Psychologists Update their Beliefs About Effect Sizes After Replication Studies

Alex D. McDiarmid^{1*}, Alexa M. Tullett^{1*}, Cassie M. Whitt¹, Simine Vazire², Paul E. Smaldino³,
Jeremy E. Stephens¹

¹Department of Psychology, University of Alabama, Tuscaloosa, AL, U.S.A

²Melbourne School of Psychological Sciences, University of Melbourne, Parkville, Vic,
Australia

³Department of Cognitive and Information Sciences, University of California, Merced, Merced,
CA, U.S.A

*Corresponding authors' email addresses: admediarmid@crimson.ua.edu, alexa.tullett@ua.edu

Abstract

Self-correction—a key feature distinguishing science from pseudoscience—requires that scientists update their beliefs in light of new evidence. However, people are often reluctant to change their beliefs. We examined belief updating in action, tracking research psychologists' beliefs in psychological effects before and after the completion of four large-scale replication projects. We found that psychologists did update their beliefs; they updated as much as they predicted they would, but not as much as our Bayesian model suggests they should, if they trust the results. We found no evidence that psychologists became more critical of replications when it would have preserved their pre-existing beliefs. We also found no evidence that personal investment or lack of expertise discouraged belief updating, but people higher on intellectual humility updated their beliefs slightly more. Overall, our results suggest that replication studies can contribute to self-correction within psychology, but psychologists may underweight their evidentiary value.

Scientific progress depends on self-correction—it rests on the assumption that supportive evidence will bolster people’s belief, whereas contradictory evidence will undermine it. This cumulative process, where theories can be supported and undermined, is a key feature distinguishing science from pseudoscience. An extensive literature on motivated reasoning and confirmation bias, however, shows that people do not always update their beliefs in a rational way, especially in the face of disconfirming evidence (Kunda, 1990; Nickerson, 1998). Considering that scientists are people and thus subject to these biases, this raises questions about the extent to which belief updating actually happens in practice.

Here, we conducted a direct test of belief updating by tracking psychologists’ beliefs in psychological effects before and after the release of evidence from rigorous, large-scale replication projects: two Registered Replication Reports (RRRs; McCarthy et al., 2018; Verschuere et al., 2018) in *Perspectives on Psychological Science*, and the Many Labs 2 (ML2; Klein et al., 2018) project. By doing so, we aimed to address three broad research questions (RQs):

RQ1: How much do psychologists update their beliefs in response to new empirical evidence?
RQ2: Do psychologists show evidence of motivated reasoning when evaluating replications?
RQ3: What predicts the extent of psychologists’ belief updating?

Investigating these questions within psychological science is important for at least two reasons. First, these inquiries are important for assessing the value of large-scale replication projects (e.g., Camerer, 2016; Camerer, 2018; Klein et al., 2018; Ebersole et al., 2016; Open Science Collaboration, 2012). If psychologists update their beliefs in light of replication evidence, this would provide evidence that replication projects can function as a mechanism of self-correction. Second, this investigation could identify shortcomings in belief updating, as well as illuminate why those shortcomings occur, and consequently, how they might be addressed. This information is necessary for ensuring that the valuable evidence provided by replication studies is translated into cumulative scientific knowledge.

Publication bias and file drawer effects limit the extent to which science can be fully self-correcting by skewing the evidence toward positive results (Ioannidis, 2012; Fanelli, 2012; McElreath & Smaldino, 2015; Nissen, Magidson, Gross, & Bergstrom, 2016). For this reason, recent critiques challenge the efficacy with which psychological science accumulates knowledge (Scheel, Schijen, & Lakens, in press; van Elk, Matzke, Gronau, Guang, Vanderkerckhove, & Wagenmakers, 2015). Bias in the publication system poses a serious boundary to self-correction; this is particularly concerning if bias also operates at the individual level.

A substantial body of psychological research suggests that self-correction should not be taken for granted. When evaluating new evidence, people often exhibit a confirmation bias, or the tendency to favor information that bolsters (rather than undermines) their pre-existing beliefs (Anderson, Lepper, & Ross, 1980; Hayden & Mischel, 1976; Kunda, 1990; Luchins, 1957; Nickerson, 1998; Rhine & Severance, 1970; Ross, Lepper, & Hubbard, 1975, c.f., Koehler, 1993; Mahoney, 1977; Hergovich, Schott, & Burger, 2010). This phenomenon is distinct from motivated reasoning that is driven by accuracy but overlaps substantially with motivated

reasoning that is driven by the desire to reach a particular conclusion; in this case, the original belief (Hart et al., 2009; Kunda, 1990).

One manifestation of confirmation bias and motivated reasoning involves selective criticism of contradictory information (Jacks & Cameron, 2003; Koehler, 1993; Lord et al., 1979; Munro, 2010). For instance, past work has shown that both laypeople and scientists evaluate the quality of research more favorably when results support their pre-existing views (Koehler, 1993; Lord et al., 1979; Munro, et al., 2004; Munro & Ditto, 1997). Given these findings, we examined whether people questioned the quality of replication studies selectively, when doing so would preserve their pre-existing beliefs.

In a similar fashion, people may also exaggerate context differences between replication studies and original studies (e.g., hidden moderator effects; Barrett, 2015; Van Bavel et al., 2016) when the replication challenges their original view. Although concerns about context differences are often justified, they also provide an opportunity for the selective evaluation of evidence. To test this second mechanism, we examined if people were more likely to identify context differences when this would protect their prior beliefs.

One important question when using replication studies to examine belief updating is: Are psychologists sometimes justified in ignoring, or downplaying, replication results? A long tradition in the philosophy of science has documented the importance of considering auxiliary hypotheses: conditions that must be met for new evidence to directly speak to a previously formed belief (Duhem, 1906; Gershman, 2019; Quine, 1951). For example, if a replication study claimed to test an original effect but used a sample of participants who could not understand the materials, a key auxiliary hypothesis (i.e., that participants understand the materials) would be unsupported and the replication evidence should be devalued. In the present project, when we specify how much participants “should” update their beliefs, we make the controversial assumption that all of the auxiliary hypotheses are true. This strong assumption is unlikely to ever be met completely, but we believe that by selecting exceptionally rigorous replication projects, high confidence in the auxiliary hypotheses is warranted. In addition, by measuring participants’ ratings of replication quality and context effects, we can directly assess these two possible auxiliary hypotheses.

Although confirmation bias appears ubiquitous, previous work has identified some factors that influence its strength. First, several studies indicate that people are more likely to demonstrate confirmation bias when they are personally invested in a theory than when they do not have a personal stake (Apsler & Sears, 1968; Hart et al., 2009; Lord, 1979; Miller et al., 1993; Nickerson, 1998; Panagiotou & Ioannidis, 2012; cf. Anderson, et al., 1980; MacCoun, 1998; Murdock, 1962; Nickerson, 1998; Nisbett & Ross, 1980). To test this, we assessed whether personal investment predicted less belief updating when the replication results cast doubt on the original finding.

Intuitively, it seems possible that people who are informed about biases in their thinking—and motivated to correct them, as scientists ideally are—might be able to combat such biases. Some evidence has demonstrated that teaching people about bias can cause them to acknowledge at least some subjectivity in their own thinking (Pronin & Kugler, 2007). Admitting bias and

changing behavior are two different things, however, and there is also evidence that people repeatedly fail to correct for biases in their decisions (Hart et al., 2009; Kahneman, 2011; Koehler, 1993; Mahoney, 1977; Hergovich et al., 2010). To test this, we examined whether expertise regarding motivated reasoning and confirmation bias predicted belief updating.

Finally, some people might be dispositionally more open to new information. Intellectual humility refers to the willingness to recognize that one's beliefs might be wrong (Krumrei-Mancuso, & Rouse, 2016; Leary et al., 2017; McElroy, Rice, & Davis, 2014). We reasoned that people who are high in intellectual humility should be more open to considering evidence that contradicts their previously held views, and thus, more likely to update their beliefs. To test this, we assessed whether intellectual humility predicted belief updating.

In the current project, we assessed psychologists' beliefs in psychological findings before (Phase I) and after (Phase II) the release of evidence from large-scale replication studies. In Phase I, all participants read about the results of three original studies, estimated the probability that the true effects were "non-trivial" (explicitly defined as a Cohen's d greater than .1), and indicated their confidence and personal investment in the effects. Participants were randomly assigned to either a prediction condition or a control condition. Participants in the prediction condition were given a brief description (excluding results, which were not yet publicly available) of the ongoing replication studies. Based on this information, they were asked to predict how they would adjust their beliefs in response to a range of potential outcomes of these replication studies. This allowed us to assess belief updating while minimizing the influence of motivated reasoning (i.e., before participants were faced with actual results). These participants also rated the quality of the replication studies and the extent of contextual differences between the replication and original study. Participants in the control condition did not learn about the replication studies during Phase I and did not make any predictions. This allowed us to assess whether merely making predictions affects subsequent belief updating, and if necessary, it enabled us to statistically control for such effects in our analyses.

In Phase II, all participants read about the results of the replication studies corresponding to the studies they viewed in Phase I. Then, they provided a new estimate of the probability that the true effects were non-trivial and indicated their confidence in the effects. Starting from participants' original beliefs, we calculated the degree of belief adjustment that would be mathematically justifiable according to a Bayesian model (assuming all auxiliary hypotheses are true). This time, all participants rated the quality of the replication studies and the extent of contextual differences from the original studies. Finally, we assessed participants' expertise regarding confirmation bias and motivated reasoning, and their intellectual humility (see Figure 1 for an overview of study procedures).

Results

All data are available on the Open Science Framework (https://osf.io/jtp4b/?view_only=63d9a7deb66e4764821b048c6f837185). All analyses were conducted using multilevel modeling and were preregistered unless otherwise noted. For all analyses we assessed assumptions of homogeneity of variance, normality of residuals, and (except for analyses with only a single binary predictor) linearity. The largest violations of

assumptions were minor enough that our MLM models were appropriate for the data without transformations. With the exception of Wald z-tests (reported only in Supplementary Methods – Random Effects), all p-values are based on two-tailed tests and. For results of additional preregister analyses, see Supplementary Data – 1-8 — the numbers are a reference to the hypothesis the additional analyses are relevant for. All reported means are those observed in our sample rather than population mean estimates from the MLM analyses. Research questions and hypotheses are summarized in Table 1.

Terminology

Actual prior: A participant's estimate of the probability that a true effect is non-trivial in size, provided before viewing the replication evidence (Phase I).

Predicted posterior: A participants' estimate of the probability than a true effect is non-trivial in size, forecast based on hypothetical replication outcomes (Phase I, prediction condition only).

Actual posterior: A participants' estimate of the probability that a true effect is non-trivial in size, provided after viewing the replication evidence (Phase II).

Bayesian posterior: A calculated probability that considers how a perfectly rational Bayesian agent would update their beliefs, incorporating the participant's *actual prior* and the strength of the replication evidence (see Figure 2 for visual example of calculating a *Bayesian posterior*).

Details on how these values are measured and/or calculated can be found in the Methods section.

Hypothesis 1. The results were consistent with Hypothesis 1. Participants adjusted their estimated probability that the effect was real (i.e., $d > .1$) in the direction indicated by our Bayesian model. When the replication results indicated that participants should adjust downward, participants' *actual posteriors* ($M = 28.06$, $SD = 24.73$) were lower than their *actual priors* ($M = 63.16$, $SD = 24.37$), $b = -35.13$, $\beta = -.582$, 95% CI $[-.609, -.555]$, $t(1851.8) = -42.77$, $p < .001$. When the replication results indicated that participants should adjust upward, participants' *actual posteriors* ($M = 79.98$, $SD = 21.15$) were higher than their *actual priors* ($M = 68.02$, $SD = 23.45$), $b = 11.98$, $\beta = .259$, 95% CI $[.224, .293]$, $t(1123.41) = 14.66$, $p < .001$, (Figure 3).

A separate analysis using Likert-scale-based confidence ratings (see Evaluation of Studies - Control Condition section) as the dependent variable, suggests a substantially smaller effect size, $b = -.400$, 95% CI $[-.429, -.364]$, $\beta = -.352$, $t(1832.8) = -23.64$, $p < .001$, when participants should adjust downward.

As unregistered analyses, we also tested professional position (e.g., graduate student or Assistant Professor) as a moderator of probability-based belief updating—separately for cases in which participants should have adjusted downward versus upward. For most categorical comparisons of professional position as a moderator, we did not find evidence that the magnitude of belief updating depended on professional position. The one exception was that, when participants should have adjusted upward, we found a significant interaction such that graduate students and postdoctoral scholars updated their beliefs less than lecturers and faculty as a combined group, $b = 6.22$, $\beta = .255$, $p = .003$ (see Supplementary Data 9 – Unregistered Moderation by Academic Position for details).

Hypothesis 2. The results were consistent with Hypothesis 2. Participants did not update their estimates as much as our Bayesian model would indicate. When the replication results indicated that participants should adjust downward, participants' *actual posteriors* ($M = 28.06$, $SD = 24.73$), were not as low as their *Bayesian posteriors* ($M = 5.16$, $SD = 17.18$), $b = 23.80$, $\beta = .482$, 95% CI [.455, .509], $t(1676.6) = 34.40$, $p < .001$. When the replication results indicated that participants should adjust upward, participants' *actual posteriors* ($M = 79.98$, $SD = 21.15$) were not as high as their *Bayesian posteriors* ($M = 97.46$, $SD = 8.04$), $b = -15.84$, $\beta = -.424$, 95% CI [-.459, -.389], $t(1041.7) = -23.70$, $p < .001$.

In an preregistered analysis, we also assessed this effect for the subset of cases when participants gave the most positive possible evaluations (i.e., provided ratings at the most extreme point on our Likert scale) of the replication study methodology during Phase II ("high quality" replication studies, and "minor contextual differences"). We conducted this analysis to examine whether the participants, for whom replication results should have most strongly influenced belief updating—because their responses indicate the most confidence that these critical auxiliary hypotheses held—updated their beliefs as strongly as they should, according to our model. Even when we limit our analyses to these 642 [629] cases in which participants should adjust downward [upward], the results still support Hypothesis 2; participants do not update as much as our Bayesian model states they should. When the replication results indicated that participants should adjust downward, participants' *actual posteriors* ($M = 16.72$, $SD = 20.67$), were still not as low as their *Bayesian posteriors* ($M = 3.96$, $SD = 15.88$), $b = 13.64$, $\beta = .340$, 95% CI [.289, .391], $t(389.9) = 13.21$, $p < .001$. When participants should have adjusted upward, their *actual posteriors* ($M = 83.16$, $SD = 19.57$) were still not as high as their *Bayesian posteriors* ($M = 98.06$, $SD = 7.82$), $b = -13.01$, $\beta = -.377$, 95% CI [-.432, -.323], $t(613.0) = -13.64$, $p < .001$.

Hypothesis 3 (prediction condition only). Hypothesis 3 was not supported. Contrary to our hypothesis, participants updated *more* than they predicted they would. We tested this hypothesis in the prediction condition only, as these were the only participants that provided information to calculate their *predicted posteriors*. When the replication results indicated that participants should adjust downward, participants' *actual posteriors* ($M = 28.06$, $SD = 24.73$) were even *lower* than their *predicted posteriors* ($M = 30.84$, $SD = 22.26$), $b = -3.02$, $\beta = -.060$, 95% CI [-.096, -.024], $t(1239.4) = -3.24$, $p = .001$. When the replication results indicated that participants should adjust upward, participants' *actual posteriors* ($M = 79.98$, $SD = 21.15$) were even *higher* than their *predicted posteriors* ($M = 70.01$, $SD = 24.47$), $b = 10.69$, $\beta = .223$, 95% CI [.184, .263], $t(898.8) = 11.04$, $p < .001$.

Hypothesis 4 (prediction condition only). The results largely failed to support our hypothesis. We observed little evidence that higher *actual priors* predicted a decrease in ratings of replication quality from Phase I to Phase II. We tested this hypothesis in the prediction condition only, as control participants did not provide ratings of replication quality in Phase I. To quantify changes in ratings of replication quality, we computed a difference score by subtracting participants' Phase I ratings from their Phase II ratings.

We ran an MLM model predicting changes in replication quality ratings from Phase I to Phase II from participants' *actual priors*. When the replication result indicated that participants should adjust downward, higher priors predicted a smaller increase in perceived replication $b = -.003$, quality, $\beta = -.094$, 95% CI $[-.181, -.006]$, $t(586.7) = -2.10$, $p = .036$. That is, for participants who should have adjusted downward, those who expressed stronger prior beliefs in an effect were somewhat less likely to increase their ratings of the perceived quality of the replication in Phase II. This result provided mild support for our hypothesis. However, in contrast, a separate non-MLM analysis (averaging across participants' responses to individual studies), failed to support our hypothesis, $\beta = -0.051$, $p = .308$.

Because the previous MLM analysis obscures whether higher *actual priors* are associated with higher Phase I ratings or lower Phase II ratings, we ran an preregistered analysis assessing whether *actual priors* predicted Phase II ratings of replication quality controlling for Phase I ratings. Our MLM model included Phase I replication quality, *actual priors* and the interaction of these terms as fixed variables with Phase II replication quality as the outcome variable. The simple effect of *actual prior* was not significant $b = -.002$, $\beta = -.068$, 95% CI $[-.146, .010]$, $t(598.0) = -1.70$, $p = .089$, nor was the interaction ($\beta = .021$, $p = .195$). Thus, when psychologists had similar Phase I ratings of replications quality, we did not observe evidence that their Phase II ratings significantly differed depending on their *actual priors*.

We realized, post-hoc, that ratings of personal investment might provide a better test of motivated reasoning than participants' *actual priors*. For this reason, we reran the previous analysis replacing *actual priors* with personal investment. The interaction between personal investment and Phase I ratings of replication quality was significant, $b = .127$, $\beta = .070$, 95% CI $[.002, .137]$, $t(538.9) = 2.03$, $p = .043$. The simple effects of personal investment on Phase II ratings of replication quality were not significant when Phase I ratings were 1 SD above the mean $b = .040$, $\beta = .037$, 95% CI $[-.076, .149]$, $t(597) = .64$, $p = .523$, or at the mean $b = -.037$, $\beta = -.033$, 95% CI $[-.111, .045]$, $t(597) = -.83$, $p = .405$. However, the simple effect of personal investment on Phase II ratings of replication quality was significant, but small, when Phase I ratings were 1 SD below the mean, $b = -.114$, $\beta = -.103$, 95% CI $[-.196, -.010]$, $t(535) = -2.17$, $p = .031$. Thus, personal investment predicted Phase II ratings of replication quality, but only when participants had relatively low Phase I ratings. Given that this analysis was not pre-registered and the 95% CI for the simple effect at 1SD below the mean – the result most relevant to our hypothesis about motivated reasoning – comes very close to including zero, we believe this result is not especially informative one way or the other regarding the effect of personal investment on belief updating.

Despite the initial MLM analysis' weak support for our hypothesis, overall, these results do not provide more than trivial support (at best) for our hypothesis. It is also noteworthy that participants rarely expressed strong personal investment in these effects (with 17.7% of ratings expressing a slight preference and 2.0% of ratings expressing a strong preference) and therefore our sample likely does not provide ideal conditions for assessing motivated reasoning.

Hypotheses 4 and 5 tested if *actual priors* predicted less favorable study evaluations (*quality* of replication and *context differences* between the replication and original study). Conceptually related (preregistered) mediation analysis tested if the (potential) negative association between

actual priors and belief updating was mediated by more negative study evaluations during Phase II. These tests provided some evidence of strikingly weak, but significant, effects in the predicted direction (see Supplementary Data 4 and 5 for detailed mediation analysis results)

Hypothesis 5 (prediction condition only). The results did not support Hypothesis 5. We did not find evidence that higher priors predicted changes in perceptions of context differences. We tested this hypothesis in the prediction condition only, as control participants did not provide ratings of context differences in Phase I. To quantify changes in perceptions of context differences, we computed a difference score by subtracting participants' Phase I ratings of context differences from their Phase II ratings of context differences. We ran an MLM model predicting the difference score from participants' *actual priors*. When the replication result indicated that participants should adjust downward, we did not observe evidence that *actual priors* are associated with changes in perceived contextual differences, $\beta = -.044$, $p = .302$.

As with Hypothesis 4, we realized that ratings of personal investment might provide a better test of motivated reasoning than participants' *actual priors*. For this reason, and to account for differences in initial ratings of contextual differences, we ran an MLM model including Phase I ratings of contextual differences, personal investment, and the interaction of these terms as fixed variables with Phase II contextual difference as the outcome variable. The simple effect of *personal investment* for participants with the average Phase I rating of contextual differences was not significant, $b = -.009$, $\beta = -.008$, 95% CI [-.084, .069], $t(595) = -.199$, $p = .843$, nor was the interaction ($\beta = -.026$, $p = .499$). Thus, when psychologists had similar Phase I ratings of contextual differences, we did not observe evidence that their Phase II ratings significantly differed depending on their *personal investment*.

Hypothesis 6. Results did not support Hypothesis 6. When replication results suggested participants should adjust their beliefs downward, people who were more personally invested in a finding adjusted *more*, not less. We ran an MLM model predicting adjustment (*actual posterior* – *actual prior*) from personal investment. When the replication result indicated that participants should adjust downward, more personal investment predicted *greater* downward adjustment, $b = -8.42$, $\beta = -.155$, 95% CI [-.207, -.103], $t(1,202.1) = -5.85$, $p < .001$. Contrary to our hypothesis, this finding suggests that psychologists who are more personally invested in a finding show a greater decrease from their priors to posteriors when replication results warrant such a change.

The previous analysis obscures whether higher personal investment is associated with higher *actual priors* or lower *actual posteriors* (the latter being more relevant to our hypothesis). Supplemental unregistered analyses revealed that higher personal investment is predictive of higher priors and predictive of slightly *higher* posteriors (see Supplementary Data 6). To more explicitly dissect the relationship between personal investment and belief updating, we ran an unregistered analysis assessing whether personal investment predicted *actual posteriors* when controlling for *actual priors*. Our MLM model included *actual prior* and personal investment, and the interaction of these terms as fixed variables with *actual posterior* as the outcome variable. The simple effect of personal investment, for the average *actual prior*, was significant, $b = 2.20$, $\beta = .052$, 95% CI [.002, .102], $t(1,104.1) = 2.02$, $p = .043$, while the interaction of personal investment and *actual prior* was not significant, $\beta = -.035$, 95% CI [-.082, .011], $p = .138$. When participants should adjust downward, these results suggest that for psychologists

with equal priors, one should predict slightly higher *actual posteriors* for psychologists with more personal investment in an effect. This result suggests a caveat to the original finding that participants with more personal investment show greater belief updating; participants who are equally credulous before evaluating replication evidence, show slightly *less* belief updating when they are more personally invested in the effect.

While our hypothesis only applied to cases in which participants should have adjusted their beliefs downward, we conducted an unpreregistered analysis, for cases in which participants should have adjusted upwards, to assess if personal investment predicted *actual posteriors* when controlling for *actual priors*. The simple effect of personal investment, for the average *actual prior*, was not significant, $\beta = .017$, $p = .568$ (see Supplementary Data 6 for details).

Hypothesis 7. Results did not support Hypothesis 7. We did not observe evidence that expertise with bias is related with the degree to which participants updated their beliefs. Our outcome variable was computed as *actual posterior* – *actual prior*, in cases when participants should have adjusted upward, and as *actual prior* – *actual posterior* in cases when participants should have adjusted downward. We ran an MLM model predicting the extent to which participants adjusted in the correct direction from expertise. The association was not significant, $\beta = .009$, 95% CI [- .034, .052], $p = .680$, indicating that expertise regarding bias was not predictive of belief updating. While the 95% confidence intervals of resamples would vary substantially, it is perhaps worth noting that not only does the observed confidence interval contain zero, but it only contains effect sizes we would consider smaller than a practically meaningful effect. .

Hypothesis 8. The results supported Hypothesis 8. Participants with greater intellectual humility showed slightly greater belief updating in the correct direction. Our outcome variable was computed as *actual posterior* – *actual prior*, in cases when participants should have adjusted upward, and as *actual prior* – *actual posterior* in cases when participants should have adjusted downward. We ran an MLM model with intellectual humility as the fixed variable predicting the extent to which participants adjusted in the correct direction. As anticipated, greater intellectual humility predicted greater belief updating in the correct direction, $b = 5.49$, $\beta = .086$, 95% CI [.043, .128], $t(978.8) = 3.93$, $p < .001$.

We followed-up on this finding with an unpreregistered analysis assessing whether intellectual humility predicted *actual posteriors* when participants started with similar priors (i.e., controlling for the effect of *actual prior* on *actual posterior*). Our MLM model included *actual prior*, intellectual humility, and the interaction of these terms as fixed variables with *actual posterior* as the outcome variable. When the replication result indicated that participants should adjust downward, the simple effect of intellectual humility was significant, $b = -6.61$, $\beta = -.121$, 95% CI [- .211, -.050], $t(411.8) = -3.19$, $p = .002$. When the replication result indicated that participants should adjust upward, the simple effect of intellectual humility was also significant, $b = 3.59$, $\beta = .085$, 95% CI [.013, .157], $t(306.1) = 2.33$, $p = .021$. Given the lack of evidence for interactions between *actual prior* and intellectual humility (regardless of the normative adjustment direction; $\beta = .034$, $p = .359$; $\beta = .019$, $p = .580$), these simple effects suggest that for psychologists with equal priors, one should predict slightly greater belief updating from psychologists with greater intellectual humility.

Discussion

The current study tracked psychologists' beliefs in psychological effects before and after they learned the results of large-scale replication studies. Our results suggest three broad conclusions. First, consistent with our predictions, psychologists did update their beliefs in light of replication evidence, but not to the extent that our Bayesian model suggested they should. Second, contrary to our predictions, we found little evidence for motivated reasoning in psychologists' evaluations of negative (or smaller than expected) results of replication studies. Third, belief updating is associated with slightly greater intellectual humility, but, contrary to our predictions, it is not associated with less personal investment and more expertise with bias.

With respect to the extent of belief updating, participants did indeed shift their beliefs in the direction of the replication evidence. We found that the extent of this shift was comparable to or, when participants should adjust their confidence upward, greater than that predicted by participants before knowing the results (i.e., during Phase I). This degree of updating is striking when compared to investigations of citation patterns, which show only modest adjustment in response to failed replications (Hardwicke et al., 2021). However, participants did not shift their beliefs as much as our Bayesian model suggested they should. This was true even when participants gave the replication the highest possible ratings on quality and similarity to the original study. These findings seem to suggest that participants update as much as they think they should, but—even when they have high confidence in two key auxiliary hypotheses—underweight the evidence relative to our Bayesian model.

We did not conduct an in-depth examination of the mechanisms underlying these effects, but we did test whether psychologists demonstrated motivated reasoning in their evaluation of replication studies by becoming more critical of replication results that challenge their prior. We found no evidence that participants shifted their ratings of replication quality and context differences in directions that would preserve their original beliefs. Overall, participants seemed to separate their evaluation of the replication from their prior beliefs. It is worth noting, however, that our participants knew their responses would not be publicly linked with their names, and so our analysis is best seen as a test of private beliefs. Incentives for privately held beliefs and public declarations of those beliefs are not always consistent (Kuran, 1997), so it is possible that *public* declarations of confidence in scientific results might still be influenced by motivated reasoning.

Finally, we tested three different individual differences that we thought may predict belief updating: personal investment, expertise with bias, and intellectual humility. Our assessment of personal investment served dual purposes in that a negative relationship with belief updating could be considered an indicator of motivated reasoning. Although we did not find strong evidence of this relationship, this doesn't necessarily mean that participants were completely unaffected by their motivations; our findings suggest that people were not particularly motivated in the first place – meaning that a key precondition for this analysis was missing. In other words, it is possible that we would have observed stronger evidence of motivated reasoning if our participants had a higher personal stake in these findings. Addressing our two remaining hypotheses, we did not observe that people with more expertise with bias updated their beliefs more, but we did observe that people who self-reported greater intellectual humility updated their

beliefs slightly more than those who claimed to be less intellectually humble. These findings call to mind work showing that people who are adept at incorporating new information – those who are willing to “treat one’s beliefs as testable hypotheses, not sacred possessions” (p. 482) – tend to be better at forecasting outcomes of politically-relevant events (Tetlock, Mellers, & Scoblic, 2017).

With respect to documenting belief updating within the field of psychology, our current sample seems quite strong. Our sample of $N = 1,096$ is relatively large and includes participants across a range of levels of training (see Participants section). Thus, there is reason to expect that our results would generalize to researchers in psychology. It remains unclear, however, how these estimates would extend to other fields or to the general public. In addition, the duration of the observed effects is unknown, and indeed, some evidence suggests that induced changes to beliefs fade over time (Jost, 1897).

One limitation of the present study is the possibility of experimenter demand effects or observer effects. Perhaps, for instance, participants exaggerated their willingness to adjust their beliefs or minimized their personal investment in original effects to look more rational or scientific (demand effects). Another possibility is that psychologists’ updating was augmented by aspects of our study, such as being asked to consider the replication study results and to report on one’s beliefs (observer effect). Either of these effects would inflate our estimate of belief updating; in the real world, psychologists might update their beliefs less than they did in our study. Ideally, these findings should be supplemented with evidence that is uncolored by demand and observer effects (e.g., behavioral indicators of belief updating that can be examined unobtrusively). Observer effects may also have impacted the results for intellectual humility. Because participants self-reported their intellectual humility immediately after they had the opportunity to update their beliefs in light of new evidence, it is possible that their responses to the belief updating questions influenced their self-perceptions (or self-presentation) in the domain of intellectual humility.

We used a Bayesian model to provide one benchmark against which to evaluate belief updating. Our model assumes that participants had complete faith in the auxiliary hypotheses—that is, that they treated the replication study as a pure test of the original effect. We attempted to address this limitation by conducting an analysis on only the subset of cases in which participants gave the maximum rating for replication quality and the minimum rating for contextual differences. For this subset, participants still did not update as much as our Bayesian model would indicate, although they were closer than the full sample. This analysis demonstrates that the underadjustment we observed in our sample probably cannot be accounted for by participants blaming these two auxiliary hypotheses. Our model also assumed that participants’ prior distributions were normally distributed, but this was likely not always accurate. More broadly, our results do not necessarily imply that participants are being “irrational;” it is possible that participants are not prioritizing knowledge goals at the exclusion of all else and that the underweighting of evidence we observed reflects rational behavior given participants’ motivational states (Longino, 1990).

Many of the conclusions drawn in the current paper are based on probability estimates provided by psychological scientists. Ultimately, however, the *behavior* of psychologists is more relevant

to the real-world process of self-correction than are their stated beliefs. We chose to examine probability estimates because they can be compared to a standard of accuracy (i.e., the Bayesian model), but this information would be of questionable value if these estimates were unrelated to behavior. An unplanned analysis showed that the relationship between participants' probability belief updating and the difference between Phase I and Phase II Likert scale-based confidence (e.g., are they confident enough in the effect to teach it to undergraduates) was $r = .50$ [.47, .53], suggesting a link between probability estimates and behavioral intentions. Nevertheless, behavioral evidence is necessary to provide a true test of this question.

For a scientific field to be self-correcting at least three conditions are necessary: (1) empirical results must be subject to replication, and members of the field must (2) evaluate those replication efforts and (3) adjust their beliefs in response to new evidence. Here, we tested this third step and observed that psychologists do update their beliefs, but they may underweight replication evidence—even when they did not report criticisms of the methodology of the replication studies. This suggests that as things currently stand, replications can function as intended, leading psychologists to adjust their beliefs in the direction of the new evidence. However, our results also reveal a discrepancy between participants' updating and that of our Bayesian model. This could mean that participants have valid reasons for deviating from Bayesian updating. It could also mean that participants are giving excessive weight to their original beliefs, even when replication evidence - deemed high quality by the participants - exists.

Arguably, science gains strength through consensus (Oreskes, 2019). Thus, it could be illuminating for future work to examine the ways in which individual belief change does, or does not, aggregate to shift the consensus in a field. Such an investigation could be particularly timely now, as the landscape of scientific discourse is rapidly being shaped by the increasing popularity of blogs, podcasts, and Twitter (Carlson & Harris, 2020).

Our study design placed replication results front and center in the attention of interested researchers. However, a persistent problem with replications is that they are rarely published, and, when published, are rarely linked to the original results being replicated (Ioannidis, 2012; Pashler & Harris, 2012; cf. scite, n.d.). This severely diminishes the potency of their impact. Our results show that replication can work when researchers are made to see their results. Communicating the results of replication studies should be more of a priority.

Methods

This research was sanctioned by the University of Alabama Institutional Review Board and the data were collected in a manner consistent with ethical standards for the treatment of human subjects. Participants provided informed consent prior to beginning the study. The preregistration (see Preregistration section for details) for this study is available here (https://osf.io/wz3ty/?view_only=16834083b7f441059be85c85db1c04b9). This Open Science Framework page has links to the original preregistration and three preregistration addendums; the webpage describes how much data had been collected and viewed at the time of each addendum.

Design.

The design was a 2 (phase: I vs. II) by 2 (condition: prediction vs. control) mixed design with phase as a within-subjects factor and condition as a between-subjects factor.

Stimuli: Original and Replication Studies.

All materials can be viewed here

(https://osf.io/jtp4b/?view_only=63d9a7deb66e4764821b048c6f837185)

The original studies used as stimuli for the present investigation were selected because we were aware that they were included in large-scale replication projects that were underway, but not complete, at the time we began Phase I. During Phase I, each participant saw three studies from a pool of 35: 23 studies from Many Labs 2 (Klein et al., 2018; see Supplementary Methods – Stimulus Exclusions for a discussion of studies from ML2 excluded from this research), nine studies from Many Labs 5 (Ebersole et al., 2019), and three studies that were the focus of RRRs (McCarthy et al., 2018; Verschuere et al., 2018).

The nine Many Labs 5 studies and one of the RRRs that were included in Phase I were excluded from Phase II because of uncertainty about their completion date. This left us with a total of 25 studies included in both phases. Participants saw the same studies in Phase II that they saw in Phase I. Thus, because we dropped some studies from Phase II, some participants saw fewer studies in Phase II than they did in Phase I.

Participants.

Participants were research psychologists—individuals with a graduate degree in psychology or completed at least some graduate level training in psychology—recruited from professional conferences and listservs. Participation in Phase I was completed between November, 2017 and March, 2018, and participation in Phase II was completed between February, 2019 and April, 2019. Participants from Phase I were re-contacted via email and were allowed at least three weeks to complete Phase II. Participants were compensated \$10 for completing Phase I and \$20 for completing Phase II.

Of the 2,309 participants who consented to participate in Phase I, 1,688 were invited to participate in Phase II. (This *N* is higher than our preregistered value; we explain our rationale in Preregistration Addendum 2-23-18). The remaining 621 participants were not invited to complete Phase II because they did not complete Phase I ($n = 574$) or completed Phase I multiple times ($n = 47$). An additional 429 were lost to attrition because they did not accept their invitation to complete Phase II ($n = 422$) or did not provide a valid email for recontact ($n = 7$). Three participants could not complete Phase II because none of the studies they viewed in Phase I were included in Phase II. Of the remaining 1,256 participants who completed Phase II, 160 were excluded from reported analyses because they evaluated considerably more than three studies due to experimenter error ($n = 14$), were undergraduate students ($n = 24$), or indicated they did not understand critical questions ($n = 122$). The first exclusion criterion was not preregistered but the latter two were preregistered. This left us with responses for 1-3 studies (M

= 2.2) from 1,096 participants ($n_{\text{control}} = 572$, $n_{\text{prediction}} = 524$) for both Phase I and Phase II (see Supplementary Methods - Power Analysis).

The majority of unexcluded participants were graduate students (53.0%), but postdoctoral scholars (12.4%), and tenure track faculty (Assistant Professors: 14.5%, Associate Professors: 6.4%, and Professors: 2.9%) also made up a substantial portion of our sample. The remaining participants identified as lectures (2.1%) or “other” (8.7%).

Data Exclusions.

We excluded Phase I responses that referred to studies excluded from Phase II; these exclusions were not preregistered. As preregistered, we excluded responses in cases where participants indicated that they knew the results of a replication of the original study before beginning Phase II of our study.

Phase I.

Participants were told that they were going to be asked about their views regarding various contemporary psychology findings. Participants were presented the informed consent and, after consenting to participate, read descriptions of, and answered questions about, three original studies randomly selected from the pool of studies.

Evaluation of studies - control condition. Participants in the control condition read the following information about each original study: citation, summary (modified from replication protocols), details about the study sample, and key finding (including effect size, test statistic, and p -value). After reading some basic guidelines regarding effect sizes (i.e., Cohen’s guidelines, and the average effect size published social psychology papers), participants estimated the probability that the true effect was non-trivial in size (i.e., Cohen’s $d > .1$). We refer to this estimate as the participant’s *actual prior* (though we also consider their prior as a probability distribution— see “Calculating Bayesian Posteriors” section). Participants also indicated their estimate of the true effect size and responded to four items regarding their confidence that the effect was non-trivial (“*The field of psychology would be justified in [funding research on this finding / hiring people who study topics related to this finding / teaching this finding / publicizing this finding]*”) using a three-point scale ([*less than/the same as/more than*] the typical psychology finding). In the instructions at the top of the scale, participants were explicitly asked to base their answers only on their impressions of the robustness of the findings, and to disregard the importance or interestingness of the effects (see Supplementary Methods - Secondary Dependent Variable for additional detail). These four items were averaged and used as a measure of participants’ confidence in the effect (for results using this variable, see Supplementary Data 1, 6, & 7).

Participants also indicated their personal investment in the finding (“*To what degree are you invested in this finding?*”) using a 5-point scale (1 = *strong preference that it be refuted*, 3 = *no preference*, 5 = *strong preference that it be supported*, $M = 3.13$, $SD = .58$). Participants completed three *yes/no* items inquiring about whether they were friends with the original authors, whether they had done research on the topic, and whether they had heard about the study before.

At the end of this section, participants were reminded of the three studies they evaluated, and asked to indicate *yes/no* to whether they knew of any study that had attempted to replicate the original finding. They also indicated their degree of comprehension (“*I am comfortable that I understood what the questions were asking and the meaning of my responses*”) using a 5-point scale (*1 = strongly disagree, to, 5 = strongly agree*).

Evaluation of studies - prediction condition. Participants in the prediction condition went through the same procedure as those in the control condition described above, but after evaluating each original study, they went on to read descriptions of, and answer questions about, the corresponding replication study.

Participants read the following information about each replication study: a reminder of the key finding from the original study, details about the study sample, and known methodological differences from the original study. Importantly, the results of the replication studies were not yet publicly available, and we did not provide any information about the replication study’s results. Participants answered two *yes/no* items inquiring about whether they knew about the results of this replication study or any other study attempting to replicate the effect. Then, based only on the description of the replication study’s methods, they rated the quality of the replication study using a 3-point scale (*1 = low quality, 2 = moderate quality, 3 = high quality*), and the extent of contextual differences between the replication and original studies using a 3-point scale (*1 = minor contextual differences, 2 = moderate contextual differences, 3 = major contextual differences*).

Participants were then asked to make predictions about how they would update their beliefs given various hypothetical results. Before making their predictions, participants were reminded of their *actual prior* and their estimate of the true effect size. They then went on to predict how those two values would change for each of 6 hypothetical replication outcomes (a negative sign indicates an effect in the opposite direction than the original study’s effect): $d = -.15$, $d = .05$, $d = .20$, $d = .60$, $d = 1.20$, and $d = 1.80$. Participants also predicted which of these six potential replication effect sizes they expected to be closest to the actual effect size found by the replication study. At the end of this section, participants indicated their degree of comprehension in the same manner as in the control condition.

Individual difference measures and demographics. At the end of the study, all participants were asked two questions regarding their general attitudes about the replicability of psychology research and if they had ever replicated a study or had someone attempt to replicate one of their studies. Then, they completed the Cognitive Reflection Test (Frederick, 2005) and demographic items assessing age, gender, racial identity, political ideology, academic position, subfield of psychology, and how they heard about our study.

Phase II.

Participants were told that they would read about the original effects that they evaluated in Phase I and then be asked to respond to the actual results of a replication of each of those effects. The procedure in Phase II was identical for participants in both conditions (prediction vs. control).

Evaluation of replication study. For each study, participants read the same information that was given to participants in the prediction condition of Phase I (i.e., details of the original study and replication study). Crucially, all participants now learned the key finding from the replication. Participants then provided ratings of replication quality and contextual differences (these were also assessed during Phase I for those in the prediction condition). These measures were included to compare attitudes about the replication study's methodology to participants' corresponding attitudes from Phase I (before they knew the results of the replication studies). Participants then provided a new estimate of the probability that the true effect was non-trivial in size (i.e., that Cohen's $d > .1$). This estimate constituted the participant's *actual posterior*. Participants also indicated their estimate of the true effect size and responded to the four Likert scale confidence items from Phase I. They then responded to two *yes/no* items, one asking whether they thought the effect replicated, and another asking whether they knew the results of this replication before completing Phase II.

Individual difference measures and demographics. Participants completed the same four items assessing general attitudes about the replicability of psychology research and experience with replication from Phase I. They also completed two self-report measures of intellectual humility (Leary et al., 2017; McElroy et al., 2014). We preregistered the choice to use the Intellectual Humility (IH) Scale created by Leary and colleagues (2017) for hypothesis testing. This scale is a 6-item measure assessing the willingness to reconsider one's pre-existing views (e.g., "*I reconsider my opinions when presented with new evidence*").

Participants completed the same demographic items as in Phase I. Finally, they responded to two items assessing their expertise with both *motivated reasoning* and *confirmation bias* using a 5-point scale ($1 = \text{much less than the average psychologist}$, $3 = \text{about the same as the average psychologist}$, $5 = \text{much more than the average psychologist}$). These two items were averaged and used as a measure of expertise with bias.

Analysis Plan.

Calculating Bayesian Posteriors. We converted each individual's responses about their priors into probability density distributions. To do this, we assumed these distributions were normally distributed. We used the participant's Phase I estimate of the population effect size as the mean of the distribution, and their stated estimate of the probability of Cohen's $d > .1$ (their *actual prior*), to calculate the standard deviation of the modeled distribution. In cases where participants' priors were incompatible with this assumption (e.g., a distribution mean of .1 with 90% of the distribution greater than .1), we did not calculate *Bayesian posteriors* (and thus these participants were excluded from analyses involving *Bayesian posteriors* - this exclusion was not preregistered). Because prior distributions cannot be defined for *actual priors* of 0% or 100%, we converted responses of "0%" and "100%" to .25% and 99.75%, respectively (detailed explanation in Supplementary Methods - Calculating Bayesian Posteriors). While our assumption that participants' prior effect size beliefs are normally distributed seems likely to approximately model participants' true beliefs, this assumption is nonetheless imprecise. For example, if participants systematically had more leptokurtic [platykurtic] bell shaped distributions, our model would systematically overestimate [underestimate] how much participants should adjust.

For each *actual prior*, we calculated a *Bayesian posterior* based on the replication evidence from Phase II, using conjugate Bayesian inference. That is, we considered how a perfectly rational Bayesian agent would update their beliefs based on new evidence, starting from participants' actual priors, and conditional on them trusting the replication results. (see formulas and spreadsheet to test different inputs here:

https://osf.io/y5n3f/?view_only=904174e2fc72413c9ecbe7b35a5507a1). To do this, we computed the weighted average of two effect-size distributions—one modeling the participant's prior belief (previously described) and one modeling the replication evidence. The replication evidence was modeled using the effect size (the distribution mean) and standard error from the replication results. Computing the weighted average of these two distributions (weighted by the ratio of the distributions' precision [the inverse of variance]) gave us a third distribution combining information from the participant's Phase I effect size estimate and the replication evidence (see Figure 3). The mean of this distribution indicated the calculated population effect size estimate, and the standard deviation was used to determine the calculated probability that the effect is greater than $d > .1$ —the *Bayesian posterior*. This *Bayesian posterior* provides one standard against which to compare participants' *actual posteriors*.

Calculating Predicted Posteriors. For each study, participants in the prediction condition were asked to forecast the posterior they thought they would provide for 6 possible outcomes (Cohen's d of $-.15$, $.05$, $.20$, $.60$, 1.20 , and 1.80). To generate a single *predicted posterior* we conducted a linear interpolation between the two posteriors provided for the possible outcomes that flanked the actual outcome (see syntax here

https://osf.io/y5n3f/?view_only=904174e2fc72413c9ecbe7b35a5507a1). For example, if a replication produced a d of $.10$ (a third of the way between possible outcomes $d = .05$ and $d = .20$), we used the value that would fall a third of the distance along a linear trend between the posterior provided for $d = .05$ and that provided for $d = .20$.

Cases When Participants Should Adjust Upward vs. Downward. Hypotheses 1 through 6 (see Table 1) make a distinction between when participants should, according to our Bayesian model, adjust their confidence (that an effect is greater than Cohen's $d = .1$) upward vs. downward. If a participant's *Bayesian posterior* was larger (smaller) than their *actual prior* for a given effect, we labeled this as a case when the participant “should” adjust upward (downward). (Throughout this manuscript, we use the phrase “should adjust upward/downward” in lieu of the more precise “would have adjusted upward/downward if acting as a Bayesian agent.”) Therefore, the direction participants should adjust was not defined solely by comparing the effect size from the original and replication studies but also by their priors. There are many cases in which participants responding to evidence from the same study should adjust their confidence in opposite directions. Nevertheless, there was a direction the overwhelming majority of participants should have adjusted for most studies. For 14 of the 25 studies, participants generally should have adjusted downward (i.e., the replication result effect size was lower than the center/mean of participants' priors). For 10 studies, participants generally should have adjusted upwards, and for one study, an approximately equal number of participants should have adjusted in each direction. Overall, participants provided 1,285 priors from which they should have adjusted downward, and 898 priors from which they should have adjusted upwards.

In some cases, the effect sizes and *actual priors* provided by participants made it impossible for us to calculate a *Bayesian posterior* (see Calculating Bayesian Posteriors section; these cases were not included in analyses for Hypothesis 2). In these cases, if the proportion of the replication evidence distribution that was above $d = .1$ was larger (smaller) than their *actual prior* for a given effect, this was classified as a case when the participants should adjust upward (downward). When both of these methods could be used, they were in agreement about adjustment direction in 99.6% of cases.

MLM Analyses. In our preregistration we specified that we would conduct both non-multilevel and multilevel analyses. We now believe that the multilevel analyses are more appropriate for our data, which is nested. We thus report only the MLM analyses in this manuscript; all other preregistered analyses are reported in Supplementary Data 1-8 (see subsequent section for a description of preregistered analyses in supplementary material and the reporting of conflicting results).

We used a cross-classified random effects model, which we deemed appropriate as responses were nested in stimulus (the studies that participants evaluated) and participant but neither stimulus nor participant were nested in one another. Therefore, the multilevel model could incorporate two random effects (stimulus and participant) with only two levels. For all hypotheses, the initial model we fit to the data included the random effect of participant on intercept and the random effect of stimulus on both intercept and slope. However, for 9 out of 11 hypotheses, model comparisons revealed that the full model was not a significantly better fit for the data (based on the Chi-square test of the log likelihood deviances) than a reduced model without the random effect of stimulus on the fixed effect slope. For both versions of Hypothesis 2 (adjust upward & adjust downward), the more complex model was a statistically significant better fit, but the improvement was minimal. When participants should have adjusted downward, the random factors in the more complex model were estimated to account for an additional .9% of residual variance (81.6% vs. 80.7%), and when participants should have adjusted upward, the more complex model were estimated to account for an additional 3.4% of residual variance (71.8% vs. 68.4%). Therefore, with few exceptions (see Supplementary Methods – Random Effects), hypotheses were tested with models that included only the random effect of stimulus and participant on intercept. The basic model can be expressed with classification notation (Browne et al. 2001) as:

$$y_i = \beta_0 + \beta_1 x_i + u^{(2)}_{\text{participant}(i)} + u^{(3)}_{\text{stimulus}(i)} + e_i$$

$$u^{(2)}_{\text{participant}(i)} \sim N(0, \sigma^2_{u(2)})$$

$$u^{(3)}_{\text{stimulus}(i)} \sim N(0, \sigma^2_{u(3)})$$

$$e_i \sim N(0, \sigma^2_e)$$

where y_i denotes the outcome variable for the i^{th} response, β_0 is the model intercept, x_i denotes the value of the predictor for the i^{th} response, β_1 is the associated slope coefficient, $u^{(2)}_{\text{participant}(i)}$ and $u^{(3)}_{\text{stimuli}(i)}$ denotes the random effect of participant and stimuli for the i^{th} response and e_i denotes the error for the i^{th} response.

Preregistration.

Deviations from preregistered plan.

1. In our preregistration, we said that we would divide analyses (for Hypotheses 1-6) based on a distinction between when the evidence was stronger versus weaker than the original evidence. However, there were cases when the replication evidence was weaker than the original evidence and still, according to our Bayesian model, warranted an upward adjustment in confidence (e.g., a precise replication effect size estimate that was smaller than that reported in the original study but larger than a trivial effect). It is the direction of adjustment that is called for by our Bayesian model, not the strength of the evidence, that is the distinction relevant to our hypotheses.
2. In our preregistration, we said that we would use absolute values to quantify belief change for Hypotheses 7 and 8. However, there were cases when participants changed their beliefs in the *opposite* direction of that indicated by our Bayesian model. Because we intended to capture the degree to which a participant's confidence was swayed by replication evidence, we instead quantified belief change as positive when it changed in the direction consistent with our Bayesian model, and negative when it changed in the opposite direction.
3. In our preregistration, we did not predict the experimenter error (showing participants all studies, instead of just 3) that led us to exclude data from 14 participants.
4. In our preregistration, we did not anticipate that we would need to drop some studies between Phases I and II, and so did not preregister that we would exclude responses in Phase I to studies that were dropped in Phase II.
5. In our preregistration, we planned to recruit 1,126 participants in Phase I but, as we were approaching this target, observed that our sample was heavily skewed toward social and cognitive psychologists. With the aim of having a more representative sample of psychologists, and, with concerns of deviating from our plan ameliorated by the fact that Phase II data was necessary for testing any of our hypotheses, we exceeded our original participant target for Phase I data collection (see Preregistration Addendum 2-23-18 for details: https://osf.io/wz3ty/?view_only=16834083b7f441059be85c85db1c04b9)

Not specified in preregistration.

1. We did not preregister which of the four possible random effects we would include in our MLM models nor the criteria for inclusion versus exclusion of random effects.
2. We did not preregister our formulas for calculating *Bayesian posteriors* and *predicted posteriors*.

All other analytical decisions for hypothesis testing were preregistered unless we explicitly note that they were not. The Supplementary Materials include some preregistered (see Supplementary Data 1 – 8) and some unpreregistered analyses (mostly in Supplementary Data 9-10 but some [clearly labeled] in other supplementary data sections), with the distinction clearly noted. The results of preregistered analyses in the Supplementary Materials—including results for a secondary dependent variable for several hypotheses (Likert-scale-based measure of confidence) and non-MLM versions of the analyses for all hypotheses (see Supplementary Methods – Calculating Variables for Non-MLM Analyses for variable calculation procedure)—generally yielded quite similar conclusions to the preregistered analyses reported here. Substantial

discrepancies—in effect size or p -value—between preregistered analyses assessing the same hypothesis are noted in the results section. We somewhat arbitrarily operationalized “substantial discrepancies” as differences in $\beta > .15$ and p -value differences that spanned both sides of $p = .001$, $p = .01$, or $p = .05$ (discrepancies not fitting these criteria are only documented in Supplementary Data 1—8).

Data availability. All data are available here (https://osf.io/jtp4b/?view_only=63d9a7deb66e4764821b048c6f837185).

Code availability. The algorithm for computing *Bayesian posteriors* is available here (https://osf.io/y5n3f/?view_only=904174e2fc72413c9ecbe7b35a5507a1).

Tables

Table 1. Research questions and hypotheses.

	<i>Prediction when participants should, according to our Bayesian model, adjust downward</i>	<i>Prediction when participants should, according to our Bayesian model, adjust upward</i>
RQ1: How much psychologists update their beliefs in response to new empirical evidence?		
H1. Participants will update their beliefs in response to new evidence	<i>Actual posteriors</i> will be lower than <i>actual priors</i>	<i>Actual posteriors</i> will be higher than <i>actual priors</i>
H2. Participants will not update as much as our Bayesian model would predict	<i>Actual posteriors</i> will be higher than <i>Bayesian posteriors</i>	<i>Actual posteriors</i> will be lower than <i>Bayesian posteriors</i>
H3. Participants will not update their beliefs as much as they predict they will	<i>Actual posteriors</i> will be higher than <i>predicted posteriors</i>	<i>Actual posteriors</i> will be lower than <i>predicted posteriors</i>
RQ2: Do psychologists show evidence of motivated reasoning when evaluating replications?		
H4. Participants will provide lower ratings of replication quality when this would preserve their pre-existing beliefs	Higher <i>actual priors</i> will be associated with greater decreases in ratings of replication quality from Phase I to Phase II	No prediction
H5. Participants will be more likely to identify context differences when this would preserve their pre-existing beliefs	Higher <i>actual priors</i> will be associated with greater increases in context difference ratings from Phase I to Phase II	No prediction
RQ3: What predicts the extent of psychologists' belief updating?		
H6. Personal investment will be negatively associated with belief updating	Participants who are more personally invested in a finding will show less of a difference between their <i>actual priors</i> and their <i>actual posteriors</i>	No prediction
H7. Expertise regarding confirmation bias and motivated reasoning will be positively associated with belief updating	Participants who have more expertise regarding confirmation bias and motivated reasoning will show more of a difference between their <i>actual priors</i> and their <i>actual posteriors</i>	
H8. Self-reported intellectual humility will be positively associated with belief updating	Participants who self-report more intellectual humility (Leary et al., 2017) will show more of a difference between their <i>actual priors</i> and their <i>actual posteriors</i>	

Note: Hypotheses 3 through 5 are only tested in the prediction condition because they use difference scores that incorporate ratings only made by participants in the prediction condition.

Acknowledgements

We received funding from Grant 1728332 from the National Science Foundation (A.M.T. & S. V.) The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript. We are grateful to Katie Finnigan and Jessie Sun for their assistance in recruiting study participants, to Jacob Miranda for helping upload and organize data files on the Open Science Framework, to Charlie Ebersole, Rick Klein, and Dan Simons for providing updates on the timelines for publication of Many Labs 2 and Many Labs 5, and to Will Hart and David McDiarmid for their feedback on statistical analyses.

Author Contributions

A.D.M., A.M.T., & S.V. designed the study. A.D.M., A.M.T., & C.M.W. developed stimuli and collected data. A.D.M. created the analytic plan and analyzed data with contributions from A.M.T., P.E.S, & E.E.S. A.D.M. & A.M.T. wrote the manuscript, and all authors edited the manuscript and gave conceptual advice.

Competing Interests

The authors declare no competing interests.

Figure Legends

Figure 1. Overview of study procedure. †Some participants saw fewer than three studies during Phase II because some studies were dropped (see main text for details).

Figure 2. A hypothetical participant's prior belief of the effect size distribution (green line) is defined by a mean of .25 and an estimate that there is an 85% probability the (population) effect size is greater than $d = .1$, both reported by the hypothetical participant. Assuming a normal distribution, the standard deviation of this distribution is calculated to be .145. The replication evidence (blue line) is defined by a mean of .05 and a standard error of .05, both reported in the replication study report. These two distributions are used to compute the hypothetical participant's Bayesian posterior distribution (purple line). For this hypothetical participant, 27.2% of their Bayesian posterior distribution would be greater than $d = .1$, and therefore, their *Bayesian posterior* would be 27.2% (the area shaded in purple).

Figure 3. Panel A displays the mean probability ratings for the prediction condition when participants should adjust their ratings downward. White bars represent participants' actual priors and posteriors, the black bar represents participants' predicted posteriors, and the striped bar represents participants' Bayesian posteriors. Panel B displays the same information as Panel A but for participants in the control condition. Here, the grey bar represents the predicted posterior provided by participants in the prediction condition, as participants in the control condition did not provide predictions. Panels C and D display the same information as A and B, respectively, but for cases when participants should adjust their ratings upward. Error bars indicate standard errors.

References

- Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology*, 39(6), 1037–1049. <https://doi.org/10.1037/h0077720>
- Apsler, R., & Sears, D. O. (1968). Warning, personal involvement, and attitude change. *Journal of Personality and Social Psychology*, 9(2, Pt.1), 162-166. <https://doi.org/10.1037/h0021248>
- Barrett, L. F. (2015, September 1). Psychology is not in crisis. *The New York Times*. <https://www.nytimes.com/2015/09/01/opinion/psychology-is-not-in-crisis.html>
- Browne, W. J., & Goldstein, H., & Rasbash, J. (2001). Multiple membership multiple classification (MMMC) models. *Statistical Modelling*, 1(2), 103-124. <https://doi.org/10.1177/1471082X0100100202>
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T. H., Huber, J., Johannesson, M., ... & Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280), 1433-1436.
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T. H., Huber, J., Johannesson, M., ... & Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637-644.
- Carlson, J., & Harris, K. (2020). Quantifying and contextualizing the impact of bioRxiv preprints through automated social media audience segmentation. *PLoS Biology*, 18(9), e3000860.
- Duhem, P. (1906). *La théorie physique, son objet et sa structure*. Chevalier et Rivière.
- Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Allen, J. M., Banks, J. B., ... & Nosek, B. A. (2016). Many Labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, 67, 68-82.
- Ebersole, C. R., Mathur, M. B., Baranski, E., Bart-Plange, D., Buttrick, N., Chartier, C. R., ... Nosek, B. A. (2019). Many Labs 5: Testing pre-data collection peer review as an intervention to increase replicability (results-blind manuscript). <https://doi.org/10.31234/osf.io/sxfrm2>
- Fanelli, D. (2012). Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90(3), 891-904.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25-42. <https://doi/pdfplus/10.1257/089533005775196732>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis*. CRC Press. <http://www.stat.columbia.edu/~gelman/book/>
- Gershman S. J. (2019). How to never be wrong. *Psychonomic Bulletin & Review*. 26(1), 13-28. <https://doi.org/10.3758/s13423-018-1488-8>
- Hardwicke, T. E., Szűcs, D., Thibault, R. T., Crüwell, S., Van den Akker, O., Nuijten, M. B., & Ioannidis, j. (2021, February 9). Post-replication citation patterns in psychology: Four case studies. <https://doi.org/10.31222/osf.io/wt5ny>
- Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: a meta-analysis of selective exposure to information. *Psychological bulletin*, 135(4), 555–588. <https://doi.org/10.1037/a0015701>

- Hayden, T., & Mischel, W. (1976). Maintaining trait consistency in the resolution of behavioral inconsistency: The wolf in sheep's clothing?. *Journal of Personality*, 44(1), 109-132. <https://doi.org/10.1111/j.1467-6494.1976.tb00587.x>
- Hergovich, A., Schott, R., & Burger, C. (2010). Biased evaluation of abstracts depending on topic and conclusion: Further evidence of a confirmation bias within scientific psychology. *Current Psychology: A Journal for Diverse Perspectives on Diverse Psychological Issues*, 29(3), 188–209. <https://doi.org/10.1007/s12144-010-9087-5>
- Ioannidis, J. P. (2012). Why science is not necessarily self-correcting. *Perspectives on Psychological Science*, 7(6), 645-654.
- Jacks, J. Z., & Cameron, K. A. (2003). Strategies for resisting persuasion. *Basic and Applied Social Psychology*, 25(2), 145–161. https://doi.org/10.1207/S15324834BASP2502_5
- Jost, A. (1897). Die Assoziationsfestigkeit in ihrer Abhängigkeit von der Verteilung der Wiederholungen [The strength of associations in their dependence on the distribution of repetitions]. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*, 16, 436–472.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., Bahník, Š., Batra, R., Berkics, M., Bernstein, M. J., Berry, D. R., Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., ... Nosek, B. A. (2018). Many Labs 2: Investigating Variation in Replicability Across Samples and Settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443–490. <https://doi.org/10.1177/2515245918810225>
- Koehler, J. J. (1993). The influence of prior beliefs on scientific judgments of evidence quality. *Organizational Behavior and Human Decision Processes*, 56(1), 28-55. <https://doi.org/10.1006/obhd.1993.1044>
- Krumrei-Mancuso, E., & Rouse, S. (2016). The development and validation of the Comprehensive Intellectual Humility Scale. *Journal of Personality Assessment*, 98(2), 209-221. <https://doi.org/10.1080/00223891.2015.1068174>
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480–498. <https://doi.org/10.1037/0033-2909.108.3.480>
- Kuran, T. (1997). *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. Harvard Univ Press.
- Leary, M. R., Diebels, K. J., Davisson, E. K., Jongman-Sereno, K. P., Isherwood, J. C., Raimi, K. T., Deffler, S. A., & Hoyle, R. H. (2017). Cognitive and interpersonal features of intellectual humility. *Personality and Social Psychology Bulletin*, 43(6), 793–813. <https://doi.org/10.1177/0146167217697695>
- Longino, H. E. (1990). *Science as Social Knowledge*. Princeton University Press.
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11), 2098–2109. <https://doi.org/10.1037/0022-3514.37.11.2098>
- Luchins, A. S. (1957). Primacy-recency in impression formation. In C. Hovland (Ed.), *The Order of Presentation in Persuasion*. Yale University Press.
- MacCoun, R. J. (1998). Biases in the interpretation and the use of research results. *Annual Review of Psychology*, 49(1), 259–287. <https://doi.org/10.1146/annurev.psych.49.1.259>
- Mahoney, M.J. (1977). Publication prejudices: An experimental study of confirmatory bias in the peer review system. *Cognitive Therapy and Research*, 1(2), 161–175. <https://doi.org/10.1007/BF01173636>

- McCarthy, R. J., Skowronski, J. J., Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., ... Yildiz, E. (2018). Registered replication report on Srull and Wyer (1979). *Advances in Methods and Practices in Psychological Science*, 1(3), 321-336. <https://doi.org/10.1177/2515245918777487>
- McElreath, R., & Smaldino, P. E. (2015). Replication, communication, and the population dynamics of scientific discovery. *PloS one*, 10(8), e0136088. <https://doi.org/10.1371/journal.pone.0136088>
- McElroy, S. E., Rice, K. G., Davis, D. E., Hook, J. N., Hill, P. C., Worthington, E. L., Jr., & Van Tongeren, D. R. (2014). Intellectual humility: Scale development and theoretical elaborations in the context of religious leadership. *Journal of Psychology and Theology*, 42(1), 19–30. <https://doi.org/10.1177/009164711404200103>
- Miller, A. G., McHoskey, J. W., Bane, C. M., & Dowd, T. G. (1993). The attitude polarization phenomenon: Role of response measure, attitude extremity, and behavioral consequences of reported attitude change. *Journal of Personality and Social Psychology*, 64, 561-574.
- Munro, G.D. (2010). The scientific impotence excuse: Discounting belief-threatening scientific abstracts. *Journal of Applied Social Psychology*, 40(3), 579-600. <https://doi.org/10.1111/j.1559-1816.2010.00588.x>
- Munro, G. D., & Ditto, P. H. (1997). Biased assimilation, attitude polarization, and affect in reactions to stereotype-relevant scientific information. *Personality and Social Psychology Bulletin*, 23(6), 636–653. <https://doi.org/10.1177/0146167297236007>
- Munro, G. D., Leary, S. P., & Lasane, T. P. (2004). Between a rock and a hard place: Biased assimilation of scientific information in the face of commitment. *North American Journal of Psychology*, 6(3), 431–444.
- Murdock, B. B., Jr. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, 64(5), 482–488. <https://doi.org/10.1037/h0045106>
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175–220. <https://doi.org/10.1037/1089-2680.2.2.175>
- Nisbett, R. E. & Ross, L. (1980). *Human Inference: Strategies and Shortcomings of Social Judgment*. Prentice-Hall.
- Nissen, S. B., Magidson, T., Gross, K. & Bergstrom, C.T. (2016). Research: Publication bias and the canonization of false facts. *Elife*, 5, e21451. <https://doi.org/10.7554/eLife.21451>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251).
- Oreskes, N. (2019). *Why trust science?*. Princeton University Press.
- Panagiotou, O. A., & Ioannidis, J. P. A. (2012). Primary study authors of significant studies are more likely to believe that a strong association exists in a heterogeneous meta-analysis compared with methodologists. *Journal of Clinical Epidemiology*, 65(7), 740-747. <https://doi.org/10.1016/j.jclinepi.2012.01.008>
- Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives on Psychological Science*, 7(6), 531-536.
- Pronin, E., & Kugler, M. B. (2007). Valuing thoughts, ignoring behavior: The introspection illusion as a source of the bias blind spot. *Journal of Experimental Social Psychology*, 43(4), 565–578. <https://doi.org/10.1016/j.jesp.2006.05.011>
- Quine, W. V. O. (1951). Two dogmas of empiricism. *Philosophical Review*, 60(1), 20–43. <https://doi.org/10.2307/2266637>

- Rhine, R. J., & Severance, L. J. (1970). Ego-involvement, discrepancy, source credibility, and attitude change. *Journal of Personality and Social Psychology*, 16(2), 175–190. <https://doi.org/10.1037/h0029832>
- Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology*, 32(5), 880–892. <https://doi.org/10.1037/0022-3514.32.5.880>
- Scheel, A. M., Schijen, M., & Lakens, D. (2021). An excess of positive results: Comparing the standard Psychology literature with Registered Reports. *Advances in Methods and Practices in Psychological Science*, 4(2), 1-12.
- Tetlock, P.E., Mellers, B. A., & Scoblic, J. P. (2017). Bringing probability judgments into policy debates via forecasting tournaments. *Science*, 355(6324), 481-483.
- Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Contextual sensitivity in scientific reproducibility. *Proceedings of the National Academy of Sciences of the United States of America*, 113(23), 6454–6459. <https://doi.org/10.1073/pnas.1521897113>
- van Elk, M., Matzke, D., Gronau, Q., Guan, M., Vandekerckhove, J., & Wagenmakers, E. J. (2015). Meta-analyses are no substitute for registered replications: A skeptical perspective on religious priming. *Frontiers in psychology*, 6, 1365.
- Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., Orthey, R., McCarthy, R. J., Skowronski, J. J., Acar, O. A., Aczel, B., Bakos, B. E., Barbosa, F., Baskin, E., Bègue, L., Ben-Shakhar, G., Birt, A. R., Blatz, L., Charman, S. D., Claesen, A., Clay, S. L., ... Yıldız, E. (2018). Registered replication report on Mazar, Amir, and Ariely (2008). *Advances in Methods and Practices in Psychological Science*, 1(3), 299–317. <https://doi.org/10.1177/2515245918781032>

All Studies