Belief updating: does the ‘good-news, bad-news’ asymmetry extend to purely financial domains?

Kai Barron¹

Received: 20 February 2018 / Revised: 17 March 2020 / Accepted: 25 March 2020 /
Published online: 13 April 2020
© The Author(s) 2020

Abstract

Bayes’ statistical rule remains the status quo for modeling belief updating in both normative and descriptive models of behavior under uncertainty. Some 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 is present in a financial decision making context (i.e. a domain that is important for understanding much economic decision making). We find no evidence of asymmetric updating in this domain. In contrast, in our experiment, belief updating is close to the Bayesian benchmark on average. However, we show that this average behavior masks substantial heterogeneity in individual updating behavior. We find no evidence in support of a sizeable subgroup of asymmetric updaters.

Keywords Economic experiments · Bayes’ rule · Belief updating · Belief measurement · Proper scoring rules · Motivated beliefs

JEL Classification C11 · C91 · D83

1 Introduction

Throughout our lives, we are constantly receiving new information about ourselves and our environment. The way that we filter, summarize and store this new information is of critical importance for the quality of our decision making. Theories of human behavior under dynamic uncertainty are therefore enriched by an

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s10683-020-09653-z>) contains supplementary material, which is available to authorized users.

✉ Kai Barron
kai.barron@wzb.eu

¹ WZB, Reichpietschufer 50, 10785 Berlin, Germany

understanding of how individuals process new information. Economists typically write down models where information is summarized in the form of probabilistic ‘beliefs’ over states of the world, and updated upon receipt of new information according to Bayes’ rule. However, the assumption that individuals process information in a statistically accurate way has increasingly been called into question, with many studies documenting systematic deviations.¹ One important strand of this literature examines whether belief formation and updating is influenced by the *affective content* of the new information, 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; Möbius et al. 2014; Coutts 2019).

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—her beliefs are completely unaffected by the prizes and punishments she would receive in different states of the world. This fundamental assumption—that people update their beliefs symmetrically—is important because it underpins the entire orthodox approach to modelling uncertainty.

In this paper, we test whether individuals update their beliefs symmetrically in response to ‘good-news’ and ‘bad-news’ when states differ only in the financial rewards they offer. To do this, we conduct a laboratory experiment in which subjects complete a series of belief updating tasks. In each task, there are two possible states of the world. Subjects are told the prior probabilities of each state, and then receive a sequence of partially informative binary signals. We elicit subjects’ beliefs after each signal. The financial rewards associated with the two states are either identical (symmetric), or different (asymmetric). This experimental design 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 the other the signal constitutes ‘bad news’. We can also conduct a similar exercise for a single individual, comparing contexts where she has identical priors and signals, but the signal constitutes ‘good news’ in one context, and ‘bad news’ in the other. These comparisons provide a clean test of the *asymmetric updating hypothesis*.

The experiment considers belief updating in two contexts. In the SYMMETRIC treatment, subjects have an equal stake in each of the underlying states. In the two ASYMMETRIC treatments, a larger bonus payment is paid in one of the two states of the world. This design permits two separate tests of the asymmetric updating hypothesis. First, we can compare how the same individual responds to ‘good-news’ and ‘bad-news’ *within* the ASYMMETRIC treatments. Second, we can conduct a *between-subject*

¹ Some notable examples include the *representativeness bias* (Grether 1978, 1980, 1992), *cognitive dissonance* (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), *correlation neglect* (Enke and Zimmermann 2019), and *selection neglect* (Espanda and Vespa 2018; Barron et al. 2019; Enke 2019).

comparison of belief updating in the SYMMETRIC treatment and 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. First, we use exogenous variation in the priors to ensure that the estimates are robust to the econometric issues that arise when a right-hand side variable (i.e. the prior) is a lagged version of the dependent variable (i.e. the posterior). Second, 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 Online Appendix C for further details). Third, 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 hedging influence—both experimentally, and econometrically. Fourth, our experimental design allows us to study belief updating from priors spanning much of the unit interval. Importantly, averaging across all subjects, the design generates a balanced distribution of ‘good’ and ‘bad’ signals.

Our results show no evidence in favor 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.² This average behavior masks substantial heterogeneity in updating behavior at the individual level, but we find no evidence in support of a sizeable subgroup of asymmetric updaters.

The evidence reported here contributes to the recent literature studying the asymmetric updating hypothesis across different contexts (e.g. Eil and Rao 2011; Ertac 2011; Grossman and Owens 2012; Mayraz 2013; Möbius et al. 2014; Kuhnen 2015; Schwarzmann and Van der Weele 2019; Gotthard-Real 2017; Charness and Dave 2017; Heger and Papageorge 2018; Buser et al. 2018; Coutts 2019). The results in this literature thus far are highly heterogeneous, with some papers finding a greater responsiveness to good-news, some a greater responsiveness to bad-news, and some no evidence of an asymmetry. The objective of the experiment discussed in this paper is not to isolate the contextual features that activate and deactivate asymmetric updating; rather the objective is to provide a clean test of the asymmetric updating hypothesis for contexts in which states differ only in their material rewards and have no direct ego-relevance. The reason for focusing on this context (where states yield different financial outcomes) is that it characterizes a large class of economically important decision problems under uncertainty—most economic models with uncertainty fit this description. This paper does not disentangle the reasons for why updating is asymmetric in some settings but

² This is in line with, e.g. Holt and Smith (2009) and Coutts (2019), who find that average posteriors across all individuals are well approximated by posteriors obtained by applying Bayes’ rule.

not in others. Instead, by providing a clean test of the asymmetric updating hypothesis for one specific domain, the paper contributes to the growing collective body of evidence pertaining to asymmetric updating across a range of contexts. There are several candidate contextual and experimental design factors that could be generating the heterogeneous results observed in the literature as a whole. Section 7 below offers a discussion of some of these candidate explanations, and asks whether the existing body of evidence can help us to detect a systematic pattern that organizes the results.

The remainder of the paper proceeds as follows. Section 2 outlines the theoretical framework, Sect. 3 details the experimental design, Sect. 4 provides some descriptive statistics, Sect. 5 presents the empirical specification, Sect. 6 discusses the related literature and Sect. 7 concludes.

2 Theoretical framework and hypotheses

The following section discusses a simple framework for belief updating that augments the standard normative Bayesian benchmark to allow for several commonly hypothesized deviations from Bayes' rule. This framework is borrowed from Möbius et al. (2014) and is commonly used for analyzing belief updating descriptively. The framework provides a basis for the empirical approach that we will use to test whether 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 the state is $\omega = A$ and $1 - \pi_t$ is the agent's belief that the state is $\omega = B$. 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, $p(a|A) = p(b|B) = q > \frac{1}{2}$. The history, H_t , is defined as the sequence of signals received by the agent in periods $1, \dots, t$, with $H_0 = \emptyset$. Therefore, the history at time t is given by $H_t = (s_1, \dots, s_t)$.

To study how individuals update their beliefs, we follow Möbius et al. (2014) in considering the following model of *augmented Bayesian updating*:

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

The parameters δ , γ_a and γ_b can be interpreted as follows. If $\delta = \gamma_a = \gamma_b = 1$ then the agent updates her beliefs according to Bayes' rule. The δ parameter captures the degree to which the agent's prior affects her updating. For example, if $\delta > 1$ then

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)	$\gamma_a < \gamma_b$

the agent displays a *confirmatory bias*³, whereby she is more responsive to information that supports her prior. In contrast, if $\delta < 1$ she is more responsive to information that contradicts her prior (i.e. *base rate neglect*⁴). The former would predict that beliefs will polarize over time, while the latter would predict that over time beliefs remain too close to 0.5.

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 information than a Bayesian. 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.

2.1.1 Affective states

In the preceding section, the *affect* or desirability of different states of the world played no role. However, in most situations in which individuals form beliefs, there are some states that yield an outcome that is preferred to the outcomes associated with other states—i.e. there are *good* and *bad* states of the world

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.⁵ To do this, assume

³ 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. $\pi_i > 0.5$, 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. $\pi_i < 0.5$.

⁴ One can think of *base rate neglect* in this context as the agent forming her beliefs as if she attenuates the influence of her prior belief when calculating her posterior—i.e. acting as if her prior was closer to 0.5 than it actually was.

⁵ To avoid ambiguity, in the discussion below, the term ‘preference’ is usually used to refer to preferences over sure outcomes—never to a preference ordering 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 realization of a state in which a *good* outcome is realized.

that each of the two states of the world is associated with a certain outcome—i.e. in state $\omega = A$, the agent receives outcome x_A , and in state $\omega = B$, she receives x_B . There are now two belief updating scenarios:

- Scenario 1 (SYMMETRIC): the agent is indifferent between outcomes (i.e. $x_A \sim x_B$); and
- Scenario 2 (ASYMMETRIC): the agent strictly prefers one of the two outcomes (i.e. $x_A \succ x_B$).

The question of interest is whether the agent will update her beliefs differently in the SYMMETRIC and ASYMMETRIC scenarios. Under the assumption that the agent's behavior is consistent with the model described above in Eq. 1, this involves asking whether the parameters δ , γ_a and γ_b , differ between the two contexts.

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

Hypothesis 1 (Bayesian updating) *Individuals update their beliefs according to Bayes' rule. Therefore, $\delta = 1$, $\gamma_a = 1$ and $\gamma_b = 1$ in both SYMMETRIC and ASYMMETRIC scenarios.*

The second benchmark that we consider is provided by the *asymmetric updating hypothesis*—that individuals respond more to 'good-news' than 'bad-news'. Here, in our simple framework there are two ways to identify *asymmetric updating*.

First, if we only consider the behavior of individuals within the ASYMMETRIC scenario, we can 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'). For example, if $\gamma_a > \gamma_b$, this would indicate that the agent updates more in response to '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 'bad-news'. We refer to such an agent as a *pessimistic updater*.

Second, if we compare behavior between the SYMMETRIC and ASYMMETRIC scenarios, we can ask whether the parameters of Eq. 1 differ according to the scenario. We use the postscript $c \in \{A, S\}$ to distinguish the parameters in the two scenarios—i.e. δ^S , γ_a^S and γ_b^S in SYMMETRIC and δ^A , γ_a^A and γ_b^A in ASYMMETRIC. In the SYMMETRIC treatment, where the agent is completely indifferent between the two states, there is no reason to expect her updating to be asymmetric. Therefore, we assume that $\gamma_a^S = \gamma_b^S = \gamma^S$. Thus, the difference $\gamma_a^A - \gamma_b^A$ 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, $\gamma_b^A - \gamma_b^S$ 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.

Hypothesis 2 (Asymmetric updating) *Individuals update their beliefs asymmetrically, responding more to good than bad news. Therefore, within the ASYMMETRIC scenario, we will observe $\gamma_a^A > \gamma_b^A$. And in a comparison between the SYMMETRIC and ASYMMETRIC scenarios, we will observe $\gamma_a^A - \gamma_a^S > 0$ and $\gamma_b^A - \gamma_b^S < 0$. Together, we can summarize the asymmetric updating hypothesis parameter predictions as follows: $\gamma_a^A > \gamma^S > \gamma_b^A$.*

2.2 Belief elicitation and incentives

To empirically test the hypotheses above using an experiment, we would like to be able to elicit our participants' true beliefs. However, eliciting beliefs when studying the relationship between preferences and beliefs presents additional challenges. In particular, one needs to account for the inherent hedging motive faced by participants who have a stake in one state of the world (see Karni and Safra 1995 for a discussion). To obtain unbiased reported beliefs, we adopt the approach developed by Offerman et al. (2009), and extended to accommodate state-dependent stakes as in Kothiyal et al. (2011).

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 reported beliefs. We therefore 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 invert this function to recover the participant's true beliefs from her reported beliefs. Our objective, therefore, 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 use in our experiment to elicit beliefs is the quadratic scoring rule (QSR).⁶ Online Appendix B.2 provides a detailed discussion of the way in which reported beliefs might be distorted under the quadratic scoring rule. In the absence of state-dependent stakes, it is well documented that under the

⁶ 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 QSR has been widely used in the literature, implying that both the *theoretical properties* and *empirical performance* are well understood (see, e.g., Armanter 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.

QSR a risk averse agent should distort their reported belief towards 50%. With state-dependant stakes, a risk averse EU maximizer will face two distortionary motives in reporting her belief: (1) she will face the motive to distort her belief towards 50% as discussed above; and (2) in addition, there is a hedging motive, which will compel a risk averse individual to *lower* her reported belief, r_i , towards 0% as the size of the exogenous stake increases.⁷

If the participants in our experiment are *risk neutral expected utility* maximizers, the reported beliefs, r_i , that we elicit under the QSR will coincide with their true beliefs, π_i , and no belief correction is necessary. However, to account for a hedging motive (e.g. due to risk aversion), we measure the size of the distortionary influence of the elicitation incentives at an individual level and correct the beliefs accordingly. The following section provides the intuition for how this correction works.

2.2.1 A Non-EU ‘truth serum’

The Offerman et al. (2009) approach proposes correcting reported beliefs for the reporting bias generated by risk aversion or non-linear probability weighting. This approach involves eliciting subjects’ reported beliefs, r , corresponding to a set of risky events where both the participant and the analyst know the objective probabilities, p (*known probability*). This is done under precisely the same QSR incentive environment in which the subjects’ subjective beliefs, π , regarding the events of interest are elicited (*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 probabilities, $p \in [0, 1]$, to the reported beliefs, r , for each individual under the relevant incentive environment. Given this map, R , we can invert the function and recover the subject’s true beliefs from her reported beliefs about events with unknown probabilities, π .

In Online Appendix B.2, we offer a detailed discussion of how the Offerman et al. (2009) method operates, describes the underlying assumptions, and demonstrate how it can be augmented (as proposed in Kothiyal et al. 2011) to allow for the scenario where there are state-contingent stakes (i.e. $x \neq 0$).

3 Experimental design

The experiment is designed to test the *asymmetric updating hypothesis* using both within-subject and between-subjects comparisons of updating behavior. The experiment consists of three treatment groups. The treatment T1:SYMMETRIC corresponds to Scenario 1 (SYMMETRIC: no exogenous state-contingent stakes) and the other two

⁷ This assumes that the state-dependent payment is associated with the 100% state, not the 0% state. This assumption is made throughout the paper and the experiment.

treatment groups, T2: COMBINED and T3: SEPARATE, correspond to Scenario 2 (ASYMMETRIC: state-contingent stakes) discussed above.

The two ASYMMETRIC treatments, T2: COMBINED and T3: SEPARATE, have identical financial incentives. The rationale for running two treatments with identical incentives was to conduct additional checks to ensure that our results were not driven by the influence of a hedging motive. To do this, we vary only the way that the incentive environment is described to participants (i.e. we only vary the framing of the incentives).

The experiment proceeds in three stages. The first stage comprises the core belief updating task in which we elicit a sequence of reported beliefs from subjects as they received a sequence of noisy signals. In the second stage we collect 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 obtain 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 is chosen at random and paid out. In addition the participants receive a fixed fee of £5 [€5] for completing Stage 3, as well as a show-up fee of £5 [€5].⁸

3.1 The belief updating task (Stage 1)

The Belief Updating Task is the primary task of the experiment, and is summarized in Fig. 1. The task consists of five rounds. In each round, participants are presented with a pair of computerized 'urns' containing blue and red colored balls. Each of these two urns represents one of the two states of the world. The composition of the two urns is always constant, with state $\omega = A$ represented by the urn containing more blue balls (5 blue and 3 red), and state $\omega = B$ represented by the urn containing more red balls (5 red and 3 blue).

3.1.1 Priors

The five rounds differ from one another only in the exogenous prior probability that $\omega = A$ is 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 is known to the participant. The order of these rounds is randomly chosen for each individual. Conditional on the prior, p_0 , one of the two urns is then chosen through the throw of a virtual die, independently for each individual in each round.

⁸ In the discussion below, I will always refer only to Pounds (£), however half of the sessions were run at the Technical University in Berlin where all the payments were made in Euros. In all cases, the payment in Pounds (£) was equivalent to the payment in Euros (€).

Fig. 1 Overview of experimental design

3.1.2 Belief updating

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

In each round, we elicit the participant's reported belief, r_t , about the likelihood that state $\omega = A$ is the correct state of the world six times (i.e. for $t = \{0, 1, 2, 3, 4, 5\}$). We first elicit her reported belief, r_0 , directly after she is informed of the exogenous prior probability, p_0 , and then after she receives each of her five signals we elicit r_t

for $t = \{1, 2, 3, 4, 5\}$. Overall, we therefore elicit 30 reported beliefs in Stage 1 from each individual (6 reported beliefs in each of 5 rounds).⁹

3.1.3 Incentives and treatment groups

The Belief Updating Task is identical across treatment groups, with the exception of the incentives faced by participants. In each treatment, a participant's payment consists of two components: (1) an exogenous state-contingent payment, and (2) an accuracy payment that depends on both their stated belief and the true state.

In treatments T2: COMBINED and T3: SEPARATE, the state-contingent payment is £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 receive an equal state-contingent payment of £0.10 in both states.

In all three treatments, participants receive nearly identical detailed written instructions describing the belief updating task as well as the two payment components. With the aim of ensuring that participants understand the incentive environment, the instructions present 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 are presented with *payment tables*, which inform them of the precise prospect they would face for each choice of r_i , 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 2.¹⁰ In order to represent all payments as integers in the instructions and payment tables, we adopt the approach of using experimental points. The exchange rate is 6000 points = £1.

Table 2 shows the difference between treatments T2: COMBINED and T3: SEPARATE. While participants in these two treatments face precisely the same incentives, the treatments differ in the salience of the hedging motive. The only difference between the two treatments is the way in which the payment information is summarised in the *payment table*. In T2: COMBINED, the *payment table* shows the combined payment from both (1) the exogenous state-contingent payment, and (2) the accuracy payment. Therefore, it summarizes the reduced form prospect associated with each reported probability choice (r_i) for subjects. In T3: SEPARATE, the *payment table* shows only the accuracy payment, so subjects need to integrate the two payments themselves to spot the hedging opportunity.

Why do we have two treatments with identical incentives, differing only in presentation? The T2: COMBINED treatment presents participants with incentives faced in the simplest and clearest way. This is desirable as it maximizes subject understanding of the incentives. In contrast, T3: SEPARATE lowers the salience of the hedging opportunity with the aim of inducing more accurate belief reporting. Implementing

⁹ Participants are not informed about the correct urn at the end of each round. They only receive feedback when their payment is calculated at the very end of the experiment.

¹⁰ 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	1770	95	71,970	1770	95	11,970	1170
90	12,480	2880	90	71,880	2880	90	11,880	2280
85	12,330	3930	85	71,730	3930	85	11,730	3330
80	12,120	4920	80	71,520	4920	80	11,520	4320
75	11,850	5850	75	71,250	5850	75	11,250	5250
70	11,520	6720	70	70,920	6720	70	10,920	6120
65	11,130	7530	65	70,530	7530	65	10,530	6930
60	10,680	8280	60	70,080	8280	60	10,080	7680
55	10,170	8970	55	69,570	8970	55	9570	8370
50	9600	9600	50	69,000	9600	50	9000	9000
45	8970	10,170	45	68,370	10,170	45	8370	9570
40	8280	10,680	40	67,680	10,680	40	7680	10,080
35	7530	11,130	35	66,120	11,130	35	6930	10,530
30	6720	11,520	30	66,120	11,520	30	6120	10,920
25	5850	11,850	25	65,250	11,850	25	5250	11,250
20	4920	12,120	20	64,320	12,120	20	4320	11,520
15	3930	12,330	15	63,330	12,330	15	3330	11,730
10	2880	12,480	10	62,280	12,480	10	2280	11,880
5	1770	12,570	5	61,170	12,570	5	1170	11,970
0	600	12,600	0	60,000	12,600	0	0	12,000

both treatments serves multiple purposes: (1) it allows us to evaluate the influence of the presentation of incentives on reported beliefs (i.e. whether a hedging motive affected reporting), (2) it provides a way to test the internal validity of the correction mechanism we use, and (3) it provides us with a form of internal replication of our results, since we can test each hypothesis more than once.

3.2 The Offerman et al. (2009) correction task (Stage 2)

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

In Stage 2, participants are 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. For T1: SYMMETRIC, this specific example 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. The twenty reported beliefs correspond to the objective probabilities 0.05, 0.1, ..., 0.95.¹¹

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. There were two sessions for each of the three Treatment groups at each location, making twelve sessions and 222 participants in total. At each location, 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 approximately 1.5 h 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 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.¹² 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 9 in the Appendices provides a check that the selection of our preferred sample has not substantially biased our treatments groups. 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.

¹¹ These objective probabilities were presented in a random order. Furthermore, in the likelihood statement used, such as the example in the main text, 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 repetition of one of the first nineteen.

¹² 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 18% of our subjects. This is comparable to the 25% 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.

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.

5.1 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 their personal belief-distortion function (this is the R function discussed in Equation 8 in Online Appendix B.2.1). We estimate this function for each individual separately in order to correct the reported beliefs at the individual level. Essentially, we are simply fitting a curve through each subject's belief elicitation distortion, for the relevant incentive environment. Online Appendix B.2.4 contains a more detailed discussion of the estimation of the belief correction function, and illustrates the main ideas through Figures 5 and 6 (for further details about this procedure, see Offerman et al. 2009 and Kothiyal et al. 2011).

In Fig. 2 we plot the CDFs of the uncorrected reported beliefs as well as the reported beliefs that have been corrected at the individual level. The left panel displays the distribution of reported beliefs, prior to correction, in each of the three treatments. One interesting feature of this figure is that, in spite of the fact that subjects in T2: COMBINED and T3: SEPARATE are offered precisely the same incentives, 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 it is more salient.

The right panel shows the corrected beliefs. It suggests that the belief correction approach was successful in removing the strong hedging influence in T2: COMBINED. After correction, the beliefs in the two treatments with identical incentives, T2: COMBINED and T3: SEPARATE, are similar. This evidence underscores the usefulness of the belief correction mechanism in reducing bias in reported beliefs when there is a salient hedging opportunity.

With the corrected beliefs in hand, we proceed to the main analysis. The analysis is done using both the corrected and uncorrected beliefs.¹³

¹³ 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.

Fig. 2 CDFs of corrected and uncorrected reported beliefs

5.2 Core estimation specifications

Our core estimation equations aim to test for systematic patterns in updating behavior, within the framework of Eq. 1. First, we examine whether there are systematic deviations from Bayes' rule in updating behavior, independently of having a stake in one of the two states of the world. Second, we assess the influence that having a stake in one of the two states of the world has by (1) testing for an asymmetry in updating *within* treatments where there is a state-contingent stake (i.e. T2: COMBINED and T3: SEPARATE); and (2) 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 Eq. 1, allowing us to test the asymmetric updating hypothesis, and also to test for other common deviations from Bayes' rule, such as a *confirmatory bias* or *base rate neglect*:

$$\tilde{\pi}_{i,j,t+1} = \delta \tilde{\pi}_{i,j,t} + \gamma_a \tilde{q} \cdot 1(s_{i,j,t+1} = a) - \gamma_b \tilde{q} \cdot 1(s_{i,j,t+1} = b) + \epsilon_{i,j,t+1} \quad (2)$$

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.

To determine the belief updating pattern *within* each incentive environment, we estimate this equation separately for each treatment. 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.

5.2.1 Endogeneity of the Lagged Belief

One potential concern with the identification of the parameters of Eq. 2 is that 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 (i.e. if $E\{\tilde{\pi}_{i,j,t}\epsilon_{i,j,t+1}\} \neq 0$).¹⁴ 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 2. 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 *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 (see Online Appendix C and earlier versions of this paper for further discussion of this issue).

6 Results

Table 3 reports the results from estimating Eq. 2 for each of the treatment groups separately. These estimates describe the updating behavior of the average individual in 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) estimations 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: (1) the similarity in the updating patterns across the three treatment groups; and (2) 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

¹⁴ 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.

Table 3 Average updating behavior across treatments

	T1: SYMMETRIC		T2: COMBINED		T3: SEPARATE	
	Reported	Corrected	Reported	Corrected	Reported	Corrected
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
OLS						
δ	0.90 (0.03)***	0.90 (0.03)***	0.86 (0.04)***	0.86 (0.04)***	0.91 (0.03)***	0.93 (0.02)***
γ_a	1.12 (0.12)	1.09 (0.11)	1.05 (0.12)	1.06 (0.12)	1.25 (0.14)*	1.16 (0.11)
γ_b	1.21 (0.13)*	1.17 (0.11)	1.19 (0.15)	1.12 (0.13)	1.20 (0.12)	1.13 (0.10)
$p(H_0 : \gamma_a = \gamma_b)$	0.30	0.32	0.20	0.48	0.58	0.73
N	1075	1075	1285	1285	1140	1140
R^2	0.71	0.73	0.73	0.74	0.81	0.84
IV						
δ	0.99 (0.03)	0.99 (0.03)	1.00 (0.02)	0.99 (0.02)	0.99 (0.03)	0.99 (0.02)
γ_a	1.12 (0.12)	1.09 (0.11)	1.08 (0.12)	1.02 (0.11)	1.22 (0.13)*	1.14 (0.11)
γ_b	1.21 (0.12)*	1.16 (0.11)	1.11 (0.13)	1.12 (0.12)	1.18 (0.12)	1.11 (0.10)
$p(H_0 : \gamma_a = \gamma_b)$	0.30	0.31	0.76	0.25	0.63	0.74
N	1075	1075	1285	1285	1140	1140
1st Stage F	61.38	84.04	107.98	107.01	76.90	95.45

(1) Standard errors in parentheses (clustered at the individual level)

(2) 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%

(3) 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)

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 was, on average, not statistically different to that of a Bayesian, since both γ_a and γ_b are not significantly different to 1 at the 5% 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 the possible sources of endogeneity discussed above using our instrumental variable strategy, the estimates are no longer indicative of base rate neglect. Since the OLS estimates may be biased, the IV estimates represent our preferred results. The first stage regression results for the IV estimation are reported in the Appendices in Table 8, indicating that we

don't have a weak instrument issue. Importantly, however, the form of endogeneity addressed by IV estimation only pertains to potential biases in the δ parameter, since it addresses endogeneity in the prior belief variable. It is therefore reassuring that for the estimates of our primary parameters of interest, namely γ_a and γ_b , the estimates are largely consistent across all the estimation specifications reported in Table 3. A reason for this is that the signals, $s_{ij,t+1}$, that subjects receive are always completely exogenous, both across rounds, and with respect to the subject's personal characteristics. Furthermore, the distribution of signals observed is also exogenous, and balanced in expectation. This helps to avoid other sources of endogeneity (see, e.g., Online Appendix C) and to alleviate the influence of other potential confounding belief updating biases (see, e.g., Coutts (2019) for a discussion of the influence of signal distributions on belief updating).

It is worth noticing that although we observe a substantial difference in the levels of the corrected and uncorrected beliefs in Fig. 2 above, the estimates for updating in Table 3 are quite similar for the corrected and uncorrected beliefs. One explanation for this apparent inconsistency is the following. If the degree of hedging by an individual is similar for both the prior and posterior belief, then the correction could have a sizeable effect on the levels of both beliefs, but not result in the large difference in the estimated updating parameters, since updating pertains to the change in the belief rather than the level.

6.1 A model free test of the asymmetric updating hypothesis

In order to alleviate the potential concern that these results are dependent on the functional form of our empirical specification, we conduct a model-free test of the *asymmetric updating hypothesis*. Perhaps the simplest and most direct test of this hypothesis is obtained by directly comparing the posterior beliefs formed in two scenarios where the information set is identical, but the rewards associated with one of the states of the world is varied. Our data is well suited for conducting this exercise.

We do this by considering a comparison of information-set-equivalent posterior beliefs after individuals have received only (1) the exogenous prior and (2) a single ball draw. 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. Our data allows us to conduct two comparisons of information-set-equivalent posterior beliefs—a *within-subject* and a *between-subject* comparison.

First, 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, e.g. $p_0 = \frac{1}{6}$, and an identical signal, e.g. $s_1 = a$ (i.e. a blue ball), across treatments.

Second, we can compare information-set-equivalent posterior beliefs *within* treatment groups. This comparison involves comparing π_1 after $\{p_0 = p, s_1 = s\}$

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

Fig. 3 Comparison of beliefs after the receipt of a single signal and an exogenous prior

with $1 - \pi_1$ after $\{p_0 = 1 - p, s_1 = s^c\}$ where s^c is the complementary signal to s .¹⁵ 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 informativeness of a red ball exactly the same as a blue ball. 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 does not rely on Bayes' rule (although it is an implication of Bayes' rule), but rather only requires symmetric updating, and therefore it provides us with a non-parametric test of the *asymmetric updating hypothesis*.

Figure 3 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. Each group is labeled on the x axis by the prior belief associated with the 'red' bars, which correspond to the information sets that 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 for these bars the x axis label corresponds to $1 - p_0$. Within each group, the first two bars represent the average posterior beliefs in T1: SYMMETRIC; the second pair of bars depict the same for T2: COMBINED; and the third pair of bars for T3: SEPARATE.

¹⁵ Therefore, if $s = a$ then $s^c = b$ and vice versa.

The results displayed in Fig. 3 show that there are no systematic differences between posterior beliefs within information-set-equivalent groups, neither *within* nor *between* treatment groups. 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 comparisons¹⁶ are significant at the 5% 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 distribution. This lends support to the results described above which indicate that we fail to find evidence in support of the *asymmetric updating hypothesis*.

6.2 Robustness exercises

In addition to this model-free test, 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 Online Appendix A.

The first exercise examines whether the results from the main specification described in Eq. 2 are robust to first differencing the dependent variable (i.e. this imposes the assumption that $\delta = 1$). The second subsection extends the main empirical specification to allow for individual-specific updating parameters. For both of these specifications, an ex post power analysis is conducted, reporting the MDE for a significance level of $\alpha = 0.05$ and a power of $\kappa = 0.8$. The third 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 Eq. 2.

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.

6.3 Heterogeneity in updating behavior

In order to investigate whether the aggregate results are masking heterogeneity in updating behavior, we estimate Specification 2 at the individual level and collect the parameters. The distributions of these individual level parameters are reported in Fig. 4. Perhaps the most conspicuous feature of this figure is the fact that all three treatment groups display such similar parameter distributions in each of the panels—i.e. for each of the parameters. Testing for differences between the underlying distributions from which the parameters are drawn in the different treatment groups fails to detect any statistically significant differences

¹⁶ For these comparisons, for each exogenous prior, we test the following binary comparisons: (1) within treatment group, we test between those that received the $s_1 = a$ and $s_1 = b$ signals (3×5 binary comparisons); (2) for those that received the same signal, s_1 , we test between treatment groups (6×5 binary comparisons).

Fig. 4 Distributions of individual level updating parameters

in any of the four panels.¹⁷ 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.

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 Fig. 4 suggests that this is not the case by plotting the distribution of the individual level difference between the γ_a and

¹⁷ More specifically, for each of the four panels in Fig. 4, we conducted three binary Mann–Whitney tests comparing each possible pair of treatments. In total, these twelve statistical tests failed to detect any significant difference in the underlying parameter distributions between the treatment groups at the 10% level. In addition, another twenty-four similar tests for differences in the mean (t test) and median (chi-squared) of the underlying distributions between treatments fail to detect any statistically significant differences at the 10% level.

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. Furthermore, this conclusion is supported by the fact there are no significant differences between the distributions of updating parameters observed across the three treatments in any of the four panels, since the motive for a ‘good-news, bad-news’ effect is switched off in the T1: SYMMETRIC treatment.

7 Discussion

7.1 Heterogeneous results observed in the asymmetric updating literature

A central question that emerges from the discussion above is why we observe no evidence of a ‘good-news, bad-news’ effect here, while some other influential contributions to this literature have found evidence for such an effect. More generally, Benjamin (2019) points out that the evidence in this nascent literature is so far very mixed. In the economics literature, three papers find evidence in favor of stronger inference from *good news*: Eil and Rao (2011), Möbius et al. (2014) and Charness and Dave (2017).¹⁸ In contrast, there are three papers that find evidence of stronger inference from *bad news*: Ertac (2011), Kuhnen (2015) and Coutts (2019). Furthermore, in addition to the current paper, there are four other papers that find no evidence in favor of a preference-biased asymmetry in belief updating: Grossman and Owens (2012), Schwardmann and Van der Weele (2019), Gotthard-Real (2017) and Buser et al. (2018).¹⁹ In the psychology literature, however, there appears to be a near-consensus arguing that there is an asymmetry in belief updating in favor of good news (see, e.g., Sharot et al. 2011, 2012; Kuzmanovic et al. 2015; Marks and Baines 2017, amongst others). A notable exception is Shah et al. (2016) who argue that many of the contributions to this literature suffer from methodological concerns; Garrett and Sharot (2017) offer a rebuttal, claiming that optimistically biased updating is robust to these concerns.

The aim of the experiment discussed in this paper is not to isolate the contextual features that generate these heterogeneous results. Rather, it is to contribute robust evidence on updating for one particular domain—belief updating when states differ in terms of the financial rewards they bring. Nevertheless, it is important to consider the body of evidence as a whole to assess whether this reveals a systematic pattern that might organize these heterogeneous results. Below, I offer a discussion of some of the candidate explanations for the heterogeneity. In general, these explanations fall into two categories: (1) the hypothesis that contextual factors mediate

¹⁸ Mayraz (2013) also presents evidence in favor of individuals forming motivated beliefs that are distorted towards more desirable states, however in his experiment one cannot calculate the Bayesian posterior, so it is less comparable to the other studies in this literature.

¹⁹ Additionally, while Eil and Rao (2011) found evidence of an asymmetry in favor of good-news in the domain of beliefs about one’s own *Beauty*, they did not find evidence of an asymmetry in the domain of beliefs about one’s own *IQ*.

asymmetric updating: belief updating *is* influenced by preferences, but this preference-biased updating is switched on or off by contextual factors. (2) the hypothesis that asymmetric updating is sometimes misidentified: belief updating *is not* actually influenced by preferences, but rather, what appears to be asymmetric updating is driven by a different cognitive bias (e.g. prior-biased inference). The majority of the explanations discussed below fall into the first category.

7.1.1 Information structure

One avenue of enquiry for attempting to reconcile the results is to consider the differences in the information structures across experiments. For example, while several of the studies adopt a two-state bookbag-and-poker-chip experimental paradigm, Ertac (2011) uses a three-state structure with signals that are perfectly informative about one state, and Eil and Rao (2011) consider a ten-state updating task with binary signals. However, this does not seem to be driving the differences in results, since we observe heterogeneous results amongst papers with similar information structures—e.g. restricting attention only to the papers with two-state structures with binary signals (e.g. the current paper, Möbius et al. 2014; Gotthard-Real 2017; Coutts 2019) yields mixed results.

7.1.2 Priors

Focusing only on two-state experiments, there is substantial variation in the average prior belief across experiments, with Coutts (2019) (by design) observing relatively low average priors in comparison to Möbius et al. (2014), for example. If belief updating is influenced by prior beliefs (e.g. a confirmatory bias), then what looks like preference-biased belief updating may be driven by a completely different cognitive deviation from Bayesian updating, namely prior-biased updating (see, e.g., Benjamin 2019, Section 8). However, if we look at the papers that find evidence for preference-biased updating, Charness and Dave (2017) do find evidence in favor of prior-biased updating, while Eil and Rao (2011) and Möbius et al. (2014) do not find evidence of prior-biased updating. This speaks against the explanation that what appears to be asymmetric belief updating due to preferences is actually driven by a confirmatory bias.

7.1.3 Ambiguity

One important dimension in the belief updating literature is whether subjects are provided with exogenous point estimate priors or update from subjectively formed prior beliefs. There are advantages and disadvantages to each approach. The former

brings increased experimental control and improved causal identification, but this comes at the expense of reduced realism and perhaps a slightly less natural setting.²⁰

This discussion highlights a key assumption that is typically made in this literature, namely that subjects are probabilistically sophisticated and therefore update as if they hold a point estimate prior. In cases where subjects must form their own subjective prior belief, this assumption is not innocuous. If, instead, the beliefs subjects hold are not precise point estimates (e.g. if subjects hold ambiguous prior beliefs over an interval), then simple Bayesian updating may no longer be the most appropriate benchmark model. First, there are several competing theoretical models of belief updating in the presence of ambiguity with differing predictions, e.g. *full Bayesian updating* (Jaffray 1989; Pires 2002) and *maximum likelihood updating* (Gilboa and Schmeidler 1993), and some recent experimental evidence testing between them (Ngangoue 2018; Liang 2019). Second, one might postulate that there is greater scope for motivated reasoning when one is updating beliefs from ambiguous priors (or ambiguous signals) in comparison to belief updating from exogenously endowed point estimate priors (and signals with clearly defined informativeness).

However, the existing evidence suggests that the presence or absence of ambiguity is not a unifying explanation for the differing results. Even within the set of papers with home-grown subjective priors (e.g. Eil and Rao 2011; Ertac 2011; Möbius et al. 2014, and Coutts 2019), the results are very mixed.

7.1.4 Domain of belief updating

Typically, we study belief formation as if it is domain-independent. However, it seems natural to consider the possibility that humans evolved to process information about their *physical environment* differently from information about their *self* and their *social environment*. For example, the mental processes involved in forming a belief about the likelihood of future rainfall may be fundamentally different to those involved in forming a strategic belief about the probability that another individual will be trustworthy in a specific scenario. Some papers in this literature have explored this question by asking whether we update differently about a given fundamental characteristic of one's own self in comparison to the same fundamental characteristic of another individual (e.g., Möbius et al. 2014; Coutts 2019). Furthermore, recent theoretical and experimental work has studied how individuals attribute outcomes to their *self* versus an external fundamental from their *physical* or *social environment* (see, e.g., Heidhues et al. 2018; Hestermann and Le Yaouanq 2020; Coutts et al. 2019).

The influence of the domain on belief updating may matter for the asymmetric updating literature because this literature considers updating scenarios pertaining to both the *environment* (with 'good' states typically represented by high monetary payments) and to the *self* (where 'good' states pertain to a desirable individual

²⁰ An interesting recent contribution by Le Yaouanq and Schwardmann (2019) proposes a methodology for studying belief updating in more natural settings, while still maintaining experimental control and permitting a comparison with Bayesian updating.

characteristic). One might posit that asymmetric updating only manifests in certain domains. While the studies focused on belief updating where only financial rewards are varied generally don't find a 'good-news, bad-news' effect (e.g. this paper, Gotthard-Real 2017; Coutts 2019), even within the group of studies considering beliefs about the self, the evidence is mixed (e.g. Eil and Rao 2011; Ertac 2011; Möbius et al. 2014, and Coutts 2019).

7.1.5 Outcomes and stake size

One caveat to the results reported here is that the financial stakes in play are not extremely large. Coutts (2019) investigates the role played by stakes, increasing the stake size up to \$80 and finds no evidence that it plays a role. This suggests that stake size may not be a pivotal concern. Nonetheless, it is worth keeping this caveat in mind when interpreting the results.

7.2 Addressing hedging

One challenge when studying belief updating in the presence of state-dependent stakes is the inherent hedging motive. Various approaches have been adopted to try to deal with it. Typically, the papers considering belief updating about an ego-related characteristic (e.g. IQ) rely on the (very reasonable) implicit assumption that hedging across the ego utility and monetary utility domains will be minimal.

Amongst the studies considering monetary state-dependent stakes, two different approaches to dealing with this challenge have been adopted. Coutts (2019) follows the method suggested in Blanco et al. (2010), which involves partitioning the world such that the participant will either be paid according to their belief *or* according to the state-dependent prize, but never both. In this paper we instead follow the method developed by Offerman et al. (2009) and Kothiyal et al. (2011). Both these approaches are theoretically valid for alleviating the influence of hedging under the assumptions they make; both approaches have advantages and drawbacks. However, neither Coutts (2019) nor the current paper finds support for a larger responsiveness to good-news. In addition, while the possibility of hedging certainly deserves attention when interpreting the results, there are several reasons why the results discussed above suggest that hedging is not a driving factor behind the absence of a good-news, bad-news effect. The observed absence of any asymmetry in updating within either of the two ASYMMETRIC treatments for either the uncorrected or corrected beliefs, as well as the consistency in updating patterns observed across all three treatments is not easily explained by a combination of a 'good-news, bad-news' effect and hedging.

7.3 Concluding remarks

The main finding of this paper is that we find no evidence for asymmetric updating when states differ only in their financial rewards. Instead, we find that the updating behavior of the average individual is approximately Bayesian, irrespective of the

presence or absence of differing financial stakes. Taken in the context of the literature as a whole, where a ‘good-news, bad-news’ effect is sometimes observed in other domains, this suggests that belief updating may be context dependent. One interesting avenue for future research, therefore, is to consider whether differing evolutionary considerations across contexts may have generated differing belief updating processes. For example, one evolutionary reason why there could be differences in belief updating about ego-relevant characteristics in comparison to belief updating about external states that differ in financial rewards is the idea that ego maintenance can yield evolutionary benefits. A positive asymmetry in updating about one’s self 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 2019; Solda et al. 2019; Coffman et al. 2020). In contrast, asymmetric updating about external states of the world would lead to overoptimism which is likely to lead to costly mistakes. Further research is needed to help enrich our understanding of precisely how context influences belief updating.

Acknowledgements Open Access funding provided by Projekt DEAL. I would like to thank Syngjoo Choi, Alex Coutts, Steffen Huck, Heather Jacklin, Terri Kneeland, Peter Schwardmann, Yves Le Yaouanq, Lukas Wenner, and the audiences at RES, ECBE, THEEM, ASFEE, ENTER Jamboree, and the ESA conferences for constructive comments and valuable discussions. I would also like to thank the anonymous referees and editor who made numerous extremely helpful suggestions. Brian Wallace, Lukas Wenner, Roberta De Filippis, and Nina Bonge provided me with valuable support in running the experiments. I gratefully acknowledge financial support from the Deutsche Forschungsgemeinschaft (through CRC TRR 190). All mistakes are my own.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Akerlof, G. A., & Dickens, W. T. (1982). The economic consequences of cognitive dissonance. *American Economic Review*, 72(3), 307–319.
- Armantier, O., & Treich, N. (2013). Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. *European Economic Review*, 62, 17–40.
- Barron, K., Huck, S., & Jehiel, P. (2019). Everyday econometricians: Selection neglect and overoptimism when learning from others. *Working paper*.
- Benabou, R., & Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics*, 117, 871–915.
- Benjamin, D. (2019). Chapter 2: Errors in probabilistic reasoning and judgment biases. In *Handbook of behavioral economics: Foundations and applications 2* (pp. 69–186). North-Holland.

- Bernardo, A. E., & Welch, I. (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy*, 10(3), 301–330.
- Blanco, M., Engelmann, D., Koch, A. K., & Normann, H.-T. (2010). Belief elicitation in experiments: Is there a hedging problem? *Experimental Economics*, 13(4), 412–438.
- Brunnermeier, M., & Parker, J. (2005). Optimal expectations. *American Economic Review*, 95(4), 1092–1118.
- Burks, S. V., Carpenter, J. P., Goette, L., & Rustichini, A. (2013). Overconfidence and social signalling. *Review of Economic Studies*, 80(3), 949–983.
- Buser, T., Gerhards, L., & Van Der Weele, J. (2018). Responsiveness to feedback as a personal trait. *Journal of Risk and Uncertainty*, 56(2), 165–192.
- Caplin, A., & Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics*, 116(1), 55–79.
- Charness, G., & Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, 104, 1–23.
- Coffman, K. B., Collis, M., & Kulkarni, L. (2020). Stereotypes and belief updating. *Working paper*.
- Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22, 369–395.
- Coutts, A., Gerhards, L., & Murad, Z. (2019). No one to blame: Biased belief updating without attribution. *Working paper*.
- Eil, D., & Rao, J. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2), 114–138.
- Enke, B. (2019). What you see is all there is. *Working paper*.
- Enke, B., & Zimmermann, F. (2019). Correlation neglect in belief formation. *Review of Economic Studies*, 86(1), 313–332.
- 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.
- Esponda, I., & Vespa, E. (2018). Endogenous sample selection: A laboratory study. *Quantitative Economics*, 9(1), 183–216.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171–178.
- Garrett, N., & Sharot, T. (2017). Optimistic update bias holds firm: Three tests of robustness following Shah et al. *Consciousness and Cognition*, 50, 12–22.
- Gilboa, I., & Schmeidler, D. (1993). Updating ambiguous beliefs. *Journal of Economic Theory*, 59(1), 33–49.
- 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.
- Grossman, Z., & Owens, D. (2012). An unlucky feeling: Overconfidence and noisy feedback. *Journal of Economic Behavior & Organization*, 84, 510–524.
- Heger, S., & Papageorge, N. (2018). We should totally open a restaurant: How optimism and overconfidence affect beliefs. *Journal of Economic Psychology*, 67, 177–190.
- Heidhues, P., Köszegi, B., & Strack, P. (2018). Unrealistic expectations and misguided learning. *Econometrica*, 86(4), 1159–1214.
- Heifetz, A., Shannon, C., & Spiegel, Y. (2007). The dynamic evolution of preferences. *Economic Theory*, 32(2), 251.
- Hestermann, N., & Le Yaouanq, Y. (2020). Experimentation with self-serving attribution bias. *American Economic Journal: Microeconomics* (forthcoming).
- Holt, C., & Smith, A. (2009). An update on Bayesian updating. *Journal of Economic Behavior & Organization*, 69(2), 125–134.

- Hossain, T., & Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies*, 80(3), 984–1001.
- Jaffray, J.-Y. (1989). Linear utility theory for belief functions. *Operations Research Letters*, 8(2), 107–112.
- Johnson, D. D., & Fowler, J. H. (2011). The evolution of overconfidence. *Nature*, 477(7364), 317–320.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237.
- Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, 77(2), 603–606.
- Karni, E., & Safra, Z. (1995). The impossibility of experimental elicitation of subjective probabilities. *Theory and Decision*, 38(3), 313–320.
- Kothiyal, A., Spinu, V., & Wakker, P. P. (2011). Comonotonic proper scoring rules to measure ambiguity and subjective beliefs. *Journal of Multi-Criteria Decision Analysis*, 17, 101–113.
- Kuhnen, C. M. (2015). Asymmetric learning from financial information. *Journal of Finance*, 70(5), 2029–2062.
- Kuzmanovic, B., Jefferson, A., & Vogeley, K. (2015). Self-specific optimism bias in belief updating is associated with high trait optimism. *Journal of Behavioral Decision Making*, 28(3), 281–293.
- Le Yaouanq, Y., & Schwardmann, P. (2019). Learning about one's self. *CESifo working paper*.
- Liang, Y. (2019). Learning from unknown information sources. *Working paper*.
- Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal*, 97, 666–684.
- Marks, J., & Baines, S. (2017). Optimistic belief updating despite inclusion of positive events. *Learning and Motivation*, 58, 88–101.
- Mayraz, G. (2013). The optimism bias: Experiment and model. *Working paper*.
- Möbius, M., Niederle, M., Niehaus, P., Rosenblat, T. (2014). Managing self-confidence. *Working paper*.
- Ngangoue, K. (2018). Learning under ambiguity: An experiment on gradual information processing. *Working paper*.
- Offerman, T., Sonnemans, J., van de Kuilen, G., & Wakker, P. (2009). A truth serum for non-bayesians: Correcting proper scoring rules for risk attitudes. *Review of Economic Studies*, 76(4), 1461–1489.
- Pires, C. P. (2002). A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2), 137–152.
- Rabin, M., & Schrag, J. (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics*, 114(1), 37–82.
- Schwardmann, P., & Van der Weele, J. (2019). Deception and self-deception. *Nature Human Behaviour*, 3(10), 1055–1061.
- Shah, P., Harris, A., Hahn, U., Catmur, C., & Bird, G. (2016). A pessimistic view of optimistic belief updating. *Cognitive Psychology*, 90, 71–127.
- Sharot, T., Kanai, R., Marston, D., Korn, C. W., Rees, G., & Dolan, R. J. (2012). Selectively altering belief formation in the human brain. *Proceedings of the National Academy of Sciences*, 109(42), 17058–17062.
- Sharot, T., Korn, C. W., & Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality. *Nature Neuroscience*, 14(11), 1475–1479.
- Solda, A., Ke, C., Page, L., & Von Hippel, W. (2019). Strategically delusional. *Experimental Economics*, <https://doi.org/10.1007/s10683-019-09636-9>.
- Trautmann, S., & van de Kuilen, G. (2015). Belief elicitation: A horse race among truth serums. *Economic Journal*, 125, 2116–2135.