A Guideline and Cautionary Note:

How to Use the Belief Update Task Correctly

Tali Sharot^{1,2} & Neil Garrett³

¹Affective Brain Lab, Department of Experimental Psychology, University College London

²The Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London

³School of Psychology, University of East Anglia

Correspondence to: t.sharot@ucl.ac.uk

Abstract: In 2011 we published a paper showing that people update their beliefs to a greater extent in response to good news (e.g., learning that the likelihood of robbery is lower than expected) than bad news (e.g., learning it is lower than expected) (Sharot et al., 2011). This phenomenon, which can lead to increased optimistic beliefs, is absent in depression. Since then, our belief update task has been used by many others to test a wide range of questions related to belief formation and optimism. Most of these studies are rigorous and well conducted. However, a small number of researchers have used the task inappropriately, inserting new confounds and failing to control for other potential ones. This has resulted in the report of false findings which have muddied the literature and caused confusion. Given these incidents and the enthusiasm for using the task across different disciplines, the need for guidelines on how to use the belief update task correctly has become apparent. The belief update task can be a helpful tool in studying beliefs in domains ranging from climate change to health, but like any other task it must be used properly if valid conclusions are to be reached. We hope this guide will be helpful for scientists who would like to use the belief update task, as well as readers, reviewers and editors who are required to evaluate studies using this task.

In 2011 (Sharot et al., 2011) we published a paper showing that healthy individuals tend to update their beliefs to a greater extent in response to unexpected positive information (e.g., learning that the likelihood of being a victim of card fraud is lower than expected) than negative information (e.g., learning it is higher than expected). This phenomenon, which can lead to increased optimistic beliefs, is absent in depression (Korn et al., 2014; Garrett et al., 2014). Since then, our task - known as the “belief update task” - has been used by scientists around the world to test a wide range of questions related to belief formation and optimism (e.g., Korn et al., 2012; Ma et al., 2016; Kuzmanovic et al., 2016; Kappes et al., 2018; Ossola et al., 2020; Oganian et al., 2015). For example, the task has been used to examine belief formation in atypical populations (Kuzmanovic et al., 2019), to examine beliefs updating about others (e.g., Kappes et al., 2018), and to test the underlying biological mechanisms of belief updating (Ma et al., 2016). The basic design has also been modified to examine belief change in response to new information about one’s own or other’s personality (e.g., Korn et al., 2012), politics (e.g., Tappin et al., 2017) and personalized genomic information (e.g., Krieger et al., 2016) among others.

All the studies cited above were carefully conducted, creative and informative. However, we have come across attempts to use the belief update task in inappropriate ways that lead to false finding (e.g., Shah et al., 2016). This is the result of authors changing the task in ways that introduce new confounds, failing to control for possible confounds and using over-lenient and inappropriate analyses. Below we explain these potential pitfalls and how to avoid them. We hope this guide will be helpful for scientists who would like to use the belief update task, as well as readers, reviewers and editors who are required to evaluate studies using this task. Below we recap the task and analysis and then list the common pitfalls and how to avoid them (if you are familiar with the task/analysis, skip to the pitfall section).

A quick recap of the belief update task:

Participants are presented with approximately 40-80 life events (e.g., ‘robbery’) and asked to estimate how likely the event is to happen to them in the future (this is referred to as the ***first estimate***). They are then presented with the base rate of the event in a demographically similar population, that is someone of a similar age, location and socio-economic background (this is referred to as ***information***). In a second session, participants are asked again to provide estimates of their likelihood of encountering the same events (this is referred to as ***second estimate***). While most studies use negative stimuli (e.g., ‘victim of card fraud’), some use positive stimuli (e.g. ‘experience a sunny day in February’). Other valid variations of the original task have been developed (Garrett and Sharot, 2017, 2014; Kappes et al., 2018; Kuzmanovic et al., 2016; Ma et al., 2016).

To assess and control for potential confounds, at the end of the task participants provide ratings for each event on a Likert scale according to their past experience with the events, familiarity with the event, how negatively or positively they believe the event to be, how vivid they imagine the event and how emotionally arousing they perceive the event (we refer to these as ***‘subjective ratings’***) as well as provide the actual probability previously presented for each event (***‘memory’***).

Figure 1: Belief Update Task. In the first session (top row), on each trial, participants are presented with a short description of a life event (Burglary in this example) and asked to estimate how likely this event is to occur to them in the future. They are then presented with the probability of that event occurring to someone from the same age, location and socio-economic background as them. After completing approximately 40-80 trials they then complete a second session (bottom row). This is the same as the first except that the average probability of the event to occur is not presented. Shown are examples of trials for which the participant's estimate of an adverse event was higher (a) or lower (b) than the statistical information provided leading to receipt of good news and bad news, respectively. Update is calculated as the difference between participants estimates in the two sessions (i.e. pre and post information). Figure adapted from Garrett et al., 2018.

A quick recap of the analysis:

Trials are divided into those in which participants received **good news** (for example, the probability presented of encountering an aversive event was lower than the subject's first estimate of their own probability; see **Figure 1a**) or **bad news** (for example, the probability presented of encountering an aversive event is higher than the subject's first estimate of their own probability; see **Figure 1b**). If using positive life events (e.g., Marks & Baines, 2017; Garrett & Sharot, 2017), good news will be trials in which the subject learns that a positive event (e.g., winning a large sum of money) is higher than the subject's first estimate and bad news is the opposite. Comparison is not between positive and negative events, but between information that is better (thus subjectively positive) or worse (thus subjectively negative) than expected.

Participants change in beliefs for each trial is referred to as the **update** and is calculated as the absolute difference between participants' second and first estimate. The update is then signed such that positive numbers indicate a move towards the new information and negative numbers indicate an update in the opposite direction of the new information.

For each subject, average update for good news and bad news trials is then calculated. The difference between these two types of trials is the **update bias**, an explicit measure of valence dependent belief updating. Positive update bias scores indicate greater updating in response to good news relative to bad news (this is an optimistic update bias) whilst negative scores indicate the reverse (a pessimistic update bias). In healthy populations, we commonly observe

the update is greater in response to good news than bad news (i.e. an optimistic update bias, **Figure 2a**).

For each trial an **estimation error** term is calculated as the absolute difference between the information (probability presented) and participant's first estimate on that trial. When examining whether the update for good news is greater than for bad news, it is vital that estimation errors are controlled for (see below).

Finally, we calculate a **learning bias**, which is the difference between the **learning parameter** for good news trials and bad news trials. A learning parameter measures to what extent participants update their beliefs in proportion to the error made. A learning bias is calculated by conducting a linear regression where the update on each trial is the dependent variable and the estimation error on each trial is the independent variable. This is done for each participant for good news trials and for bad news trials separately, which provides two beta coefficients for each participant. The difference in this beta coefficient is the learning bias. This learning bias captures differences in the extent to which people update beliefs in proportion to the evidence when receiving good versus bad news. We commonly observe the learning parameter is greater in response to good news than bad news (**Figure 2b**).

Other valid and useful analysis have also been developed (Garrett and Sharot, 2017, 2014; Kappes et al., 2018; Kuzmanovic et al., 2016; Ma et al., 2016).

Fig. 2. Typical Results. Results typically show that participants **(a)** update their beliefs more in response to “good news” than “bad news” and **(b)** that greater weight is put on the estimation error when updating in response to “good news” than “bad news”. Bar shows the mean, error bars SEM and dots represent participants (overlaid on top of each other). Data displayed in **(a)** is taken from Garrett and Sharot, 2014. Data displayed in **(b)** taken from Garrett et al., 2018.

Six common pitfalls and how to avoid them

1. **Skewing Event Probabilities** – Consider a scientist who wishes to use the belief update task and includes events that all have a true probability of 20% (i.e. far to the left of

the midpoint, 50%). Let's imagine that all these events are aversive (like 'leukaemia'). Any event for which a participant's first estimate is greater than 20% will be considered a 'good news' trial. For example, if a participant indicates that they believe their likelihood of leukaemia is 50% and then learns it is only 20%, that is good news. They will then likely provide a new estimate that is lower. Let's say 28%. In contrast, any event in which a participant's first estimate is lower than 20% will be considered a 'bad news' trial. For example, if a participant indicates that they believe their likelihood of leukaemia is 10% and learns it is 20%, that is bad news. They will then likely provide a new estimate that is higher. Let's say 18%. But here is the problem – over the course of the experiment, there will be more room for participants to update when receiving good news and less room to update when receiving bad news. The closer the true probabilities are towards zero, the more this will be the case (Garrett and Sharot, 2017). As we have shown using simulations, this will create an artificial update bias even when none exists (Garrett and Sharot, 2017, **Figure 3**).

For example, a previous study that overlooked this pitfall (Shah et al., 2016) falsely reported a pessimistic bias in belief updating about positive stimuli. The effect failed to replicate (Garrett and Sharot, 2017; Marks and Baines, 2017). To correct for this problem, scientists need to select stimuli with probabilities with a mean at the centre of the scale that is available to participants. These events should be either normally or uniformly distributed around the mean. For example, if using a scale between 3% and 77% the mean should be 40%. In cases where this is not possible, all analysis must control for estimation errors. Since estimation errors quantify the difference between first estimates and information provided they are, by definition, the "room" available for updating. However, readers should be very sceptical of studies that severely skew stimuli, such as allowing a scale between 0% and 100% with a mean probability of 30%.

2. **Not Using an Appropriate Scale** – In our original task, very rare or very common events are not included with all event probabilities lying between 10% and 70%. Participants were told that the range of probabilities lay between 3% and 77% and were only permitted to enter estimates within this range. This is done for two reasons. First, it is known that people's perception of very low and very large probabilities is distorted (Allais, 1953; Kahneman and Tversky, 1979; Luce, 1999). Second, it is important to ensure that the range of possible overestimation is equal to the range of possible underestimation. That is, if all event probabilities lie between 5% and 80% and participants are allowed to enter numbers between 0% and 100% then by design they will not be able to update upwards as much as downwards.
3. **Not Controlling for Estimation Errors** – Even when the mean probabilities of the events are at the centre of the scale, estimation errors still need to be controlled for. This is because participants may have a certain tendency to provide low (or high) first estimates, which will lead to an artificial difference in the "room" participants have to

update beliefs upwards or downwards. Any study that does not control for estimation errors (or alternatively for first estimates) is at high risk of reporting false findings.

4. **Not Controlling for Possible Confounds** – An optimistic update bias in belief updating may be due to differences in remembering good and bad news. To control for this, the normal protocol is to ask participants to recall the information provided (after the Belief Update Task has finished) and then use this to calculate *memory errors* - the absolute difference between information presented during the task and participants' recollection of that information after the task. These memory errors can then be added as covariates in the analysis. In addition, differences in past experience with the events, familiarity with them, how negativity/positively they are perceived to be, how vividly they are imagined and how emotionally arousing they are can emerge between events that end up in the "good news" bin and the "bad news" bin. This is why it is important to collect these subjective ratings and add them as covariates. Whilst we and others have shown many times that these differences are usually null and do not account for the optimistic update bias, they may emerge as an effect if using new stimuli, testing atypical populations or using modified versions of the original task.
5. **Using a Low Number of Trials** – We recommend using at least 40 trials and preferably up to 80 trials. This is because in the analysis, trials are subsequently divided into "good news" trials and "bad news" trials (based on participants responses), a process in which some trials may also need to be omitted – for example if estimates are equal to the information provided they cannot be partitioned into either category. It is important therefore to use at least 40 trials so that a decent number of trials end up in each bin for analysis. Failure to do this will increase the noise in the update bias and learning bias scores thereby inflating the chance of false positives.
6. **Degrees of Freedom** – Some researchers opt to use a mixed linear model to analyse the data. This is a valid choice if these models include random intercepts and slopes for each participant (and control for estimation errors in the model as discussed above). Failure to do this (e.g., not including any random effects or only including random intercepts) can inflate degrees of freedom even 40 fold thereby increase Type-1 error rates theoretically by up to 100% (Barr et al., 2013; Judd et al., 2012; Murayama et al., 2014). Given this, if a model does not converge with random slopes, we recommend using a more conservative approach such as an ANOVA (Sharot et al., 2011). We advise readers to be very cautious when evaluating findings using the update task, which do not adequately account for random effects and inflated degree of freedom.

Fig. 3. Example of Pitfall 1. (a) Example of stimuli used in a past study (Shah et al., 2016) in which probabilities for positive and negative life events are skewed towards the low end of the scale. (b) Example of stimuli used in a past study (Garrett and Sharot, 2017) where the distribution of base rates are normally distributed around the centre of the scale used for positive and negative life events. (c) Using the base rates of Shah et al., simulations published by us (Garrett and Sharot, 2017) reveal an optimistic update bias in updating for negative events (that is more updating for good news than bad news) and a pessimistic update bias for positive events (that is more updating for bad news than good news). (e) This pattern closely matches the human data reported by Shah et al. (d) Simulations using normally distributed base rates reveal no bias in updating. (f) Human participants nonetheless show an optimistic bias in updating for both positive and negative life events using these same base rates (Garrett and Sharot, 2017). (g) An illustration of how providing participants with event probabilities towards the edge of the response scales will artificially bias their ability to update beliefs. Error bars represent SEM.

Conclusion

Despite many examples of rigorous and insightful use of the update bias task (e.g., Korn et al., 2012; Ma et al., 2016; Kuzmanovic et al., 2016; Kuzmanovic et al., 2019; Kappes et al., 2018; Oganian et al., 2019; Moutsiana et al., 2014) a small number of researchers have used the task inappropriately (e.g., Shah et al., 2016), committing all or part of the pitfalls listed above (e.g., skewing event base rates, not controlling for potential confounds etc.). As a result they

report impossible findings, muddy the published literature and caused confusion. The belief update task can be a helpful tool in studying belief update in domains ranging from climate change (Sunstein et al., 2016) to health (Kreiger et al., 2016; Ossola et al., 2020), but like any other task it must be used correctly if valid conclusions are to be reached.

References

- Allais M. 1953. Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* **21**:503–546. doi:10.2307/1907921
- Barr DJ, Levy R, Scheepers C, Tily HJ. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *J Mem Lang* **68**:10.1016/j.jml.2012.11.001. doi:10.1016/j.jml.2012.11.001
- Garrett N, González-Garzón AM, Foulkes L, Levita L, Sharot T. 2018. Updating beliefs under perceived threat. *Journal of Neuroscience* **38**:7901–7911.
- Garrett N, Sharot T. 2017. Optimistic update bias holds firm: Three tests of robustness following Shah et al. *Consciousness and cognition* **50**:12–22.
- Garrett N, Sharot T. 2014. How robust is the optimistic update bias for estimating self-risk and population base rates? *PLoS One* **9**:e98848.
- Garrett N, Sharot T, Faulkner P, Korn CW, Roiser JP, Dolan RJ. 2014. Losing the rose tinted glasses: neural substrates of unbiased belief updating in depression. *Frontiers in human neuroscience* **8**:639.
- Judd CM, Westfall J, Kenny DA. 2012. Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology* **103**:54–69. doi:10.1037/a0028347
- Kahneman D, Tversky A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* **47**:263–291. doi:10.2307/1914185
- Kappes A, Faber NS, Kahane G, Savulescu J, Crockett MJ. 2018. Concern for Others Leads to Vicarious Optimism. *Psychol Sci* **29**:379–389. doi:10.1177/0956797617737129
- Kappes A, Harvey AH, Lohrenz T, Montague PR, Sharot T. 2020. Confirmation bias in the utilization of others' opinion strength. *Nature Neuroscience* **23**:130–137. doi:10.1038/s41593-019-0549-2
- Korn, C. W., Sharot, T., Walter, H., Heekeren, H. R., & Dolan, R. J. (2014). Depression is related to an absence of optimistically biased belief updating about future life events. *Psychological medicine*, *44*(3), 579-592.
- Korn, C. W., Prehn, K., Park, S. Q., Walter, H., & Heekeren, H. R. (2012). Positively biased processing of self-relevant social feedback. *Journal of Neuroscience*, *32*(47), 16832-16844.
- Korn, C. W., La Rosée, L., Heekeren, H. R., & Roepke, S. (2016). Social feedback processing in borderline personality disorder. *Psychological Medicine*, *46*(3), 575-587.
- Krieger JL, Murray F, Roberts JS, Green RC. 2016. The impact of personal genomics on risk perceptions and medical decision-making. *Nat Biotechnol* **34**:912–918. doi:10.1038/nbt.3661
- Kuzmanovic B, Jefferson A, Vogeley K. 2016. The role of the neural reward circuitry in self-referential optimistic belief updates. *NeuroImage* **133**:151–162. doi:10.1016/j.neuroimage.2016.02.014
- Kuzmanovic B, Rigoux L, Vogeley K. 2019. Brief Report: Reduced Optimism Bias in Self-Referential Belief Updating in High-Functioning Autism. *J Autism Dev Disord* **49**:2990–2998. doi:10.1007/s10803-016-2940-0
- Luce RD. 1999. Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches. New York: Psychology Press. doi:10.4324/9781410602831
- Ma Y, Li S, Wang C, Liu Y, Li W, Yan X, Chen Q, Han S. 2016. Distinct oxytocin effects on belief updating in response to desirable and undesirable feedback. *Proceedings of the National Academy of Sciences of the United States of America* **113**:9256–61. doi:10.1073/pnas.1604285113
- Marks J, Baines S. 2017. Optimistic belief updating despite inclusion of positive events. *Learning and Motivation* **58**:88–101. doi:10.1016/j.lmot.2017.05.001

- Moutsiana C, Charpentier CJ, Garrett N, Cohen MX, Sharot T. 2015. Human Frontal-Subcortical Circuit and Asymmetric Belief Updating. *J Neurosci* **35**:14077–14085. doi:10.1523/JNEUROSCI.1120-15.2015
- Murayama K, Sakaki M, Yan VX, Smith GM. 2014. Type I error inflation in the traditional by-participant analysis to metamemory accuracy: A generalized mixed-effects model perspective. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **40**:1287–1306. doi:10.1037/a0036914
- Oganian, Y., Heekeren, H. R., & Korn, C. W. (2019). Low foreign language proficiency reduces optimism about the personal future. *Quarterly Journal of Experimental Psychology*, 72(1), 60-75.
- Ossola, P., Garrett, N., Sharot, T., & Marchesi, C. (2020). Belief updating in bipolar disorder predicts time of recurrence. *Elife*, 9, e58891.
- Shah P, Harris AJ, Bird G, Catmur C, Hahn U. 2016. A pessimistic view of optimistic belief updating. *Cognitive Psychology* **90**:71–127.
- Sharot T, Korn CW, Dolan RJ. 2011. How unrealistic optimism is maintained in the face of reality. *Nature neuroscience* **14**:1475–1479. doi:10.1038/nn.2949
- Tappin BM, Van Der Leer L, McKay RT. 2017. The heart trumps the head: Desirability bias in political belief revision. *Journal of Experimental Psychology: General* **146**:1143.

Acknowledgments. The authors thank Chris Korn for comments on an earlier version of this manuscript. Tali Sharot is funded by a Wellcome Trust Senior Researcher Fellowship. Neil Garrett is funded by a Sir Henry Wellcome Fellowship (209108/Z/17/Z).