

# A descriptive Bayesian account of optimism in belief revision

Derek Powell (dmpowell@asu.edu)  
School of Social and Behavioral Sciences  
Arizona State University

## Abstract

A number of findings suggest that people's expectations about the future are unrealistically optimistic (e.g. Sharot et al., 2011). This bias is thought to result from the "motivational modulation" of evidence, driven by the desire to feel positively about one's own future (Sharot, 2011). However, evaluating "bias" in belief revision requires careful comparison against a rational standard, and recent arguments and findings (Shah et al., 2016) give reason to doubt much of the evidence for optimism bias. Descriptive Bayesian models allow for a direct comparison of human belief updating against the Bayesian rational standard (Tauber et al., 2017). Here, these analyses indicate widespread "conservatism," or weaker-than-rational belief revision. However, in contrast to the widely-reported "optimism bias," participants more commonly displayed pessimism than optimism in their belief revision. Both effects were marked by significant heterogeneity, with a sizable fraction of participants engaging in largely rational updating.

**Keywords:** Cognitive Psychology; Belief updating; Motivated reasoning; Bayesian computational models

"The truth hurts" or so you may have variously heard from poets, singers, and philosophers—from Twain, Nietzsche, and Lizzo alike.

The formation and maintenance of veridical beliefs in light of new experiences is a key function of cognition. Accurate beliefs are essential to inform adaptive action in the present and plans for the future. Nevertheless, some truths can sting: It can be hard to accept that a leader has betrayed us, that our actions have hurt others, or that our hopes for the future are unlikely to be realized. In such cases, truth-seeking might compete with other motivational factors: for instance, people might seek to maintain face within their ingroup (e.g. Kahan, 2013), preserve their positive self-concept (e.g. Dunning et al., 1989; Sanitioso et al., 1990), or feel good about the future (Sharot et al., 2011; e.g. Sharot, 2011).

Consistent with the idea that such motivated reasoning is pervasive (Kunda, 1990), a large body of findings suggest that people's expectations about the future are unreasonably rosy (Chowdhury et al., 2014; Sharot et al., 2011; e.g. Sharot, 2011; Weinstein, 1980). This "unrealistic optimism" is typically thought to reflect a self-serving bias to feel positively about one's own future. This desire affects how new experiences and information are interpreted, leading to the "motivational modulation" of evidence (Sharot, 2011) and the discounting of "bad news."

In a now-classic study, Sharot and colleagues (2011) asked participants to estimate their own personal risk for a number

of different positive and negative life events. Then, they presented those participants with information on the true base rate risk for those events in the population and asked them to again estimate their own personal risk. Many studies report that the differences between participants' initial and updated estimates are greater (in magnitude) when the base rate was lower than participants' own estimates for negative events and higher for positive events. From this, researchers have concluded that people are optimistically biased in their belief updating (e.g. Chowdhury et al., 2014; Kuzmanovic et al., 2015; Sharot et al., 2011).

However, Shah and colleagues (2016) identified serious flaws in the logic of these belief updating experiments. Most fundamentally, these original studies do not provide sufficient information to determine whether the information being presented to participants constitutes good or bad news. This is because people's perceptions of their own personal risk and baseline risk are not necessarily the same. Alice might reasonably believe that the risk for skin cancer in the general population is 3%, but that her own risk is 10% given her complexion, family history, and so forth. And Bob might reasonably believe that the prevalence of skin cancer in the general population is 10%, but that his own risk is 5% given his own personal characteristics. Suppose these two are both informed that the true risk is 6%. Intuitively, it should be clear this is bad news for Alice (being higher than they expected) and good news for Bob, but the procedures used by previous research (e.g. Sharot et al., 2011) would reverse these categorizations.

Determining whether people's belief updating is biased and could therefore constitute evidence for motivated reasoning requires a comparison against a rational standard. More broadly, there are a number of examples of apparent biases in belief revision and formation that can be explained by alternative rational accounts (e.g. Austerweil & Griffiths, 2011; Chater et al., 2020; Jern et al., 2014; Navarro & Perfors, 2011).

So, is there an optimism bias? And can such a bias be demonstrated in clear contrast to any competing rational accounts?

## A rational model for belief revision

Bayes rule provides a normative, rational standard for belief updating. According to Bayes' rule, the posterior probabil-

ity of a hypothesis given some evidence can be calculated from the probability of the evidence given the hypothesis (the likelihood) and the prior probability of the hypothesis in the absence of this evidence.

$$P(h|d) = \frac{P(d|h)P(h)}{P(d)} \quad (1)$$

To judge their own personal risk, a rational reasoner would integrate any knowledge or evidence they possess about their own relative personal risk with their estimation of the prior or average risk. This can be expressed simply using the log-odds form of Bayes Rule, whereby the log odds of the posterior is the sum of the prior odds and the log of what is sometimes called an “evidence ratio” (ER).

$$\text{Posterior log-odds} = \text{Prior log-odds} + \log \text{ER} \quad (2)$$

The evidence ratio is the ratio of the likelihoods of the data given the hypothesis,  $P(D|H)/P(D|\neg H)$  (hence it is also sometimes called a likelihood ratio). This evidence ratio, or the evidence an individual possesses (or believes themselves to possess), can thus be inferred from their estimates of the base rates and their own risk. After being informed of the true base rate, a rational reasoner should then update their beliefs about their own personal risk accordingly, substituting in the new base rate. With some algebra, the log-odds of the posterior belief can be calculated from the prior base rate estimate, the prior personal risk estimate, and the true base rate provided. With all terms transformed to log-odds, the schematic form of the equation is:

$$\text{Posterior risk} = \text{prior risk} + (\text{true base rate} - \text{prior perceived base rate}) \quad (3)$$

Shah and colleagues (2016) use this rational analysis to derive predictions for rational belief updating and attempted to design studies that could demonstrate evidence for optimism as compared with a true rational standard. Their results actually provided some evidence for *pessimism* rather than optimism. However, they note several challenges that complicate their attempts to address this research question: most fundamentally, issues posed by the boundedness of the probability scale for the use of additive statistical models.

Fortunately these concerns are relatively straightforwardly addressed by recent computational advances that simplify the estimation of Bayesian generalized linear models with non-normal response distributions. As described in the following section, this makes Shah and colleagues’ (2016) paradigm an attractive test case for exploring belief updating and motivated reasoning generally. Much of the original foundational work on Bayesian belief updating concerned reasoning tasks about entirely artificial contexts, such as flipping coins or drawing differently-colored balls from urns (e.g. Edwards, 1968). These artificial tasks made it easy to define what is “rational.” Sharot and colleagues’ (2011) base rate task, as refined by Shah and colleagues (2016), similarly offers a clear

definition of rational belief updating, but in the context of more ecologically-valid and meaningful beliefs.

### Descriptive Bayesian models of belief revision

The methodological challenges raised by Shah and colleagues (2016) can be addressed by adopting a descriptive Bayesian approach that estimates deviations from optimality as part of a Beta regression model predicting posterior beliefs.

The equations above provide an optimal rational model for belief revision in light of base rate information. But this optimal model can also be extended to form a variety of descriptive Bayesian models that can be used to examine possible deviations from optimality—including biases plausibly evidencing motivated reasoning. Comparisons of these descriptive models promise a test of motivated versus rational belief updating that is valid by the Bayesian perspective’s own lights.

The equation below expresses the rational model’s predictions for the posterior judgment ( $\mu$ ) given a specific prior risk judgment ( $x$ ), the true base rate ( $b$ ) and the initial perceived base rate ( $\tilde{b}$ ) for a participant  $i$  and item  $j$ . Grouping the last two terms of equation 4 together with parenthesis should make its appearance familiar: together these last two terms form the log evidence ratio representing the evidence provided by the true base rate for a person’s personal risk estimate.

$$\text{logit}(\mu_{ij}) = \text{logit}(x_{ij}) + (\text{logit}(b_j) - \text{logit}(\tilde{b}_{ij})) \quad (4)$$

The regression equation above defines a generalized linear model, where inputs are transformed by the logit function and responses on the probability scale are connected to the linear model with a logit link function. This regression equation can be combined with a Beta likelihood to model people’s posterior probability judgments ( $y_{ij}$ ). The Beta distribution is a continuous probability distribution for probabilities or proportions, that is naturally bounded in  $[0, 1]$ . In the context of Beta regression, it is parameterized by its central tendency  $\mu$  and precision  $k$ .

$$y_{ij} \sim \text{Beta}(\mu_{ij}k, (1 - \mu_{ij})k) \quad (5)$$

Under the rational model, the predicted posterior probability is fixed by the equation as written. But under a descriptive Bayesian model, additional terms can be added to capture how belief revision might deviate from optimality.

Over 50 years of cognitive psychological research has found that people often do not update their beliefs as they ought to according to Bayes rule, a phenomenon known as “conservatism” (Edwards, 1968; Erev et al., 1994; Fischhoff & Beyth-Marom, 1983). These findings are quite robust, although there is considerable contention about their source (Erev et al., 1994). They may owe to a general “anchoring” response-bias (e.g. Slovic & Lichtenstein, 1971) or to imperfect representation of the evidence (Edwards, 1968). Or, apparent “conservatism” may simply reflect an imperfect trust of the information being provided in cognitive psychological

experiments (Corner et al., 2010). In any case, the degree to which participants incorporate new base rate evidence when updating their beliefs can be captured by adding the multiplicative parameter  $\alpha$  to the regression equation.<sup>1</sup>

$$\text{logit}(\mu_{ij}) = \text{logit}(x_{ij}) + \alpha_i(\text{logit}(b_j) - \text{logit}(\tilde{b}_{ij})) \quad (6)$$

Extending this descriptive Bayesian approach further allows for a test of optimism in belief updating. The multiplicative term can be expanded into a linear equation, with a new coefficient  $\beta$  and a new variable  $g_{ij}$ , a binary variable representing whether the information being presented represents good news (1) or bad news (0). This binary variable is coded 1 for negative events where participants perceived the base rates to be higher than the true base rates, 1 for positive events where participants perceived the base rates to be lower than the true base rates, and zero otherwise. The  $\beta$  coefficient thus offers a test of optimism: if  $\beta$  is greater than zero, this indicates optimism (greater incorporation of evidence for good rather than bad news). If it is less than zero, this indicates pessimism.

$$\text{logit}(\mu_{ij}) = \text{logit}(x_{ij}) + (\alpha_i + \beta_i g_{ij})(\text{logit}(b_j) - \text{logit}(\tilde{b}_{ij})) \quad (7)$$

## Methods

This study was preregistered (<https://osf.io/k4am6>). All data and analysis code are available at <https://osf.io/crp98/>.

### Participants

A total of 197 participants were recruited through the CloudResearch survey recruitment platform, which provides access to a group of prescreened workers from Amazon’s mechanical Turk work distribution website. All participants were at least 18 years old and were located in the United States. Participants were compensated \$2.25 for approximately 10 to 15 minutes of participation. Participants who failed a simple attention check question were excluded from analysis, leaving a final sample of 186 participants (76 female, median age 36).

### Materials and procedures

Twenty potential life events were drawn from a larger pool of items used by Shah et al. (2016). Ten of the events were positive (e.g. living to be at least 90 years old), and ten were negative (e.g. being diagnosed with cancer).

<sup>1</sup>If “mistrust” of the experimentally-provided evidence is the source of conservatism (as suggested by Corner and colleagues (2010), then some belief-updating contexts might be better represented as a case of Jeffrey-conditionalization (Jeffrey, 1983). In this case, an analogous  $\alpha$  parameter would represent the probability people assign to the facts presented. However, given that the base rate information being provided here is a continuous quantity, this may not be a complete representation of the evidence’s interpretation.

For each event, participants estimated their own probability of experiencing the event and the average person’s probability of experiencing the event (order counterbalanced across participants). Then, they were informed of the true base rates for the event. Finally, they rated their own personal estimate of their risk after having seen this information. All questions were presented on individual screens and numerical responses were freely typed.

After completing the primary task, participants were also asked to respond to the Life Orientation Test-Revised (LOT-R, Scheier et al., 1994), a measure of personality trait optimism.

## Results

All models were implemented as fully Bayesian mixed-effects models using the non-linear syntax functions of the brms R package (Bürkner, 2017). Model posteriors were estimated using the No-U-Turn Markov Chain Monte Carlo (MCMC) sampler implemented in Stan. Four MCMC chains were run, with 2000 samples (1000 burn-in) drawn from each. Chains were assessed for convergence with  $\hat{R}$  and the total estimated effective sample size was verified to be greater than 1000 for all parameters (Gelman et al., 2014).

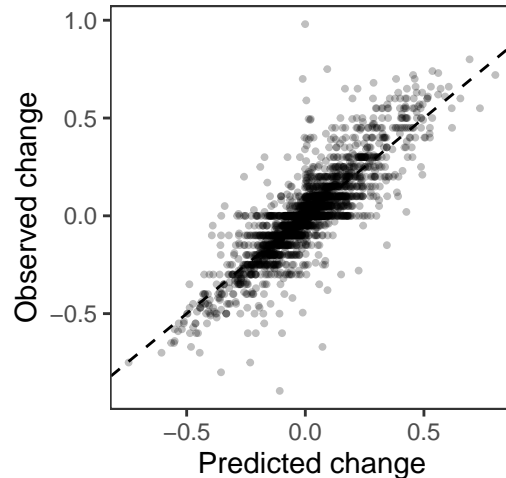


Figure 1: Predicted versus observed changes in probability judgments under the descriptive Bayesian model with  $\alpha$  parameter across all responses.

### Rational versus descriptive Bayesian models

A rational Bayesian model of belief updating can be used to predict posterior (i.e. posttest) beliefs a priori, without requiring any parameters be estimated. In contrast, descriptive Bayesian models must be fit to data.

First, I compared a rational Bayesian model’s predictions (equation 4) against a descriptive Bayesian model incorporating a multiplicative parameter for each individual participant that estimates the degree to which they incorporate new evidence into their beliefs (equation 6). This was a




Figure 2: Forest plots of participant-level  $\alpha$  (left) and  $\beta$  (right) parameters from descriptive Bayesian models of belief revision. Individual estimates for  $\alpha$  that do not credibly differ from 1 are highlighted, indicating responses consistent with rational responding. On the right, estimates for  $\beta$  that are credibly above or below zero are highlighted to indicate optimism and pessimism.

Bayesian hierarchical Beta regression model that also included population-level parameters for the expectation and spread of this update multiplier.

Both models can be treated as Beta regressions, where participants' posterior (posttest) beliefs are distributed under a Beta distribution parameterized according to its expectation  $\mu$  and spread  $k$  (equation 5).

These models both predict posterior beliefs, but it is of particular interest how well they capture changes in people's beliefs. The descriptive Bayesian model generates substantially better predictions for these changes ( $R^2 = 0.724$ ) than did the rational model ( $R^2 = 0.62$ ).

Of course the descriptive Bayesian model is far more complex in its parameterization. The models can also be compared formally by  $ELPD_{loo}$ , an estimate of the posterior probability of new unseen data given the model and the posterior distribution of its parameters (Vehtari et al., 2017). This measure, an estimate of leave-one-out cross validation performance, provides an estimate of a model's ability to generalize to new data that accounts for differences in model complexity.<sup>2</sup> Again, the descriptive Bayesian model is very strongly preferred ( $\Delta ELPD_{loo} = -1843$ ,  $SE = 136.1$ ).

On average, participants did not update their beliefs as much as Bayes' rule prescribes, indicated by the estimated population-level multiplier parameter,  $\alpha_{pop} = .60$ , 95% CI [.57, .64]. This is consistent with a large body of prior findings demonstrating conservatism in belief revision (Edwards, 1968; e.g. Erev et al., 1994; Fischhoff & Beyth-Marom, 1983). Figure 2 presents a forest plot showing the posterior distribution of participants' multiplier parameters. As can be seen, the clear majority of these estimates are credibly below one, indicating that these participants shifted their beliefs less than predicted by Bayes rule. In addition, a number of participants' estimated multiplier parameters are near zero,

<sup>2</sup>Note that this does require one parameter to be fit for the rational model, capturing the spread of the Beta distribution,  $k$ .

indicating they hardly updated their beliefs at all.

Nevertheless, there is substantial heterogeneity across individuals. It is especially worth emphasizing that a substantial portion of participants' estimated multiplier parameters do not credibly differ from 1. Correspondingly, a good number of participants' posterior beliefs and belief changes are well-captured by the rational Bayesian model ( $R^2 > .90$  for 30 of 186 participants). Figure 3 shows predicted belief changes under the rational model against observed belief changes for a sampling of participants with different multiplier estimates.

There were some participants who did not substantively update their beliefs at all in this task. These failures seem qualitatively different from the under-appreciation of evidence consistent with "conservatism." One possibility is that these participants were entirely distrustful of the experimental context. Another possibility is that they are exhibiting a form of base rate neglect. Base rate neglect is most commonly observed when people's judgments are *overly* sensitive to diagnostic information, as they apparently neglect prior probability information (a sort of anti-conservatism, Kahneman & Tversky, 1973; Tversky & Kahneman, 1974). For instance, people might judge that a "shy student who enjoys math" is more likely to be a physics major than a psychology major, neglecting to consider that many more students study psychology than physics.

The base rate updating task here requires some active understanding of Bayesian inference: participants must appreciate that base rates (priors) are relevant to their own personal risk (posteriors). Some participants may fail to reason about this appropriately, thereby failing to see the relevance of the base rate information and neglecting to update their beliefs.

### Optimism in belief updating

Optimism or pessimism in belief updating can be examined by extending the descriptive Bayesian model with an additional parameter sensitive to the valence of the evidence be-

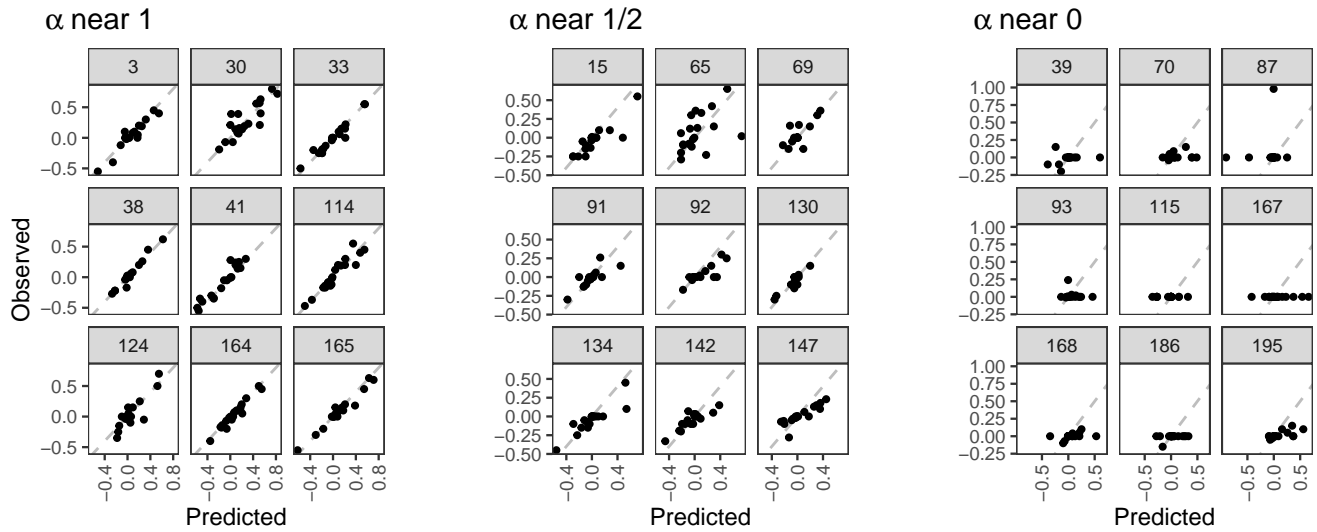


Figure 3: Illustration of participant-level variation in belief updating. Random participants were sampled from among participants with varying  $\alpha$  parameters. Faceted plots show predicted versus observed belief changes under the rational Bayesian model. Belief updating for participants with  $\alpha$  near 1 is largely consistent with the rational model’s predictions. Participants with  $\alpha$  near zero generally fail to update their beliefs. Dashed lines indicate perfect correlation.

ing provided (captured in equation 7 above). This model adds a predictor indicating whether the news is classified as good (1) or bad (0), and a multiplicative coefficient that codes for the difference in updating for good versus bad news. The addition of this parameter substantially improves the quality of the model fit, as indicated by  $ELPD_{loo}$  ( $\Delta ELPD_{loo} = -72.8$ ,  $SE = 26.5$ ).

Optimism can be tested by examining the  $\beta$  coefficient in this model, which codes for the difference in the use of evidence constituting “good news” versus “bad news.” Fitting the model reveals that participants displayed modest pessimism on average, with  $\beta$ ’s posterior expectation equal to  $-.10$  and a 95% credible interval ranging from  $-.14$  to  $-.06$ .

Figure 2 (right) presents a forest plot showing the posterior distribution of participants’  $\beta$  coefficients. As can be seen, there is substantial heterogeneity across individual participants. Consistent with the average estimates, a greater number (25 of 186) of participants’  $\beta$  coefficient estimates reliably indicate pessimism at the individual-level, and only a handful (3 of 186) reliably evinced optimism (by 95% credible interval).

Consistent with this heterogeneity, a number of studies have explored the potential for individual differences in optimistic belief updating. For instance, Kuzmanovic and colleagues (Kuzmanovic et al., 2015) report that “optimism” in belief revision was correlated with trait optimism as measured by the Life Orientation Test - Revised (LOT-R) (Scheier et al., 1994). In addition, Chowdhury and colleagues (2014) report finding greater optimism among older adults as compared with younger adults, over and above the effects of trait optimism.

However, neither age nor trait optimism appears to be cor-

related with “optimism” in belief updating as captured by the descriptive Bayesian model. Figure 4 shows the inferred optimism parameters (with 95% credible intervals) for all participants against their age and LOT-R scores. As is plain from the figure, neither are meaningfully correlated with optimism. One possibility is that prior findings were artifacts of the flawed measurement paradigm used in these studies. However, despite recruiting many more participants than these prior studies, the present study is somewhat lacking in its recruitment of older individuals. Thus further investigation may be warranted, including studies that recruit participants more uniformly across the age range and that utilize more experimental trials so as to better estimate individual-level optimism parameters.

## Discussion

This study tested optimism in belief revision through descriptive Bayesian modeling. Using Beta regression and a descriptive model inspired by rational analysis, I found evidence for substantial deviation from normative updating. In particular, most participants updated their beliefs less than would be expected from the direct application of Bayes’ rule. However, in contrast to widespread claims of “optimism” in belief updating, I instead found greater evidence for pessimism.

There was substantial individual heterogeneity in both the conservatism and pessimism effects that must qualify the population-average findings. A sizable subset of participants revised their beliefs in ways consistent with an entirely rational model of Bayesian belief updating. These participants did not exhibit any substantial conservatism, nor was their updating affected by the valence of the evidence provided. In addition, a small number of participants did exhibit optimism as could be reliably inferred from the model’s parameters.




Figure 4: Scatterplot and least-squares regression line showing the association between age (top) and LOT-R responses (bottom) and inferred  $\beta$  optimism parameters for individual participants. Error bars indicate 95% credible intervals of parameter estimates.

Motivational and Bayesian accounts of belief revision may be at odds, but there is no reason to expect these theories to be mutually exclusive. Instead, as these findings indicate, there is liable to be substantial individual and contextual variation: some people might be pessimists, perhaps fewer might be optimists, and many might be relatively even-keeled. Some people might reason rationally—only to show bias in another context.

Bearing this in mind, the development and application of Bayesian versus motivational accounts should be informed by the prevalence and relative magnitude of deviations from rationality in belief revision. On average, the findings here suggest conservatism is a more prevalent and stronger bias than pessimism. Consider, for example, public health officials trying to persuade people of the benefits of vaccines: these findings suggest they would be better served identifying and countering the sources of conservatism rather than pessimism. Moreover, any overall account of biases in belief revision should be tempered by the understanding that substantial numbers of people are apparently capable of and do engage in normatively rational belief revision.

Nevertheless, the impacts of affective valence on belief revision remains a worthy area of study with several interesting avenues for further research.

Pessimism presents some challenges to current motivational accounts of belief revision. According to such accounts, the pursuit of accurate beliefs is but one among many different motivations that might guide how people form and revise their beliefs. Perhaps the paradigmatic motivation is pleasure-seeking—to feel good about ourselves and the future. This is the sort of motivation thought to be served by optimistic belief updating biases. In addition, some have argued that optimistic biases might even serve some adaptive functions, such as supporting exploration by reducing anxiety and stress (Scheier et al., 1989; Sharot et al., 2011; Taylor et al., 2000). However, the present findings suggest that optimistic biases are quite rare and instead pessimistic biases are more prevalent.

Pessimism runs directly counter to the most immediate and paradigmatic pleasure-seeking motivations. An alternate motivational account of pessimism could be that lowering one's expectations guards against *future* disappointments, thereby serving as a regulatory coping mechanism (see e.g. Gul, 1991; Mellers et al., 1997; Powell & Horne, 2018). It could be further argued that a pessimistic outlook serves adaptive functions by encouraging defensive planning and preparation that helps people to weather hardships. If these proposals seem at all persuasive, this should make clear that current motivational accounts are underspecified, as they can apparently offer justifications for either optimistic or pessimistic biases.

One potential explanation is that optimistic and pessimistic motivations compete: that there is an immediate motivation for optimism but also a longer-term motivation for pessimism. Empirical support for this sort of motivational account might be found by quantifying the subjective value of present anticipation or dread over time against the anticipated dampening of future negative emotional experiences. Under a motivational account, people should be optimistic when the positive feelings associated with the formation of a desirable belief (and then experienced over time) outweigh the expected increase in negative feelings if a negative outcome should actually eventually occur. And people should be pessimistic if the negative feelings associated with the formation and maintenance of an undesirable belief are outweighed by the expected decrease in negative feelings if the negative outcome should occur.

Altogether, the present findings call for a re-evaluation of “optimism” in belief revision and of existing theories of motivated reasoning. It may be possible to develop a motivational account of the present findings of pessimism, but further nuance and investigation are needed to advance such an account.

## References

- Austerweil, J. L., & Griffiths, T. L. (2011). Seeking Confirmation Is Rational for Deterministic Hypotheses. *Cognitive Science*, 35(3), 499–526. <https://doi.org/10>

- .1111/j.1551-6709.2010.01161.x
- Bürkner, P.-C. (2017). Brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Chater, N., Zhu, J.-Q., Spicer, J., Sundh, J., León-Villagrà, P., & Sanborn, A. (2020). Probabilistic Biases Meet the Bayesian Brain. *Current Directions in Psychological Science*, 29(5), 506–512. <https://doi.org/10.1177/0963721420954801>
- Chowdhury, R., Sharot, T., Wolfe, T., Düzel, E., & Dolan, R. J. (2014). Optimistic update bias increases in older age. *Psychological Medicine*, 44(9), 2003–2012. <https://doi.org/10.1017/S0033291713002602>
- Corner, A., Harris, A. J. L., & Hahn, U. (2010). Conservatism in Belief Revision and Participant Skepticism. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, 6.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, 57(6), 1082–1090. <https://doi.org/10.1037/0022-3514.57.6.1082>
- Edwards, W. (1968). Conservatism in Human Information Processing. In B. Kleinmuntz (Ed.), *Formal representation of human judgment* (pp. 17–52). Wiley.
- Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101(3), 519–527. <https://doi.org/http://dx.doi.org/10.1037/0033-295X.101.3.519>
- Fischhoff, B., & Beyth-Marom, R. (1983). Hypothesis evaluation from a Bayesian perspective. *Psychological Review*, 90(3), 239–260. <https://doi.org/10.1037/0033-295X.90.3.239>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (Third edition). CRC Press.
- Gul, F. (1991). A Theory of Disappointment Aversion. *Econometrica*, 59(3), 667. <https://doi.org/10.2307/2938223>
- Jeffrey, R. C. (1983). *The logic of decision* (2. ed). University of Chicago Press.
- Jern, A., Chang, K. K., & Kemp, C. (2014). Belief polarization is not always irrational. *Psychological Review*, 121(2), 206–224. <https://doi.org/10.1037/a0035941>
- Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, 8(4), 19.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237–251. <https://doi.org/10.1037/h0034747>
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480–498. <https://doi.org/10.1037/0033-2909.108.3.480>
- 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. <https://doi.org/10.1002/bdm.1849>
- Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision Affect Theory: Emotional Reactions to the Outcomes of Risky Options. *Psychological Science*, 8(6), 423–429.
- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, 118(1), 120–134. <https://doi.org/http://dx.doi.org/10.1037/a0021110>
- Powell, D., & Horne, Z. (2018). Expectations bias judgments of harm against others. *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, 6.
- Sanitioso, R., Kunda, Z., & Fong, G. T. (1990). Motivated recruitment of autobiographical memories. *Journal of Personality and Social Psychology*, 59(2), 229–241. <https://doi.org/10.1037/0022-3514.59.2.229>
- Scheier, M. F., Carver, C. S., & Bridges, M. W. (1994). Distinguishing optimism from neuroticism (and trait anxiety, self-mastery, and self-esteem): A reevaluation of the Life Orientation Test. *Journal of Personality and Social Psychology*, 67(6), 1063–1078. <https://doi.org/10.1037//0022-3514.67.6.1063>
- Scheier, M. F., Matthews, K. A., Owens, J. F., Magovern, G. J., Lefebvre, R. C., Abbott, R. A., & Carver, C. S. (1989). Dispositional optimism and recovery from coronary artery bypass surgery: The beneficial effects on physical and psychological well-being. *Journal of Personality and Social Psychology*, 57(6), 1024–1040. <https://doi.org/10.1037//0022-3514.57.6.1024>
- Shah, P., Harris, A. J. L., Bird, G., Catmur, C., & Hahn, U. (2016). A pessimistic view of optimistic belief updating. *Cognitive Psychology*, 90, 71–127. <https://doi.org/10.1016/j.cogpsych.2016.05.004>
- Sharot, T. (2011). The optimism bias. *Current Biology*, 21(23), R941–R945. <https://doi.org/10.1016/j.cub.2011.10.030>
- 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. <https://doi.org/10.1038/nn.2949>
- Slovic, A., & Lichtenstein, S. (1971). *Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment*. 96.
- Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological Review*, 124(4), 410–441. <https://doi.org/10.1037/rev0000052>
- Taylor, S. E., Kemeny, M. E., Reed, G. M., Bower, J. E., & Gruenewald, T. L. (2000). Psychological resources,

- positive illusions, and health. *The American Psychologist*, 55(1), 99–109. <https://doi.org/10.1037//0003-066x.55.1.99>
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1413–1432. <https://doi.org/10.1007/s11222-016-9696-4>
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 806–820.