Errors in probabilistic reasoning and judgment biases: Online appendix

Tushar Kundu

Edited by: Daniel J. Benjamin

First draft: April 3, 2019

This draft: April 3, 2019

Contents

1 Introduction

This online appendix covers all empirical analyses conducted for the chapter in the *Handbook of Behavioral Economics*, "Errors in probabilistic reasoning and judgment biases." This includes two main analyses: a meta-analysis of *simultaneous* and *sequential* updating problems, and a meta-analysis of regression estimates for biased use of likelihoods and priors. Sections 2-5 of this online appendix will focus on the meta-analysis of simultaneous and sequential updating problems, with section 2 describing the data, and sections 3-5 detailing the analyses. Section 6 will cover the meta-analysis of regression estimates from other papers. For code that accompanies the analyses described in sections 3-5, please refer to `beliefupdatemeta_public.R`, which may be downloaded on Daniel Benjamin's website (as of the date of this document: `https://www.danieljbenjamin.com/handbook-chapter`).

2 Data on updating problems

2.1 Simultaneous samples

The meta-analysis sample for simultaneous problems (see chapter section 4.2) includes 455 observations drawn from 16 papers. The papers are listed below in Table 1.

Table 1: Papers with simultaneous updating problems

Paper	Number of observations	Incentivized?
Bar-Hillel (1980)	4	N
Beach, Wise and Barclay (1970)	4	N
Camerer (1987)	8	N
Donnell and DuCharme (1975)	2	Y
Gettys and Manley (1968)	50	N
Green, Halbert, and Robinson (1965)	83	Y
Grether (1992)	91	Y/N*
Griffin and Tversky (1992)	45	N
Holt and Smith (2009)	30	Y
Kahneman and Tversky (1972)	10	N
Kraemer and Weber (2004)	4	Y
Marks and Clarkson (1972)	23	N
Nelson, Bloomfield, Hales, and Libby (2001)	5	Y
Peterson and Miller (1965)	72	N
Peterson, Schneider, and Miller (1965)	13	N
Sasaki and Kawagoe (2007)	1	N

*Grether (1992) ran several experiments, some of which had financial incentives. We include observations from both types of experiments, and indicate which observations are incentivized.

All data is contained in the file `sim_data.csv`. Columns in the dataset are defined as:

- `Source`: Name of the paper that reported the original observation.
- `SourceNo`: Internally assigned source id number. Note that "Source" for Grether (1992) is split into "Grether (1992)" and "Grether (1992) Not Incentivized." However, SourceNo = 1 for all Grether (1992) observations.
- `thetaA`: Probability of observing an a signal in state A (see def. of θ_A in chapter p. 104)
- `thetaB`: Probability of observing an a signal in state B (θ_B)
- `pA`: Prior probability of state A being the true state ($p(A)$)
- `pB`: Prior probability of state B being the true state ($p(B)$)
- `N`: Total number of signals (N)
- `Na`: Total number of a signals (N_a)
- `ObsPostMedian`: Sample median of posterior beliefs that state A is the true state
- `ObsPostMean`: Sample mean of posterior beliefs that state A is the true state
- `ObsPostMode`: Sample mode of posterior beliefs that state A is the true state

2.2 Sequential samples

The meta-analysis sample for sequential problems (see chapter section 4.3) includes 127 observations drawn from 5 papers. The papers are listed below in Table 2.

Table 2: Papers with sequential updating problems

Paper	Number of observations	Incentivized?
Beach, Wise, and Barclay (1970)	54	N
Dave and Wolfe (2003)	50	Y
Kraemer and Weber (2004)	3	Y
Peterson, Schneider, and Miller (1965)	13	N
Sasaki and Kawagoe (2007)	7	N

All data is contained in the file `seq_data.csv`. Columns in the dataset are the same as those defined above in section 2.1.

3 Participants' log-posterior-odds versus Bayesian log-posterior-odds

3.1 Calculating log-posterior-odds

The first step in conducting the meta-analysis of simultaneous and sequential updating problems is to calculate the true and empirical (participants') log-posterior-odds for each updating problem. To calculate participants' log-posterior-odds, $\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)$, we take posterior probabilities for state A directly from each paper, using either the median, mean, or mode reported belief. For simultaneous updating problems, when multiple statistics are reported, we use this preference ordering to choose which statistic to use: first the median (145 observations), then the mean (296 observations), and then the mode (4 observations). For sequential updating problems, we use the mean posterior belief, which is reported for all 127 observations. Participants' log-posterior-odds are calculated as $\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)$, for which $\pi(A|S)$ is the reported posterior probability, and $\pi(B|S) = 1 - \pi(A|S)$.

To calculate the true log-posterior-odds, $\ln\left(\frac{p(A|S)}{p(B|S)}\right)$, we use Bayes' Theorem:

$$\ln\left(\frac{p(A|S)}{p(B|S)}\right) = \ln\left(\frac{\frac{p(S|A)p(A)}{p(S)}}{\frac{p(S|B)p(B)}{p(S)}}\right) = \ln\left(\frac{p(S|A)p(A)}{p(S|B)p(B)}\right). \tag{A1}$$

Recall that all observations are for two-state, binomial updating problems. Let θ_A be the probability of observing an a signal in state A, θ_B be the probability of observing an a signal in state B, N be the total number of signals, N_a be the total number of a signals, and N_b be the total number of b signals ($N = N_a + N_B$). The likelihood of observing a set of signals, S, conditional on being in state A is $p(S|A) = (\theta_A)^{N_a} \times (1-\theta_A)^{N-N_a}$. Similarly, conditional on being in state B, the likelihood of observing S is $p(S|B) = (\theta_B)^{N_a} \times (1-\theta_B)^{N-N_a}$. We use the features of each updating problem to calculate these likelihood values, and calculate true log-posterior-odds by plugging the likelihood values and given prior probabilities into Eq. (A1).

3.2 Figure 2 and Table 2

To examine biased inference, we first restrict the sample to observations that satisfy three conditions: the problem uses (i) equal prior probabilities for states A and B, (ii) symmetric signal probabilities, i.e., $\theta_A + \theta_B = 1$, and (iii) an unequal number of a and b signals. Restricting the sample to problems with equal prior probabilities (condition (i)) ensures that the bias in posterior beliefs is entirely driven by biased inferences (see chapter p. 103). Under conditions (i) and (ii), the posterior odds for all problems with equal numbers of a and b signals is equal to 1. Since "overinference" and "underinference" are not well defined in those cases, we impose condition (iii).

For simultaneous updating problems, condition (i) drops 225 observations, condition (ii) drops an additional 68 observations, and condition (iii) drops 5 further observations, leaving us with 147 observations. For sequential updating problems, conditions (i) and (ii) drop zero observations, and condition (iii) drops 16 observations, leaving us with 111 observations.

Figure 2 plots participants' log-posterior-odds, $\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)$, against the true log-posterior-odds, $\ln\left(\frac{p(A|S)}{p(B|S)}\right)$, for the 147 simultaneous observations (panel A) and the 111 sequential observations (panel B). Each plot includes a linear regression line and local regression curve with gray bands indicating 95% confidence intervals. The LOESS curve uses a span of 0.75, and uses locally linear fits (degree = 1). Both options are the default in R.

Table 2 reports the results from regressing participants' log-posterior-odds on true log-posterior-odds. The second column in each panel contains the results after restricting the sample to observations for which participants were incentivized. Heteroskedasticity-robust standard errors are reported in parentheses.

4 Inference based on features of the observed sample

The results shown in Figure 2 and reported in Table 2 show the amount of biased inference as measured by $\hat{c} = \ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)/\ln\left(\frac{p(A|S)}{p(B|S)}\right)$. Next, we ask how biased inference is related to features of the observed sample, namely sample size (N) and diagnosticity (θ).

4.1 Figures 3 and 4

To produce Figure 3, we take the ratio of the participants' log-posterior-odds and the true log-posterior-odds, and plot this against the number of signals (N) for each observation. We include the same 147 simultaneous observations (panel A) and 111 sequential observations (panel B). Similar to Figure 2, we plot linear regression lines and local regression curves with gray bands indicating 95% confidence intervals.

In footnote 27, we report the mean of the inference measure, $\hat{c}$, for specific sample sizes, N. We calculate these as the mean ratio of participants' log-posterior-odds and the true log-posterior-odds among the observations at the given sample size. We estimate the standard error as the standard deviation of these ratios divided by the square root of the number of observations.

Figure 4 is analogous to Figure 3 with diagnosticity, θ, replacing the number of signals, N, on the x-axis. Note that $\theta = \max\{\theta_A, \theta_B\}$. Hence, higher values of θ necessarily imply higher diagnosticity for a given set of signals. We include the same 147 simultaneous observations (panel A) and 111 sequential observations (panel B). We plot the linear regression lines in each panel, but omit local regression curves as they would lie almost directly on top of the linear regression lines.

4.2 Table 3

To examine updating based on features of the observed sample, we run the regression model defined in chapter Eq. (4.13):

$$\ln\left(\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)\right) = \alpha_0 + \alpha_1 \ln\left(\frac{N_a - N_b}{N}\right) + \alpha_2 \ln(N) + \alpha_3 \ln\left(\ln\left(\frac{\theta}{1-\theta}\right)\right) + \epsilon \tag{4.13}$$

Note that this equation assumes we have defined the diagnosticity parameter, θ, as $\theta \equiv \theta_A = 1 - \theta_B$ (see p. 36 in the chapter). Under this assumption, Eq. (4.13) is only well-defined for inference problems such that $\theta_A > 0.5$ and $N_a > N_b$. This holds for 94 out of 147 simultaneous observations and 96 out of 111 sequential observations. However, by relabeling the rates and proportions used in the model, we are able to use the full sample of 147 simultaneous observations and 111 sequential observations that satisfy the three conditions specified above in section 3.2. For inference problems such that $\theta_A > 0.5$ and $N_a < N_b$, we express Bayes' Theorem as $\frac{p(B|S)}{p(A|S)} = \left(\frac{\theta_B}{1-\theta_B}\right)^{\frac{N_a-N_b}{N}\times N} = \left(\frac{\theta_A}{1-\theta_A}\right)^{\frac{N_b-N_a}{N}\times N}$ so that Eq. (4.13) can be rewritten as:

$$\ln\left(\ln\left(\frac{\pi(B|S)}{\pi(A|S)}\right)\right) = \alpha_0 + \alpha_1 \ln\left(\frac{N_b - N_a}{N}\right) + \alpha_2 \ln(N) + \alpha_3 \ln\left(\ln\left(\frac{\theta_A}{1-\theta_A}\right)\right) + \epsilon. \tag{A2}$$

This allows us to include 49 additional simultaneous observations and 15 additional sequential observations. For problems such that $\theta_A < 0.5$ and $N_a < N_b$ we express Bayes' Formula as $\frac{p(A|S)}{p(B|S)} = \left(\frac{\theta_A}{1-\theta_A}\right)^{\frac{N_a-N_b}{N}\times N} = \left(\frac{1-\theta_A}{\theta_A}\right)^{\frac{N_b-N_a}{N}\times N}$ so that the model can be rewritten as:

$$\ln\left(\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)\right) = \alpha_0 + \alpha_1 \ln\left(\frac{N_b - N_a}{N}\right) + \alpha_2 \ln(N) + \alpha_3 \ln\left(\ln\left(\frac{1-\theta_A}{\theta_A}\right)\right) + \epsilon. \tag{A3}$$

This allows us to include the remaining 4 simultaneous observations. To summarize:

For observations such that $\theta_A > 0.5$ and $N_a > N_b$:

- log-log-posterior-odds: $\ln\left(\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)\right)$
- sample-proportion term: $\ln\left(\frac{N_a-N_b}{N}\right)$
- sample-size term: $\ln(N)$
- diagnosticity term: $\ln\left(\ln\left(\frac{\theta_A}{1-\theta_A}\right)\right)$

For observations such that $\theta_A > 0.5$ and $N_a < N_b$:

- log-log-posterior-odds: $\ln\left(\ln\left(\frac{\pi(B|S)}{\pi(A|S)}\right)\right)$
- sample-proportion term: $\ln\left(\frac{N_b - N_a}{N}\right)$
- sample-size term: $\ln(N)$
- diagnosticity term: $\ln\left(\ln\left(\frac{\theta_A}{1-\theta_A}\right)\right)$

For observations such that $\theta_A < 0.5$ and $N_a < N_b$:

- log-log-posterior-odds: $\ln\left(\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)\right)$
- sample-proportion term: $\ln\left(\frac{N_b - N_a}{N}\right)$
- sample-size term: $\ln(N)$
- diagnosticity term: $\ln\left(\ln\left(\frac{1-\theta_A}{\theta_A}\right)\right)$.

The results in columns A1 and B1 of Table 3 are derived from regressing log-log-posterior-odds on the sample-proportion term, sample-size term, and diagnosticity term as defined above for the three classes of observations. Columns A3 and B3 show estimates for the same specification, but with the sample restricted to incentivized observations. To assess evidence regarding "exact representativeness," we include an additional independent variable: an indicator for whether the proportion of a signals, a/N, is equal to θ. The results are shown in columns A2 and B2.

5 Base-rate neglect

To assess biased use of priors, we no longer restrict the sample to observations with equal prior probabilities for states A and B (condition (ii) from section 3.2 above). This increases the number of simultaneous observations from 147 to 296. No sequential observations have unequal prior probabilities, so we can only conduct this analysis in the simultaneous data.

Recall chapter Eq. (4.17):

$$\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right) - \ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right) = d \ln\left(\frac{p(A)}{p(B)}\right). \tag{4.17}$$

To estimate the log of participants' subjective likelihood ratios, $\ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right)$, we first calculate predicted values of their log-log subjective likelihood ratios, $\ln\left(\ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right)\right)$. To do this, we exploit the fact that we estimated biased inferences in regression Eq. (4.13). The fitted values estimate posterior odds for problems with equal priors. Therefore, we treat these predicted values as participants' biased subjective likelihood ratios. For each of the 296 observations, we calculate each term on the right-hand side of Eq. (4.13) (sample-proportion term, sample-size term, and diagnosticity term). To ensure that all the terms are well defined, we need to consider each combination of $\theta_A \gtrless 0.5$ and $N_a \gtrless N_b$. For three of the four possible combinations, the well-defined expression for each term is reported above in section 4.2. For the final combination, $\theta_A < 0.5$ and $N_a > N_b$, we use the expressions:

- sample-proportion term: $\ln\left(\frac{N_a - N_b}{N}\right)$
- sample-size term: $\ln(N)$
- diagnosticity term: $\ln\left(\ln\left(\frac{1-\theta_A}{\theta_A}\right)\right)$.

After calculating these terms, we use the coefficient estimates reported in Table 3 column A1 to predict log-log subjective likelihood ratios. Specifically:

$$\ln\left(\ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right)\right) = -0.052 + 0.848\,\ln\left(\frac{N_a - N_b}{N}\right) + 0.411\,\ln(N) + 0.394\,\ln\left(\ln\left(\frac{\theta_A}{1-\theta_A}\right)\right) \tag{A4}$$

When we restrict the sample to incentivized observations, we use the estimates reported in Table 3 column A3 (estimated only using incentivized observations) to predict log-log subjective likelihood ratios:

$$\ln\left(\ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right)\right) = -0.120 + 0.870\,\ln\left(\frac{N_a - N_b}{N}\right) + 0.462\,\ln(N) + 0.515\,\ln\left(\ln\left(\frac{\theta_A}{1-\theta_A}\right)\right) \tag{A5}$$

Next, we use our estimates of log-log subjective likelihood ratios to derive estimates for log subjective likelihood ratios. Footnote 39 (chapter p. 119) mistakenly describes how this was done in an earlier version of the chapter:

> Simply exponentiating the estimate $\widehat{\ln\left(\ln\left(\frac{\pi(S|A)}{\pi(S|B)}\right)\right)}$ is not a consistent estimator for $\ln\left(\ln\left(\frac{\pi(S|A)}{\pi(S|B)}\right)\right)$ due to Jensen's inequality. Therefore, we generate an estimate of $\ln\left(\ln\left(\frac{\pi(S|A)}{\pi(S|B)}\right)\right)$ by calculating $e^{\hat{\mu}+\frac{1}{2}\hat{\sigma}^2}$ where $\hat{\mu} = \ln\left(\ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right)\right)$ and $\hat{\sigma}^2$ is the estimated variance of the residual from Eq. (4.13). This estimator is consistent under the assumption that $\ln\left(\ln\left(\frac{\pi(S|A)}{\pi(S|B)}\right)\right)$ is normally distributed.

Since the true distribution of $\ln\left(\frac{\pi(S|A)}{\pi(S|B)}\right)$ is unlikely to be normally distributed, for the results reported in the published version of the chapter, we followed a different procedure that does not rely on the normality assumption:

1. Exponentiate the estimates: $e^{\ln\left(\ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right)\right)}$.
2. Calculate the mean of these values: $\overline{e^{\ln\left(\ln\left(\widehat{\frac{\pi(S|A)}{\pi(S|B)}}\right)\right)}}$.

3. Subtract the mean calculated in step 2 from the estimates in step 1. We use the resulting values as estimates for $\ln\left(\frac{\pi(S|A)}{\pi(S|B)}\right)$.

Next, we calculate participants' log-posterior-odds and true log-prior-odds for each observation. To ensure all terms are well defined, we calculate log-posterior-odds as:

- $\ln\left(\frac{\pi(A|S)}{\pi(B|S)}\right)$ if ($\theta_A > 0.5$ and $N_a > N_b$) or ($\theta_A < 0.5$ and $N_a < N_b$)
- $\ln\left(\frac{\pi(B|S)}{\pi(A|S)}\right)$ if ($\theta_A > 0.5$ and $N_a < N_b$) or ($\theta_A < 0.5$ and $N_a > N_b$),

and we calculate log-prior-odds as:

- $\ln\left(\frac{p(A)}{p(B)}\right)$ if ($\theta_A > 0.5$ and $N_a > N_b$) or ($\theta_A < 0.5$ and $N_a < N_b$)
- $\ln\left(\frac{p(B)}{p(A)}\right)$ if ($\theta_A > 0.5$ and $N_a < N_b$) or ($\theta_A < 0.5$ and $N_a > N_b$).

5.1 Figure 5 and Table 4

To produce Figure 5, we subtract participants' predicted log-likelihood-ratio from their log-posterior-odds and plot this against the log-prior-odds. We include the 296 simultaneous observations and plot the linear regression line and local regression curve with gray bands indicating 95% confidence intervals.

We calculate the estimates reported in Table 3 by regressing participants' (log-posterior-odds − predicted log-likelihood-ratio) on the true log-prior-odds. Column 2 restricts the sample to observations for which $p(A)$ is not equal to $p(B)$, and column 3 restricts the sample to incentivized observations.

6 Meta-analysis of sequential-sample updating estimates from other papers

6.1 Data

The second piece of analysis involves meta-analyzing regression estimates (as opposed to posteriors from individual updating problems) for parameters that capture the biased use of likelihoods, c, and the biased use of priors, d (see chapter section 4.3, pp. 124-125). Estimates are drawn from 8 papers:

- Barron (2016)
- Buser, Gerhards, and Van der Weele (2018)
- Charness and Dave (2017)
- Coutts (2017)
- Gotthard-Real (2017)

- Grether (1992)
- Holt and Smith (2009)
- Möbius, Niederle, Niehaus, and Rosenblat (2014)

The estimates are contained in the file `updating_regression_estimates.xlsx`. In the first tab titled "Updating_regression_estimates," columns are defined as:

- `Source`: Name of the paper that reported the original estimates.
- `Original Model`: Specification for the regression.
- `Method of estimation`: Approach for estimating the coefficients in the specification.
- `alpha_`x`_hat`: Reported estimate for the coefficient alpha_x.
- `SE(alpha_`x`)`: Reported standard error for alpha_x.
- `Notes`: Location of estimates within the paper.

The second tab, titled "Meta-analysis," contains two tables. The first table corresponds to estimates for β_1, and the second table corresponds to estimates for β_2 as defined in chapter Eq. (4.21) (p. 124). Columns are defined as:

- `Source`: Name of the paper that reported the original estimates.
- β_x: $\hat{\beta}_x$.
- `SE(`β_x`)`: Standard error for $\hat{\beta}_x$.
- β_x `Weight`: $1/(SE(\beta\_x))^2$ (Inverse-variance weight).
- β_x `Weight *` β_x: Inverse-variance weight × estimate.
- `sum(`β_x `Weight *` β_x`)/sum(`β_x `Weight)`: Inverse-variance-weighted mean of $\hat{\beta}_x$ estimates.
- `SE of weighted mean for` β_x: Standard error of inverse-variance-weighted mean of $\hat{\beta}_x$ estimates.

6.2 Analysis

This section describes our procedure for conducting a meta-analysis of regression estimates from papers that examine biased inference and biased use of priors in sequential updating problems. We list the eight papers used in the meta-analysis above in section 6.1. All papers run a regression similar to the specification in chapter Eq. (4.21), reproduced here:

$$\ln\left(\frac{\pi(A|s_1, s_2, ..., s_t)}{\pi(B|s_1, s_2, ..., s_t)}\right) = \beta_0 + \beta_1 \ln\left(\frac{p(s_t|A)}{p(s_t|B)}\right) + \beta_2 \ln\left(\frac{\pi(A|s_1, s_2, ..., s_{t-1})}{\pi(B|s_1, s_2, ..., s_{t-1})}\right) + \eta_t. \tag{4.21}$$

In this model, β_1 captures the biased use of likelihoods, and β_2 captures the biased use of priors.

The estimates we use are reported in the tab titled "Updating_regression_estimates," in the spreadsheet `updating_regression_estimates.xlsx`, which can be downloaded on Daniel Benjamin's website (as of the date of this document: `https://www.danieljbenjamin.com/handbook-chapter`). Note that we use estimates from experiments II and III in Grether (1992). We do not include the estimates from Charness and Dave (2017) in the meta-analysis because they do not report standard errors.

Five papers do not report estimates in the exact form of Eq. (4.21), so we first transform the regression outputs from these papers. The tab titled "Updating_regression_estimates" lists the estimates as reported in Barron (2016), Buser et al. (2017), Coutts (2017), Gotthard-Real (2017), and Mobius et al. (2014). In those papers, the likelihood term is split into two indicator variables: one for an affirmative signal (signal that favors state A), and one for a negative signal (favors state B). Since we are interested in a single measure of biased inference, we combine these into one estimate by computing the inverse-variance-weighted average of the two coefficients. As an example, the reported estimates for Coutts (2017) are 0.576 (SE = 0.051) for the affirmative signal indicator and 0.812 (SE = 0.051) for the negative signal indicator. We calculate the combined estimate as

$$\frac{\frac{\hat{\beta}_{\text{aff}}}{\sigma^2_{\text{aff}}} + \frac{\hat{\beta}_{\text{neg}}}{\sigma^2_{\text{neg}}}}{\frac{1}{\sigma^2_{\text{aff}}} + \frac{1}{\sigma^2_{\text{neg}}}} = \frac{\frac{0.576}{0.051^2} + \frac{0.812}{0.051^2}}{\frac{1}{0.051^2} + \frac{1}{0.051^2}} = 0.694,$$

and the standard error for the combined estimate is calculated as

$$\sqrt{\frac{1}{\frac{1}{\sigma^2_{\text{aff}}} + \frac{1}{\sigma^2_{\text{neg}}}}} = \sqrt{\frac{1}{\frac{1}{0.051^2} + \frac{1}{0.051^2}}} = 0.036.$$

In my initial calculations, I mistakenly calculated incorrect standard errors for the combined estimates. Specifically, I used σ_{aff} and σ_{neg} instead of σ^2_{aff} and σ^2_{neg} in the standard error formula above (although the point estimates were calculated correctly). This inflated the standard errors for the combined estimates; for example, using the results from Coutts (2017): $\sqrt{\frac{1}{\frac{1}{\sigma_{\text{aff}}} + \frac{1}{\sigma_{\text{neg}}}}} = \sqrt{\frac{1}{\frac{1}{0.051} + \frac{1}{0.051}}} = 0.160 > 0.036$. Hence, the combined estimates from Barron (2016), Buser et al. (2017), Coutts (2017), Gotthard-Real (2017), and Mobius et al. (2014) were underweighted (due to their inflated standard errors) relative to the estimates from the two remaining papers in the meta-analysis. The incorrect meta-analysis estimates are reported in the chapter, although the qualitative conclusion is the same. We report the correct meta-analysis estimates below and in the spreadsheet `updating_regression_estimates.xlsx`. The tab titled "Meta-analysis" in the spreadsheet also contains the correct combined estimates and standard errors for the five papers that used affirmative and negative signals.

After deriving a single estimate for biased inference ($\hat{\beta}_1$) and a single estimate for biased use of priors ($\hat{\beta}_2$) from each paper, the next step is to meta-analyze the 8 estimates for each parameter. We meta-analyzed the estimates for each parameter by calculating the inverse-variance-weighted mean. This is analogous to how we combined the estimates for affirmative and negative signals. Weights for each estimate are shown in column D of the tab titled "Meta-analysis." The final meta-analyzed estimates are shown in column F, and standard errors are reported in column G. In the chapter, we report that the inverse-variance-weighted mean of the $\hat{\beta}_1$ estimates is 0.53 (SE = 0.012). However, this is incorrect due to the error described above. The correct inverse-variance-weighted mean of the $\hat{\beta}_1$ estimates is 0.38 (SE = 0.007).