

Belief Updating: An experimental test of Bayes rule and the ‘good-news, bad-news’ asymmetry

Kai Barron, *University College London*

EXTENDED ABSTRACT

Beliefs are important for all economic situations where an agent makes decisions, while facing uncertainty over the past, present or future states of the world. In many of these situations, the agent will prefer one state of the world to another. In thinking about the way in which we form beliefs over these states, the norm within economics has been to model agents’ information processing as being unaffected by these preferences over states. In other words, belief updating is treated as being orthogonal to preferences, with agents represented as dispassionate Bayesians. Recently, however, a growing literature has questioned this separability between belief formation and preferences. For example, several theoretical papers have explored the channels through which preferences may distort beliefs¹ or have taken payoff-distorted beliefs as a premise of the model and explored the implications for behaviour in a specific context². Similarly, there is an increasing trend in applied work to use payoff-distorted beliefs to explain empirical regularities that are difficult to reconcile with the standard rational Bayesian model³.

It is clear that there is a demand for robust empirical evidence that provides a better understanding, firstly, of whether beliefs are distorted by payoffs in certain contexts, and secondly, for evidence that helps to understand the mechanisms driving this distortion. While survey data can provide suggestive evidence in this direction, it is extremely difficult to separately identify the role of preferences and beliefs. Therefore, these questions provide a clear mandate for the experimental

¹See, for example: Brunnermeier and Parker (2005), Mayraz (2013), Bracha and Brown (2014), Benabou and Tirole (2002), Benoît and Dubra (2011), Compte and Postlewaite (2004)

²E.g. Eliaz and Spiegler (2008), De la Rosa (2011), Bridet and Schwardmann (2013), and Kőszegi (2006).

³Spinnewijn (2008, 2013) documents that job-seekers substantially overestimate how quickly they will find work. He examines the implications of this systematic bias for optimal unemployment insurance design. Oster, Shoulson and Dorsey (2013) examine the behaviour of individuals at high risk of Huntington’s disease, and reach the conclusion that their behaviour is most consistent with the Brunnermeier and Parker (2005) “Optimal Expectations” model for optimism. Malmendier and Tate (2008) classify CEO’s according to their confidence level and show that CEO’s that they identify as overconfident tend to overestimate their ability to generate returns from a company and undertake mergers that destroy value. Grubb (2009) provides evidence that cell phone companies in the USA adopt pricing plans that take advantage of consumers’ overconfidence regarding the precision of their demand forecasts.

approach to have a voice in this conversation. Unsurprisingly, the last few years has seen several important contributions from the experimental literature aimed precisely at addressing these questions. These papers can be broadly taxonomised into two groups that differ along two dimensions: (i) the domain over which beliefs are formed (e.g over attributes relating to *self-image* or events associated with *monetary outcomes*); and (ii) whether they examine single shot *static belief formation* or *dynamic belief updating*.

The first group of papers elicit beliefs over some personal characteristic or attribute that relates to the participant's *self-image* (e.g. the participant's intelligence, ability or attractiveness). Möbius et al. (2014) examine updating over beliefs regarding students' relative performance on an IQ test, while in a closely related paper Eil and Rao (2011) look at belief updating over both intelligence and attractiveness to members of the opposite sex, relative to the other participants in the session. These papers both examine *dynamic belief updating* in response to a sequence of information signals.

In contrast, the second group of papers, including Coutts (2014), Mayraz (2013) and Heger and Papageorge (2013), elicit beliefs over states that are associated with *monetary outcomes* of differing magnitude. This class of papers examine *static belief formation* situations in the sense that participants form a single belief over an event, E , that is associated with a monetary outcome.

In this paper, we bridge the gap between these two strands of the literature by examining *dynamic belief updating* with varying *monetary state-dependent stakes*. We conduct a laboratory experiment that allows us to robustly detect whether belief updating in response to a sequence of signals is affected by how much the participants *want* a certain state to be realised. We do this by comparing belief updating in two contexts: (i) a context in which neither state is preferred to the other; (ii) a context where one state of the world is clearly preferred. We use monetary incentives to vary the relative desirability of the states and examine how individuals update their beliefs in response to a sequence of noisy information signals in a balls-and-urns experiment. This paper contributes to the literature on belief updating in the presence of state-dependent stakes in several ways.

Firstly, this setup allows for prior beliefs to be exogenously assigned to participants. This allows us to avoid the issues that arise when the subject enters the lab with a history of information regarding a given belief (e.g. an individual's belief regarding her intelligence⁴). Secondly, in our experiment we are able to examine updating behaviour from priors that span most of the unit interval⁵. This allows us to distinguish belief distortions due to payoffs from other factors, such as probability weighting (see Wakker (2010) for a review). Thirdly, while beliefs about one's own intelligence and attractiveness are certainly important for a wide range of economic situations, there are a vast set of economic contexts in which individuals process information relating to the

⁴In this case it is difficult for the experimenter to fully control for the agent's information set; or to observe whether this prior belief is correlated with other factors that would influence how she updates her beliefs

⁵We exogenously assign priors that range from $p_0 = \frac{1}{6}$ to a prior of $p_0 = \frac{5}{6}$. Each individual carries out 5 rounds of updating from priors spanning this interval.

likelihood that different monetary outcomes will be realised. It is unclear that the results from beliefs regarding ones ability will translate to this context. It is therefore important to test whether these results are transferable, which is one of the primary objectives of the current paper. In short, in our setting the agent's entire information set as well as the payoffs associated with each state are exogenously determined. To the best of our knowledge, this is the first paper that examines the influence of preferences on dynamic belief updating where the priors are exogenously assigned.

In addition, this paper contributes to the discussion surrounding the major technical challenge faced by this literature in that it inherently requires the measurement of beliefs in contexts where agents have stakes in the possible states. This measurement problem⁶ arises when there are state-dependent stakes, since under any belief elicitation incentives linked to states, risk averse participants will face a hedging motive that generates difficulties in accurately measuring participants' true beliefs⁷. We offer a solution to this hedging problem by augmenting the Offerman et al. (2009) 'truth serum elicitation' method to elicit unbiased beliefs, while allowing participants to have state-dependent stakes.

In contrast to earlier papers in this literature, we find no evidence of an asymmetry in updating towards (or away from) more desirable states of the world that one would expect from an optimist (pessimist). Interestingly, we find that on aggregate participants in our experiment are surprisingly close to Bayesian in their belief updating. This is in sharp contrast to, for example, Möbius et al. (2014) who find that subjects are approximately $\frac{1}{3}$ as responsive to new information about their own ability than a Bayesian would be. In order to test whether the aggregate results are masking underlying individual heterogeneity, we conduct a finite mixture model analysis of updating, allowing for three types in each treatment group. The striking result of this analysis is that in each treatment group we observe the same three distinct types of updaters. The first group, comprising around 20-25% of each treatment group, are largely unresponsive to new information. The second group, also comprising around 20 - 25%, are (i) extremely responsive to new information, treating a single signal as a Bayesian would treat 3 similar signals, and (ii) tend to update more away from the the state favoured by their prior, than towards it. The third group, comprising the majority of each treatment group, are less responsive to new information than a Bayesian, updating $\frac{1}{2}$ to $\frac{2}{3}$ as much as a Bayesian would to a new piece of information. However, none of these groups exhibit asymmetric updating. Furthermore, it is informative to highlight that while aggregate updating behaviour in our experiment is close to Bayesian, this hides the fact that each of the sub-groups exhibit updating behaviour that is distinctly different from Bayesian.

⁶This measurement problem is precisely the subject of Karni and Safra (1995) who show that when decision makers have stakes in events, the experimental elicitation of unbiased subjective probabilities is impossible.

⁷For a good recent review of this literature, see Schotter and Trevino (2014) or see Trautmann and van de Kuilen (2014) for a horse race between elicitation methods.

References

- Benabou, R and Tirole, J (2002), ‘Self-confidence and personal motivation’, *The Quarterly Journal of Economics*, 117:3, pp.871–915.
- Benoît, JP and Dubra, J (2011), ‘Apparent overconfidence’, *Econometrica*, 79:5, pp. 1591–1625.
- Bracha, A. and Brown, D. (2014), ‘Affective Utilities: A Rational Theory of Optimistic Bias in Asset Markets’,.
- Bridet, L. and Schwardmann, P. (2013), ‘Selling dreams: Endogenous overoptimism and collateral use in financial contracts’, *working paper*.
- Brunnermeier, M. and Parker, J. (2005), ‘Optimal expectations’, *AER*.
- Compte, O. and Postlewaite, A. (2004), ‘Confidence-enhanced performance’, *American Economic Review*, pp.1536–1557.
- Coutts, A. (2014), ‘Testing Models of Belief Bias: An Experiment’, *working paper*.
- Eil, D. and Rao, J. (2011), ‘The good news-bad news effect: asymmetric processing of objective information about yourself’, *American Economic Journal: Microeconomics*, 3:2, pp. 114–138.
- Eliasz, K and Spiegler, R (2008), ‘Consumer optimism and price discrimination’, *Theoretical Economics*, 3:4, pp. 459–497.
- Grubb, M. (2009), ‘Selling to overconfident consumers’, *The American Economic Review*, 99:5, pp. 1770–1807.
- Heger, S. and Papageorge, N. (2013), ‘We should totally open a restaurant: performance uncertainty and optimistic beliefs’, *Available at SSRN 2348410*.
- Kőszegi, B. (2006), ‘Emotional agency’, *The Quarterly Journal of Economics*, pp. 121–155.
- Malmendier, U. and Tate, G. (2008), ‘Who makes acquisitions? CEO overconfidence and the market’s reaction’, *Journal of Financial Economics*, 89:1, pp.20–43.
- Mayraz, G. (2013), ‘The optimism bias: experiment and model’, *working paper*.
- Möbius, M. et al. (2014), ‘Managing Self-Confidence’, *working paper*.
- Offerman, T. et al. (2009), ‘A truth serum for non-bayesians: Correcting proper scoring rules for risk attitudes’, *The Review of Economic Studies*, 76:4, pp.1461–1489.
- Oster, E, Shoulson, I and Dorsey, E (2013), ‘Optimal expectations and limited medical testing: evidence from Huntington disease’, *AER*.

- Rosa, L. De la (2011), 'Overconfidence and moral hazard', *Games and Economic Behavior*, 73:2, pp. 429–451.
- Schotter, A. and Trevino, I. (2014), 'Belief Elicitation in the Laboratory', *Annual Review of Economics*:0.
- Spinnewijn, J. (2008), 'Unemployed but optimistic: Optimal insurance design with biased beliefs', *Available at SSRN 1291566*.
- Spinnewijn, J. (2013), 'Insurance and perceptions: how to screen optimists and pessimists', *The Economic Journal*.
- Trautmann, S. and Kuilen, G. van de (2014), 'Belief elicitation: A horse race among truth serums', *The Economic Journal*.
- Wakker, P. (2010), *Prospect theory: For risk and ambiguity*, (Cambridge University Press).