.

Figure E1 presents the evolution of beliefs for different values of the prior (first reported beliefs). Updating appears conservative for low values of the prior, well calibrated for moderate values, and too responsive for high values of the prior. These patterns are suggestive that some of the differences in updating observed across events are driven by differences in average values of the prior, rather than differences in the events themselves. In particular, elicited priors for the two dice events and the quiz (other) event are significantly lower than those for the weather event and the quiz (self) event, and additionally exhibit substantially greater levels of conservatism.

Table E1 examines the primary results of Table 1, but restricting priors to lie between 0.4 and 0.6. From the table, one is able to see that the results are very similar. Negative signals continue to be weighted significantly more than positive signals, and one cannot reject that the difference in any of the parameters in the valenced contexts (good/bad news) are the same as those found in the neutral (just news) contexts. Analogously, Table E2 presents the results of Table 2 with the same prior restrictions, finding no differences in the patterns of updating.

Table E3 splits the sample into priors less than one half or greater than one-half. There, one can see similar asymmetries across the two subsamples. One finding of note is that the value of δ is very close to one in Column (2), when priors are greater than one-half. In fact, this has implications for interpreting differences between patterns in the empirical framework and the raw data. It implies that an over-weighting of priors less than one half occurs (since $\delta < 1$ in this case), but no corresponding under-weighting of priors greater than one-half is occurring (since δ is approximately 1 in this case).

Figure E1: Evolution of Beliefs By Value of the Prior

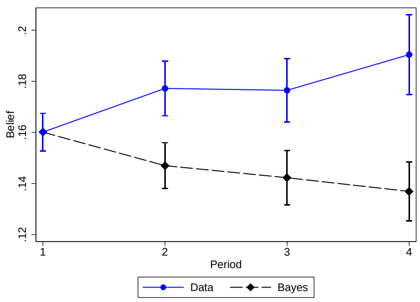

(a) Prior $\leq$ 40%

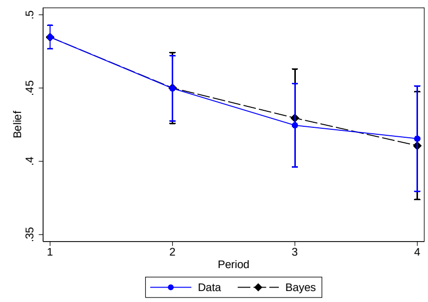

(b) 40% $\leq$ Prior $\leq$ 60%

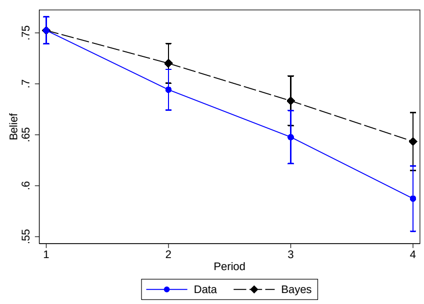

(c) 60% $\leq$ Prior

All events and all stake conditions, sample restricted to individuals updating with indicated prior. The path of beliefs starting from the prior (period 1), and after each sequential signal (periods 2 through 4). Average individual responses are the blue solid line, the Bayesian benchmark is marked as the black dashed line. Error bands represent 95% confidence intervals. $N = \{798, 185, 297\}$ average per round, respectively for (a)-(c).

Table E1: Updating Beliefs for All Events: 40% $\leq$ Prior $\leq$ 60%

Dependent Variable: Logit Posterior Belief			
Regressor	(1) Good/Bad News	(2) Just News	(3) All
δ^V	0.897		
	(0.091)		
β_1^V	0.586***		
	(0.085)		
β_0^V	0.777***		
	(0.078)		
δ^N		0.918	
		(0.055)	
β_1^N		0.576***	
		(0.102)	
β_0^N		0.755**	
		(0.107)	
δ			0.908*
			(0.052)
β_1			0.581***
			(0.067)
β_0			0.770***
			(0.066)
P-Value ($\delta = 1$)	0.2598	0.1444	0.0790
P-Value ($\beta_1 = 1$)	0.0000	0.0001	0.0000
P-Value ($\beta_0 = 1$)	0.0053	0.0255	0.0006
Diff ($\beta_1 - \beta_0$)	-0.191	-0.179	-0.188
P-Value ($\beta_1 = \beta_0$)	0.0219	0.2074	0.0118
R^2	0.59	0.65	0.61
Observations	297	183	480
P-Value [Chow-test] for $\delta^V = \delta^N$			0.8400
P-Value [Chow-test] for $\beta_1^V = \beta_1^N$			0.9367
P-Value [Chow-test] for $\beta_0^V = \beta_0^N$			0.8578
P-Value [Chow-test] for $(\beta_1^V - \beta_0^V) - (\beta_1^N - \beta_0^N)$			0.9403

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

Table E2: Updating Beliefs Within Events: 40% $\leq$ Prior $\leq$ 60%

Dependent Variable: Logit Posterior Belief					
Regressor	(1) Easy Dice	(2) Hard Dice	(3) Weather	(4) Quiz (S)	(5) Quiz (O)
δ	0.828**	0.832**	1.036	0.768	1.069
	(0.077)	(0.076)	(0.120)	(0.181)	(0.093)
β_1	0.311***	0.337***	0.740**	0.639***	0.552**
	(0.156)	(0.093)	(0.111)	(0.122)	(0.217)
β_0	0.934	0.886	0.782**	0.694***	0.505***
	(0.260)	(0.178)	(0.088)	(0.110)	(0.126)
P-Value ($\delta = 1$)	0.0366	0.0344	0.7676	0.2093	0.4682
P-Value ($\beta_1 = 1$)	0.0002	0.0000	0.0216	0.0052	0.0596
P-Value ($\beta_0 = 1$)	0.8019	0.5254	0.0159	0.0084	0.0018
Diff ($\beta_1 - \beta_0$)	−0.623	−0.549	−0.042	−0.055	0.047
P-Value ($\beta_1 = \beta_0$)	0.0674	0.0032	0.6916	0.6816	0.8545
R^2	0.67	0.62	0.61	0.47	0.80
Observations	64	81	189	111	35

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

Table E3: Priors Greater or Less than One Half

Dependent Variable: Logit Posterior Belief			
Regressor	(1) Prior $> \frac{1}{2}$	(2) Prior $< \frac{1}{2}$	(3) All
δ	0.899***	0.987	0.914***
	(0.016)	(0.036)	(0.009)
β_1	0.542***	0.466***	0.588***
	(0.049)	(0.064)	(0.034)
β_0	0.819***	0.888	0.793***
	(0.058)	(0.069)	(0.038)
P-Value ($\delta = 1$)	0.0000	0.7085	0.0000
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_0 = 1$)	0.0021	0.1072	0.0000
Diff ($\beta_1 - \beta_0$)	-0.277	-0.422	-0.205
P-Value ($\beta_1 = \beta_0$)	0.0011	0.0001	0.0000
R^2	0.69	0.65	0.84
Observations	2253	927	3360

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant. First column includes updates in direction predicted by Bayes' rule. Second column replaces boundary probabilities with 0.01 or 0.99 respectively. Third column is entire sample.

E.2 Additional Results for Investigating Signal Structure

Table E5 presents the analogous analysis to Table 3 in the paper, but omitting the Quiz (Self) event, since signals regarding the quiz event depend on ability, a potential confound. Regarding Columns 1 to 4, the results are by and large unchanged. Regarding Columns 5 and 6, if anything, the results present even stronger evidence of differential asymmetry between those who received exactly the same sequence of signals (1 affirmative and 1 negative), but only differed in the order these were received. In Column 5 the negative asymmetry is significant at the 5% level, while in Column 6 the positive asymmetry is not significant at conventional levels. However, the difference in the asymmetry is statistically significant at the 5% level (Chow Test).

The result that levels of asymmetry in Columns 5 and 6 are different, solely based on the *order* of signals received is highly surprising. This result is not driven by differences in the average or even the distribution of prior beliefs for these individuals. Table E4 presents tests of equality for the prior beliefs used in Tables 3 and E5 (note these are updated beliefs after receiving two rounds of signals). From these tests, one can see that prior beliefs are quite similar, as one would expect given individuals who received identical signals in the past. Excluding the quiz, in fact leads to slightly improved balance across the two groups.

Table E4: Comparing Beliefs for Individuals in Columns 5 and 6 in Tables 3 and E5

Equality tests of prior beliefs		1st '−'; 2nd '+'	1st '+'; 2nd '−'	Difference
Incl. Quiz (Table 3)	Mean	0.346	.3733	-0.027
	Median	0.250	0.260	-0.010
	Std. Dev.	0.274	0.300	
	Observations	289	270	
	Wilcoxon rank-sum (p-value)			0.450
	Kolmogorov-Smirnov (p-value)			0.365
Excl. Quiz (Table E5)	Mean	0.323	.334	-0.011
	Median	0.250	0.200	-0.050
	Std. Dev.	0.266	0.290	
	Observations	243	215	
	Wilcoxon rank-sum (p-value)			0.880
	Kolmogorov-Smirnov (p-value)			0.841

Table E6 presents additional specifications intended to examine the observed bias related to signal structure discussed in Section 4.4. Columns 1 and 2 present updating in the second round, after individuals had received two signals in total. It separates those who received a negative ($-$) signal as their previous (first) signal (Column 1), with those who received previously an affirmative (+) signal (Column 2). Columns 3 and 4 present analogous regressions for updating in the third and final round, given the previous (second) signal.

Confirming earlier observed patterns, and contrary to the Bayesian prediction, Table E6 shows a significant negative asymmetry in Column 1 (following a negative signal), and a positive (though not significant) asymmetry in Column 2 (following an affirmative signal). In Columns 3 and 4 the asymmetry is negative in both cases, though it is worth noting that the difference in asymmetry between the two regressions is of similar magnitude.[44]

Table E5: Updating Beliefs in Final Round By Distribution of Signals Received (Excluding Quiz (self))

Dependent Variable: Logit Posterior Belief						
Regressor	(1) 0 '+' Signals	(2) 1 '+' Signal	(3) 2 '+' Signals	(4) 3 '+' Signals	(5) 1st '−'; 2nd '+'	(6) 1st '+'; 2nd '−'
δ	0.899***	0.890***	0.915**	0.988	0.860***	0.906***
	(0.033)	(0.025)	(0.038)	(0.077)	(0.035)	(0.036)
β_1		0.323***	0.903	1.244	0.750**	0.973
		(0.103)	(0.090)	(0.171)	(0.117)	(0.158)
β_0	1.106	0.920	0.659***		1.090	0.757**
	(0.118)	(0.081)	(0.124)		(0.113)	(0.102)
P-Value ($\delta = 1$)	0.0023	0.0000	0.0261	0.8820	0.0001	0.0097
P-Value ($\beta_1 = 1$)		0.0000	0.2817	0.1588	0.0339	0.8644
P-Value ($\beta_0 = 1$)	0.3719	0.3273	0.0064		0.4264	0.0183
Diff ($\beta_1 - \beta_0$)		-0.597	0.244		-0.340	0.216
P-Value ($\beta_1 = \beta_0$)		0.0001	0.1053		0.0439	0.2502
R^2	0.77	0.81	0.76	0.76	0.82	0.81
Observations	206	380	220	60	225	199

Analysis uses OLS regression. Columns (1)-(4): K '+ Signals' refers to K affirmative signals, out of a possible maximum of 3. Columns (5)-(6): Compares individuals who received exactly 1 affirmative and 1 negative signal, only differing in the order these signals were received. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

[44] Given the patterns observed here and in Section 4.4, more negative asymmetry is to be expected in the third rather than second round, as the average proportion of negative signals received is greater.

Table E6: Updating Beliefs by Sequences of Signals Received

	Dependent Variable: Logit Posterior Belief			
	After 2nd Signal (Round 2)		After 3rd Signal (Round 3)	
Regressor	(1) 1st Signal '−'	(2) 1st Signal '+'	(3) 2nd Signal '−'	(4) 2nd Signal '+'
δ	0.910***	0.913***	0.890***	0.878***
	(0.019)	(0.023)	(0.026)	(0.025)
β_1	0.372***	0.746***	0.328***	0.811**
	(0.053)	(0.082)	(0.091)	(0.083)
β_0	0.887*	0.679***	1.142	0.971
	(0.060)	(0.071)	(0.094)	(0.079)
P-Value ($\delta = 1$)	0.0000	0.0002	0.0000	0.0000
P-Value ($\beta_1 = 1$)	0.0000	0.0022	0.0000	0.0234
P-Value ($\beta_0 = 1$)	0.0590	0.0000	0.1332	0.7139
Diff ($\beta_1 - \beta_0$)	−0.514	0.067	−0.814	−0.160
P-Value ($\beta_1 = \beta_0$)	0.0000	0.5398	0.0000	0.1499
R^2	0.83	0.83	0.78	0.81
Observations	700	435	377	515

Analysis uses OLS regression. '+' refers to affirmative signal, '−' to negative. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

E.3 Excluding the Last Round

Table E7 presents the analogous analysis to Table 1 in the main paper, excluding the last round. The motivation for excluding the last round of updating is to understand how much of the asymmetry in the data could be explained by representativeness bias. In the final round, subjects have potentially observed one of the two sets of "representative" signal sequences - i.e. sequences that exactly match the strength of signals. Two negative and one affirmative, exactly matches the expected number of signals for an event that did not occur, while two affirmative and one negative exactly matches the expected number of signals for an event that did occur. The representativeness bias would be to bias the posterior towards 0 in the first case, and towards 1 in the second case. In the framework this could be manifested as an exaggerated response to signals that go in the majority direction, and a conservative response to signals that go against the majority. By eliminating the final round, subjects cannot make use of the representativeness heuristic.

From Table E7 it is possible to see that the observed negative asymmetry persists in earlier updating rounds. Thus, while representativeness bias may play a role, it cannot explain the negative asymmetry observed in the data.

Table E7: Updating Beliefs for All Events: Rounds 1-2 only

Dependent Variable: Logit Posterior Belief			
Regressor	(1) Good/Bad News	(2) Just News	(3) All
δ^V	0.910***		
	(0.014)		
β_1^V	0.528***		
	(0.042)		
β_0^V	0.722***		
	(0.045)		
δ^N		0.895***	
		(0.017)	
β_1^N		0.454***	
		(0.050)	
β_0^N		0.762***	
		(0.060)	
δ			0.904***
			(0.011)
β_1			0.500***
			(0.035)
β_0			0.736***
			(0.041)
P-Value ($\delta = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_0 = 1$)	0.0000	0.0001	0.0000
Diff ($\beta_1 - \beta_0$)	-0.194	-0.308	-0.236
P-Value ($\beta_1 = \beta_0$)	0.0002	0.0000	0.0000
R^2	0.85	0.83	0.84
Observations	1343	972	2315
P-Value [Chow-test] for $\delta^V = \delta^N$			0.4444
P-Value [Chow-test] for $\beta_1^V = \beta_1^N$			0.2195
P-Value [Chow-test] for $\beta_0^V = \beta_0^N$			0.5313
P-Value [Chow-test] for $(\beta_1^V - \beta_0^V) - (\beta_1^N - \beta_0^N)$			0.1709

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

F Sampling Robustness Checks

F.1 Sample Restrictions

Table F1 examines the impact of how restricting the sample alters updating estimates in the main framework. The first column presents the main analysis (Column 3 in Table 1), but includes observations where belief updates go in the opposite direction that Bayes' rule predicts. The second column replaces boundary observations of 0 or 1 with 0.01 or 0.99 respectively. In Table 1 these were dropped. Finally the third column also truncates boundary observations, and includes updates in the wrong direction. The third column thus presents the full data, with no exclusions.

Table F1: Relaxing Sample Restrictions and Full Sample

Dependent Variable: Logit Posterior Belief			
Regressor	(1) Include Wrong Dir.	(2) Include Boundary	(3) Include All
δ	0.910***	0.914***	0.914***
	(0.010)	(0.010)	(0.011)
β_1	0.506***	0.727***	0.649***
	(0.033)	(0.045)	(0.045)
β_0	0.714***	0.903**	0.805***
	(0.038)	(0.045)	(0.045)
P-Value ($\delta = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_1 = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_0 = 1$)	0.0000	0.0306	0.0000
Diff ($\beta_1 - \beta_0$)	-0.208	-0.176	-0.156
P-Value ($\beta_1 = \beta_0$)	0.0000	0.0003	0.0022
R^2	0.81	0.81	0.79
Observations	3537	3654	3840

Analysis uses OLS regression. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant. First column includes updates in direction predicted by Bayes' rule. Second column replaces boundary probabilities with 0.01 or 0.99 respectively. Third column is entire sample.

F.2 Restricting to Active Updates

Table F2 presents the analysis of Table 1, but restricting the sample to only active updates. The results show that subjects appear to suffer from the opposite bias of conservatism, as they are over-responsive to information. This is largely drive by response to a negative signal, but does not appear to differ between good or bad news, versus just news. As such, symmetry can be rejected at the 1% level.

Table F2: Active Updates: Reponse to Contemporaneous Signal

Dependent Variable: Logit Posterior Belief			
Regressor	(1) Good/Bad News	(2) Just News	(3) All
δ	0.882***	0.863***	0.873***
	(0.018)	(0.022)	(0.014)
β_1	1.060	1.092	1.074
	(0.052)	(0.071)	(0.047)
β_0	1.295***	1.323***	1.305***
	(0.050)	(0.071)	(0.046)
P-Value ($\delta = 1$)	0.0000	0.0000	0.0000
P-Value ($\beta_1 = 1$)	0.2529	0.1962	0.1138
P-Value ($\beta_0 = 1$)	0.0000	0.0000	0.0000
Diff ($\beta_1 - \beta_0$)	-0.235	-0.232	-0.231
P-Value ($\beta_1 = \beta_0$)	0.0002	0.0121	0.0000
R^2	0.81	0.79	0.80
Observations	1121	799	1920

Analysis uses OLS regression. Includes only active updates. Difference is *significant from 1* at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at individual level. R^2 corrected for no-constant.

G Aggregate Updating by Event/Stake/Accuracy Payment

In this section I examine patterns in updating behavior for different events and financial stake conditions. Recall that the lump sum payment used for the lottery method was randomized at the session level, and was either \$3, \$10, or, \$20. The financial stake was randomized at the individual-event level, and was either \$0 or \$80 with 50% probability respectively. The financial stake was an amount of money that would be gifted to the subject if the event occurred and had been randomly selected for payment.

In Figure G1 I examine the analog to Figure 4, for each of the two financial stake conditions ($0 and $80), as well as each of the three accuracy payment conditions ($3, $10, $20). While different values of the accuracy payment do not affect whether news is good or bad, note that having an $80 stake in an event necessitates that signals contain either good or bad news.

From Figure G1 there does not appear to be any sizeable differences in updating behavior across these different payment conditions. The results on differences between a stake of $0 versus $80 are consistent with Barron (2016), who does not find evidence of asymmetry when individuals have a financial stake in an event. Note also that the prior varies slightly by payment conditions; updating patterns by prior are presented in Figure E1 below.

Next, in Figure G2 I present the analogous analysis for each of the four events, with the quiz event split into the self and other treatments. For the two dice events, which involved the probability that particular outcomes from rolls of either two or four dice had occurred, updating appears to be more conservative than the aggregate. The pattern is also seen when individuals estimate the probability that another randomly selected, anonymous individual in the room had scored in the top 15% on the earlier taken quiz (quiz: other performance).

For the quiz (self performance) event, which involved the probability that the individual believed they scored in the top 15% of quiz takers, updating appears to adhere more closely to the Bayesian prediction. This is also true for the weather event, which occurred when subjects had correctly estimated the mean temperature $\pm$ 5 degrees F in New York City on a randomly selected day in the previous calendar year. In the aggregate, updating about own performance does not appear to deviate much from the Bayesian prediction.

Figure G1: Evolution of Beliefs By Stake and Accuracy Conditions

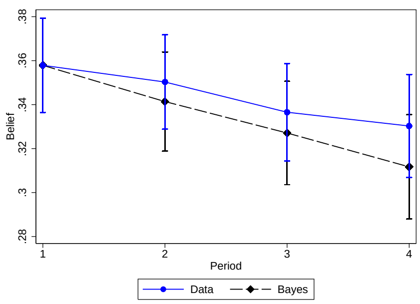

(a) Stake = \$0

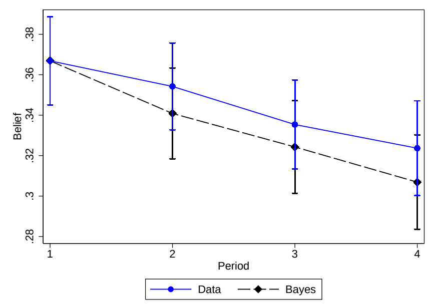

(b) Stake = \$80

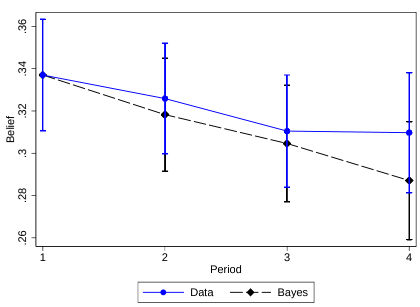

(c) Accuracy Payment = \$3

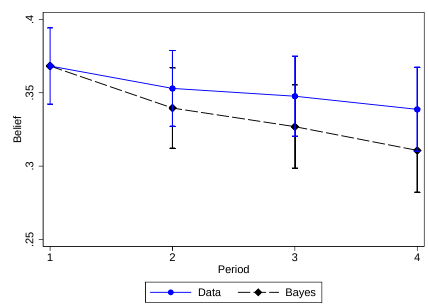

(d) Accuracy Payment = \$10

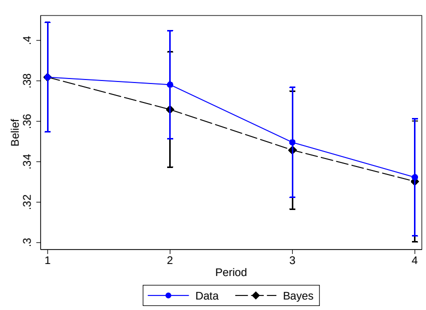

(e) Accuracy Payment = \$20

The path of beliefs starting from the prior (period 1), and after each sequential signal (periods 2 through 4). Average individual responses are the blue solid line, the Bayesian benchmark is marked as the black dashed line. Bayesian benchmark takes prior beliefs, and subsequently uses Bayes' rule to update beliefs. Error bands represent 95% confidence intervals. Note the potential difference in the range of prior beliefs, on the vertical axis. $N = \{646, 634, 424, 436, 420\}$ per round, respectively for (a)-(e).

Figure G2: Evolution of Beliefs: By Event

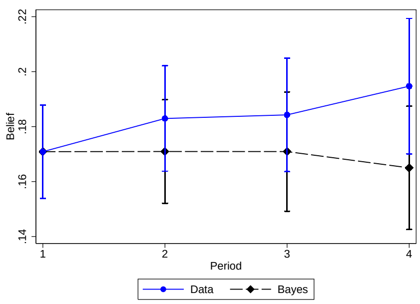

(a) Easy Dice

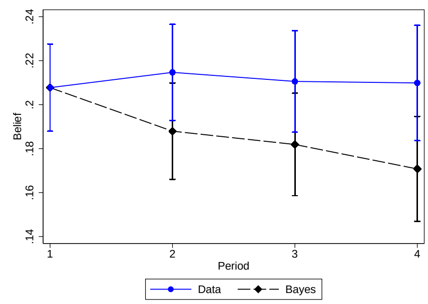

(b) Hard Dice

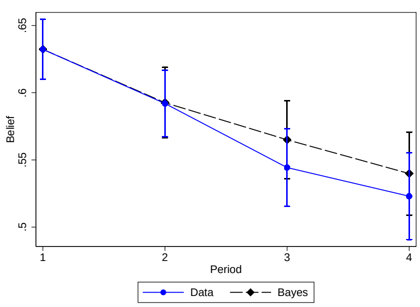

(c) Weather Event

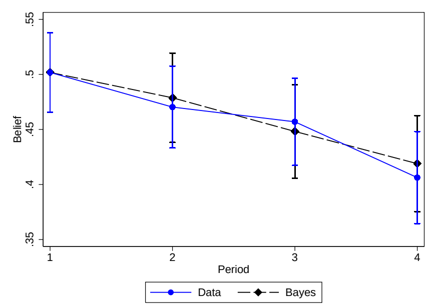

(d) Quiz Event (self performance)

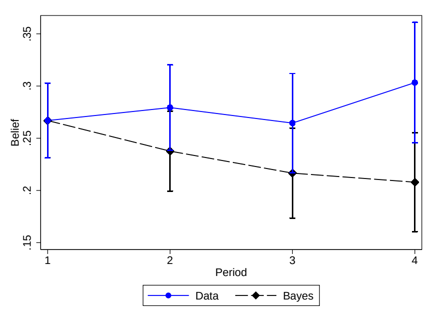

(e) Quiz Event (other performance)

The path of beliefs starting from the prior (period 1), and after each sequential signal (periods 2 through 4). Average individual responses are the blue solid line, the Bayesian benchmark is marked as the black dashed line. Bayesian benchmark takes prior beliefs, and subsequently uses Bayes' rule to update beliefs. Error bands represent 95% confidence intervals. Note the difference in the range of prior beliefs, on the vertical axis. $N = \{318, 318, 326, 223, 95\}$ per round, respectively for (a)-(e).

H Experiment Instructions

Note: Accuracy payments were randomized at the session level $a \in \$\{3, 10, 20\}$. Instructions show \$20 for exposition only.

Instructions (Section 1)

Thank you for your participation in this experiment! This experiment will last approximately 80 minutes. This experiment is about how likely you think an uncertain event is to have occurred. You will consider four such separate events today, which will be presented one at a time. For these events, we want you to think in terms of the percent chance out of 100 that they occurred. For example, you may believe that there is 50% chance that when flipping a coin it will come up TAILS. This experiment has been designed so that you have the greatest chance of earning the most money when you carefully and accurately think about the percent chance of such an event occurring.

You will be awarded a $10 show-up fee for your participation until the end, in addition to anything you may earn during the experiment. Please also note the following during the experiment:

- Please put away any cell phones/devices. Outside communication or accessing the internet during this experiment is forbidden. Violators will not receive payment and will be blacklisted from the lab.
- Please do not communicate with others in the lab, except to ask questions
- If you have a question please do not hesitate to ask! Questions are encouraged!

We will now introduce the experiment through Instructions 1-3 and three short practice sessions that go with each set of instructions. The practice sessions are to help you get familiar with the experiment's components that will ALL be combined when doing the final experiment for money.

The "Main Event"

In this experiment you are estimating the percent chance that a "main event" occurred. An example of a "main event" is: the average temperature in the contiguous USA was warmer in 2013 than 2012. Your earnings are in part based on the accuracy of your predictions of whether the "main event" occurred. Think about the following: What is the probability the average temperature in the USA was warmer in 2013 than 2012?

How will I record my percent chance estimate?

First we introduce a gumball machine with 100 green and black gumballs. For example, suppose there are 40 green and 60 black gumballs. Most people would agree that the probability of drawing a green gumball is exactly 40%. Now think back to the "main event" about the weather being warmer in 2013 than 2012 in the US. We next give you $20. But this $20 must be wagered on one of two scenarios.

1. The “gumball event”: Drawing a green gumball from a machine with 40 out of 100 green, OR

2. The “main event”: the average US temperature in 2013 was warmer than it was in 2012.

You have to decide if you think the chance that the weather was warmer in 2013 is greater than 40%, or less than 40%. If you decide to wager the $20 on the “gumball event”, the computer randomly draws a gumball from the machine with 40 green (60 black) gumballs. If it’s green you win the $20. If black, you get nothing. If you decided to go with the “main event”: the climate being warmer in 2013, we check the statistics. If it was warmer, you win the $20. If it was colder, you get nothing.

Consider different numbers of green gumballs:

If the gumball machine has only 2 green gumballs (98 black) would you prefer to wager $20 on the “gumball event” or the “main event”? Most of you probably think the climate being warmer in 2013 than 2012 is more likely than 2% and prefer to wager the $20 on the “main event”.

What if the gumball machine has 25 green gumballs? Those who think the “main event” is more likely than 25% would want to wager on the “main event”. Now, what if the gumball machine has 90 green gumballs? The “gumball event” now pays off with 90% chance. Probably, almost everyone will prefer to wager the $20 on the gumball machine, except for those that think there is a greater than 90% chance that the weather was warmer in 2013.

Example – You think there is a 35% chance the weather is warmer in 2013 than 2012.

- Case 1: Whenever you see a gumball machine with 34 or less green gumballs, to earn the most money you would want to wager the $20 on the “main event”. E.g. if there were 5 green gumballs, 5% is a lower chance than 35% of earning the $20.

- Case 2: If you see a gumball machine with 36 or more green gumballs, you would prefer to wager the $20 on the “gumball event”. E.g. If there were 60 green gumballs, this is a 60% chance of drawing green – better than the 35% chance you think the weather would be warmer.

- If there are exactly 35 green gumballs, you probably don’t care whether to wager your $20 on the “gumball event” or the “main event”. Both give you a 35% chance of earning the $20.

The “Slider”

In this experiment you are going to indicate on a “slider” exactly how many gumballs need to be green before you prefer to wager $20 on the “gumball event” instead of some other “main event”. In other words, you will indicate the minimum number of gumballs that have to be green, before you prefer to wager \$20 on the gumball machine. To make sure it is in your best financial interest to do this, after you have made your slider choice we are going to randomly fill a gumball machine with 0 to 100 green gumballs and the rest black. Each possible number of green gumballs is equally likely – and your slider choice has no effect on the number chosen. Based on your slider choice, we will then make the \$20 wager for you. If there happen to be less green gumballs than the minimum you chose, your \$20 is wagered on whatever main event you are predicting. If there happen to be more (or the same) green gumballs than the minimum you indicated in the slider, we will wager your \$20 on drawing a green gumball from this machine we randomly filled.

If this is a little confusing, you can just remember, to have the highest chance of earning money, your slider choice should be exactly the probability out of 100 you think the event has of occurring.

Summary of Section 1

- Make selection on the "Slider" for your estimate of the "main event"
- Computer randomly generates an amount (out of 100) of "green gumballs"
- The amount of green gumballs determines how the \$20 is wagered in your best interest. 1) The "main event" or 2) The "gumball event". The outcome of the \$20 wager is then revealed.

Are there any questions?

Instructions (Section 2) – "Feedback"

Now we're going to make things more interesting. Suppose now the "Main event" is that the average temperature in 1998 was warmer than 1997 in the contiguous USA.

Please note – these events are used for practice. The real events may (and will) be different.

You will again adjust the slider to indicate how likely you believe this is to be true. But now, after you adjust the "Slider" the first time, you are going to get some "feedback" about whether or not 1998 was in fact warmer than 1997.

What is "Feedback"?

"Feedback" is information about the main event that gives you additional clues to help you make your selection. Please note that you are provided three rounds of this "feedback" – however each time you are presented with this "feedback" it may or may not be telling you the truth. For our experiment we use gremlins to provide the three rounds of feedback when making your selection. For each round, two gremlins always tell the truth while one of them, Larry, always lies. You will not know which gremlin is talking and after you get this "feedback", you can adjust your prediction on the 'Slider" if you choose to use their information. Note: The gremlins are randomly chosen “with replacement”, meaning that every time you get “feedback” it is true with 2/3 probability. This means, that it’s even possible (though unlikely) that all three rounds of feedback come from the gremlin that lied!

Remember: All 3 gremlins always know whether the event happened or not. It’s just that only 2 of these 3 tell the truth. When we determine your earnings, before filling the gumball machine we are going to randomly only pick one of these four slider choices. Are there any questions at this point? Next we proceed to the second practice. In this example please note two additional tools for your use.

1. Calculate Fraction: Pulls up a calculator in case you want to transform a fraction to a decimal.

2. Show History: Shows you your history of feedback from gremlins AND your past slider choices.

Instructions (Section 3) – Payment groups

The last component explains how you might earn additional money during this experiment. This is very important to understand when conducting the final experiment. You will all be in one of two payment groups: “red” or “blue”. NOTE: You will not know which payment group (red or blue) you are in when you make your slider choices. Suppose now the ‘main event” is whether the climate in the USA was warmer in 1990 than 1980.

“The Red Group”

Half of you are going to be in the “red” group. In the “red” group, your payment at the end looks exactly like how we have been practicing so far. We will pick one of your four slider choices incorporating the “feedback”, and then fill a gumball machine with a random number of green gumballs. Based on your selection, if the $20 is wagered on the “gumball event” then a gumball would be drawn – if green you earn the $20. If the $20 is wagered on the “main event”, then if that event occurred you earn the $20.

“The Blue Group”

The other half of you will be in the “blue” group. The “blue” group automatically gets $20, just for being blue. In this group, the slider choices previously selected do not matter for payment. Instead payment depends on a “blue bonus chip” provided that pays out only if the event you are predicting actually occurs. Taking the example of climate, if 1990 was warmer than 1980, and if you are in the blue group, you would receive $20 automatically, plus whatever amount is on the “blue bonus chip”. The amount on the chip is either $0 or $80. Each is equally likely. Example: If you’re in the “blue” group you would automatically earn $20, and if the main event you are predicting occurs you would also earn the amount on the blue bonus chip ($0 or $80): for a maximum earnings of $100.

"Blue Bonus chip"

Everyone will get a "blue bonus chip" prior to knowing which group you are in and prior to each of the four events. The experiment coordinator will fill a bag with half \$0 chips and half \$80 chips. Then, each of you will draw one of these chips from the bag. Note that having a "blue bonus chip" is only significant when you end up in the "blue" group and indicates how much is earned if the event happens AND if you are in the "blue" group.

Each of you has a fair, 50% chance of drawing an \$80 bonus chip. There is no advantage to drawing a chip earlier or later, everyone in this room has the same 50% chance. Even if you are the last to draw, and there is only one chip left, that one chip is \$0 with 50% chance and \$80 with 50% chance. Since you don't know if you're "red" or "blue" until all slider choices have been made, in order to have the best chance of earning the most money, it pays to be as accurate as possible when making slider choices.

Are there any questions at this point? Next we proceed to the final practice. Note that your "blue bonus chip" has an 8-digit code that you are required to enter into the computer. Your "blue bonus chip" does not affect in any way the event that you will be predicting. The event is the same if you pick a \$0 chip or an \$80 chip. Forget about the gremlins or "feedback" for this practice, yet they will be in the main experiment.

Summary for the Final experiment

Now we are ready to put ALL the pieces together for the final experiment! There are going to be four main events, however only one will be picked at random for payment.

1. The coordinator will come around with a bag that contains a 50/50 mix of \$0 and \$80 "blue bonus chips" for the upcoming event.

2. Make a note of your "blue bonus chip" amount. This is what you could earn if the event happens AND if you also happen to be in the blue group.

3. The event will be described to you. Next, indicate on the "Slider" the probability you believe the event occurred. Your slider choice does not affect how many green gumballs the random gumball machine will have nor does it affect the chances of the "main event".

4. You'll get "Feedback" three times from a random gremlin. Remember there is a 2/3 chance the feedback is true. You can choose to use this information if you want to reassess the probability by indicating this on the slider after each "Feedback".

5. Steps 1 to 4 are repeated for each of the four events.

After making all of your slider choices:

1. The coordinator will come with two bags. The color bag contains 50/50 mix of blue and red chips. The chip you draw determines if your payment group is red or blue. If it is red, the slider choice (1-4) is indicated on the chip.

2. The event bag contains an equal amount of Event #1, #2, #3 and #4 chips. The number on the chip determines what event will be paid.

Suppose you picked the chip for Event #1.

1. IF draw RED: The chip indicates the slider choice. A gumball machine is filled with a random number of green gumballs. Based on your slider choice, $20 is wagered on gumball machine or Event #1, as we practiced.

2. IF draw BLUE: The outcome of Event #1 is revealed. If the event occurred you earn $20 + the amount on your event #1 bonus chip, $80 or $0. If the event did not occur you just earn the $20. After your payment is determined, we will reveal the outcomes of the other three events. This is for your information only, and it does not affect your payment.

Important Notes:

The procedures that will occur today have been approved by the University Committee on Activities Involving Human Subjects (UCAIHS). This experiment complies with UCAIHS requirements (HS# 10-8117), in particular, not to engage in any deception or misinformation about the probabilities presented today.

- When you encounter random chance off the computer (e.g. when drawing chips from the bag) we make every effort to ensure that this is transparent and legitimate. If we state there is a 50-50 chance of drawing a particular chip, we will have at least one participant verify that this is indeed the case. (any participant may ask to verify the bag contents before the draws begin)
- When you encounter random chance on the computer (e.g. drawing a gumball from a hypothetical machine) the computer has been programmed to perform the randomization exactly as is stated in this experiment. For example, if you are told that there are 30 green gumballs and 70 black, the computer is programmed to randomly select a green gumball with exactly 30 chances out of 100.

Before moving forward to the next main event, the computer will wait for everyone to finish the current event. There is no advantage to finishing quickly, as you will end up waiting for other participants.