Belief updating: Does the 'good-news, bad-news' asymmetry extend to purely financial domains?

Barron, Kai

WZB

20 February 2018

Belief Updating: Does the 'Good-news, Bad-news' Asymmetry Extend to Purely Financial Domains?*

Kai Barron†

Abstract

Bayes' statistical rule remains the status quo for modeling belief updating in both normative and descriptive models of behavior under uncertainty. Recent research has questioned the use of Bayes' rule in descriptive models of behavior, presenting evidence that people overweight 'good news' relative to 'bad news' when updating ego-relevant beliefs. In this paper, we present experimental evidence testing whether this 'good-news, bad-news' effect extends to belief updating in the domain of financial decision making, i.e. the domain of most applied economic decision making. We find no evidence of asymmetric updating in this domain. In contrast, the average participant in our experiment is strikingly close to Bayesian in her belief updating. However, we show that this average behavior masks substantial heterogeneity in updating behavior, but we find no evidence in support of a sizeable subgroup of asymmetric updators.

JEL Classification: C11, C91, D83

Keywords: economic experiments, Bayes' rule, belief updating, belief measurement, proper scoring rules, subjective probability, motivated beliefs.

*I would like to thank Syngjoo Choi, Alex Coutts, Steffen Huck, Terri Kneeland, Yves Le Yaouanq, Daniel Martin, Lukas Wenner, and the audiences at RES, ECBE, THEEM, ASFEE, ENTER Jamboree, and the ESA conferences for comments and helpful discussions, as well as Brian Wallace, Lukas Wenner, Roberta De Filippis, and Nina Bonge for their help and support in running the experiments. All mistakes are my own.

†University College London and WZB Berlin: kai.barron@wzb.eu.

1 Introduction

In recent decades, the subjective expected utility (SEU) theory has provided the foundation for the majority of economic thinking regarding choice under uncertainty. Much of its appeal derives from its elegance and simplicity in representing a wide range of choice behavior using two flexible constructs: (i) subjective 'beliefs' and (ii) a utility function that reflects preferences over certain outcomes. However, it is a purely *paramorphic* ('as if') model that makes no claims to capture the underlying psychological processes driving choices. Consequently, the behavioral economics literature was born with the objective of developing *homeomorphic* models that accurately reflect choice behavior as well as the underlying psychological processes. Much of this literature has focused on the second construct of SEU, the utility function, by developing and testing models that move away from SEU. These models typically consider alternative representations of preferences, including classical examples such as *reference dependence* (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Kőszegi and Rabin, 2006) and *social preferences* (Kahneman et al., 1986; Rabin, 1993; Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000).

In contrast, the current paper focuses on the first construct of SEU, beliefs, and examines the question of whether individuals update their beliefs symmetrically in response to 'good-news' and 'bad-news' in the domain of purely financial decision making. This paper therefore joins a rapidly growing body of literature[1] that draws inspiration from the hypothetical 'beliefs' of the SEU model, instead taking the view that beliefs and expectations are real measurable objects that guide decision making. Much of this literature is devoted to studying the formation and dynamic evolution of beliefs, taking the statistical rule for how conditional probabilities are determined when new information arrives (i.e. Bayes' rule[2]) as a benchmark for how beliefs evolve. Many empirical studies in this literature document that people update their beliefs in a way that departs systematically from Bayesian updating. Consequently, a lively theoretical literature has developed with the aim of organizing the patterns observed in the empirical evidence through means of psychologically accurate theories of belief updating, with noteworthy contributions on the wide range of departures

[1] This literature is divided into several branches examining: (a) *the formation of beliefs* (the area of focus of the current paper); (b) *the measurement of beliefs* (for informative reviews, see Manski (2004); Attanasio (2009); Hurd (2009); Delavande (2014); Schotter and Trevino (2014)); and (c) *the mapping from beliefs to actions* (e.g. probability weighting in Kahneman and Tversky (1979) and experimental evidence on the belief-action relationship in Nyarko and Schotter (2002) and Costa-Gomes et al. (2014), amongst others).

[2] Bayes' rule is a statistical rule that represents the normative optimum for processing new information most efficiently. One argument in favor of the use of Bayes' rule in descriptive psychological models is that evolution should therefore have selected individuals who were able to process new information efficiently according to Bayes' rule over individuals who could not, as the Bayesian individuals would have more accurate beliefs guiding their decision making.

from Bayes' rule including the *representativeness bias* (Grether, 1978, 1980, 1992), *cognitive dissonance*[3] (Akerlof and Dickens, 1982), *anticipatory utility* (Loewenstein, 1987; Caplin and Leahy, 2001; Brunnermeier and Parker, 2005), *base rate neglect* (Kahneman and Tversky, 1973; Holt and Smith, 2009), *confirmatory bias* (Rabin and Schrag, 1999), *motivated belief formation* (Benabou and Tirole, 2002), *gambler's and hot-hand fallacy* (Rabin and Vayanos, 2010; Croson and Sundali, 2005; Ayton and Fischer, 2004), and *correlation neglect* (Enke and Zimmermann, 2015).

One important strand of this literature examines whether belief formation and updating is influenced by the *affective content* of the new information[4], i.e. whether individuals update their beliefs symmetrically in response to 'good-news' and 'bad-news' (see, for example, Eil and Rao (2011); Ertac (2011); Mayraz (2013); Möbius et al. (2014); Coutts (2016); Heger and Papageorge (2016); Gotthard-Real (2017)). Essentially, this literature tests an implicit assumption of Bayesian updating, namely that the only object that is relevant for predicting an individual's belief is her information set, and therefore her beliefs are completely unaffected by the prizes and punishments she will receive in different states of the world. This fundamental assumption - that people update their beliefs symmetrically - is of paramount importance because it underpins the results of a wide range of theoretical studies within economics, including all research in which agents receive new information and form rational expectations. As noted by Brunnermeier and Parker (2005), since Muth (1960, 1961) and Lucas (1976), the vast majority of economic research involving uncertainty has built on the rational expectations assumption, with diverse applications including the vast literatures on *social learning* (see, for example, Chamley (2004) for a review), *capital markets* (e.g. Fama (1970, 1976)), *intertemporal portfolio choice problems* (e.g. Mossin (1968); Merton (1969); Samuelson (1969); Lewellen and Shanken (2002)) and *consumption savings problems* (e.g. Friedman (1957); Hall (1978)). Therefore, it is essential to empirically test whether the assumption that people form beliefs symmetrically in response to 'good-news' and 'bad-news', such that their posterior beliefs are completely unaffected by their hopes and wishes, is consistent with how belief updating actually works.

With this objective in mind, in this paper we employ a laboratory experiment to study how individuals update their beliefs from exogenously assigned prior beliefs in a two-state world when

[3]This paper by Akerlof and Dickens (1982) is closely related and a precursor to the anticipatory utility models discussed below in this paragraph (in particular, Brunnermeier and Parker (2005)) as their model allows the agent to choose her belief by trading off her anticipatory preferences over states against the instrumental value of this information. However, these chosen beliefs are persistent and the model shows that this can lead to surprising behavioral patterns grouped under the heading of *cognitive dissonance* in the psychology literature.

[4]As an indication of the importance of this strand of the literature, in their recent paper taking stock of the current state of the motivated reasoning literature, Bénabou and Tirole (2016) point to *asymmetric updating* as one of the main testable implications of motivated reasoning.

they receive a sequence of partially informative binary signals. In particular, we vary the financial rewards associated with each of the two states of the world in order to test whether these state-contingent financial rewards influence belief updating. A nice feature of this experimental design is that it allows us to compare posterior beliefs in situations where the entire information set is held constant, but the rewards associated with the states of the world are varied. For example, we can compare how two groups of individuals revise their beliefs when both groups share the same prior belief and receive an equally informative signal, but for one group of individuals the signal is 'good news' and for others the signal constitutes 'bad news'. Furthermore, we can conduct a similar exercise for a single individual, by comparing two situations in which she has the same information set, and receives the same signal, but the signal consistutes 'good news' in one instance, and 'bad news' in the other. Our experimental design therefore permits a clean test of the *asymmetric updating hypothesis*.

More specifically, in our experiment we study belief updating in two contexts. In our SYMMETRIC treatment, we examine how subjects update their beliefs when they have an equal stake in each of the underlying states and therefore are indifferent about which state is realized. We compare this with updating behavior in our two ASYMMETRIC treatments, in which a large bonus payment will be paid if one state of the world is realized. Here, one would expect that subjects prefer that the state with the bonus payment is realized. In all treatments, subjects are told the true prior probability with which each of the two states of the world will be realized, and then update their beliefs upon receiving a sequence of informative but noisy signals regarding the true realized state.[5] We elicit this sequence of beliefs. In our experiment, we can therefore make two comparisons. Firstly, we can compare how the same individual responds to 'good-news' and 'bad-news' *within* the ASYMMETRIC treatments. Secondly, we can compare the belief updating of two groups facing different incentive environments in a *between-subject* comparison of the SYMMETRIC treatment and the ASYMMETRIC treatments. Each individual in our experiment faces only one incentive environment. However, since we exogenously endow participants with a prior over the states of the world, we are able to repeat the exercise several times for each individual and study how they update from each of five different priors, p_0, chosen from the set $\{\frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}\}$.

The experimental design and analysis aim to address several challenges that are present when studying belief updating in the presence of state dependent stakes. Firstly, we use exogenous variation in the priors to ensure that the estimates are robust to the econometric issues that arise when

[5] The ASYMMETRIC treatments correspond to the setup studied in Möbius et al. (2014) and Eil and Rao (2011), where a 'good' underlying state is being higher in the IQ or attractiveness distribution in relation to the other participants in the experiment.

a right-hand side variable (i.e. the prior) is a lagged version of the dependant variable (i.e. the posterior). Secondly, we avoid a second type of endogeneity issue, which arises when the underlying states are defined as a function of some personal characteristic of the individual (e.g. her relative IQ) that might also be related to how she updates (see Appendix C for further details). Thirdly, we measure the influence that hedging has on belief elicitation when there are state-dependent stakes. Furthermore, we conduct several exercises to correct our estimates for this heding influence - both experimentally, and econometrically. Fourth, our experimental design allows us to study belief updating from priors spanning much of the unit interval. Furthermore, averaging across all subjects, the design generates a balanced distribution of 'good' and 'bad' signals.

The empirical strategy employed permits testing for several commonly hypothesised deviations from Bayesian updating, including confirmatory bias and base rate neglect, however the focus of both the experimental design, and the analysis, is on testing for the presence of an asymmetry in updating. Our results show no evidence in favour of asymmetric updating in response to 'good-news' in comparison to 'bad-news' in the domain of financial outcomes. Several robustness exercises are carried out in support of this conclusion. Furthermore, we find that average updating behavior is well approximated by Bayes' rule.[6] However, this average behavior masks substantial heterogeneity in updating behavior, but we find no evidence in support of a sizeable subgroup of asymmetric updators.

The absence of a 'good-news, bad-news' asymmetry in our context is in contrast to the evidence reported in studies considering belief updating in the domain of self-image, which have typically reported evidence in favour of asymmetric updating that is either positively (Eil and Rao, 2011; Sharot et al., 2011; Möbius et al., 2014) or negatively (Ertac, 2011) skewed.

However, this result is consistent with complementary work by Gotthard-Real (2017) and Coutts (2016) that also studies belief updating in the domain of financial updating. Gotthard-Real (2017) uses a similar belief updating task as the current paper, with all subjects holding a \$7 stake in one of two possible states of the world and updating from a prior of 0.5. This corresponds to belief updating from one of the five priors we consider, under one of our three treatment conditions. In line with this paper, Gotthard-Real (2017) finds no evidence of asymmetric updating. Coutts (2016) offers an important contribution to this literature with an experiment that considers belief updating within both the financial and self-image domains. He finds that the average posterior belief is well approximated by the Bayesian posterior. Furthermore, while he finds evidence for asymmetric

[6] This is in line with Holt and Smith (2009) and Coutts (2016), who find that average posteriors across all individuals are well approximated by posteriors obtained by applying Bayes' rule.

updating, it appears to be unlinked to the desirability of the states. Rather, Coutts (2016) argues that it is driven by the low objective probability of the good state and consequent unbalanced ratio of positive to negative signals generated by his experimental design. Therefore, together with the current paper, this evidence suggests that the asymmetry in updating in the self-image domain may not extend to the domain of financial stakes.

One potential explanation for the difference in belief updating between the domains of self-image and financial decision making is the idea that ego maintainance could yield evolutionary benefits. In particular, a positive asymmetry in updating about one's self-image would lead to overconfident beliefs, and several authors have posited that maintaining a high self-confidence may be associated with evolutionary advantages (see, e.g., Bernardo and Welch (2001); Heifetz et al. (2007); Johnson and Fowler (2011); Burks et al. (2013); Schwardmann and Van der Weele (2016)). In contrast, asymmetric updating about external states of the world would lead to overoptimism which is likely to lead to costly mistakes.

The remainder of the paper proceeds as follows. Section 2 outlines the theoretical framework, Section 3 details the experimental design, Section 4 provides some descriptive statistics, Section 5 presents the empirical specification, and Section 6 concludes.

2 Theoretical Framework

In the following section, we discuss a simple framework for belief updating that augments the standard benchmark of Bayesian updating by allowing for several of the deviations from Bayes' rule commonly discussed in the psychology and related economics literature. The basic idea is that, while Bayes' rule captures the normative benchmark for how we might think a rational agent *should* update her beliefs[7], it has been argued that, descriptively, people may update their beliefs in ways that depart systematically from Bayes' rule. The framework below serves to facilitate a discussion of these different potential deviations, and motivates the empirical approach that we

[7] This statement is not uncontentious. One argument in favor of Bayes' rule as a normative benchmark for belief updating is that Bayes' rule is a statistical theorem that captures the objective relationship between a prior probability and a posterior probability, given new information. Therefore, if an individual updates her beliefs according to Bayes' rule, she will always hold beliefs that are as accurate as possible, given her information set. This will allow her to use these optimally accurate beliefs to guide her decision making and make decisions. This implies that the decision maker can make decisions that are as informed as possible. However, this argument rests on the assumption that beliefs serve only an *instrumental* role in guiding decision making. If, for example, we relax this assumption and allow beliefs to yield *intrinsic* utility, then this argument no longer holds as it may be optimal for the decision maker to hold beliefs that are distorted away from the Bayesian posterior.

will use to test whether they are observed in our data.[8]

In short, below we describe a model that embeds the normative Bayesian benchmark, but also allows for commonly discussed belief updating distortions. The aim will be to then make use of this model to test whether subjects update their beliefs like a Bayesian automaton or display some systematic deviations from this Bayesian benchmark. Most importantly, the model permits a discussion of what observed updating behavior should look like if agents update their beliefs asymmetrically in response to 'good-news' and 'bad-news'.

2.1 A Simple Model of Belief Formation

We consider a single agent who forms a belief over two states of the world, $\omega \in \{A, B\}$, at each point in time, t. One of these states of the world is selected by nature as the 'correct' (or 'realized') state, where state $\omega = A$ is chosen with prior probability $\bar{p_0}$ (known to the agent). The agent's belief at time t is denoted by $\pi_t \in [0, 1]$, where π_t is the agent's belief regarding the likelihood that $\omega = A$ and $1 - \pi_t$ is the agent's belief that $\omega = B$. In each period, the agent receives a signal, $s_t \in \{a, b\}$, regarding the state of the world, which is correct with probability $q \in (\frac{1}{2}, 1)$. In other words, $p(a|A) = p(b|B) = q > \frac{1}{2}$. Furthermore, the history, H_t, is defined as the sequence of signals received by the agent in periods $1, \ldots, t$, with $H_0 = \emptyset$. Therefore, the history at time t is given by $H_t = (s_1, ..., s_t)$.

We are interested in how the agent's beliefs evolve over time in response to the sequence of information signals, s_t, she receives. As discussed above, Bayes' rule provides a useful benchmark by describing the objective evolution of the probabilities, given the agent's information set at each point in time. Therefore, we let $p_t \equiv P(\omega = A|H_t, \bar{p_0})$ be the Bayesian posterior probability at time t, given the agent's information set $\{H_t, \bar{p_0}\}$. Since, $p_0 \equiv \bar{p_0}$[9], we will use p_0 below to refer to the initial objective prior probability of state A being the correct state. With these definitions in hand, we can apply Bayes' rule to show how objective probabilities evolve in response to the signal $s_{t+1} = a$

$$p_{t+1}|_{s_{t+1}=a} \equiv P(\omega = A|H_t, p_0, s_{t+1} = a) = \frac{q \cdot p_t}{q \cdot p_t + (1-q) \cdot (1-p_t)} \tag{1}$$

and similarly, in response to the signal, $s_{t+1} = b$, the Bayesian posterior is given by:

[8] This framework is very similar to the one used in Möbius et al. (2014), amongst others.

[9] Notice that at $t = 0$, we have $p_0 \equiv P(\omega = A|H_0 = \emptyset, p_0) = P(\omega = A|p_0) = \bar{p_0}$.

$$p_{t+1}|_{s_{t+1}=b} \equiv P(\omega = A|H_t, p_0, s_{t+1} = b) = \frac{(1-q)\cdot p_t}{(1-q)\cdot p_t + q\cdot(1-p_t)} \tag{2}$$

Therefore, if our agent is approximately Bayesian in her updating, her beliefs, π_{t+1}, regarding the likelihood that state $\omega = A$ is the correct state will evolve in a similar way to the statistical rule described in Equations 1 and 2 above.[10] Since we are dealing with binary states (as in our experiment), a more convenient expression of Bayes' rule can be written in terms of the logit function:

$$\text{logit}(\pi_{t+1}) = \text{logit}(\pi_t) + \log(\frac{q}{1-q}) \cdot 1(s_{t+1} = a) - \log(\frac{q}{1-q}) \cdot 1(s_{t+1} = b) \tag{5}$$

where $1(s_{t+1} = a)$ [or $1(s_{t+1} = b)$] is an indicator function that takes a value of 1 if the signal at time t is a [b], $\log(\frac{q}{1-q})$ is the log likelihood ratio for the signal $s_{t+1} = a$, and $\log(\frac{1-q}{q}) = -\log(\frac{q}{1-q})$ is the log likelihood ratio for the signal $s_{t+1} = b$. This formulation of Bayes' rule provides a useful benchmark for testing for the presence of commonly hypothesized systematic deviations from Bayes' rule. In order to do this, we follow Möbius et al. (2014) in defining a model of *augmented Bayesian updating*[11]:

$$\text{logit}(\pi_{t+1}) = \delta\text{logit}(\pi_t) + \gamma_a \log(\frac{q}{1-q}) \cdot 1(s_{t+1} = a) - \gamma_b \log(\frac{q}{1-q}) \cdot 1(s_{t+1} = b) \tag{6}$$

where this model is identical to Equation 5, with the exception of the additional parameters δ, γ_a and γ_b. These parameters serve, firstly, to provide structure for a discussion of different ways in which our agent may depart from Bayesian updating, and secondly, to provide a clear prescription for how to test for these departures in our empirical analysis below.

[10] i.e. in Equations 1 and 2, Bayes' rule dictated how objective probabilities, p_{t+1}, should evolve if they efficiently incorporate all new information into the posterior, and similarly, if an individual updates her beliefs, π_{t+1}, in a statistically efficient way, then they should evolve according to:

$$\pi_{t+1}(s_{t+1} = a) = \frac{q\cdot\pi_t}{q\cdot\pi_t + (1-q)\cdot(1-\pi_t)} \tag{3}$$

and likewise, for $s_{t+1} = b$:

$$\pi_{t+1}(s_{t+1} = b) = \frac{(1-q)\cdot\pi_t}{(1-q)\cdot\pi_t + q\cdot(1-\pi_t)} \tag{4}$$

[11] For a discussion of core properties underlying Bayes' rule, see Appendix B.1. This discussion serves as a motivation for the approach taken in augmenting Bayes' rule in Equation 6.

If $\delta = \gamma_a = \gamma_b = 1$ then the agent updates her beliefs according to Bayes' rule. However, if we consider deviations from this benchmark, we see that δ captures the degree to which the magnitude of the agent's prior affects her updating. For example, if $\delta > 1$ then this suggests that the agent displays a *confirmatory bias*[12], whereby she is more responsive to information that supports her prior. In contrast, $\delta < 1$ suggests she is more responsive to information that contradicts her prior (i.e. *base rate neglect*[13]). The former would predict that beliefs will polarize over time, while the latter would predict that over time beliefs remain closer to 0.5 in a two-state world than Bayes' rule would predict.

The parameters, γ_a and γ_b capture the agent's responsiveness to information. If $\gamma_a = \gamma_b < 1$ then the agent is less responsive to the information that she receives than a Bayesian updater would be. And if $\gamma_a = \gamma_b > 1$, then she is more responsive than a Bayesian. For example, if $\gamma_a = 2$, then whenever the agent receives a signal $s_t = a$, she updates her belief exactly as much as a Bayesian would if he received two a signals, $s_t = \{a, a\}$. The interpretation of the parameters is summarized in the first five rows of Table 1 below.

Table 1: Interpretation of Parameters: A Summary

Belief Updating Distortion	**Parameter Values**
Bayesian Updating	$\delta = 1$, $\gamma_a = 1$ and $\gamma_b = 1$
Confirmatory Bias	$\delta > 1$
Base Rate Neglect	$\delta < 1$
Conservatism	$\gamma_j < 1$ for $\forall j \in \{a, b\}$
Overresponsiveness	$\gamma_j > 1$ for $\forall j \in \{a, b\}$
Optimistic updating (in ASYMMETRIC)	$\gamma_a > \gamma_b$
Pessimistic updating (in ASYMMETRIC)	$\gamma_a < \gamma_b$

Affective States

The simple belief updating framework developed in the preceding section has so far focused on purely cognitive deviations from Bayes' rule. The *affect* or desirability of different states of the world has played no role. However, in most situations in which individuals form beliefs, there are

[12]For a detailed discussion of the *confirmatory bias*, see Rabin and Schrag (1999). Essentially, it is the tendency to weight information that supports one's priors more heavily than information that opposes one's priors. In this case, when one's prior regarding state $\omega = A$ is greater than 0.5, i.e. $\pi_t > 0.5$, a participant who is prone to the *confirmatory bias* weights signals that support state $\omega = A$ more heavily than signals that support state $\omega = B$; and *vice versa* when her prior suggests state $\omega = B$ is more likely, i.e. $\pi_t < 0.5$.

[13]One can think of this as the agent attenuating the influence of her prior belief in calculating her posterior - i.e. acting as if her prior was closer to 0.5 than it actually was.

some states that yield an outcome that is preferred to the outcome in other states - i.e. there are *good* and *bad* states of the world. For example, an individual would generally prefer to be more intelligent rather than less intelligent, to win a lottery rather than lose, and for the price of assets in her possession to increase rather than decrease. This implies that new information is also often either *good-news* or *bad-news*.

In order to allow for the possibility that individuals update their beliefs differently in response to *good-news* in comparison to *bad-news*, we relax the assumption that belief updating is orthogonal to the *affect* of the information.[14] To do this, assume that each of the two states of the world is associated with a single certain outcome - i.e. in state $\omega = A$, the agent receives outcome x_A, and in state $\omega = B$, she receives x_B. We therefore consider two cases:

- Case 1 (SYMMETRIC): the agent is indifferent between outcomes (i.e. $x_A \sim x_B$); and
- Case 2 (ASYMMETRIC): the agent strictly prefers one of the two outcomes. Without loss of generality, we assume that x_A is preferred (i.e. $x_A \succ x_B$).

Now, the question we wish to consider is whether the agent will update her beliefs differently in the SYMMETRIC and ASYMMETRIC contexts. Under the assumption that the agent's behavior is consistent with the model described above in Equation 6, this involves asking whether the parameters δ,γ_a and γ_b, differ between the two contexts.

In order to guide our discussion of the differences between the parameters in the two cases, we consider the following two benchmarks. The first natural benchmark is Bayes' rule. Bayes' rule prescribes that all three parameters equal 1 in both the SYMMETRIC and ASYMMETRIC contexts as the statistically efficient updating of probabilities is unaffected by the state-dependent rewards and punishments. According to Bayes' rule, news is news, independent of its affective content.

The second benchmark that we consider is that individuals respond more to 'good-news' than 'bad-news'. This benchmark derives from the literature which has presented evidence in favor of a 'good-news, bad-news' effect in belief updating, specifically in the context of ego-relevant beliefs where individuals may wish to preserve a positive *self-image*. Here, we test this *asymmetric updating hypothesis* in two ways.

[14] In this paper, we are following Wakker's (2010) *homeomorphic* approach to modeling behavior (i.e. where underlying parameters have psychological interpretations). In the discussion below, the term 'preferences' is used to refer to preferences over sure outcomes, but not preferences over lotteries. We will also sometimes refer to 'preferring' one state of the world to another. This simply captures the idea that an individual prefers the realisation of a state in which a *good* outcome is realised.

Firstly, we focus only on subjects **within** the ASYMMETRIC context (i.e. $x_A \succ x_B$). For these subjects, we ask whether there is an *asymmetry in updating* after signals that favor the more desirable state $\omega = A$ ('good-news'), relative to signals that favor the less desirable state $\omega = B$ ('bad-news')[15]. For example, if $\gamma_a > \gamma_b$, this would indicate that the agent updates more in response to signals that support the preferred state, $\omega = A$ ('good-news'). We refer to such an agent as an *optimistic updater*. Conversely, if we have $\gamma_a < \gamma_b$ then the agent updates more in response to signals supporting the less desirable state, $\omega = B$ ('bad-news'). We refer to such an agent as a *pessimistic updater*[16].

Secondly, we can ask whether the parameters of Equation 6 differ **between** the SYMMETRIC (δ^1,γ_a^1 and γ_b^1) and ASYMMETRIC (δ^2,γ_a^2 and γ_b^2) contexts, where we use the postscript $c \in \{1, 2\}$ to distinguish the parameters in the two cases. In the SYMMETRIC treatment, where the agent is completely indifferent between the two states, we would expect her updating to be symmetric, with $\gamma_a^1 = \gamma_b^1$ (i.e. responding equally to the signals $s_t = a$ and $s_t = b$). Therefore, the difference $\gamma_a^2 - \gamma_a^1$ reflects a measure of the increase in the agent's responsiveness when information is desirable, relative to when information is neutral in terms of its affect. Similarly, $\gamma_b^2 - \gamma_b^1$ is a measure of the increase in the agent's responsiveness when information is undesirable, relative to the case in which information is neutral in affect.

In our experiment, we will provide evidence on both of these questions. Firstly, we will examine whether there is an asymmetry in updating behavior within the ASYMMETRIC context; and secondly, we will examine whether the parameters of the updating process differ between the SYMMETRIC and the ASYMMETRIC contexts, in a between subjects design. Furthermore, our experiment will allow us to test for other systematic deviations from Bayes' rule, such as those mentioned in the discussion above. Table 1 summarizes the interpretations of the different values that the belief updating parameters may take.

[15] In order to reduce the verbosity of my discussion, in the text below I will use the phrases *desirable [undesirable] information* or *good [bad] news* to refer to "information that supports a desirable [undesirable] state of the world".

[16] Notice that these definitions of optimistic and pessimistic updating relate only to the asymmetry in responsiveness to signals supporting different states, but depend on how responsive the agent is to the signals she receives relative to Bayes' rule. Therefore, under the definition given, we can have an optimistic updater who is less responsive to signals supporting the desirable state than a Bayesian (i.e. $1 > \gamma_a > \gamma_b$). This agent is less responsive than a Bayesian to both desirable and undesirable information, but more responsive to desirable information relative to undesirable information.

2.2 Belief Elicitation and Incentives

In order to empirically test the features of the model using an experiment, we would like to be able to elicit our participants' true beliefs. Belief elicitation is now fairly well established, with a growing literature debating the merits of several belief elicitation techniques. Many of these approaches assume that participants use a particular model to evaluate prospects (e.g. expected value is often used) and show that under the incentives of the elicitation technique, a participant following the assumed model should report her true belief (i.e. under the model, truthful revelation is *incentive compatible*).

However, in the context of studying the relationship between preferences and beliefs, this approach to the elicitation of beliefs is problematic for two reasons: firstly, the assumption that participants follow expected utility involves a stronger assumption than we might wish to make; and secondly, the inherent hedging motive faced by participants who have a stake in one state of the world poses an additional challenge (see Karni and Safra (1995) for a discussion). Therefore we adopt the approach developed by Offerman et al. (2009).

The central idea behind this approach is to acknowledge that the incentive environment within which we elicit beliefs in the laboratory may exert a distortionary influence on the beliefs which some participants report, relative to the beliefs they actually hold, and then measure this distortionary influence of the incentive environment in a separate part of the experiment. Once we have constructed a mapping from true beliefs to reported beliefs within the relevant incentive environment, we can use this function to recover the participant's true beliefs from her reported beliefs. In other words, our objective is to recover the function that each individual uses to map her true beliefs to the beliefs that she reports within the given incentive environment.

The incentive environment that we will use in our experiment to elicit beliefs is the quadratic scoring rule (QSR).[17]

[17] There are several reasons for adopting this approach: firstly, the QSR has the advantage that it ensures that the decision environment is clear and simple for the participants - essentially they are making a single choice from a list of binary prospects; secondly, the quadratic scoring rule has been commonly used in the literature, with both the *theoretical properties* and *empirical performance* having been studied in detail (see, e.g., Armantier and Treich (2013)); thirdly, in a horse race between elicitation methods, Trautmann and van de Kuilen (2015) show that there is no improvement in the empirical performance of more complex elicitation methods over the Offerman et al. (2009) method, neither in terms of internal validity, nor in terms of behavior prediction. Out of the set of alternative elicitation techniques, the two that are most theoretically attractive are, the *binarized scoring rule*, proposed by Hossain and Okui (2013), and the *probability matching mechanism,* described by Grether (1992) and Karni (2009). However, in the context of the current paper, we viewed neither of these approaches as being preferable to the Offerman et al. (2009) technique, since both of these approaches introduce an additional layer of probabilities and in the study of probability bias, this is an undesirable attribute of the elicitation strategy.

Belief Elicitation Incentives

In this section, we discuss how beliefs reported under the QSR might be distorted, and how we address this challenge. Consider the binary event, denoted by E_ω, where $\omega \in \{A, B\}$. Therefore, E_A refers to the event that state $\omega = A$ is realized. The object that we would like to elicit is the participant's belief, $\pi_t = P(E_A) = P(\omega = A)$, regarding the likelihood that state $\omega = A$ is the correct state at time t. However, the object that we will observe is the participant's reported belief, r_t, at each point in time under the incentives prescribed by the quadratic scoring rule. The *Quadratic Scoring Rule* at time t is defined by:

$$S_A(r_t) = 1 - (1 - r_t)^2 \tag{7}$$

$$S_B(r_t) = 1 - r_t^2 \tag{8}$$

where r_t is the reported probability of event E_A occurring; $S_A(r_t)$ is the payment if the state $\omega = A$ is realized; $S_B(r_t)$ is the payment if the state $\omega = B$ is realized. Therefore, the QSR essentially involves a single choice from a list of binary prospects, $(1 - (1 - r_t)^2)_{E_A}(1 - r_t^2)$. The QSR is a '*proper*' scoring rule since, if the agent is a risk neutral EU maximizer then she is incentivized to truthfully reveal her belief, π_t:

$$\pi_t = arg \max_{r_t \epsilon [0,1]} \pi_t S_A(r_t) + (1 - \pi_t) S_B(r_t)$$

However, the QSR is no longer incentive compatible once we allow for (i) *risk aversion / loving* and (ii) participants who have exogenous stakes in the state of the world. The reasons for this are the following. Firstly, it has been well documented theoretically that, if the participant is *risk averse,* then the QSR leads to reporting of beliefs, r_t, that are distorted towards 0.5, away from her true belief, π_t, when the participant has no exogenous stakes in the realized state.[18] This distortion has been observed in experimental data (Offerman et al., 2009; Armantier and Treich, 2013). Secondly, in our experiment, we will also be interested in eliciting beliefs when participants have an exogenous stake associated with one of the two states. More precisely, we will be interested in recovering the participant's true belief when she receives an exogenous payment, x, if state $\omega =$

[18] i.e. if $\pi_t > 0.5$ then $\pi_t > r_t > 0.5$, and if $\pi_t < 0.5$ then $\pi_t < r_t < 0.5$ for a risk averse individual reporting her beliefs under QSR incentives.

A is realized. This payment, x, is in addition to the payment she receives from the QSR. In other words, she will choose from a menu of binary prospects of the form: $(x+1-(1-r_t)^2)_{E_A}(1-r_t^2)$.

In the context of state-dependant stakes, a risk averse EU maximizer[19] will face two distortionary motives in reporting her belief: (i) she will face the motive to distort her belief towards 0.5 as discussed above; and (ii) in addition, there is a hedging motive, which will compel a risk averse individual to *lower* her reported belief, r_t, towards zero as x increases.

If the participants in our experiment are *risk neutral expected utility* maximizers, the reported beliefs, r_t, that we elicit under the QSR will coincide with their true beliefs, π_t. However, in order to allow for choice behaviors consistent with a wider range of decision models, we will measure the size of the distortionary influence of the elicitation incentives at an individual level and correct the beliefs accordingly. This approach is valid under the weak assumption that individuals evaluate binary prospects according to the *biseparable preferences*[20] model and are *probabilistically sophisticated*.[21] This restriction on behavior is very weak and includes individuals who behave according to EU with any risk preferences as well as the majority of commonly used NEU models.[22]

[19]A participant who is a risk averse EU maximizer will choose her reported belief r_t by solving the following maximization problem:

$$\max_{r_t \epsilon [0,1]} \pi_t U(x+1-(1-r_t)^2) + (1-\pi_t)U(1-r_t^2)$$

[20]The *biseparable preference* model holds if the preference ordering, $\succsim$, over prospects of the form, $y_E z$, can be represented by:

$$y_E z \rightarrow W(E)U(y) + (1-W(E))U(z)$$

where U is a real-valued function unique up to level and unit; and W is a unique weighting function, satisfying $W(\varnothing) = 0$, $W(S) = 1$ and $W(E) \leq W(F)$ if $E \subseteq F$. S is the set of all states and events are subsets of the full set of states: i.e. $E, F \subseteq S$. In this paper, we only consider two-state prospects, where the state-space is partitioned into two parts by an event, E and its complement E^c. Making the further assumption that the decision maker is *probabilistically sophisticated* gives the following refinement:

$$y_E z \rightarrow w(P(E))U(y) + (1-w(P(E)))U(z)$$

[21]*Probabilistic sophistication* is the assumption that we can model that individual's preferences over prospects as if the individual's beliefs over states can be summarized by a probability measure, P. In other words, probabilistic sophistication implies that we can model the individual's belief regarding the likelihood of an event E as being completely summarized by a single probability judgment, $P(E_A)$.

[22]Amongst the models subsumed within the biseparable preferences model are EU, Choquet expected utility (Schmeidler, 1989), maxmin expected utility (Gilboa and Schmeidler, 1989), prospect theory (Tversky and Kahneman, 1992), and α-maxmin expected utility (Ghirardato et al., 2004). See Offerman et al. (2009) for a discussion.

A Non-EU 'Truth Serum'

The discussion above has highlighted how beliefs might be distorted under QSR incentives. The Offerman et al. (2009) approach proposes correcting the reported beliefs for the risk aversion caused by the curvature of the utility function or by non-linear probability weighting. This approach involves eliciting participants' reported belief parameter, r, for a set of risky events where they know the objective probability, p (*known probability)*. This is done under precisely the same QSR incentive environment in which we elicit the participants' subjective beliefs, π, regarding the events of interest (where they don't know the objective probability: *unknown probability*). If a subject's reported beliefs, r, differ from the known objective probabilities, p, this indicates that the subject is distorting her beliefs due to the incentive environment (e.g. due to risk aversion). The objective of the correction mechanism is therefore to construct a map, R, from the objective beliefs, $p \in [0, 1]$, to the reported beliefs, r, for each individual under the relevant incentive environment.

Offerman et al. (2009) show that under the assumption that individuals evaluate prospects in a way that is consistent with the weak assumptions of the *biseparable preferences* model, then in the case where there are no state-contingent stakes (i.e. $x = 0$), individuals evaluate the QSR menu of prospects $(1-(1-r_t)^2)_{E_A}(1-r_t^2)$ according to $w(P(E_A))U(1-(1-r_t)^2) + (1 - w(P(E_A)))U(1-r_t^2)$ for $r_t \geq 0.5$ and therefore the inverse of the map from objective probabilities to reported probabilities, R, is given by:

$$p = R^{-1}(r) = w^{-1}\left(\frac{r}{r + (1-r)\frac{U'(1-(1-r)^2)}{U'(1-r^2)}}\right) \tag{9}$$

In Appendix B.2, we provide a derivation for this equation, as well as augmenting the Offerman et al. (2009) approach to allow for the case where there are state-contingent stakes (i.e. $x \neq 0$). In our empirical analysis below, we will discuss how we use Equation 9 to recover the function, R, for each individual and thereby recover their beliefs, π_t, from their reported beliefs, r_t.

3 Experimental Design

The experiment was designed to test the *asymmetric updating hypothesis* using both within-subject and between-subjects comparisons of updating behavior. The experiment consisted of three treatment groups. The treatment T1.SYMMETRIC corresponds to Case 1 (SYMMETRIC: no exogenous state-contingent stakes) and the other two treatment groups, T2.COMBINED and T3.SEPARATE, correspond to Case 2 (ASYMMETRIC: state-contingent stakes) discussed above.

In T1.SYMMETRIC, one would expect the participants to be indifferent between each of the two states being realized, while in T2.COMBINED and T3.SEPARATE, the larger payment associated with state $\omega = A$ should imply that the participants prefer that this state be realized. T2.COMBINED and T3.SEPARATE are identical in terms of the financial incentives. The rationale for running two treatments with identical incentives was to examine the influence of the hedging motive discussed above. In order to examine this, we varied the way in which the information regarding the incentive environment was presented to participants (i.e. we only varied the framing of the incentives). In T2.COMBINED, this information was summarized in a way that made it much easier for participants to notice the hedging opportunity in comparison to T3.SEPARATE. This allowed us to assess whether a hedging motive influenced the beliefs elicited, and also provided an opportunity to assess the validity of the mechanism we use to correct reported beliefs. The difference between these two treatments is discussed in more detail in the '*Incentives and Treatment Groups*' section below.

The experiment proceeded in three stages. The first stage comprised the core belief updating task in which we elicited a sequence of reported beliefs from subjects as they received a sequence of noisy signals regarding the true state of the world, updating from an exogenously provided prior. In the second stage we collected the reported probabilities associated with known objective probabilities on the interval $[0, 1]$ required for the Offerman et al. (2009) correction approach, as well as data on risk preferences. In the third stage, we obtained data on several demographic characteristics as well as some further non-incentivized measures. In each of the first two stages, one of the subject's choices was chosen at random and paid out. In addition the participants received a fixed fee of £5 [€5] for completing Stage 3, as well as a show-up fee of £5 [€5].[23]

The Belief Updating Task (Stage 1)

The Belief Updating Task was the primary task of the experiment. We used this task to collect data about participants' belief updating behaviour. The experimental design for this task is summarized in Figure 1 and described in the following discussion.

[23] In the discussion below, I will always refer only to Pounds (£), however in the experiments run at the Technical University in Berlin all the payments were made in Euros. In all cases, the payment in Pounds (£) was equivalent to the payment in Euros (€). In other words, £1 was replaced with €1. In terms of the cost of living in the two locations, this seemed more appropriate than using the actual exchange rate. In addition, it had the benefit of keeping all the quantities constant across locations.

The Belief Updating Task consisted of five rounds. In each round, participants were presented with a pair of computerized 'urns' containing blue and red colored balls, with each of these two urns representing one of the two states of the world. The composition of the two urns was always constant, with the state $\omega = A$ represented by the urn containing more blue balls (5 blue and 3 red), while the state $\omega = B$ was represented by the urn containing more red balls (5 red and 3 blue).

Figure 1: Overview of Experimental Design

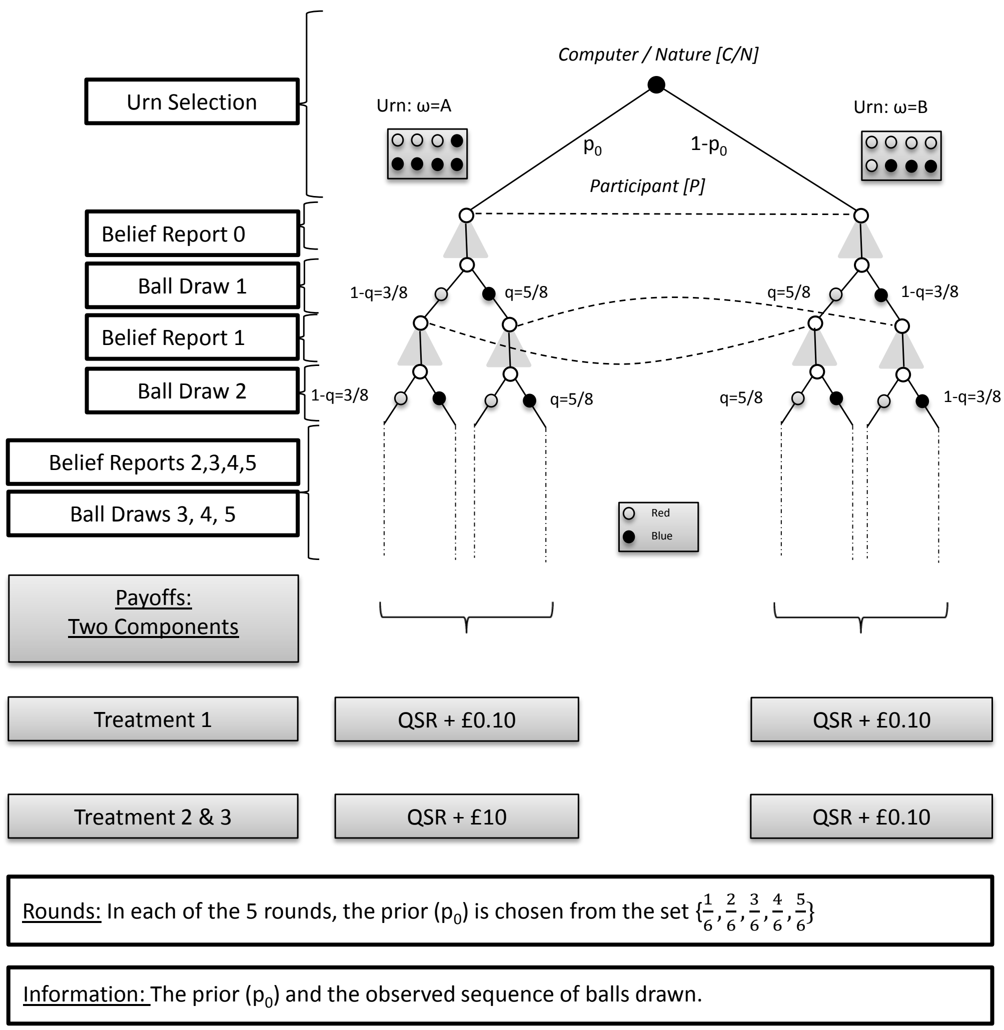

Priors

The five rounds differed from one another only in the exogenous prior probability that $\omega = A$ was the true state, with this prior, p_0, chosen from the set $\{\frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}\}$. In each round, this prior was known to the participant. The order of these rounds was randomly chosen for each individual. Conditional on the prior, p_0, one of the two urns was then chosen through the throw of a virtual die, independently for each individual, in each round.

Belief Updating

In each round, after being informed of this prior probability, p_0, the participant received a sequence of five partially informative signals, s_t, for $t = \{1, 2, 3, 4, 5\}$. These signals consisted of draws, with replacement, from the urn chosen for that round. Therefore, if the state of the world in a specific round was $\omega = A$ then the chance of drawing a red ball was $\frac{3}{8}$ and the chance of drawing a blue ball was $\frac{5}{8}$ for each of the draws in that round (see Figure 1).

In each round, we elicited the participant's reported belief, r_t, about the likelihood that state $\omega = A$ was the correct state of the world, six times (i.e. for $t = \{0, 1, 2, 3, 4, 5\}$). We first elicited her reported belief, r_0, directly after she was informed of the exogenous prior probability, p_0, and then after she received each of her five signals we elicited r_t for $t = \{1, 2, 3, 4, 5\}$. Overall, we therefore elicited 30 reported beliefs in Stage 1 from each individual (6 reported beliefs in each of 5 rounds).[24]

Incentives and Treatment Groups

The Belief Updating Task was identical across treatment groups with the exception of the incentives faced by participants. In each treatment, participants' payment consisted of two components: (i) an exogenous state-contingent payment[25], and (ii) an accuracy payment that depended their stated belief[26] and the true state (i.e. the QSR payment described in equations 7 and 8 above). In treatments T2.COMBINED and T3.SEPARATE, the state-contingent payment was substantially higher at £10 in state $\omega = A$ in comparison to £0.10 in state $\omega = B$, making $\omega = A$ the more attractive state of the world. In T1.SYMMETRIC, participants simply received an equal state-contingent payment of £0.10 in both states, $\omega = A$ and $\omega = B$, implying neither was preferable.

In all three treatments, participants received nearly identical detailed written instructions describing the belief updating task as well as the two payment components. In order to further simplify the task faced by participants and try to ensure that they understood the incentive environment they faced, we presented the QSR as a choice from a list of lotteries (this approach is also used, for example, by Armantier and Treich (2013) and Offerman et al. (2009)). To this effect, subjects

[24] Participants were not informed about the correct urn at the end of each round. They only received feedback when their payment was calculated at the very end of the experiment.

[25] This was called the "urn bonus" in the experiment.

[26] In the experiment, we refer to the object that individuals reported about the likelihood that one of the two urns was the chosen one as their "probability judgment". This corresponds to the reported belief in the discussion above.

were presented with *payment tables*, which informed them of the precise prospect they would face for each choice of r_t, in increments of 0.01. An abbreviated version of the three payment tables associated with each of the three treatment groups is presented in Table 2.[27] In order to represent all payments as integers in the instructions and payment tables, we adopted the approach of using experimental points. At the end of the experiment, these experimental points were converted to money using an experimental exchange rate of 6000 points = £1.

Table 2 highlights the difference between treatments T2.COMBINED and T3.SEPARATE. While, participants in these two treatments faced precisely the same incentives, the treatments differed in terms the salience of the hedging motive. In particular, the only difference between the two treatments was the way in which the payment information was summarised in the *payment table*. As shown in Table 2, in T2.COMBINED, the *payment table* showed the combined payment from both (i) the exogenous state-contingent payment, and (ii) the accuracy payment, together. Therefore, it summarised the reduced form prospect associated with each reported probability choice (r_t) for subjects.

The motivation behind having two treatments with identical incentives, but a different presentation of the incentives, was following. The rationale for T2.COMBINED was that presenting the incentives in a combined form is the simplest and clearest way of relaying the true incentives faced to participants. The rationale for T3.SEPARATE was that if participants "narrowly bracket", then the separate presentation of incentives could reduce the influence of the hedging motive, and therefore would induce more accurate belief reporting. By implementing both treatments, we were able to evaluate the influence of the presentation of the incentives on the reported beliefs. As we will see below, this influence was substantial and corresponds to the theoretical predictions for how a risk averse individual would act if she were hedging more when the hedging opportunity was made salient.

Furthermore, an additional benefit of running both treatments was that it provided us with a way to test the internal validity of the correction mechanism we use. We will see below that, while the uncorrected distribution of beliefs differ substantially between the two treatments, the distributions of the corrected beliefs are very similar to one another.

[27] Note, these payment tables are abbreviated in comparison to the tables presented to participants. The only substantive difference is that Table 2 contains 21 rows, one for each 5% increment in the reported belief. In contrast, the participants received payment tables that contained 101 rows, one for each 1% increase in the reported belief.

Table 2: Comparison of Incentive Summary Tables between Treatment Groups

T1 SYMMETRIC			**T2 COMBINED**			**T3 SEPARATE**		
Probability judgment for Urn A	**True Urn**		**Probability judgment for Urn A**	**True Urn**		**Probability judgment for Urn A**	**True Urn**	
	Urn A	**Urn B**		**Urn A**	**Urn B**		**Urn A**	**Urn B**
100	12 600	600	100	72 000	600	100	12 000	0
95	12 570	1 770	95	71 970	1 770	95	11 970	1 170
90	12 480	2 880	90	71 880	2 880	90	11 880	2 280
85	12 330	3 930	85	71 730	3 930	85	11 730	3 330
80	12 120	4 920	80	71 520	4 920	80	11 520	4 320
75	11 850	5 850	75	71 250	5 850	75	11 250	5 250
70	11 520	6 720	70	70 920	6 720	70	10 920	6 120
65	11 130	7 530	65	70 530	7 530	65	10 530	6 930
60	10 680	8 280	60	70 080	8 280	60	10 080	7 680
55	10 170	8 970	55	69 570	8 970	55	9 570	8 370
50	9 600	9 600	50	69 000	9 600	50	9 000	9 000
45	8 970	10 170	45	68 370	10 170	45	8 370	9 570
40	8 280	10 680	40	67 680	10 680	40	7 680	10 080
35	7 530	11 130	35	66 120	11 130	35	6 930	10 530
30	6 720	11 520	30	66 120	11 520	30	6 120	10 920
25	5 850	11 850	25	65 250	11 850	25	5 250	11 250
20	4 920	12 120	20	64 320	12 120	20	4 320	11 520
15	3 930	12 330	15	63 330	12 330	15	3 330	11 730
10	2 880	12 480	10	62 280	12 480	10	2 280	11 880
5	1 770	12 570	5	61 170	12 570	5	1 170	11 970
0	600	12 600	0	60 000	12 600	0	0	12 000

The Offerman et al. (2009) Correction Task (Stage 2)

In the second stage of the experiment we elicited the twenty reported beliefs, r, for events with known objective probabilities required to estimate the incentive distortion function, R, for each individual. In each of the three treatments, we estimate the function R using belief from Stage 2 elicited under the same incentive environment as in the Belief Updating Task in Stage 1.

In Stage 2, participants were asked to report their probability judgment regarding the likelihood that statements of the form: "the number the computer chooses will be between 1 and 75" (i.e. $p = 0.75$), were true, after being told that the computer would randomly choose a number between 1 and 100, with each number equally probable. For T1.SYMMETRIC, this specific example of the probability of the randomly chosen number being in the interval between 1 and 75 essentially involves choosing r from the list of prospects defined by $1-(1-r)^2{}_{0.75}(1-r^2)$. For T2.COMBINED and T3.SEPARATE, this example would involve choosing r from the list of prospects defined by $x+1-(1-r)^2{}_{0.75}(1-r^2)$. As in Stage 1, in each of the treatments, the Stage 2 payment table summarized the relevant payment information. For each treatment, this payment table contained identical values in Stage 1 and Stage 2. Therefore, Table 2 above also provides a summary of the Stage 2 payment tables.

The twenty reported beliefs corresponded to the objective probabilities $0.05, 0.1, \ldots, 0.95$.[28] At the end of the experiment, one choice from Stage 2 was randomly chosen to contribute to each participant's final payment.

4 Data and Descriptive Evidence

The experiment was conducted at the UCL-ELSE experimental laboratory in London as well as at the WZB-TU laboratory in Berlin, with two sessions for each of the three Treatment groups at each location, making twelve sessions and 222 participants in total.[29] At both locations, participants were solicited through an online database using ORSEE (Greiner, 2015) and the experiment was run using the experimental software, z-Tree (Fischbacher, 2007). On average, sessions lasted

[28] These objective probabilities were presented in a random order. Furthermore, in the likelihood statement used, such as the example in the main test, the interval given started at a randomly chosen lower bound (from the feasible set of lower bounds for that specific objective probability). Nineteen of the twenty reported beliefs were unique, and the twentieth was a randomly chosen repitition of one of the first nineteen.

[29] The experiment was therefore completely symmetric between the two locations, with approximately the same number of participants in each treatment group from each location.

about 1.5 hours and the average participant earned £19.7 in London and €20.3 in Berlin. Realized payments ranged between £11 [€11] and £34 [€34].

One challenge faced by belief updating studies is ensuring that subjects understand and engage with the task. In order to facilitate this, we were careful to ensure that the instructions received were as clear and simple as possible. Nonetheless, there remained a non-trivial fraction of participants who took decisions to update in the 'incorrect' direction[30] upon receiving new information. In order to ensure that the behavior we are studying is reflective of actual updating behavior of individuals who understood and engaged with the task, we restrict our sample for our main analysis by removing rounds where an individual updates in the incorrect direction[31]. However, we also estimate all the main results on the full sample, and the general patterns of behavior are similar.

While randomization to treatment group should ensure that the samples are balanced on observable and unobservable characteristics, Table 10 provides a check that the selection of our preferred sample has not substantially biased our treatments groups by reporting the sample means of a set of individual characteristics for each treatment group. Overall, the treatments appear to be balanced, with the exception that individuals in T3.SEPARATE are more likely to speak English at home than individuals in T2.COMBINED.

Figure 2 provides a first overview of the reported belief data by reporting the cumulative density functions of reported beliefs for each of the three treatment groups. One interesting feature of this figure is that, in spite of the fact that the T2.COMBINED and T3.SEPARATE treatments are offering subjects precisely the same incentives with only the way in which they are presented being slightly varied, the distributions of reported beliefs of these two groups are significantly different from one another (Mann-Whitney rank-sum test, $p < 0.01$). The larger mass of reported beliefs to the left of 50 in T2.COMBINED suggests that individuals are more likely to respond to the hedging opportunity when the incentives are presented in the simpler reduced form lottery format. The correction mechanism we use helps to remove this distortive effect of the incentive environment on the reported beliefs.

[30]'Incorrect' in the sense that the participant received a signal that increased the likelihood that one state was the true state, but the individual updated her reported belief in the opposite direction. For example, if the individual drew a red ball, but increased her reported belief regarding the probability that the state was $\omega = A$, then we would label this as an update in the incorrect direction.

[31]More specifically, for our preferred sample, we remove all individuals who make more than five out of twenty five updating decisions in the incorrect direction. This comprised eighteen percent of our subjects. This is comparable to the twenty five percent of individuals who make at least one mistake out of four decisions in Möbius et al. (2014). For those who updated in the wrong direction five or fewer times, we remove only the round of decisions in which they made a mistake. This removes 210 out of the remaining 910 rounds. We estimate our results on the full sample and all the main results remain the same.

Figure 2: CDF of Reported Beliefs

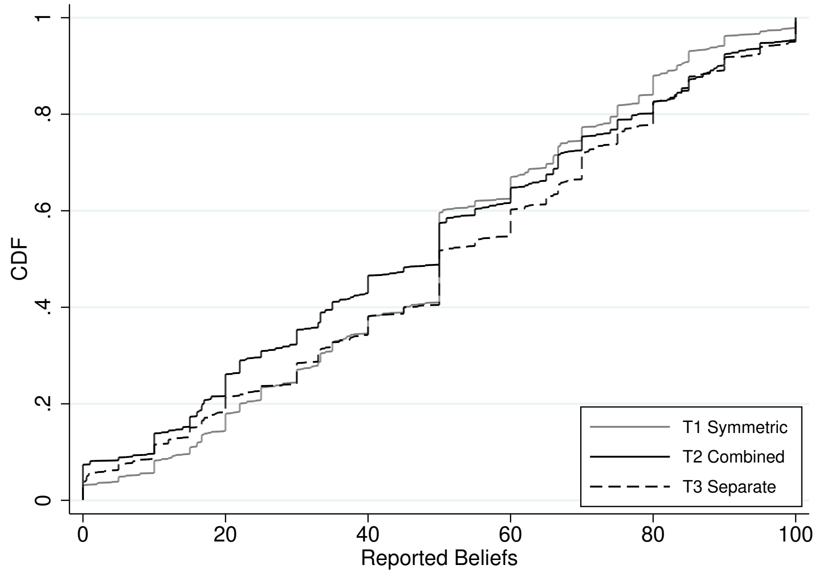

5 Empirical Specification

In this section, we first discuss the calibration exercise used to correct the reported beliefs. Second, we describe the core estimation equations used in our analysis. These build on the work by Möbius et al. (2014) as their data has a similar structure to ours. One key difference in our data is the exogenous assignment of the participants' entire information set. We exploit this feature of our data to address endogeneity issues that can arise when studying belief updating.

Calibration of the Belief Correction Procedure

The belief correction procedure that we adopt involves assuming a flexible parametric form for the participants' utility and probability weighting functions in order to estimate the R function discussed in Equation 9 above. We estimate this function for each individual separately in order to correct the reported beliefs at the individual level. In addition, we estimate this function at the aggregate level for each of the treatment groups in order to obtain a measure of the average distortion of the incentive environment faced in each of the treatment groups. A detailed discussion of the mechanics of the Belief Correction Procedure we use is relegated to Appendix B.3. Essentially, we are simply fitting a curve through each subject's belief elicitation incentive distortion.

Figure 3: Average Correction Functions across Treatments.

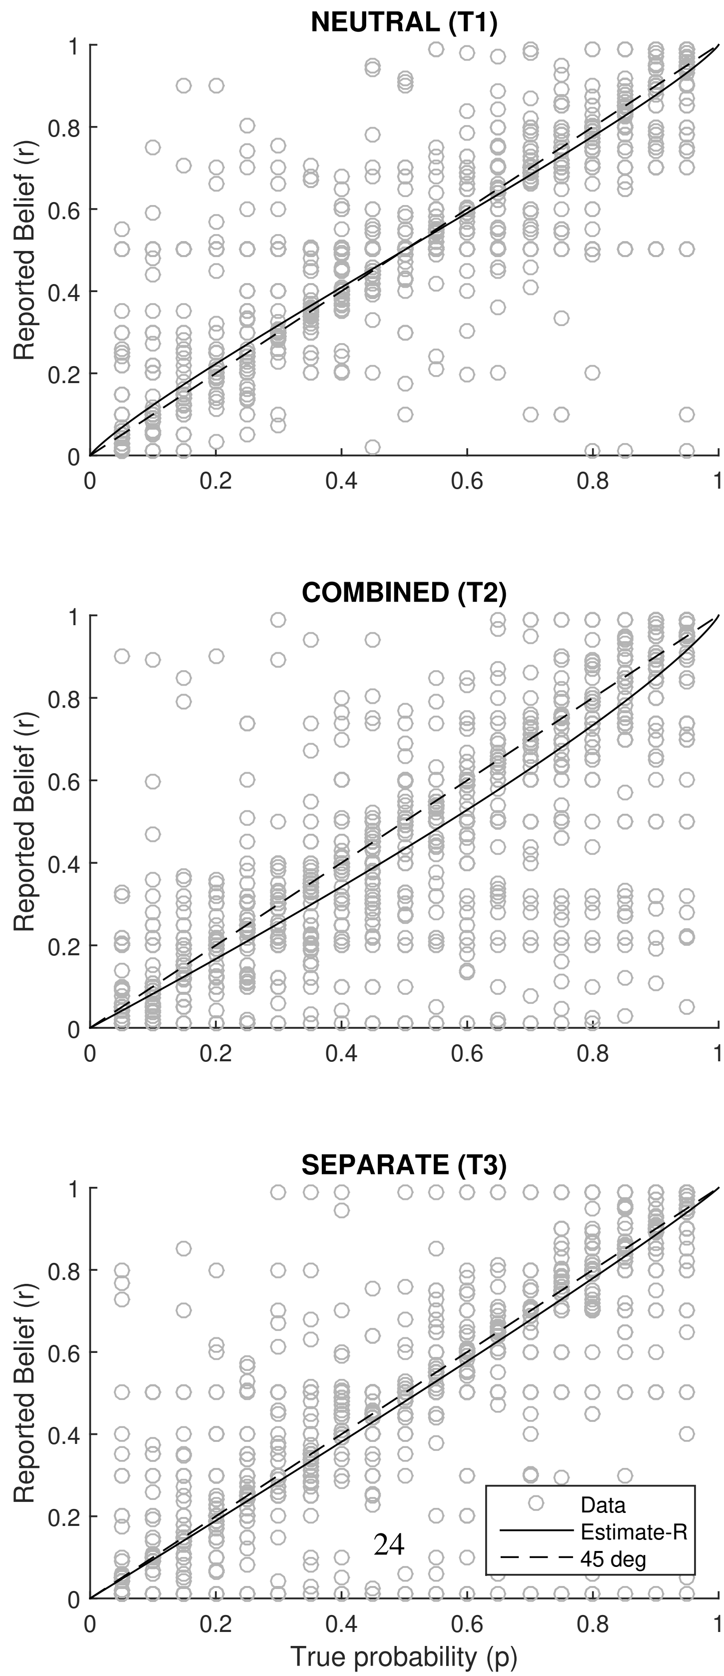

Figure 4: Individual Level Estimates of the Incentive Correction Function.

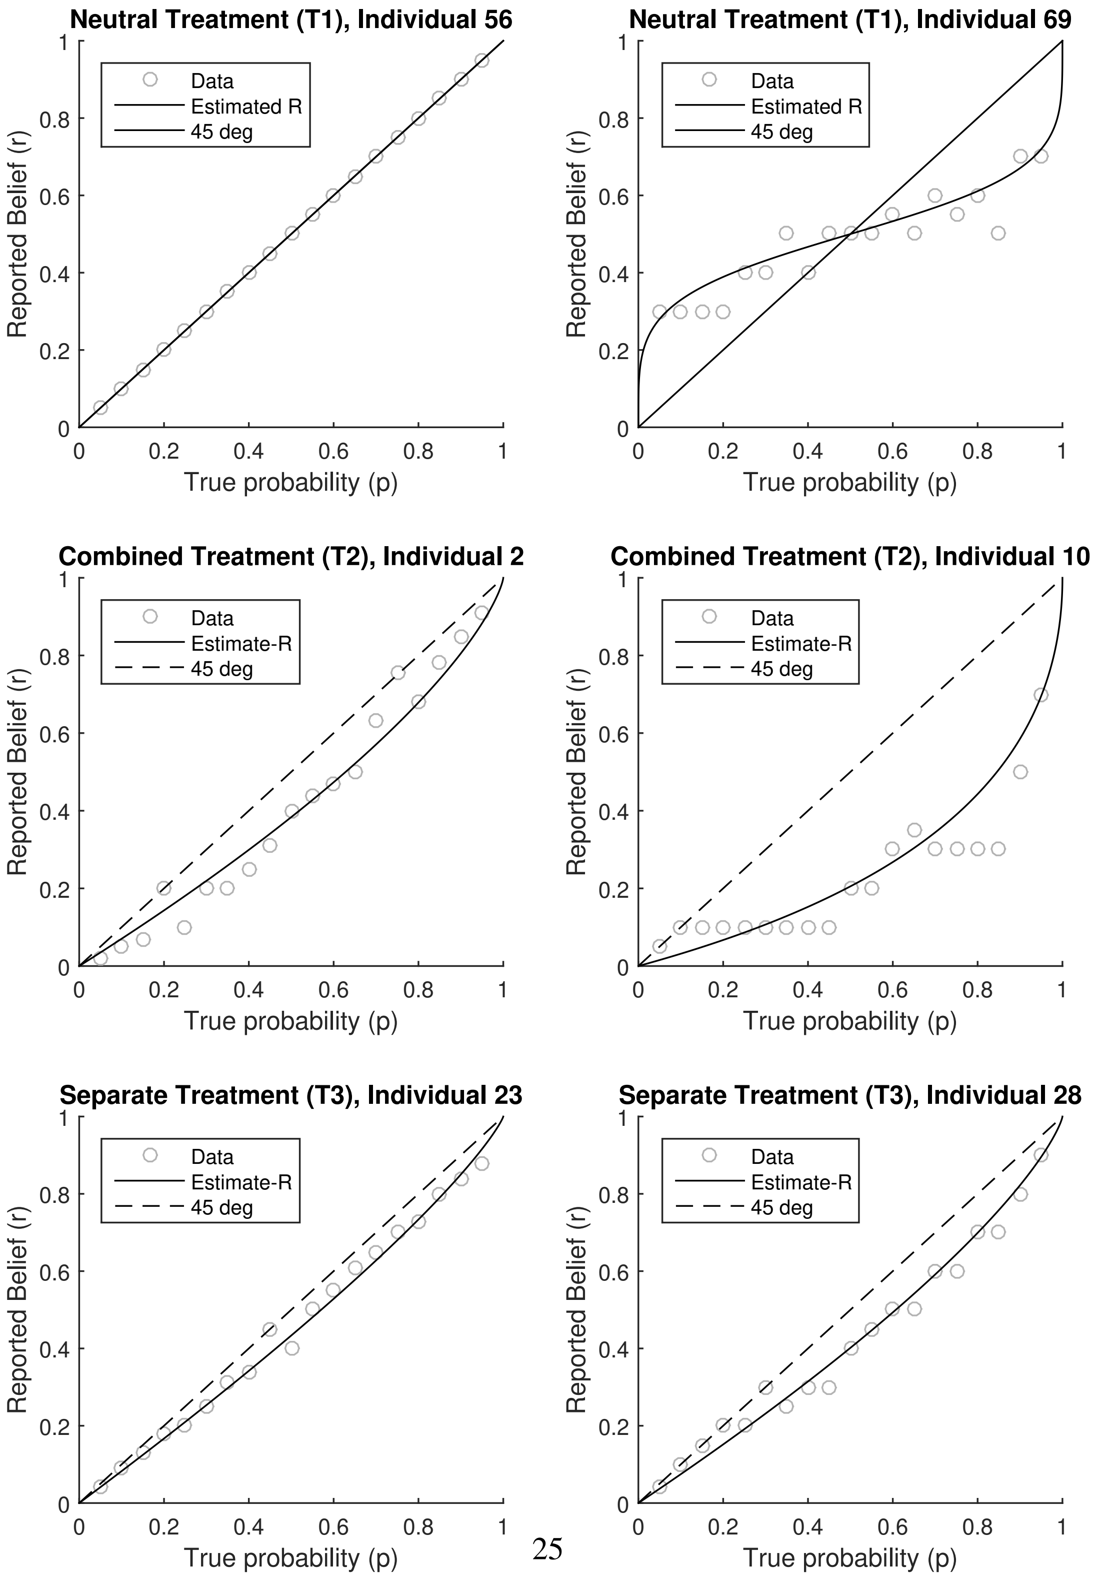

Figure 3 displays the average correction curves for each of the treatment groups. Comparing the three subgraphs, we see that the average individual distorts the beliefs she reports in a way that is consistent with what risk aversion under EU would predict, with the inverse-S shape distortion in the T1.SYMMETRIC stakes treatment and the strong distortion downwards (away from the more desirable state) in both of the ASYMMETRIC stakes treatments. Furthermore, we see that the different ways of framing the same incentives in the two ASYMMETRIC treatments has a clear influence on behavior. In T2.COMBINED, the participants hedge far more when choosing their reported beliefs in comparison to those in the T3.SEPARATE group. This is in spite of the fact that the incentives are identical in these two treatments. This indicates that the reported beliefs in the T3.SEPARATE treatment are closer to the participants' true beliefs and motivates this presentation of incentives as preferable for future work that calls for the elicitation of beliefs when there are exogenous state-contingent payments.

At the individual level, there is a large degree of heterogeneity in the degree to which individuals distorted their reported belief away from their actual belief, given the incentive environment. Figure 4 displays the correction curves estimated for two individuals from each treatment group. It is clear from this figure that some individuals responded very strongly to the incentive environment in which their belief was elicited, while others reported their belief more accurately. The belief correction procedure is therefore very helpful for recovering the true beliefs of participants who responded strongly to the incentive environment. In cases where the individual simply reported their belief accurately, the corrected beliefs and the reported beliefs are exactly the same.

In Figure 5 we plot the CDFs of the uncorrected reported beliefs as well as the reported beliefs that have been corrected at the individual level. This figure suggests that the belief correction approach that we use was successful in removing the strong hedging influence in the COMBINED treatment as the correction procedure removes the majority of the difference between the distributions of beliefs for the two treatments with identical incentives, COMBINED and SEPARATE. This evidence provides support for the hypothesis that combining the payment incentives into a single summary table made it easier for the subjects to identify the hedging opportunity and distort the belief they reported away from the belief they held in the COMBINED treatment. Furthermore, this evidence underscores the usefulness of the belief correction mechanism in reducing the bias in the reported beliefs.

Figure 5: CDFs of Corrected and Uncorrected Reported Beliefs

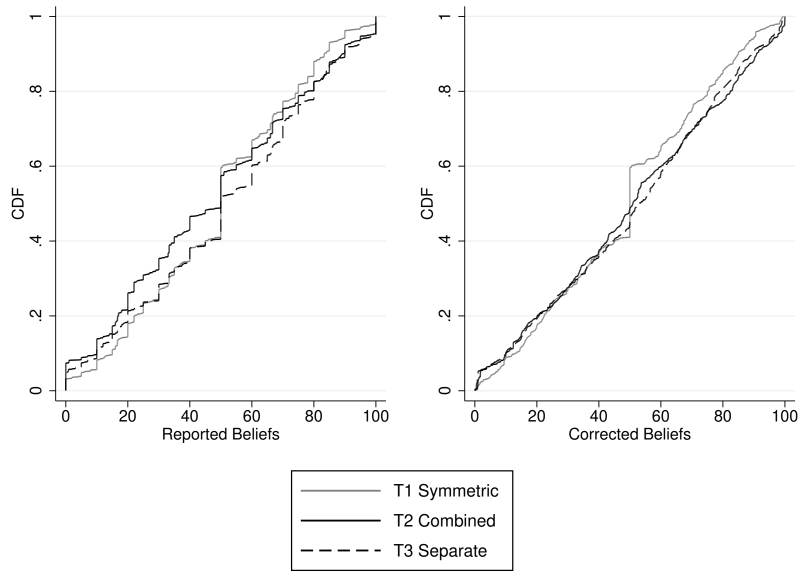

With the corrected beliefs in hand, we can proceed to the main analysis of belief updating in our sample. The analysis is done using both the corrected and uncorrected beliefs.[32]

Core Estimation Specifications

Our core estimation equations aim to test for systematic patterns in updating behavior, within the framework developed in Equation 6. Firstly, we examine whether there are systematic deviations from Bayes' rule in updating behavior, independent of having a stake in one of the two states of the world. Secondly, we assess the influence that having a stake in one of the two states of the world has by (i) testing for an asymmetry in updating *within* treatments where there is a state-contingent stake (i.e. T2.COMBINED and T3.SEPARATE); and (ii) testing whether there are differences in updating behavior *between* the treatments with and without a state-contingent stake.

The first estimation equation follows directly from Equation 6, allowing us to test the asymmetric updating hypothesis, and also to test for other common deviations from Bayes' rule, such as a

[32]The analysis using the uncorrected beliefs is essentially an analysis using beliefs elicited using the QSR, which is by far the most common approach used in the belief elicitation literature to date. Therefore, these results have greater comparability with much of the experimental literature. However, in line with the discussion above, our preferred specification uses the corrected beliefs as these should be closer to the subjects' true beliefs.

confirmatory bias or *base rate neglect*[33]:

$$\tilde{\pi}_{i,j,t+1} = \delta\tilde{\pi}_{i,j,t} + \gamma_a\tilde{q}\cdot 1(s_{i,j,t+1} = a) - \gamma_b\tilde{q}\cdot 1(s_{i,j,t+1} = b) + \epsilon_{i,j,t+1} \quad (10)$$

where $\tilde{\pi}_{i,j,t} = \text{logit}(\pi_{i,j,t})$ and $\tilde{q} = \log(\frac{q}{1-q})$; j refers to a round of decisions; and the errors, ϵ_{ijt}, are clustered at the individual (i) level.

In order to determine the belief updating pattern *within* each incentive environment, we estimate this equation separately for each treatment; and then to test for significant differences *between* the coefficients in different incentive environments, we pool our sample and interact the treatment variable with all three of the coefficients in this equation (i.e. δ,γ_a, and γ_b). This provides us with a test of whether the parameters differ between either of the two ASYMMETRIC treatments and the SYMMETRIC treatment.

Endogeneity of the Lagged Belief

One potential concern with the identification of the parameters of Equation 10 is the common issue where the right hand side contains lagged versions of the dependent variable. This implies that there is a possible endogeneity of the lagged beliefs, $\pi_{i,j,t}$, if they are correlated with the error term[34], $E\{\tilde{\pi}_{i,j,t}\epsilon_{i,j,t+1}\} \neq 0$. If this is the case, it can result in biased and inconsistent estimators for the parameters of the regression. Our experiment was designed to avoid this issue by virtue of exogenously assigning the subjects' entire information set. This allows us to use the exogenously assigned prior probability of state $\omega = A$ being the true state, $p_{i,j,t=0}$, as well as the sequence of signals observed, s_t, to construct an instrument for the lagged belief, $\pi_{i,j,t}$, in Specification 10.[35] We do this by calculating the objective Bayesian posterior, given the agent's information set at time t, and using this as an instrument for her belief, $\pi_{i,j,t}$.

The approach used here also avoids a second type of endogeneity issue that can arise when studying belief updating when the states of the world are a functions of *personal characteristics* (e.g. when examining beliefs regarding individual attributes, such as one's own skills, IQ, or beauty) or

[33]Here, base rate neglect is the tendency to underweight or partially disregard one's prior regarding which urn is more likely to be generating the signal. This implies that when the individual receives a new signal, she weights information that supports her prior less heavily than information that opposes it, relative to a Bayesian. In this case, when one's prior regarding state $\omega = A$ is greater than 0.5, $\pi_{t-1} > 0.5$, a participant who is prone to *base rate neglect* weights signals that support state $\omega = A$ less heavily than signals that support state $\omega = B$; and vice versa when her prior is $\pi_{t-1} < 0.5$, using Bayesian updating as the benchmark.

[34]For example, this would be the case if there is individual heterogeneity in the way individuals respond to information. We provide evidence below that this individual heterogeneity is present.

[35]More precisely, we instrument using the logit of the accurate Bayesian posterior, logit($p_{i,j,t}$), as an instrument for the logit of the lagged belief, $\tilde{\pi}_{i,j,t} = \text{logit}(\pi_{i,j,t})$ in Specification 10.

personal choices. When this is the case, the conditional probability of observing a specific signal depends on the state of the world, and therefore can be correlated with personal characteristics.[36]

6 Results

Table 3 reports the results from estimating Equation 10 for each of the treatment groups separately. These estimates provide an overview of the updating behavior of the average individual within each of the three treatment groups. Within each treatment group, we report the results for both the OLS (top panel) and the IV (bottom panel) estimates discussed above. Columns (1a), (2a) and (3a) use the uncorrected reported beliefs, while columns (1b), (2b) and (3b) use the corrected beliefs. Every coefficient in the table is statistically different from 0 at the 1% level. Since our primary interest is in testing whether the coefficients are different from 1, in this table, we use asterisks to reflect the significance of a t-test of whether a coefficient is statistically different from 1.

Perhaps the most striking features of this table are: (i) the similarity in the updating patterns across the three treatment groups; and (ii) that for the average individual, the observed updating behavior is close to Bayesian in all three treatment groups. The p-values from the test of the null hypothesis, $H_0 : \gamma_a = \gamma_b$, show that in none of the three treatment groups do we observe a statistically significant difference (at the 5% level) between the responsiveness to the signals in favor of $\omega = A$ and $\omega = B$ (i.e. we don't observe asymmetric updating).

Both the OLS results in the top panel, and the IV results in the bottom panel, indicate that the responsiveness to new information of the average participant was not statistically different to that of a Bayesian, since both γ_a and γ_b are not significantly different to 1 at the 5 percent level. The primary difference between the OLS results and the IV results is that, while the OLS estimates suggest a small degree of base rate neglect across all three treatments ($\delta < 1$), once we control for

[36]Consider the following toy example to illustrate this. Assume there are two states of the world - being TALL and SHORT - and being tall is viewed as *"good"*. Individuals are initially uncertain about their height group (because only relative height is important). Now, assume that TALL individuals are more responsive to new information than SHORT individuals. For example, assume TALL individuals update exactly according to Bayes' rule, but SHORT individuals don't update their prior beliefs at all in response to new information. Now, let subjects receive a sequence of noisy signals that are informative (in the sense that, for each type, the signal they receive is correct more than fifty percent of the time). In this toy example, both types respond symmetrically to 'good-news' and 'bad-news', but because TALL individuals are always more likely to receive 'good-news', and SHORT individuals are always more likely to receive 'bad-news', the resulting aggregate behavior will appear to display the "good-news, bad news effect". This is due to the fact that, on average, a good signal was more likely received by a TALL individual, who updates like a Bayesian, and a bad signal is more likely to go to a SHORT individual who doesn't update at all. Furthermore, updating would display conservatism on average. See Appendix C for further details.

the possible sources of endogeneity discussed above using our instrumental variable strategy, the estimates no longer indicate base rate neglect. Interestingly, although we observed a substantial difference in the levels of the corrected and uncorrected beliefs in Figure 5 above, the estimates for updating in Table 3 are quite similar.

Table 3: Average Updating Behavior across Treatments

	T1: SYMMETRIC		T2: COMBINED		T3: SEPARATE	
	Reported (1a)	Corrected (1b)	Reported (2a)	Corrected (2b)	Reported (3a)	Corrected (3b)
OLS						
δ	0.90	0.90	0.86	0.86	0.91	0.93
	(0.03)***	(0.03)***	(0.04)***	(0.04)***	(0.03)***	(0.02)***
γ_a	1.12	1.09	1.05	1.06	1.25	1.16
	(0.12)	(0.11)	(0.12)	(0.12)	(0.14)*	(0.11)
γ_b	1.21	1.17	1.19	1.12	1.20	1.13
	(0.13)*	(0.11)	(0.15)	(0.13)	(0.12)	(0.10)
p $(H_0 : \gamma_a = \gamma_b)$	0.30	0.32	0.20	0.48	0.58	0.73
N	1,075	1,075	1,285	1,285	1,140	1,140
R^2	0.71	0.73	0.73	0.74	0.81	0.84
IV						
δ	0.99	0.99	1.00	0.99	0.99	0.99
	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.02)
γ_a	1.12	1.09	1.08	1.02	1.22	1.14
	(0.12)	(0.11)	(0.12)	(0.11)	(0.13)*	(0.11)
γ_b	1.21	1.16	1.11	1.12	1.18	1.11
	(0.12)*	(0.11)	(0.13)	(0.12)	(0.12)	(0.10)
p $(H_0 : \gamma_a = \gamma_b)$	0.30	0.31	0.76	0.25	0.63	0.74
N	1,075	1,075	1,285	1,285	1,140	1,140
1st Stage F	61.38	84.04	107.98	107.01	76.90	95.45

(i) Standard errors in parentheses (clustered at the individual level).
(ii) All coefficients are significantly different from 0 at the 1% level. Therefore, t-tests of the null hypothesis (H_0: Coefficient $=$ 1) are reported: * = 10%, ** = 5%, *** = 1%.
(iii) The rows corresponding to p $(H_0 : \gamma_a = \gamma_b)$ report the p-statistic from a t-test of the equality of the coefficients γ_a and γ_b (i.e. a test of the asymmetric updating hypothesis).

Robustness Exercises

To check for the robustness of the belief updating results for the average individual presented in Table 3, we conducted several robustness exercises. These exercises, and their corresponding results, are discussed in detail in Appendix A.

The first exercise considers a comparison of information-set-equivalent posterior beliefs after individuals have received only (i) the exogenous prior and (ii) a single ball draw. This is an informative exercise as it allows for a model-free test of the *asymmetric updating hypothesis* in situations where the exogenous information set is kept constant, and the only factor that differs is the payment associated with different states. The second subsection examines whether the results from the main specification described in Equation 10 are robust to first differencing the dependent variable (i.e. this considers how new information shifts the change in beliefs, imposing the assumption that $\delta = 1$). The third subsection extends the main empirical specification to allow for individual-specific updating parameters. The fourth subsection pools all the observations across the three treatments together, and then tests whether the average updating parameters differ across treatments by interacting treatment group dummies with the regressors of the main specification described in Equation 10.

The results from all of these exercises are highly consistent with those in Table 3 and fail to provide any evidence in favor of an asymmetry in updating.

Heterogeneity in Updating Behavior

In order to investigate whether the aggregate results are masking heterogeneity in updating behavior, we estimate Specification 10 at the individual level and collect the parameters. The distributions of these individual level parameters are reported in Figure 6. Perhaps the most conspicuous feature of this figure is the fact that all three treatment groups display such similar distributions in each of the panels - i.e. for each of the parameters. The upper-right panel shows that the majority of individuals have an estimated δ parameter in the interval $[0.6, 1.1]$, with a large proportion of these concentrated around 1 in all three treatment groups. The two left-hand panels show that there is substantially more individual heterogeneity in the estimated γ_a and γ_b parameters, which are dispersed over the interval $[0, 3.5]$ in all three treatments.

Figure 6: Distributions of Individual Level Updating Parameters.

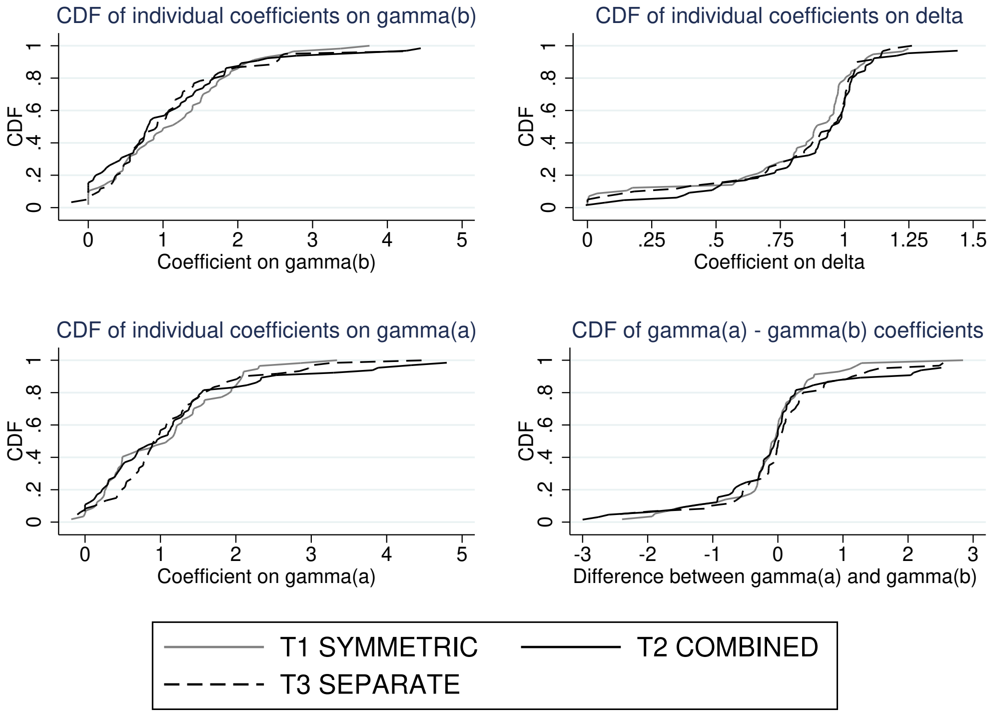

With such a large degree of variation in the individual level parameters, a natural conjecture to make is that, while we do not observe asymmetric updating at the aggregate level, it is entirely plausible that there may be a subsample of individuals who are optimistic updaters and another subsample of individuals who are pessimistic updaters. If these two subsamples are of a similar size and their bias is of a similar magnitude, we would observe no asymmetry at the aggregate level. The lower-right panel of Figure 6 suggests that this is not the case by plotting the distribution of the individual level difference between the γ_a and γ_b parameters. The majority of the distribution is concentrated in a narrow interval around 0 for all three treatment groups, suggesting that there is no asymmetry for any sizable subsample. In order to further investigate whether there exists a subgroup of asymmetric updators, we estimate a a mixture model, which uses maximum likelihood to classify the data into the predominant 'types' of updaters in the sample, where the 'types' are determined endogenously. The results are reported in Appendix D and also show no support for the existence of a subgroup of asymmetric updators in any of the treatments groups.

7 Concluding Discussion

The objective of this paper was to study belief updating when an individual prefers one state of the world to another. In particular, the experiment was designed to test the *asymmetric updating hypothesis* in the domain of financial decision making, and contribute to the body of work that is constructing a descriptive understanding of how individuals form beliefs.

The main finding of the current paper is that, in contrast to the literature considering belief updating over ego-related attributes, we find no evidence for asymmetric updating when there are financial state-dependent stakes. Instead, we find that the updating behavior of the average individual is approximately Bayesian, irrespective of the presence or absence of financial stakes. This is highly consistent with the results reported in Coutts (2016) and Gotthard-Real (2017), who also consider belief updating in the presence of financial rewards. The current paper complements this contemporaneous work in several ways. Firstly, we consider belief updating from a wider range of prior beliefs. Secondly, our experimental design ensures that the distribution of realised signals observed is balanced in terms of the frequency of 'good' and 'bad' signals. This is useful as it removes the potential influence of the signal distribution that is documented by Coutts (2016). Thirdly, our data allows us to conduct both *within subject* and *between subject* tests, as well as several robustness exercises in support of the main result.

Exploring heterogeneity in updating behavior, we find that the Bayesian behavior of the average individual masks a large degree of heterogeneity in individual's responsiveness to new information, however, we don't find any evidence in support of a sizeable subgroup of asymmetric updators.

Overall, the results suggest that the average individual updates her beliefs symmetrically in contexts of *structured belief updating,*[37] even if one of the states yields a larger financial reward.

One caveat to the results described here is that the financial stakes in play are not extremely large, and therefore it is feasible that the results would change in the presence of larger financial stakes. While small stake sizes are a standard caveat to most laboratory experiments, it is perhaps a larger concern here, since the asymmetric updating hypothesis may rely on the degree to which individuals desire the occurance of the good state of the world. However, Coutts (2016) investigates the role played by stakes, increasing the stake size up to \$80 and finds no evidence that it plays a role.

[37] It is useful to consider a taxonomy of the papers in this literature along two dimensions: (i) the domain over which beliefs are formed (e.g over attributes relating to *self-image* or events associated with *financial outcomes*); and (ii) whether they examine *structured belief updating* (updating from a prior belief known to the analyst) or *unstructured belief formation* (forming a belief where priors are unknown to the analyst, or on the basis of ambiguous information).

This suggests that stake size may not be a pivotal concern. Nonetheless, it is worth keeping this caveat in mind when interpreting the results.

The aim of this paper was to contribute towards a descriptive understanding how individuals update their beliefs in the presence of state-contingent stakes. This is instructive as there is a large class of economically important situations in which individuals form beliefs, preferring one state to another, ranging from capital markets to intertemporal portfolio choice problems to consumption savings problems. An appealing conclusion of this paper is that, while standard economic models may not be appropriate when an individual's ego is involved, in a large class of situations when an individual's ego is not involved, standard theory, which assumes Bayesian updating, may do just fine.

References

Akerlof, G. A. and W. T. Dickens (1982). The economic consequences of cognitive dissonance. *American Economic Review 72*(3), 307–319.

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

Attanasio, O. P. (2009). Expectations and perceptions in developing countries: Their measurement and their use. *American Economic Review 99*(2), 87–92.

Ayton, P. and I. Fischer (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory and Cognition 32*(8), 1369–1378.

Benabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics 117*, 871–915.

Bénabou, R. and J. Tirole (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives 30*(3), 141–164.

Bernardo, A. E. and I. Welch (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy 10*(3), 301–330.

Bolton, G. E. and A. Ockenfels (2000). Erc: A theory of equity, reciprocity, and competition. *American Economic Review 90*(1), 166–193.

Brunnermeier, M. and J. Parker (2005). Optimal expectations. *American Economic Review 95*(4), 1092–1118.

Burks, S. V., J. P. Carpenter, L. Goette, and A. Rustichini (2013). Overconfidence and social signalling. *Review of Economic Studies 80*(3), 949–983.

Caplin, A. and J. Leahy (2001). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics 116*(1), 55–79.

Chamley, C. (2004). *Rational Herds: Economic Models of Social Learning*. Cambridge University Press.

Costa-Gomes, M. A., S. Huck, and G. Weizsäcker (2014). Beliefs and actions in the trust game: Creating instrumental variables to estimate the causal effect. *Games and Economic Behavior 88*, 298–309.

Coutts, A. (2016). Good news and bad news are still news: Experimental evidence on belief updating. *Working Paper*.

Croson, R. and J. Sundali (2005). The gambler's fallacy and the hot hand: Empirical data from casinos. *Journal of Risk and Uncertainty 30*(3), 195–209.

Delavande, A. (2014). Probabilistic expectations in developing countries. *Annual Review of Economics 6*(1), 1–20.

Eil, D. and J. Rao (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics 3*(2), 114–138.

Enke, B. and F. Zimmermann (2015). Correlation neglect in belief formation. *Working Paper*.

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.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance 25*(2), 383–417.

Fama, E. F. (1976). Efficient capital markets: reply. *Journal of Finance 31*(1), 143–145.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics 114*(3), 817–868.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton University Press.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory 118*(2), 133–173.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.

Gotthard-Real, A. (2017). Desirability and information processing: An experimental study. *Economics Letters 152*, 96–99.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association 1:1*, 114–125.

Grether, D. M. (1978). Recent psychological studies of behavior under uncertainty. *American Economic Review 68*(2), 70–74.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *Quarterly Journal of Economics 95*, 537–557.

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

Hall, R. (1978). Stochastic implications of the life-cycle permanent income hypothesis. *Journal of Political Economy 86*, 971–987.

Heger, S. and N. Papageorge (2016). We should totally open a restaurant: How optimism and overconfidence affect beliefs. *Working Paper*.

Heifetz, A., C. Shannon, and Y. Spiegel (2007). The dynamic evolution of preferences. *Economic Theory 32*(2), 251.

Holt, C. and A. Smith (2009). An update on bayesian updating. *Journal of Economic Behavior & Organization 69*(2), 125–134.

Hossain, T. and R. Okui (2013). The binarized scoring rule. *Review of Economic Studies 80*(3), 984–1001.

Hurd, M. D. (2009). Subjective probabilities in household surveys. *Annual Review of Economics 1*, 543–562.

Johnson, D. D. and J. H. Fowler (2011). The evolution of overconfidence. *Nature 477*(7364), 317–320.

Kahneman, D., J. L. Knetsch, and R. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review 76*, 728–741.

Kahneman, D. and A. Tversky (1973). On the psychology of prediction. *Psychological Review 80*(4), 237.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–291.

Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica 77*(2), 603–606.

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

Kőszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics 121*(4), 1133–1165.

Lewellen, J. and J. Shanken (2002). Learning, asset-pricing tests, and market efficiency. *Journal of Finance 57*(3), 1113–1145.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal 97*, 666–684.

Lucas, R. E. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 1, pp. 19–46. North-Holland.

Manski, C. F. (2004). Measuring expectations. *Econometrica 72*(5), 1329–1376.

Mayraz, G. (2013). The optimism bias: Experiment and model. *Working Paper*.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics 51*, 247–257.

Möbius, M., M. Niederle, P. Niehaus, and T. Rosenblat (2014). Managing self-confidence. *Working Paper*.

Mossin, J. (1968). Optimal multiperiod portfolio policies. *Journal of Business 41*(2), 215–229.

Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association 55*, 299–306.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica 29*, 315–335.

Nyarko, Y. and A. Schotter (2002). An experimental study of belief learning using elicited beliefs. *Econometrica 70*(3), 971–1005.

Offerman, T., J. Sonnemans, G. van de Kuilen, and P. Wakker (2009). A truth serum for non-bayesians: Correcting proper scoring rules for risk attitudes. *Review of Economic Studies 76*(4), 1461–1489.

Palfrey, T. R. and S. W. Wang (2012). Speculative overpricing in asset markets with information flows. *Econometrica 80*(5), 1937–1976.

Prelec, D. (1998). The probability weighting function. *Econometrica 66*(3), 497–528.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review 83*(5), 1281–1302.

Rabin, M. and J. Schrag (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics 114*(1), 37–82.

Rabin, M. and D. Vayanos (2010). The gambler's and hot-hand fallacies: Theory and applications. *Review of Economic Studies 77*(2), 730–778.

Ramaswamy, V., W. S. DeSarbo, D. J. Reibstein, and W. T. Robinson (1993). An empirical pooling approach for estimating marketing mix elasticities with pims data. *Marketing Science 12*(1), 103–124.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics 51*, 239–246.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica 57*, 571–587.

Schotter, A. and I. Trevino (2014). Belief elicitation in the laboratory. *Annual Review of Economics 6*(1), 103–128.

Schwardmann, P. and J. J. Van der Weele (2016). Deception and self-deception.

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

Trautmann, S. and G. van de Kuilen (2015). Belief elicitation: A horse race among truth serums. *Economic Journal 125*, 2116–2135.

Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty 5*(4), 297–323.

Wakker, P. (2010). *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press.

APPENDICES

Appendix A: Robustness Checks

The empirical specification used in the main text of this paper assumes that updating follows the flexible parametric process described in Equation 6. This specification allows for a wide range of deviations from Bayes' rule, as discussed in Section 2. However, as one of the primary objectives of this paper is to test the *asymmetric updating hypothesis,* this section provides further evidence assessing this hypothesis, while relaxing the assumptions of the parametric process described in Equation 6.

The first subsection considers a comparison of information-set-equivalent posterior beliefs after individuals have received only (i) the exogenous prior and (ii) a single ball draw. This is an informative exercise as it allows for a model-free test of the *asymmetric updating hypothesis* in situations where the exogenous information set is kept constant and the only factor that differs is the payment associated with different states. The second subsection examines whether the results from the main specification described in Equation 10 are robust to first differencing the dependent variable (i.e. this considers how new information moves the change in beliefs, imposing the assumption that $\delta = 1$). The third subsection extends the main empirical specification to allow for individual-specific updating parameters. The fourth subsection pools all the observations across the three treatments together, and then tests whether the average updating parameters differ across treatments, by interacting treatment group dummies with the regressors of the main specification described in Equation 10.

Robustness Check 1: Model Free Test of the *Asymmetric Updating Hypothesis*

Perhaps the simplest and most direct test of the *asymmetric updating hypothesis* is obtained by comparing the posterior beliefs formed in two contexts where the information set is identical, but the rewards associated with one of the states of the world are varied. This allows us to test the *asymmetric updating hypothesis* while remaining agnostic regarding the process that guides belief updating, testing only whether it is symmetric. In the context of our experiment, we can achieve this model-free comparison of posteriors with identical information sets by comparing the posterior beliefs formed after a single ball draw (i.e. the information set consists of the exogenously assigned prior, p_0, and a single signal, $s_1 \in \{a, b\}$). We can then conduct two comparisons of information-set-equivalent posterior beliefs.

Firstly, we can compare posterior beliefs, π_1, formed with identical information sets $\{p_0, s_1\}$ *between* treatment groups, where the payments associated with states of the world differ. For example, we can compare the average posterior formed after an identical prior $p_0 = \frac{1}{6}$ and an identical signal $s_1 = a$ (i.e. a blue ball) differs between treatments.

Secondly, we can compare information-set-equivalent posterior beliefs *within* treatment groups. This comparison involves comparing π_1 after $\{p_0 = p, s_1 = s\}$ with $1-\pi_1$ after $\{p_0 = 1-p, s_1 = s^c\}$ where s^c is the complementary signal to s.[38] For example, we can compare the posterior, π_1, formed after a prior of $p_0 = \frac{1}{6}$ and the signal $s_1 = a$ (i.e. a blue ball), with $1 - \pi_1$ after a prior of $p_0 = \frac{5}{6}$ and the signal $s_1 = b$ (i.e. a red ball). To see why this comparison involves a comparison of information-set-equivalent posterior beliefs, recall that the experiment is designed to be completely symmetric in terms of information, with the information content of a red ball exactly the same as a blue ball, except in support of the other state of the world. Therefore, if an individual updates symmetrically, then $\pi_1|_{\{p_0=p,s_1=s\}} = 1 - \pi_1|_{\{p_0=1-p,s_1=s^c\}}$. This prediction results directly from assuming that the individual updates symmetrically and does not rely on Bayes' rule (although it is an implication of Bayes' rule), therefore it provides us with a non-parametric test of the *asymmetric updating hypothesis.*

Figure 7 depicts both of these comparisons, with each group of six bars collecting together the relevant information-set-equivalent groups. Each bar presents the mean posterior belief for that group, as well as a 95% confidence interval around the mean.[39] Each group is labeled on the x-axis by the prior belief associated with the red bars, which are the information sets which include a red ball as a signal (i.e. $s_1 = b$). The blue bars report the mean of $1 - \pi_1$, for information sets containing a blue ball (i.e. $s_1 = a$) and the x-axis label corresponds to $1 - p_0$. Within each group, the first two bars represent the average posterior beliefs in Treatment group 1 (SYMMETRIC) after a single red ball draw ('bad-news' in the ASYMMETRIC treatments) or blue ball draw ('good-news' in the ASYMMETRIC treatments); the second pair of bars depict the same for Treatment group 2 (ASYMMETRIC COMBINED); and the third pair of bars for Treatment group 3 (ASYMMETRIC SEPARATE).

The results depicted in Figure 7 show that there are no systematic differences between posterior beliefs within information-set-equivalent groups, neither *within* nor *between* treatment groups.

[38]Therefore, if $s = a$ then $s^c = b$ and vice versa.

[39]The calculation of these confidence interval assumes a t-distribution.

Furthermore, when testing non-parametrically whether there are differences *within* or *between* treatment groups for information-set-equivalent groups, none of the 45 relevant binary comparisons[40] are significant at the 5 percent significance level under a Mann-Whitney test, suggesting that we cannot reject the hypothesis that the posterior beliefs within information-set-equivalent groups are drawn from the same population. This lends support to the results described in the main text, which indicate that we fail to find evidence in support of the *asymmetric updating hypothesis*.

Figure 7: Comparison of Beliefs after the Receipt of a Single Signal and an Exogenous Prior

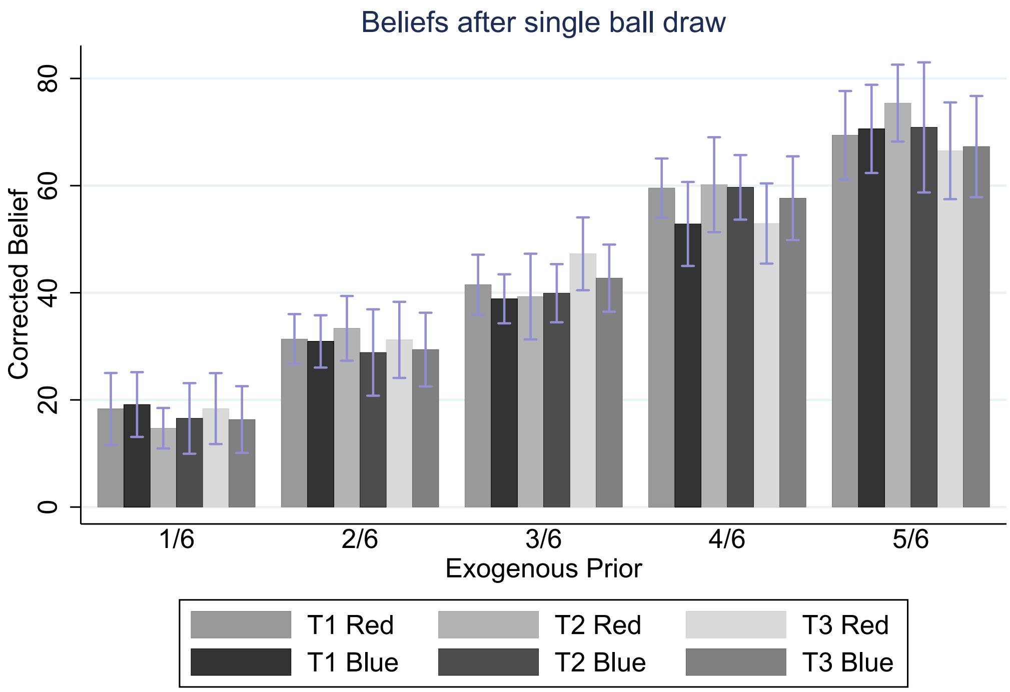

Notes: (i) Bar denotes 95 percent confidence interval around the mean; (ii) Beliefs after blue balls reflect (100 - belief).

Robustness Check 2: First-Differences Specification and Power Calculation

This section of the robustness checks serves two purposes. The first purpose is to check for the robustness of the results from the core empirical specification to the use of a first differences specification (DIFF), which essentially involves imposing the assumption that $\delta = 1$. The second purpose of this section is to report the size of the minimum detectable effect (MDE) from power calculations for both our main OLS and IV empirical specification, and the DIFF specification.

[40] For these comparisons, for each exogenous prior, we test the following binary comparisons: (i) within treatment group, we test between those that received the $s_1 = a$ and $s_1 = b$ signals (3×5 binary comparisons); (ii) for those that received the same signal, s_1, we test between treatment groups (6×5 binary comparisons).

One of the challenges in carrying out a statistical analysis of belief updating behavior is that an individual's current posterior belief necessarily depends upon her prior belief, which in turn is the result of updating in response to past information. Therefore, when estimating a parametric belief updating function, one concern is that the individual's prior belief is correlated with unobservables. In the main text, we devoted substantial space to discussing how the experiment was designed explicitly to address this concern by generating a completely exogenous information set, facilitating a natural instrumental variables (IV) approach to estimation. The first differences specification results presented here serve to further complement the IV analysis, since the DIFF specification avoids the potential endogeneity issue by removing the lagged belief from the set of dependent variables in the regression.

In columns (#a) and (#b), Table 4 repeats the OLS and IV results from Table 3 for the corrected beliefs, with one minor change to the core specification in Equation 10. Here we report, instead, the results for the equivalent specification:

$$\tilde{\pi}_{i,j,t+1} = \delta\tilde{\pi}_{i,j,t} + \gamma_a\widehat{q} - (\gamma_b - \gamma_a)\tilde{q} \cdot 1(s_{i,j,t+1} = b) + \epsilon_{i,j,t+1} \tag{11}$$

where $\tilde{\pi}_{i,j,t} = \text{logit}(\pi_{i,j,t})$ and $\widehat{q} = \log(\frac{q}{1-q}) \cdot [1(s_{i,j,t} = a) - 1(s_{i,j,t} = b)]$; while as above, $\tilde{q} = \log(\frac{q}{1-q})$; j refers to a round of decisions; t counts the decision numbers within a round, and the errors ϵ_{ijt+1} are clustered at the individual (i) level. The difference $\gamma_b - \gamma_a$ denotes a single parameter estimated in the regression, but is denoted as the difference between γ_b and γ_a as this is the natural way to think about this parameter in the context of the discussion above (i.e. the difference between how subjects update in response to 'bad news' and 'good news').

The reason for the rearrangement of the equation is that, while it is equivalent[41] to the specification in Equation 10, it displays the test of the difference between γ_a and γ_b more clearly (i.e. the test of the *asymmetric updating hypothesis*), and thereby also facilitates calculating the MDE. In Table 3, we have presented the MDE for a power of $\kappa = 0.8$.

[41]Notice that the regression coefficients and standard errors on δ and γ_a are the same in Tables 3 and 4 (where we are only considering the corrected beliefs). Furthermore, we can see the equivalence from the following simple rearrangement:

$$\begin{aligned}\tilde{\pi}_{i,j,t+1} &= \delta\tilde{\pi}_{i,j,t} + \gamma_a\tilde{q} \cdot 1(s_{i,j,t+1} = a) - \gamma_b\tilde{q} \cdot 1(s_{i,j,t+1} = b) \\ &= \delta\tilde{\pi}_{i,j,t} + \gamma_a\tilde{q} \cdot 1(s_{i,j,t+1} = a) - \gamma_a\tilde{q} \cdot 1(s_{i,j,t+1} = b) + \gamma_a\tilde{q} \cdot 1(s_{i,j,t+1} = b) - \gamma_b\tilde{q} \cdot 1(s_{i,j,t+1} = b) \\ &= \delta\tilde{\pi}_{i,j,t} + \gamma_a\tilde{q} \cdot [1(s_{i,j,t+1} = a) - \tilde{q} \cdot 1(s_{i,j,t+1} = b)] + [\gamma_a - \gamma_b] \cdot \tilde{q} \cdot 1(s_{i,j,t+1} = b)\end{aligned}$$

Columns (#c) report the results for the first difference specification, which imposes the restriction that $\delta = 1$:

$$\Delta\tilde{\pi}_{i,j,t+1} = \gamma_a \widehat{q} - (\gamma_b - \gamma_a)\tilde{q} \cdot 1(s_{i,j,t+1} = b) + \epsilon_{i,j,t+1} \tag{12}$$

where $\Delta\tilde{\pi}_{i,j,t+1} = \text{logit}(\pi_{i,j,t+1}) - \text{logit}(\pi_{i,j,t})$ and $\widehat{q} = \log(\frac{q}{1-q}) \cdot [1(s_{i,j,t} = a) - 1(s_{i,j,t} = b)]$; j refers to a round of decisions; t counts the decision numbers within a round, and the errors ϵ_{ijt+1} are clustered at the individual (i) level.

Table 4: First Difference Specification and Power Calculations

	T1: SYMMETRIC			T2: COMBINED			T3: SEPARATE		
	OLS	IV	DIFF	OLS	IV	DIFF	OLS	IV	DIFF
	(1a)	(1b)	(1c)	(2a)	(2b)	(2c)	(3a)	(3b)	(3c)
δ	0.90	0.99		0.86	0.99		0.93	0.99	
	(0.03)***	(0.03)		(0.04)***	(0.02)		(0.02)***	(0.02)	
γ_a	1.09	1.09	1.09	1.06	1.02	1.01	1.16	1.14	1.13
	(0.11)	(0.11)	(0.11)	(0.12)	(0.11)	(0.11)	(0.11)	(0.11)	(0.10)
$\gamma_b - \gamma_a$	0.08	0.08	0.08	0.07	0.10	0.11	-0.03	-0.03	-0.02
	(0.08)	(0.08)	(0.08)	(0.10)	(0.09)	(0.09)	(0.08)	(0.08)	(0.08)
p($\gamma_a = \gamma_b$)	0.32	0.31	0.32	0.48	0.25	0.25	0.73	0.74	0.75
MDE ($\kappa = 0.8$)	0.24	0.22	0.22	0.27	0.25	0.26	0.24	0.22	0.22
R^2	0.73		0.31	0.74		0.21	0.84		0.34
1st Stage F		84.04			107.01			95.45	
N	1,075	1,075	1,075	1,285	1,285	1,285	1,140	1,140	1,140

(i) Standard errors in parentheses (clustered at the individual level)
(ii) T-tests of H_0: $\delta = 1$; $\gamma_a = 1$; $\gamma_b - \gamma_a = 0$ indicated by * = 10%, ** = 5%, *** = 1%
(iii) MDE reports the minimum detectable effect size for a power of κ.

The results indicate that the $\gamma_b - \gamma_a$ parameter is robust to the different empirical specifications adopted, and also doesn't vary substantially across treatment groups. In all treatment groups, and for each of the empirical specifications considered, we cannot reject the null hypothesis that this parameter is equal to zero, which implies that we do not find support for the *asymmetric updating hypothesis.* Furthermore, we calculate the MDE for each specification, considering a significance level of $\alpha = 0.05$ and a power of $\kappa = 0.8$. Under these assumptions, the MDE for the difference between the γ_b and γ_a parameters in each of the regressions considered in isolation ranges between 0.22 and 0.27. As a result, we cannot conclusively reject the possibility that there exists a small asymmetry in updating; however none of our results provide any support for this conclusion.

Robustness Check 3: Allowing for Individual-Specific Updating Parameters

As discussed above, one reason we might think that endogeneity of the lagged belief could lead to biased estimates is if there is heterogeneity in individual updating behavior and this leads to a correlation between the unobserved error term and the lagged belief variable amongst the regressors. We have tried to address this issue above using, firstly, an instrumental variable approach, and secondly, a first differences empirical specification. However, since the data were collected in the form of a panel of belief updates for each individual, the data lends itself to controlling for individual-specific behavior through exploiting the panel. A typical fixed effects model is not appropriate here, as it is not the *level* of the regression that shifts from individual to individual. However, we can include individual-specific updating parameters to control for the *slope* to shift at the individual level. This allows us to extract the individual heterogeneity in how responsive individuals are to their prior belief, and to new information in general, and reduce the possible bias in the main parameter of interest, the average difference in responsiveness to 'bad news' and 'good news' : $\gamma_b - \gamma_a$. With this in mind, our third robustness check involves estimating the following empirical specification:

$$\tilde{\pi}_{i,j,t+1} = \delta_i \tilde{\pi}_{i,j,t} + \gamma_i \widehat{q} - (\gamma_b - \gamma_a)\tilde{q} \cdot 1(s_{i,j,t+1} = b) + \epsilon_{i,j,t+1} \tag{13}$$

where δ_i and γ_i are estimated at the individual level, and the remaining parameters and variables are defined as above. The results from this exercise using the corrected beliefs are reported in Table 5.

These results are very consistent with the estimates from the core specification, as well as from the DIFF specification in Robustness Check 2. In summary, all the empirical estimates provide support for the same underlying story that the data collected in this experiment provide no support for the *asymmetric updating hypothesis* in this context.

Table 5: Allowing for Individual-Specific Updating Parameters.

	T1 SYMMETRIC (1)	T2 COMBINED (2)	T3 SEPARATE (3)	T2+T3 (4)
$\gamma_b - \gamma_a$	0.09	0.14	-0.04	0.06
	(0.10)	(0.10)	(0.09)	(0.07)
p($\gamma_a = \gamma_b$)	0.35	0.18	0.67	0.41
MDE ($\kappa = 0.8$)	0.28	0.29	0.26	0.20
MDE ($\kappa = 0.9$)	0.33	0.33	0.30	0.23
N	1,075	1,285	1,140	2,425
R^2	0.80	0.84	0.89	0.86

(i) Standard errors in parentheses

(ii) T-tests of H_0: Coefficient = 0 reported: * = 10%, ** = 5%, *** = 1%

(iii) MDE reports the minimum detectable effect size for a power of κ.

Robustness Check 4: Between Treatments Comparison of Updating Parameters

This section tests whether the belief updating parameters in our core specification are significantly different *between* the three treatment groups. This is done by pooling together the three treatment groups and estimating Equation 10, but with the inclusion of treatment dummies interacted with the updating coefficients. This provides us with a test of whether the parameters differ between either of the two ASYMMETRIC treatments and SYMMETRIC.

More specifically, this involves estimating the following equation:

$$\begin{aligned}\tilde{\pi}_{i,j,t+1} =& \delta\tilde{\pi}_{i,j,t} + \gamma_a \tilde{q} \cdot 1(s_{i,j,t+1} = a) - \gamma_b \tilde{q} \cdot 1(s_{i,j,t+1} = b) + \\ & \sum_{k=2}^{3} [\delta^k \tilde{\pi}_{i,j,t} \cdot T^k_{i,j,t} + \gamma^k_a \tilde{q} \cdot 1(s_{i,j,t+1} = a) \cdot T^k_{i,j,t} - \gamma^k_b \tilde{q} \cdot 1(s_{i,j,t+1} = b) \cdot T^k_{i,j,t}] + \epsilon_{i,j,t+1}\end{aligned}$$

where $T^k_{i,j,t}$ is an indicator variable for treatment k [i.e. $T^k_{i,j,t} = 1(T_{i,j,t} = k)$], with $T_{i,j,t}$ a treatment variable taking the values $\{1, 2, 3\}$ corresponding to the three treatment groups. The coefficients δ,γ_a, and γ_b reflect the baseline parameters without the influence of state-contingent stakes and the parametersδ^k, γ_a^k and γ_b^k estimate the movement from these parameters for each of the two state-contingent stake treatments, $k \in \{2, 3\}$.

The results from this exercise are presented in Table 6. The results show that, for the average individual, there are no systematic differences in the updating parameters across treatment groups. This implies that the differences in exogenous state-contingent incentives do not exert a strong influence on how individuals update their beliefs in the different treatments.

Table 6: Testing for Differences in Average Updating Behavior *between* Treatment Groups.

	Belief 1 (1)	Belief 2 (2)	Belief 3 (3)	Belief 4 (4)	Belief 5 (5)	Pooled (6)	Full Sample (7)
Priors							
δ	0.92	0.92	0.99	1.01	1.03	0.99	0.99
	(0.04)	(0.03)	(0.03)	(0.03)	(0.04)	(0.02)	(0.02)
δ * T2	-0.01	-0.00	0.01	0.01	0.00	-0.00	-0.04
	(0.06)	(0.07)	(0.06)	(0.06)	(0.06)	(0.03)	(0.03)
δ * T3	0.03	0.04	-0.07	0.08	-0.05	0.00	0.01
	(0.07)	(0.05)	(0.05)	(0.05)	(0.08)	(0.03)	(0.03)
Signal: Blue ($s = a$)							
γ_a	0.85	0.91	0.94	1.10	1.36	1.04	0.78
	(0.10)	(0.12)	(0.11)	(0.13)	(0.16)	(0.08)	(0.06)
γ_a * T2	0.05	0.04	-0.07	0.13	-0.24	-0.02	0.05
	(0.19)	(0.22)	(0.18)	(0.30)	(0.26)	(0.13)	(0.11)
γ_a * T3	0.07	0.05	-0.01	0.28	0.12	0.10	0.09
	(0.19)	(0.17)	(0.15)	(0.22)	(0.28)	(0.13)	(0.11)
Signal: Red ($s = b$)							
γ_b	0.83	0.85	0.99	1.20	1.45	1.06	0.78
	(0.10)	(0.11)	(0.11)	(0.13)	(0.15)	(0.07)	(0.06)
γ_b * T2	0.08	-0.00	-0.13	-0.01	0.48	0.06	0.10
	(0.19)	(0.19)	(0.20)	(0.22)	(0.36)	(0.15)	(0.13)
γ_b * T3	0.17	-0.02	0.05	-0.17	0.29	0.05	0.10
	(0.18)	(0.16)	(0.16)	(0.19)	(0.27)	(0.12)	(0.10)
Observations	894	894	894	894	894	4,470	7,175
Kleibergen-Paap F	46.78	47.64	56.84	60.68	64.26	93.33	53.95

(i) Robust standard errors in parentheses (clustered at the individual level).
(ii) Estimates use the corrected beliefs and are instrumented using the correct lagged Bayesian posterior.
(iii) All of the non-interacted coefficients are significantly different from 0 at the 1% level. None of the interaction coefficients are significantly different from zero at the 10% level.

Appendix B.1: Core Properties of Bayes' Rule

Möbius et al. (2014) argue that the core structure of Bayesian updating is captured by the following three properties:

1. *invariance,* whereby the difference in logit beliefs between t and $t+1$ depends only on the history of signals, H_{t+1}, and the initial prior, p_0 (i.e. on the agent's information set). An updating process is *invariant* if we can find a function g_t such that:

   $$\text{logit}(\pi_{t+1}) - \text{logit}(\pi_t) = g_t(s_{t+1}, s_t, ..., s_1; p_0)$$

   If an individual displayed base rate neglect or confirmatory bias, this would constitute a violation of invariance (i.e. in the context of the model outlined in the main text, this assumption stipulates that $\delta = 1$).

2. π_t is a *sufficient statistic* for all information received at time t or earlier, such that the change in logit beliefs depends only on the new information in time $t+1$: $\text{logit}(\pi_{t+1}) - \text{logit}(\pi_t) = g_t(s_{t+1})$

3. *stability* of the updating process over time. This property is satisfied if $g_t = g$ for all t.

Under the assumption that these properties are satisfied, the authors note that the class of updating processes that remain can be fully described by the two parameter function, $g(s_t)$, where:

$$g(s_t) = \log(\frac{q}{1-q}) \cdot 1(s_{t+1} = a) - \log(\frac{q}{1-q}) \cdot 1(s_{t+1} = b)$$

This serves to motivate the model described in Equation 6.

Appendix B.2: Augmenting the Offerman et al. (2009) 'Truth Serum' Approach to Include Stakes

In the case where $x \neq 0$, the main text discussed how participants who face the quadratic scoring rule incentives, along with the non-zero state-contingent bonus x, essentially face a choice from a menu of lotteries denoted by $(x + 1 - (1 - r_t)^2)_{E_A}(1 - r_t^2)$. An individual who satisfies the *biseparable preferences* model and is probabilistically sophisticated will evaluate this prospect using the following Equations[42]:

For $x \geq 1$ or $r_t \geq 0.5$:

$$w(P(E_A))U(x+1-(1-r_t)^2)+(1-w(P(E_A)))U(1-r_t^2) \tag{14}$$

and similarly,

For $x = 0$ & $r_t < 0.5$:

$$(1-w(P(E_A^c)))U(x+1-(1-r)_t^2)+w(P(E_A^c))U(1-r_t^2) \tag{15}$$

The reason for the two separate conditions is due to the way in which many NEU models, subsumed within *biseparable preferences* model, allow the probability weighting function, $w(.)$, over events to be influenced by the ordinal ranking over the associated outcomes, from best to worst[43]. Since the case where $x = 0$ is discussed extensively in Offerman et al. (2009), we will focus on the case where $x \geq 1$ in the discussion that follows. This case only requires a very minor adjustment to their discussion. The key results are the following (adjusted to include the influence of x):

Result 1: *Under NEU with* ***known probabilities****, p, the optimal reported probability, $r = R_x(p)$ satisfies:*

$$\text{If } x \geq 1, \text{then } p = R_x^{-1}(r_t) = w^{-1}\left(\frac{r_t}{r_t+(1-r_t)\frac{U'(x+1-(1-r_t)^2)}{U'(1-r_t^2)}}\right) \tag{16}$$

Result 2: *Under NEU with* ***unknown probabilities****, the optimal reported probability, r, satisfies:*

$$\text{If } x \geq 1, \text{then } P(E) = w^{-1}\left(\frac{r_t}{r_t+(1-r_t)\frac{U'(x+1-(1-r)_t^2)}{U'(1-r_t^2)}}\right) \tag{17}$$

This motivates the simple strategy for recovering the agent's subjective beliefs from her reported

[42]For expositional simplicity, we don't consider $x\epsilon(0,1)$. The discussion below is easily extended to these cases, but they are irrelevant for the purposes of this paper. This case is slightly different due to the fact that the probability weights on events or states may depend on their ordinal ranking according to preferences in this model.

[43]When $x = 0$ and $r_t < 0.5$, then $1-r_t^2 > 1-(1-r_t)^2$ (i.e. in this case, E_A^c becomes the preferred event, rather than E_A, and therefore the probability weighting function is reversed).

beliefs under the specific incentive environment that she faces. Since the RHS of (16) and (17) agree, we have:

$$P(E) = R_x^{-1}(r) \tag{18}$$

which implies that if we can recover the function R_x^{-1} then we can map the reported beliefs to the participant's subjective beliefs. Equation 16 shows that we can recover this R_x^{-1} function in the same way here, with the bonus payment of x, as in the case where $x = 0$ considered in Offerman et al. (2009). Essentially, we provide the participant with prospects over known probabilities, p, and ask them for their belief regarding the likelihood that one state will be realized. In order to ensure the incentives to distort one's reported beliefs are kept constant, we do this exercise under precisely the same incentive environment as in the main belief updating task. By eliciting these reported beliefs associated with known probabilities spanning the whole unit interval, we can use these (p, r_t) pairs to estimate $R_x(p)$ for each individual for the relevant incentive environment created by the belief elicitation. Having estimated $R_x(p)$, we can calculate its inverse, $R_x^{-1}(r)$.

We can take any beliefs reported by the participant under the same belief elicitation incentives and then use this estimated $R_x^{-1}(r)$ to recover her true beliefs. In particular, we can use this estimated function to recover her true beliefs from her reported beliefs in the belief updating task that is the focus of this paper. Essentially, we are using this procedure to remove any misreporting effect that the belief elicitation incentive environment may have. It allows us to correct for the possibility that individuals may hold some belief $P(E)$ or π, but instead report a different belief, r.

If the incentive environment does not cause the participant to report a belief different from her true belief, then this procedure is unnecessary, but applying the procedure to her reported beliefs will not have any effect. In this case, the corrected beliefs will be the same as the reported beliefs.

Appendix B.3: Calibration of the Belief Correction Procedure

It is clear from the discussion in the main text and Equation 9, that we could recover $R(.)$ *non-parametrically* for each individual if we were to collect a large number of (p, r) pairs from participants, such that the interval between the known probabilities, p, is sufficiently small. However, since it is not practical here to elicit such a large number of observations from each participant, we instead impose a parametric structure similar to the one used by Offerman et al. (2009).

For the utility function, $U(.)$, we use the *constant relative risk aversion (CRRA)*[44] functional form:

$$U(x) = \begin{cases} x^\rho & \text{if } \rho > 0 \\ \ln x & \text{if } \rho = 0 \\ -x^\rho & \text{if } \rho < 0 \end{cases} \tag{19}$$

For the probability weighting function, $w(.)$, we adopt Prelec's (1998), one-parameter family[45]:

$$w(p) = \exp[-(-\ln(p))^\alpha] \tag{20}$$

Substituting these parametric functional form specifications into Equation 9 for the case where $x = 0$ gives:

$$p = R^{-1}(r_t) = \exp\left(-\left[-\ln\left(\frac{r_t(2r_t - r_t^2)^{1-\rho}}{r_t(2r_t - r_t^2)^{1-\rho} + (1 - r_t)(1 - r_t^2)^{1-\rho}}\right)\right]^{\frac{1}{\alpha}}\right) \tag{21}$$

For the the case where $x \neq 0$, substituting these parametric functional form specifications described in Equations 19 and 20 into Equation 16 gives:

$$p = R_x^{-1}(r_t) = \exp\left(-\left[-\ln\left(\frac{r_t(1 - (1 - r_t)^2 + x)^{1-\rho}}{r_t(1 - (1 - r_t)^2 + x)^{1-\rho} + (1 - r_t)(1 - r_t^2)^{1-\rho}}\right)\right]^{\frac{1}{\alpha}}\right) \tag{22}$$

We therefore use this adapted specification for our correction mechanism for the COMBINED and SEPARATE treatment groups.

[44] It is worth noting that the name "constant relative risk aversion (CRRA)" is not appropriate here as risk attitudes are also captured in the probability weighting function.

[45] This is a special case of Prelec's two-parameter family of weighting functions:

$$w(p) = \exp[-\beta(-\ln(p))^\alpha]$$

For the purposes of the current context, the two-parameter family is not practically suitable due to the limited data we use at the individual level. This functional form permits the standard inverse-S shaped probability weighting function that has been found to be consistent with the majority of the existing empirical evidence. When $\beta = 1$ in the one-parameter family, the α parameter captures the degree of curvature of the inverse-S shape but the point at which $w(p)$ intersects the 45 degree line is predetermined. Adding the second parameter, β, extends the one-parameter specification by allowing this fixed point to vary.

In our core analysis, for our individual level reported belief corrections, we will make the simplifying assumption that $\alpha = 1$, such that risk aversion is captured only through the curvature of the utility function and not through the probability weighting function. The results are similar when we use Prelec's one parameter weighting function. Furthermore, it is substantially easier to interpret the risk aversion parameter estimates when we estimate ρ alone, due to the strong relationship between the ρ and α estimates .

We therefore estimate the following model, for each participant, in order to acquire a numerical estimate for the inverse of this function, $R(.)$:

$$\text{logit}[R(j/20)] = \text{logit}[h(j/20, \alpha, \rho)] + u_j \tag{23}$$

where $R(j/20)$ is the probability reported by the individual that corresponds to true known probability, $p = \frac{j}{20}$ where $1 \leq j \leq 19$[46]. As discussed above, α is the parameter of the probability weighting function; ρ gives the curvature of the utility function. The function $h(.)$ is the inverse of R^{-1}. We estimate this function, $h(.)$ numerically at each step[47] within the maximum likelihood estimation. The error terms, u_j, are independently and identically distributed across participants and choices and are drawn from a normal distribution. Essentially, here we are using each participant's 20 (r, p) pairs in order to estimate an R function that reflects the distortion in her reported beliefs due to the particular quadratic scoring rule incentive structure that she is subject to. Notice, that this structure varies across treatments as x varies and therefore the same subject would require a different adjustment curve if she were reassigned to a different treatment.

Using these estimates at an individual level allows us to recover the participants' true subjective beliefs, π_t, from the first stage of the experiment in which they report their beliefs, r_t, regarding the likelihood of $\omega = A$ being the true state. In Figure 4 in the main text, we graph the individual level correction curve estimates for two individuals in each treatment group. It is clear from these examples, firstly, that individuals in the sample are distorting their reported beliefs substantially relative to the known probabilities, and secondly, that the estimated correction curves are sufficiently flexible to fit different types of belief distortion behavior reasonably well. Furthermore, importantly, the graph in the top-left panel of the figure shows that, when an individual accurately reports her beliefs, then the correction mechanism has no harmful effect.

At the aggregate level, for each treatment group, $T \in \{1, 2, 3\}$, we estimate:

[46] In other words, for known probabilities, p, between 0.05 and 0.95 at intervals of 0.05.
[47] i.e. given the current parameter guesses.

$$\text{logit}[R_i(j/20)] = \text{logit}[h_i(j/20, \alpha, \rho)] + u_{i,j} \tag{24}$$

where j indexes the 20 reported probabilities of individual i. This specification allows us to examine the distortion caused by the incentive environment to the average individual in each of the three treatment groups.

Appendix C: Endogeneity: An Illustrative Example

This section provides a simple illustrative example of why it is important to exercise a little caution in addressing endogeneity when studying belief updating. When studying belief updating we are often interested in considering beliefs about the *self*. This is the case for much of the nascent belief updating literature (understandably, since beliefs about the *self* are extremely interesting and important). However, the way this belief updating has typically been studied implies that distribution of "exogenous" noisy signals that an individual observes is related to individual characteristics, and to her prior belief.

The following discussion has the simple objective of highlighting the importance of paying attention to the endogeneity of the signal distribution. In particular, I use a very simple "toy" simulation to demonstrate that ignoring this issue can (in principal) lead to mistakenly find evidence for asymmetric updating when all individuals update symmetrically. It is important to point out that I am *not* suggesting that this is the explanation for asymmetric updating results observed in the literature[48] - I am simply highlighting a potential issue that should be addressed in this literature going forwards. In this regard, I also suggest a simple solution for dealing with the issue.

Much of the belief updating literature considers situations that resembles the following basic structure: consider an agent who updates about two states of the world, $\omega \in \{High, Low\}$, and receives a sequence of noisy signals, $s_t \in \{UP, DOWN\}$. This also reflects the setup considered in the current paper. However, since it is important to also understand how we form beliefs about *self,* in some studies the states are determined by personal characteristics of the individual (e.g. IQ). This means that states are essentially equivalent to personal types (i.e. ***states = types***). The implication of this is that if signals are informative about the state of the world, then *High* types are more likely than *Low* types to receive *Up* signals (and vice versa for *Down* signals). If *High* types update their beliefs differently from *Low* types, this can (in principal) lead to finding evidence that the average individual updates asymmetrically when no individual actually does.

In order to show this, I conduct a very simple simulation exercise. I construct a population of 10 000 individuals who are randomly assigned to one of two types, $\omega \in \{High, Low\}$. Within each type, the agents' prior beliefs about the likelihood of being the *High* type are assigned randomly

[48] For example, Figure 3 in Möbius et al. (2014) suggests that this is probably not a major concern for their main results. However, it is still important to control for this potential endogeneity issue as a robustness check.

using a uniform distribution, distributed between zero and one[49]. *High* types receive an *Up* signal with probability $q = \frac{5}{8}$ and *Low* types receive a *Down* signal with probability $q = \frac{5}{8}$. Using a seed of 1000 in STATA, the empirical observed distribution of signals across types is given by the following figure.

Figure 8: Frequencies of Signals by Type

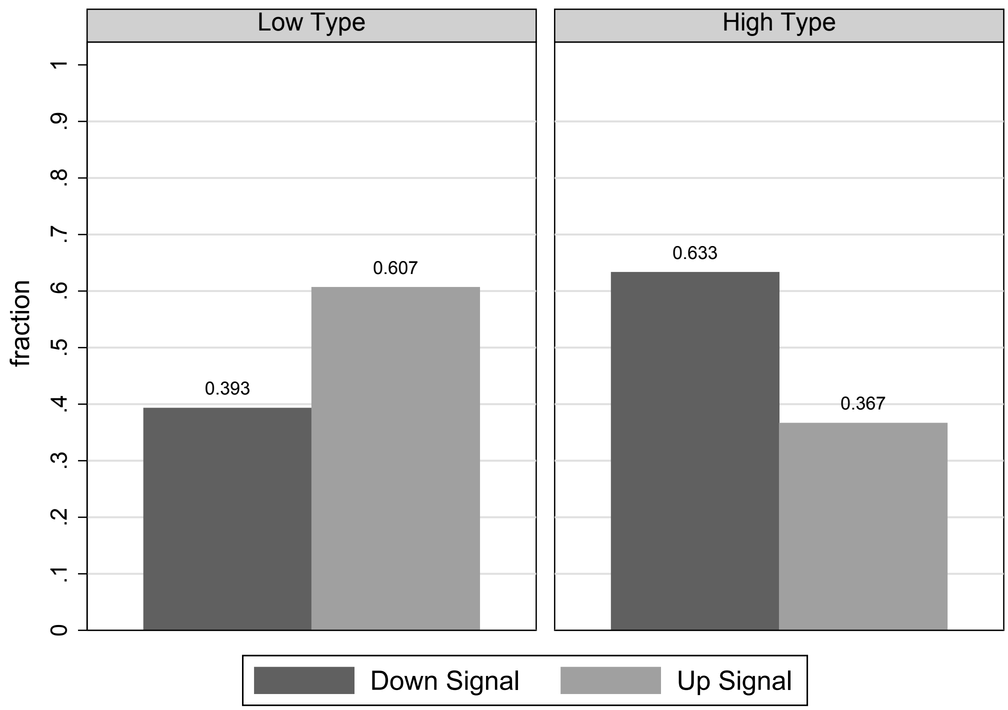

Now, the important part of this story is that belief updating may (in principal) be related to the underlying characteristic of interest. For example, it is conceivable that high IQ individuals process information and update their beliefs differently from low IQ individuals.[50]

[49]Note, this is an unrealistic assumption. In general, prior beliefs are related to the true state of the world. For example, beliefs about one's rank in an IQ distribution tend to be correlated with one's actual rank. However, constructing type and prior belief to be orthogonal allows us to isolate only the effect of the endogeneity of types and signals (with exogenous priors).

[50]Note, even if one doesn't find the story that the two *types* might update their beliefs differently compelling, a very similar pattern could also be generated if there is a relationship between *prior beliefs* and updating. This follows because: (i) prior beliefs are related to actual types (e.g. higher IQ individuals believe they are more likely to have a high IQ), (ii) types are mechanically related to the distribution of signals in the class of experiments we're considering, and therefore (iii) prior beliefs are related to the distribution of signals observed. Therefore, the story described in this section is worth paying attention to if one is not willing to assume both:

Here, we consider two types that use different, *but always symmetric*, updating rules. In particular, we consider a *High* type that is perfectly Bayesian, and a *Low* type that is not very responsive to new information (but otherwise very well behaved her belief updating).[51]

The *High* type updates according to the following rule ($\delta = 1$, $\gamma_{UP} = 1$, $\gamma_{DOWN} = 1$):

$$\text{logit}(\pi_{t+1}) = 1 \cdot \text{logit}(\pi_t) + 1 \cdot \log(\frac{5}{3}) \cdot 1(s_{t+1} = UP) - 1 \cdot \log(\frac{5}{3}) \cdot 1(s_{t+1} = DOWN) \quad (25)$$

The *Low* type updates according to the following rule ($\delta = 1$, $\gamma_{UP} = 0.2$, $\gamma_{DOWN} = 0.2$):

$$\text{logit}(\pi_{t+1}) = 1 \cdot \text{logit}(\pi_t) + (0.2) \cdot \log(\frac{5}{3}) \cdot 1(s_{t+1} = UP) - (0.2) \cdot \log(\frac{5}{3}) \cdot 1(s_{t+1} = DOWN) \quad (26)$$

However, if we as the analyst neglects the possibility that the two types update their beliefs differently, then we can obtain biased parameters. This is illustrated by the regression estimates presented in column 1 and 2 of Table 7 below. These columns reflect the estimates from the standard specification used in this literature (i.e. equation 10). These parameter estimates, along with the true population averages are summarised as follows:

	True Parameter Values (Population Ave.)	Estimates
δ	1	1
γ_{UP}	0.6	0.7
γ_{DOWN}	0.6	0.5

It is clear from this that in spite of the fact that there is not a single individual in this population who updates asymmetrically that the estimated parameters suggest that there is an asymmetry. Notice, the standard errors are small and the adjusted R^2 suggests a good model fit (see Table 7). As mentioned above, this is simply an illustration of why it is important to pay attention to the relationship between the distribution of signals and the types. Furthermore, it is important to point out that I could have made an equivalent argument to show that even if the majority of individuals

(1) types $\perp$ belief updating *and* (2) priors $\perp$ belief updating

[51] Note, neither of the types makes any errors in their belief updating. They both follow their updating rule perfectly. This exercise therefore rules out several other channels that can make life challenging for the analyst (e.g. errors related to priors or types).

in the population are asymmetric updaters that the neglect of this relationship between signals and types could (in principal) generate estimates that suggest completely symmetric updating.

Fortunately, this particular endogeneity issue is easy to deal with by simply considering updating behavior within each type (e.g. interacting the RHS variables of equation 10 with the Type dummy variable). This is illustrated in column 3 of Table 7. Notice, also, that simply including the Type dummy variable in the regression does not solve the problem (see column 2).

Table 7: Estimates of Simulated Data Parameters

	Model 1 (1)	Model 2 (2)	Model 3 (3)
δ	1.002 (0.001)	1.002 (0.001)	1.000 (.)
γ_{UP}	0.696 (0.005)	0.683 (0.007)	0.200 (.)
γ_{DOWN}	0.504 (0.006)	0.513 (0.006)	0.200 (.)
High Type (=1)		0.0114 (0.004)	
High Type (=1) * δ			0.000 (.)
High Type (=1) * γ_{UP}			0.800 (.)
High Type (=1) * γ_{DOWN}			0.800 (.)
N	10000	10000	10000
Adjusted R^2	0.99	0.99	1.00

(i) Standard errors in parentheses
(ii) Note: Std errors in column 3 missing due to perfect fit.

Appendix D: Heterogeneity using a Finite Mixture Model

To complement the heterogeneity results presented in the main text, we adopt a second approach to examining the individual level heterogeneity in the sample by conducting a finite mixture model analysis within each treatment group. This allows us to examine whether we can find distinct subgroups that differ substantially in terms of updating behavior. In particular, we can test whether there is a subgroup of individuals who are asymmetric in their updating behavior.

To address this question, we consider a mixture model extension to our core empirical specification in Equation 10.

Empirical Specification

In order to examine individual heterogeneity in updating behavior by means of a mixture model, we use maximum likelihood to estimate the specification below, with $f(\cdot)$ denoting the likelihood contribution of a particular observation:

$$f(\tilde{\pi}_{i,j,t+1}|\tilde{\pi}_{i,j,t}, s_{i,j,t+1}; \theta_1, \theta_2, \theta_3; \mu_1, \mu_2, \mu_3) = \sum_{m=1}^{M} \mu_m f_m(\tilde{\pi}_{i,j,t+1}|\tilde{\pi}_{i,j,t}, s_{i,j,t+1}; \theta_m) \tag{27}$$

where:

$$f_m(\tilde{\pi}_{i,j,t+1}|\tilde{\pi}_{i,j,t}, s_{i,j,t+1}; \theta_m) = \left\{ \frac{1}{\sqrt{2\pi\sigma_m^2}} \exp\left\{ -\frac{1}{2} \frac{[\tilde{\pi}_{i,j,t+1} - g(\tilde{\pi}_{i,j,t}, s_{i,j,t+1}; \theta_m)]^2}{\sigma_m^2} \right\} \right\} \tag{28}$$

and:

$$g(\tilde{\pi}_{i,j,t}, s_{i,j,t+1}; \theta_m) = \delta_m \tilde{\pi}_{i,j,t} + \gamma_{a,m} \tilde{q} * 1(s_{i,j,t+1} = a) - \gamma_{b,m} \tilde{q} * 1(s_{i,j,t+1} = b) \tag{29}$$

such that the following conditions are satisfied: $0 < \mu_m < 1$ and $\sum_{m=1}^{3} \mu_m = 1$. The parameters μ_m denote the mixing weights. Therefore, we estimate the parameters of Equation 10, δ_m, $\gamma_{a,m}$, $\gamma_{b,m}$ and σ_m^2 for each type, $m \in \{1, 2, \ldots M\}$. As indicated by Equation 28, we are assuming a mixture over M Gaussian distributions, each with variance denoted by σ_m^2. In addition, we are assuming that observations are independently distributed across individuals, but allow for errors to be correlated for decisions made by a single individual.

In essence, the estimation procedure involves using maximum likelihood to estimate the three parameters of the core model (δ_m, $\gamma_{a,m}$ and $\gamma_{b,m}$) for each of the m updating types in the mixture model. We refer to each 'type' as a mixture class. In addition to estimating these three parameters for each mixture class, the mixture model procedure estimates two further parameters for each class: (i) a weighting parameter, which is referred to as the prior probability, μ_m, of an observation being in that mixture class; and (ii) the variance term, σ_m^2, for each class. Each group m therefore has distinct δ_m, $\gamma_{a,m}$, $\gamma_{b,m}$, μ_m and σ_m^2 parameters.

This empirical approach allows the data to dictate how it should best be split into subgroups in terms of updating behavior, choosing the combination of subgroups that best[52] explains updating in the sample as a whole. For example, this would allow us to detect if there were substantial subgroups of optimistic updaters and pessimistic updaters in the sample. Table 8 reports the results from the finite mixture model (fmm) analysis, allowing for three types[53]. The corresponding distributions of the posterior beliefs are reported in Table 9.

[52] Here, we refer to "best" in the sense of the parameters that maximize the likelihood function.

[53] As in any mixture model analysis, the choice of the number of mixture classes requires substantial consideration. Here, this issue is further complicated by the fact that there are three treatment groups and we wanted to estimate the mixture model separately for each treatment group, allowing for the possibility of distinct updating types (classes) across treatment groups. Given this context, the reasons for choosing a mixture model with three classes are the following. We estimated the mixture model for each treatment group with two mixture classes, three mixture classes and four mixture classes. In moving from two classes to three classes, the AIC and BIC were substantially reduced, suggesting an improved model fit, in two of the three treatment groups (COMBINED and SEPARATE). In moving from three classes to four classes, the extra parameters didn't seem justified as in one treatment group (COMBINED), one mixture class had zero observations assigned to it according to posterior probabilities, and another had a substantially worse AIC and BIC (SEPARATE). Furthermore, in each of the treatments, the mixture classes under four types were generally all similar to those under three types (i.e. two of the mixture classes had similar parameter values to one another). Therefore, the choice of three mixture classes followed from these considerations.

Table 8: Mixture Model Results

	T1 SYMMETRIC (1)	T2 COMBINED (2)	T3 SEPARATE (3)
Class 1			
δ	1.00	1.00	1.00
	(0.00)	(0.00)	(0.00)
γ_a	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
γ_b	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
Class 2			
δ	1.00	0.96	0.99
	(0.01)	(0.01)	(0.01)
γ_a	0.58	0.77	0.70
	(0.08)	(0.09)	(0.08)
γ_b	0.56	0.84	0.78
	(0.09)	(0.07)	(0.07)
Class 3			
δ	0.87	0.55	0.87
	(0.06)	(0.10)	(0.05)
γ_a	3.03	4.47	3.32
	(0.26)	(0.46)	(0.29)
γ_b	3.26	4.32	3.00
	(0.27)	(0.51)	(0.27)
Error Variance			
σ_1	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
σ_2	0.25	0.28	0.23
	(0.02)	(0.03)	(0.02)
σ_3	1.06	1.55	1.06
	(0.07)	(0.13)	(0.10)
Prior Probability			
μ_1	0.22	0.35	0.24
	(0.03)	(0.04)	(0.03)
μ_2	0.52	0.50	0.54
	(0.04)	(0.04)	(0.04)
μ_3	0.26	0.14	0.23
	(0.03)	(0.02)	(0.03)
Entropy	0.83	0.90	0.84
Model Fit			
AIC	-48,833.39	-109,702.65	-59,983.90
BIC	-48,763.67	-109,630.43	-59,913.35

(i) With the exception of γ_a and γ_b from Class 1, all δ, γ_a and γ_b coefficients are significantly different from 0 at the $p < 0.1$ level.

(ii) Entropy gives a measure of the quality of a classification (Ramaswamy et al. (1993)).

(iii) There are no significant differences between the γ_a and γ_b parameters within a mixture class.

The striking feature of this table is the similarity in the fmm patterns observed in each of the treatments. Firstly, in each treatment, there is one group (Class 1) which didn't update their beliefs at all ($\delta = 1; \gamma_a = \gamma_b = 0$). The prior probability of being classified in this updating group ranges from $\mu_1 = 0.22$ in SYMMETRIC to $\mu_1 = 0.35$ in COMBINED. Secondly, there is a group (Class 3) in each treatment group who are (i) highly responsive to new information ($3.0 < \gamma_a, \gamma_b < 4.5$), and (ii) tend to update more *away* from the the state favored by their prior, than towards it ($\delta < 1$). To get a sense of the magnitude of the responsiveness of these agents to new information, notice that $\gamma_a = 3$ implies that upon seeing a single signal in support of the state $\omega = A$, the subject updates as a Bayesian would in response to three consecutive signals of this type. The prior probability of this group ranges from $\mu_1 = 0.14$ in COMBINED to $\mu_1 = 0.26$ in SYMMETRIC. Thirdly, there is a group that contains the majority of observations (Class 2) and updates in a manner that is less responsive to information than a Bayesian ($0.56 < \gamma_a, \gamma_b < 0.84$). The prior probability of this group ranges from $\mu_1 = 0.50$ in COMBINED to $\mu_1 = 0.54$ in SEPARATE.

Since one of the primary objectives of this paper is to test for the existence of an asymmetry in updating behavior, it is interesting to observe that in none of the treatments do we observe a subgroup whose updating displays a statistically significant asymmetry in updating behavior. Rather, the overall pattern of updating is reasonably similar across the three treatments groups, although there is substantial heterogeneity in updating between the different types within each treatment group.[54] This suggests that the variation in state-contingent rewards across the treatment groups did not influence the way in which subjects updated their beliefs.

Posterior Probabilities

Table 9 reports the average distribution of posterior probabilities for each mixture model class. Essentially, it shows that, when an observation is assigned to a particular mixture model class, then on average it has a posterior probability of at least 0.9 associated with that mixture class, as opposed to all other mixture classes. This is comforting as it suggests that observations are generally well classified into one of the three mixture classes. The last column of the Table also reports the proportion of observations classified to each of the mixture classes according to the posterior probabilities. In contrast to the prior probabilities reported in Table 8, which reflect

[54] Interestingly, Palfrey and Wang (2012) demonstrate that such heterogeneity in belief updating can lead to asymmetric fluctuations in prices in an asset market, with the price responding more to 'good-news' than 'bad-news'. It is important to note that this asymmetry is different from the belief updating asymmetry we test for in this paper as the price-response asymmetry that they document can result from completely symmetric, but heterogenous, belief updating (such as the heterogeneity demonstrated in our results). Therefore, the results of the two papers are consistent.

weights placed on each class for all observations, these proportions give a sense of how well observations fit into a given class.

Table 9: Average Posterior Probabilities from the Mixture Model

	Posterior 1	Posterior 2	Posterior 3	Proportions
Treatment 1				
Class 1	1.00	0.00	0.00	0.22
Class 2	0.00	0.90	0.10	0.56
Class 3	0.00	0.04	0.96	0.21
Treatment 2				
Class 1	1.00	0.00	0.00	0.36
Class 2	0.00	0.95	0.05	0.52
Class 3	0.00	0.06	0.94	0.12
Treatment 3				
Class 1	1.00	0.00	0.00	0.25
Class 2	0.00	0.91	0.09	0.57
Class 3	0.00	0.05	0.95	0.18

(i) Posteriors reflect average posterior probabilites by fmm class.
(i) Proportions report the fraction of observations assigned to each fmm class.

Appendix E: Supplementary Figures, Results and Experimental Instructions

Figure 9: Comparison of Initial (Period 0) Belief with the Exogenous Prior

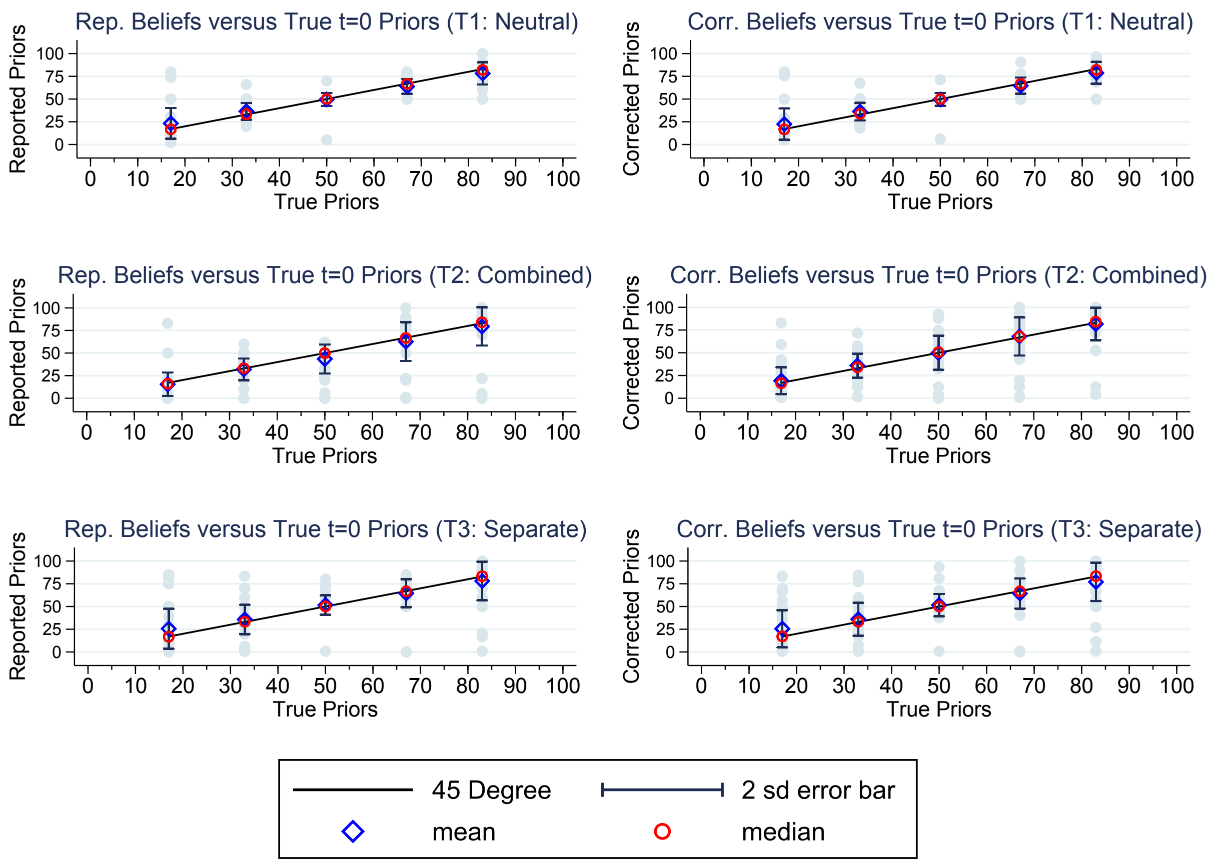

Table 10: Testing for Sample Balance Across Treatment Groups

	T1 SYMMETRIC	T2 COMBINED	T3 SEPARATE	D(2,1)	D(3,1)	D(3,2)
Gender (Male = 1)	0.56	0.58	0.58	0.02	0.02	0
	(0.50)	(0.50)	(0.50)			
Age	22.88	23.02	21.88	0.14	-0.99	-1.13
	(3.46)	(4.47)	(3.11)			
Location (London = 1)	0.51	0.48	0.52	-0.03	0.01	0.04
	(0.50)	(0.50)	(0.50)			
Economics Class	0.44	0.54	0.42	0.10	-0.02	-0.12
	(0.50)	(0.50)	(0.50)			
Home Language German	0.42	0.45	0.42	0.03	-0.00	-0.03
	(0.50)	(0.50)	(0.50)			
Home Language English	0.28	0.20	0.38	-0.08	0.10	0.18**
	(0.45)	(0.40)	(0.49)			
Cognitive Reflection Score	2.07	2.02	2.05	-0.05	-0.02	0.03
	(0.94)	(0.93)	(1.05)			
N	57	65	60			

(i) Standard deviations in parentheses

(ii) T-test for difference being significant: * = 10%, ** = 5%, *** = 1%

Table 11: Average Updating Behavior across Treatments (Full Sample)

	T1: SYMMETRIC		T2: COMBINED		T3: SEPARATE	
	Reported	Corrected	Reported	Corrected	Reported	Corrected
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
OLS						
δ	0.74	0.73	0.80	0.80	0.84	0.86
	(0.05)***	(0.06)***	(0.03)***	(0.03)***	(0.04)***	(0.03)***
γ_a	0.95	0.92	0.86	0.91	1.00	0.93
	(0.10)	(0.09)	(0.10)	(0.10)	(0.11)	(0.10)
γ_b	0.83	0.78	0.99	0.90	0.96	0.91
	(0.10)*	(0.10)**	(0.14)	(0.12)	(0.10)	(0.09)
p ($H_0 : \gamma_a = \gamma_b$)	0.26	0.19	0.25	0.90	0.64	0.80
N	1,875	1,875	1,850	1,850	1,825	1,825
R^2	0.52	0.51	0.62	0.64	0.67	0.70
IV						
δ	0.99	1.00	0.96	0.95	0.99	1.00
	(0.03)	(0.03)	(0.02)	(0.02)**	(0.02)	(0.02)
γ_a	0.87	0.84	0.87	0.83	0.95	0.88
	(0.09)	(0.09)*	(0.10)	(0.09)*	(0.11)	(0.09)
γ_b	0.82	0.78	0.89	0.89	0.92	0.88
	(0.09)*	(0.09)**	(0.12)	(0.12)	(0.09)	(0.08)
p ($H_0 : \gamma_a = \gamma_b$)	0.47	0.44	0.83	0.51	0.74	0.98
N	1,875	1,875	1,850	1,850	1,825	1,825
1st Stage F	70.89	82.63	82.83	82.45	54.72	59.90

(i) Standard errors in parentheses (clustered at the individual level).
(ii) All coefficients are significantly different from 0 at the 1% level. Therefore, t-tests of the null hypothesis (H_0: Coefficient = 1) are reported: * = 10%, ** = 5%, *** = 1%.
(iii) The rows corresponding to p ($H_0 : \gamma_a = \gamma_b$) report the p-statistic from a t-test of the equality of the coefficients γ_a and γ_b (i.e. a test of the asymmetric updating hypothesis).