How do scientists update their beliefs? An investigation of **N scientists engaged in replication research**

John Protzko*, University of California, Santa Barbara
Alex Holcombe, University of Sydney
Daniel J. Simons, University of Illinois

***Corresponding Author:** protzko@gmail.com

Abstract

How, if at all, do scientists adjust their beliefs based on observed evidence? Few studies have examined how beliefs change over the course of a study, from developing the protocol to collecting and analyzing original data. **N** scientists who were contributing to one or more of six multi-lab replication projects were asked to estimate their degree of belief in the original effect and to estimate the true effect size at three timepoints: before data collection, after learning their own results, and after learning the results of all of the other laboratories contributing replication studies. We examine how beliefs changed in response to data, whether scientists are disproportionately influenced by their own study results relative to those of other labs, and whether people optimally update their prior beliefs according to the observed evidence.

Keywords: Metascience, Beliefs; Scientists' beliefs; Conditionalization; Replications; Open Science

How do scientists update their beliefs? An investigation of **N** scientists engaged in replication research

Walt Whitman described the scientific spirit as “the holding off, the being sure but not too sure, the willingness to surrender ideas when the evidence is against them” (Traubel et al., 1915; p. 101). This view of the scientific enterprise—evidence steering beliefs—runs headlong into an apparently contradictory view: Scientists find evidence confirming their beliefs. A believer works hard to find consistent evidence, but a skeptic subtly tweaks their experimental choices to contradict the hypothesis. If belief follows from evidence, but the evidence produced by a scientist conforms to that scientist’s beliefs, to what extent does Whitman’s ideal of the scientific spirit match reality in scientific practice?

Most astronomers, in one study, barely budged in their beliefs about the formation of the moon despite evidence from the moon rocks supplied by the Apollo moon landings (and, any observed changes in belief dissipated within a few months; Mitroff, 1974; Mitroff & Mason, 1974). Perhaps unsurprisingly, *prominent* astronomers were even more likely to stick with their prior beliefs and were *less* influenced by the evidence. Similarly, the initial claims that a meteor strike precipitated the mass extinction of the dinosaurs, based on evidence of anomalous concentrations of an extremely rare substance in clay from the time of the extinction, was met with derision and largely ignored (Raup, 1986; Clemens, 1994).

If prior beliefs are sticky in the face of evidence, perhaps researcher beliefs influence not just how they interpret evidence but the evidence they report. Even the measurement of physical constants can change gradually rather than abruptly in the presence of measurement noise and uncertainty. In a review of issues in the measurement of physical constants, experimental results

at a certain time period tend to cluster together; with the following explanation (reportedly provided by E. O. Lawrence, cited in Birge, 1957):

“In any highly precise experimental arrangement there are initially many instrumental difficulties that lead to numerical results far from the accepted value of the quantity being measured...the investigator searches for the source or sources of such errors, and continues to search until he gets a result close to the accepted value. Then he stops! ... In this way one can account for the close agreement of several different results and also for the possibility that all of them are in error by an unexpectedly large amount.” (Birge, 1957, p. 51)

In other words, researchers do not necessarily update their beliefs strictly in accordance with Bayesian principles. Instead, their observations can be theory-laden (Brewer & Lambert, 2001), with discrepant results discounted unless they are close enough to the expected ones.

One study exploring how different analysis strategies affect the interpretation of the same data set also examined belief updating among the participating researchers. Twenty-nine teams of researchers were asked to report their belief (on a 1 to 5 scale where 1 is very unlikely and 5 is very likely) that soccer players with darker skin tones receive more penalties in matches. They did so at the start of the study, after they submitted their proposed analysis strategy, after they had conducted their own analysis, and again after learning about the other teams' analysis strategies and results (Silberzahn et al., 2018). Overall, initial subjective beliefs declined after teams submitted their analysis plans, returned to initial levels after teams had conducted their planned analyses, and then converged somewhat after learning about and discussing the approaches and results of other teams. After discussion, and after seeing the results from a variety of analyses, researchers came closer to agreement about the hypothesis.

What is unclear, however, is how this process of belief updating plays out when researchers are engaged in original data collection. Each sample of data is drawn from a larger

population and provides an estimate of the “True” effect. But samples vary. Even when holding the analysis strategy constant, different researchers running a study using the same protocol will observe different effect sizes due to sampling (and potentially to other uncontrolled differences between the samples or setting; i.e., true heterogeneity). Do researchers update their beliefs based on their observed results? Do they appropriately incorporate the possible noisiness of their estimate when faced with additional data? In this paper, we investigate a number of open questions about how originally collected data and beliefs affect one another among practicing scientists:

Question 1: Are prior beliefs related to the size of the effects a researcher finds in their own research? By recording researchers’ initial beliefs before data collection, we can test whether belief in an effect relates to the strength of the effect they observe (e.g., see Alcock, 2003).

Question 2: Do scientists adhere to Bayes’ rule when they update their beliefs in response to their own results? Bayes rule specifies how an ideal researcher *should* update their prior beliefs in response to new data. We will assess how well Bayes rule captures actual changes in belief. If they do not update according to Bayes rule, how do they update their beliefs?

Question 3: Once researchers have updated their beliefs based on their own results, do they further update their beliefs according to Bayes’ rule after learning the results of similar studies from other researchers? That is, do researchers treat their own evidence as privileged and weight it more heavily than Bayes rule requires when faced with additional data? A failure to do so might reflect a misunderstanding of sampling error when estimating effects or a belief that the researcher’s own data is superior to data from other researchers.

Question 4: What is the covariance among researchers' degrees of belief, belief that the true effect size is non-zero, and belief in the magnitude of the true effect size? Beliefs can be measured in a continuous, "Bayesian" way using scale of subjective beliefs in a hypothesis (0-100%), as a judgment of whether an effect is 'real' or 'non-zero', or as an estimate of the size of the "true" effect (Good, 1971; Gigerenzer & Marewski, 2015). By collecting different measures of belief, we can understand how they covary and how their covariance changes with the new evidence.

Experimental replication research provides a rich source of information about belief updating in the face of evidence. Starting in 2013, psychologists began conducting multi-lab direct replications of prominent experimental findings using a new article type: Registered Replication Reports (RRR; Simons et al., 2014). For these projects, a lead team of researchers develops a protocol to replicate an original finding, and that plan is vetted either by the original researchers or by surrogates who share their theoretical perspective. Then, many other laboratories—at least 10 and as many as 45—follow that vetted protocol, conducting independent replications using the same procedures, with all results published collectively regardless of the outcome. The individual contributing laboratories remain blind to the outcomes from the other laboratories until all data collection has ceased and all individual laboratories have analyzed their own results (using provided analysis scripts). Thus, researchers have an opportunity to adjust their beliefs after learning their own result and again after learning the results from many other laboratories conducting the same study, without the explicit interference that might come from knowing what other researchers believe.

Scientists joining these projects vary in their beliefs about the original effects; they are not all 'true believers' or "skeptics." Consequently, the RRRs provide a fertile testbed to study

how practicing scientists update their beliefs during the process of evidence collection. The use of a standardized protocol allows a test of the idea that prior beliefs alone will moderate the size of the observed effect, with proponents more likely than skeptics to observe a larger effect (Question 1 above). Such experimenter effects have been proposed to explain differential results in fringe fields of research (see Alcock, 2003 for an overview) and occasionally for more mainstream replication research (e.g. Wilson, 2014), but direct evidence for this form of moderation is lacking.

By surveying researchers about their belief in the original hypothesis before data collection, after analyzing their own data and again after learning the meta-analytic results from all contributing laboratories, we can measure both the effects of beliefs on evidence and the effects of evidence on beliefs.

[NOTE: THIS STAGE-1, PRE-DATA REPORT IS WRITTEN AS IF DATA HAD ALREADY BEEN ANALYZED SO THAT THE PROSE WILL BE THE SAME AFTER DATA ANALYSIS. THE RESULTS ARE DESCRIBED IN PAST TENSE WITH CONDITIONAL LANGUAGE OPTIONS FOR DIFFERENT POSSIBLE OUTCOMES. WE USE COMMENTS IN SQUARE BRACKETS WITH ALL-CAPS, RED TEXT (LIKE THIS NOTE) TO FLAG NOTES ABOUT THE PROSE THAT WILL NOT BE IN THE FINAL MANUSCRIPT.]

Methods

Participants

We contacted all researchers involved in conducting replications in six RRR projects (Table 1) at three separate points: before data collection, after they had completed data collection for their own study, and after they learned the meta-analytic results from all contributing labs. At

each timepoint, each participant reported their degree-of-belief in the hypothesis (using a 0-100% scale, conforming to classical Bayesian measurement), their confidence that the true difference is non-zero (using a 0-100% scale), and their belief about the size of the ‘true’ effect (using units specific to the study; see Table 1). Although all six RRR projects are now in print, we did not inspect or analyze the belief judgments and ratings prior to provisional acceptance of the stage-1 registered report version of this manuscript. **[WE HAVE WRITTEN THIS SENTENCE SO THAT IT WILL NOT NEED TO CHANGE BETWEEN STAGE 1 AND STAGE 2]**

All participants were practicing scientists ranging in experience from tenured professors to undergraduate students. These scientists were invited by the editor of the RRR projects (DJS) to participate and were informed that it was an additional study, independent of the RRR, and that participation was optional. In some cases, multiple scientists from the same team provided responses.

Data collection for RRR5 and RRR6 was bundled as part of a battery of tests in a single session to match the testing setting and conditions used for the original study under investigation in RRR5. RRR6 was added as the first task in that battery. We treat them separately here, although the data collection link was the same for all participants, meaning that the participants in our study were the same for RRR5 and RRR6.¹

Materials

Included RRR Studies.

¹ Some researchers contributed to multiple RRRs under study here. RRR5 and RRR6 have completely overlapping samples, but the overlap in participating researchers for the other RRRs is partial. This overlap of participants (i.e., non-independence) is handled analytically by using a robust maximum likelihood estimator for all analyses and dummy variables for the different surveys, given the low number of overlapping researchers (see Muthén & Muthén, 2017).

RRR1: Piovesan et al., 2017

When asked how much money they wished to contribute to a public pot (that ostensibly would be split evenly among four contributors), people contributed more when responding quickly than when forced to delay their response (Rand et al., 2012). 21 labs attempted to replicate this finding.

RRR2: Wagenmakers et al., 2016

When people were asked to respond using a pen held in their mouth, they rate cartoons as funnier if they gripped the end of the pen in their teeth (forcing a smile) than when they held it with their lips (forcing more of a pout; Strack et al., 1988). 17 labs attempted to replicate this finding.

RRR3: O'Donnell et al., 2017

Imagining daily life as a professor led to better performance on a general knowledge trivia task than did imagining life as a soccer hooligan (Dijksterhuis & van Knippenberg, 1998). 40 labs attempted to replicate this effect.

In refining the protocol for this study, the original author conducted additional pilot studies and reported that the effects were present only for men. Participants in our study of belief updating were informed that: “Follow-up studies conducted based on the original 1998 protocol reported sex differences in the effects of priming with professor or hooligan. Whereas men showed better performance on a trivia test when primed with professor than when primed with hooligan—women did not. Below we will ask you to report your beliefs about the effects of priming with professor and hooligan, and we ask for separate responses for men and for women. If you believe men and women will show the same effect, please give the same answers for men and for women.” Participants then gave their responses about the effects for women and for men in random order.

RRR4: Cheung et al., 2016

Participants reacted less negatively to hypothetical betrayals if they wrote open-ended responses to prompts emphasizing their dependence on and commitment to their partner (high commitment) than if they had written about their independence and lack of commitment to their partner; Finkel et al., 2002). 16 labs attempted to replicate this result.

This study included four dependent variables: exit responses (e.g., trying to get even); neglect responses (e.g., giving their partner the cold shoulder); voice responses (e.g., discussing the betrayal with their partner to understand why it happened); and loyalty responses (e.g., maintaining the view that their partner is a good person or believing that their partner's behavior was unintentional). Participants were asked for their beliefs about all four DVs in random order. Participants read: "The following three questions are about [exit responses (e.g., seeking vengeance)/neglect responses (e.g., giving the cold shoulder)/voice responses (e.g., partner talking about it)/loyalty responses (e.g., continuing to support the partner despite dissatisfaction)]".

RRR5: Verschuere et al., 2018

Being asked to recall the 10 commandments resulted in less cheating on a problem-solving test than did recalling 10 books they read in high school (Mazar et al., 2008). 25 labs attempted to replicate this finding.

RRR6 McCarthy et al., 2018

Unscrambling scrambled words that included hostile actions (e.g. "break his nose") caused people to interpret ambiguous behaviors (e.g. insisting a waitress replace all the silverware because it was dirty) and ambiguous descriptions of an actor as more hostile (Srull & Wyer, 1979). There were two dependent variables, how hostile the actor is and how hostile the

behaviors are. We asked participants about their beliefs for both dependent variables. 26 labs attempted to replicate these effects.

Procedures

The questions about beliefs for each RRR were tailored to the study content (see Table 1). Prior to reporting their belief, each participant provided their name and university (so that we could track their responses across timepoints) and answered two questions to verify their knowledge of the study results: “Have you begun collecting data for the Registered Replication Report (RRR) on the effects of ...?” [No/Yes]. And [“Have you analyzed the data?” [No/Yes]. Before providing their judgments, each participant was told: “As part of this Registered Replication Report (RRR) we would like to know your beliefs about the research you will be conducting.”

Participants were then asked about their degree-of-belief (e.g. how much do you believe the hypothesis that ____), their belief that the ‘true’ effect size is non-zero, and their belief about the size of the ‘true’ effect size in study-specific units (e.g. number of additional questions solved, % of additional money donated etc.). The question wording was adapted for each study. This study was considered to be exempt from the need for IRB approval by the UCSB Office of Research on Human Subjects.

Testable Hypotheses

This investigation allows us to answer a number of questions based on the study outcomes and each respondent’s degree of belief in the effect, their belief that the true difference is greater than zero, and their belief about the true effect size. All hypotheses and analysis code were pre-registered prior to inspecting the data (<https://osf.io/dq38r/>).

[AS NOTED ABOVE, DESCRIPTIONS ARE IN PAST TENSE TO MINIMIZE CHANGES BETWEEN STAGE 1 AND STAGE 2 SHOULD THE STAGE 1 PAPER BE PROVISIONALLY ACCEPTED]

Question 1: Do priors influence the observed effects?

The idea that researchers who are skeptical of a finding might be less likely to replicate it is often invoked in response to replication failures. We used structural equation modeling to test whether reported prior beliefs (all three measures) predicted the observed effect size they obtained in the RRR study. The three belief measures were allowed to covary in the model, and we included five dummy variables for the six RRRs to account for differences in the estimates. We predict there will be no relationship between any of the three prior beliefs and the observed effect size $< p = .1$. This p -value is chosen because any p -value $\geq .1$, given our sample size, would represent evidence in favor of the null when converted into a Bayes Factor. We do, however, predict that researchers will hold a prior beliefs that the ES will be smaller than the originally reported ES. This corresponds to researchers believing that effect sizes reported in the published literature are inflated.

Question 2: Do scientists update their priors in response to their own results according to Bayes' rule?

Bayes' rule and conditionalization give rise to what are referred to as posteriors in response to learning new evidence. In this study, all participants provided their priors before data collection, then indicated their first posteriors after analyzing their own data. We are thus able to tell whether and how scientists are updating their beliefs in response to their own evidence collection.

We test this hypothesis by constructing a model of the ‘ideal’ Bayesian updater whose posterior matches their prior multiplied by the Bayes Factor of their own research. As all of these variables are known quantities, the observed posterior beliefs can be tested against what their beliefs ‘should’ be given Bayes’ rule, although this figure is only correct if the subjective probability of all the auxiliary hypotheses are unchanged. We subtracted participants’ first posteriors (after their own data collection) from what their posteriors would be given their priors and the Bayes Factor of the evidence they obtained. This should result in a distribution of residual posteriors with a mean of 0 (indicating ideal updating on average) and variance σ (which would correspond, given ideal updating, only to measurement error), given ‘ideal’ updating. A simple regression of residual posteriors with no predictors can test whether the residual posteriors conform to this ‘ideal’ updater.

Three possibilities exist when observing these data: 1) residual posteriors being negative (corresponding to an over-correction in instances of null-favoring evidence) and non-0 (statistically significantly different from 0 at $p < .05$); 2) residual posteriors being positive (corresponding to an under-correction in instances of null-favoring evidence) and non-0 (statistically significantly different from 0 at $p < .05$); and 3) residual posteriors being statistically not different from 0.

Over-correction would occur when practicing scientists are not ideal updaters in light of evidence (e.g., observing a non-significant effect and shifting to disbelief in the hypothesis rather than updating priors). Posterior matching, where people abandon their priors in light of evidence and simply match their posteriors to what is observed (see Icard, 2017), also would lead to a pattern of overcorrection. Those who had higher priors, upon seeing negative evidence, would place their posteriors closer to the data than ideal updating would predict. Under-correction

would occur when researchers engage auxiliary hypotheses to retain their beliefs in the light of negative evidence (Duhem, 1954, see also Strevens, 2001). This could involve attributing negative evidence to something outside the bounds of the target hypothesis (e.g., characteristics of the researchers, context of administration, subtle shifts in the materials). We predict the data will show a pattern of over-correction for first priors.

Question 3: Once researchers have updated their beliefs based on their own results, how do they further update their beliefs after learning the results from all contributed studies according to Bayes' rule?

This analysis asks the same question as Question 3, using participants' first posteriors as their new priors before they learn the results from all of the other labs. For this analysis, the meta-analytic results are re-computed individually for each lab while excluding that lab's effect size. We test this hypothesis by calculating the Bayes Factor of all of the labs *except for* the home labs results (to prevent double updating on their own data), creating a residual second posterior as in the analysis for Question 3. We then analyze whether these second posteriors (residuals) are zero (which assumes an 'ideal' updater).

Question 4: What is the covariance among degrees of belief, belief that the true effect size is greater than zero, and belief in the true effect size?

We collected three variables about beliefs: (degree-of-belief, true difference > 0 , effect size). To address this question, we assess the covariance among these three belief variables by covarying them in a model for just the first timepoint (prior). We also test this covariance for posteriors after replicators analyze their own data and again after they learn the results of all replications. This will be done by comparing the overlap of 95% CI of the covariance estimates across the three timepoints. These analyses are descriptive.

References

- Alcock, J. (2003). Give the null hypothesis a chance: Reasons to remain doubtful about the existence of psi. *Journal of Consciousness Studies*, 10(6-7), 29-50.
- Anderson, M. S., Martinson, B. C., & De Vries, R. (2007). Normative dissonance in science: Results from a national survey of US scientists. *Journal of Empirical Research on Human Research Ethics*, 2(4), 3-14.
- Azoulay, P., Fons-Rosen, C., & Zivin, J. S. G. (2015). Does science advance one funeral at a time? (No. w21788). *National Bureau of Economic Research*.
- Benjamin, D., Mandel, D. R., & Kimmelman, J. (2017). Can cancer researchers accurately judge whether preclinical reports will reproduce?. *PLoS biology*, 15(6), e2002212.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B (Methodological)*, 57(1), 289-300.
- Birge, R. T. (1957). A survey of the systematic evaluation of the universal physical constants. II. *Nuovo Cimento*, 6(S1), 39–67. doi:10.1007/bf02724770
- Brewer, W. F., & Lambert, B. L. (2001). The theory-ladenness of observation and the theory-ladenness of the rest of the scientific process. *Philosophy of Science*, 68(S3), S176-S186.
- Carlson, S. M., Shoda, Y., Ayduk, O., Aber, L., Schaefer, C., Sethi, A., ... & Mischel, W. (2018). Cohort effects in children's delay of gratification. *Developmental psychology*.
- Chang, H. (2012). *Is water H2O?: Evidence, realism and pluralism* (Vol. 293). Springer Science & Business Media.

- Cheung, I., Campbell, L., LeBel, E. P., Ackerman, R. A., Aykutoğlu, B., Bahník, Š., ... & Carcedo, R. J. (2016). Registered Replication Report: Study 1 from Finkel, Rusbult, Kumashiro, & Hannon (2002). *Perspectives on Psychological Science*, 11(5), 750-764.
- Clemens, E.S. (1994). The impact hypothesis and popular science: Conditions and consequences of interdisciplinary debate. In W. Glen (Ed.), *The mass-extinction debates: How science works in a crisis* (pp. 92–120). Stanford, CA: Stanford University Press.
- Cunningham, M. R., & Baumeister, R. F. (2016). How to Make Nothing Out of Something: Analyses of the Impact of Study Sampling and Statistical Interpretation in Misleading Meta-Analytic Conclusions. *Frontiers in Psychology*, 7.
- DellaVigna, S., & Pope, D. (2018). Predicting experimental results: who knows what?. *Journal of Political Economy*, 126(6), 2410-2456.
- Dijksterhuis, A. & van Knippenberg, A. (1998). The relation between perception and behavior, or how to win a game of trivial pursuit. *Journal of Personality and Social Psychology*, 74(4), 865-877.
- Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., ... & Johannesson, M. (2015). Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences*, 112(50), 15343-15347.
- Duhem, P. (1954). *The Aim and Structure of Physical Theory*. Translated by P. P. Wiener. Princeton University Press, Princeton, NJ.
- Ebersole, C. R., Axt, J. R., & Nosek, B. A. (2016). Scientists' reputations are based on getting it right, not being right. *PLoS biology*, 14(5), e1002460.

- Finkel, E. J., Rusbult, C. E., Kumashiro, M., & Hannon, P. A. (2002). Dealing with betrayal in close relationships: Does commitment promote forgiveness? *Journal of Personality and Social Psychology*, 82, 956–974.
- Fiske, S. T. (2016). Mob rule or wisdom of crowds. *APS Observer*, 1-2.
- Forsell, E., Viganola, D., Pfeiffer, T., Almenberg, J., Wilson, B., Chen, Y., ... & Dreber, A. (2018). Predicting replication outcomes in the Many Labs 2 study. *Journal of Economic Psychology*.
- Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). A Response to the Reply to Our Technical Comment on ‘Estimating the Reproducibility of Psychological Science.’
http://projects.iq.harvard.edu/files/psychologyreplications/files/gkpw_response_to_osc_rebutal.pdf?m=1457031863
- Gigerenzer, G. and Marewski, J. N. (2015) Surrogate science: The idol of a universal method for scientific inference. *Journal of Management*, 41(2), 421–440.
- Good, I. J. (1971). 46656 varieties of Bayesianism. *American Statistician*, 25: 62-63.
- Icard, T. F. (2018). Bayes, bounds, and rational analysis. *Philosophy of Science*, 85(1), 79-101.
- 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.
- Mazar, N., Amir, O., & Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, 45, 633–644.
- McCarthy, R. J., Skowronski, J. J., Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., ... & Barbosa, F. (2018). Registered replication report on Srull and Wyer (1979). *Advances in Methods and Practices in Psychological Science*, 1(3), 321-336.

- McEvoy, J. G. (1997). Positivism, whiggism, and the Chemical Revolution: A study in the historiography of chemistry. *History of Science*, 35(1), 1-33.
- Mitroff, I. (1974). *The Subjective Side of Science*. Amsterdam: Elsevier
- Mitroff, I. I., & Mason, R. O. (1974). On evaluating the scientific contribution of the Apollo moon missions via information theory: A study of the scientist-scientist relationship. *Management Science*, 20(12), 1501-1513.
- Muthén, L.K. & Muthén, B.O. (2017). *Mplus User's Guide. Eighth Edition*. Los Angeles, CA: Muthén & Muthén.
- O'Donnell, M., Nelson, L. D., Ackermann, E., Aczel, B., Akhtar, A., Aldrovandi, S., & Bialobrzeska, O. (2017). Registered Replication Report: Dijksterhuis & van Knippenberg (1998) Multilab direct replication of: A variant of study 4 from Dijksterhuis, A. & van Knippenberg, A.(1998). The relation between perception and behavior, or how to win a game of, 1-38.
- Piovesan, M., Døssing, F. S., & Wengström, E. R. (2017). Registered Replication Report: Rand, Greene & Nowak (2012). *Perspectives on Psychological Science*.
- Pritschet, L., Powell, D., & Horne, Z. (2016). Marginally significant effects as evidence for hypotheses: Changing attitudes over four decades. *Psychological Science*, 27(7), 1036-1042.
- Protzko, J. (2017, September 13). Kids these days pre-print. Retrieved from psyarxiv.com/rs9hd
- Raup, D.M. (1986). *The Nemesis Affair: A story of the death of dinosaurs and the ways of science*. New York: W.W. Norton.
- Schmeidler, G. R. (1945). Separating the sheep from the goats. *Journal of the American Society for Psychical Research*, 39, 47-49.

- Shapin, S., & Schaffer, S. (2011). *Leviathan and the air-pump: Hobbes, Boyle, and the experimental life* (New in paper) (Vol. 32). Princeton University Press.
- Sharot, T., & Garrett, N. (2016). Forming beliefs: Why valence matters. *Trends in cognitive sciences*, 20(1), 25-33.
- Shea, W. R., & Artigas, M. (2003). *Galileo in Rome: The rise and fall of a troublesome genius*. Oxford University Press.
- Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., . . . Nosek, B. A. (2018). Many analysts, one data set: Making transparent how variations in analytic choices affect results. *Advances in Methods and Practices in Psychological Science*, 1, 337–356. doi:10.1177/2515245917747646.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological science*, 22(11), 1359-1366.
- Simons, D. J., Holcombe, A. O., & Spellman, B. A. (2014). An introduction to registered replication reports at perspectives on psychological science. *Perspectives on Psychological Science*, 9(5), 552-555.
- Srull, T. K., & Wyer, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, 37, 1660–1672.
- Strack, F. (2017). From Data to Truth in Psychological Science. A Personal Perspective. *Frontiers in psychology*, 8, 702.

- Strack, F., Martin, L. L., & Stepper, S. (1988). Inhibiting and facilitating conditions of the human smile: a nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology*, 54, 768-777.
- Strevens, M. (2001). The Bayesian treatment of auxiliary hypotheses. *The British Journal for the Philosophy of Science*, 52(3), 515-537.
- Veldkamp, C. L., Hartgerink, C. H., van Assen, M. A., & Wicherts, J. M. (2017). Who believes in the storybook image of the scientist?. *Accountability in research*, 24(3), 127-151.
- Verschuere, B., Meijer, E. H., Jim, A., Hoogesteyn, K., Orthey, R., McCarthy, R. J., ... & Barbosa, F. (2018). Registered replication report on Mazar, Amir, and Ariely (2008). *Advances in Methods and Practices in Psychological Science*, 1(3), 299-317.
- Wagenmakers, E. J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., Adams Jr, R. B., ... & Bulnes, L. C. (2016). Registered replication report: Strack, Martin, & Stepper (1988). *Perspectives on Psychological Science*, 11(6), 917-928.
- Wilson, T. (2014, June 15). *Is There a Crisis of False Negatives in Psychology?* Retrieved from <https://timwilsonredirect.wordpress.com/2014/06/15/is-there-a-crisis-of-false-negatives-in-psychology/>
- Wood, G. (1978). The knew-it-all-along effect. *Journal of Experimental Psychology: Human Perception and Performance*, 4(2), 345.