THEORETICAL REVIEW

The frequency of excess success for articles in *Psychological Science*

Gregory Francis

Abstract Recent controversies have questioned the quality of scientific practice in the field of psychology, but these concerns are often based on anecdotes and seemingly isolated cases. To gain a broader perspective, this article applies an objective test for excess success to a large set of articles published in the journal *Psychological Science* between 2009 and 2012. When empirical studies succeed at a rate much higher than is appropriate for the estimated effects and sample sizes, readers should suspect that unsuccessful findings have been suppressed, the experiments or analyses were improper, or the theory does not properly account for the data. In total, problems appeared for 82 % (36 out of 44) of the articles in *Psychological Science* that had four or more experiments and could be analyzed.

Keywords Statistical inference · Statistics · Probabilistic reasoning

It is widely recognized that a bias exists *across* articles in the field of psychology (Fanelli, 2010; Sterling, 1959; Sterling, Rosenbaum, & Weinkam, 1995). These studies have noted that approximately 90 % of published experiments are reported to be successful, which suggests that many unsuccessful experiments remained unpublished. However, it is not clear what such a bias means with regard to believing a specific reported experimental finding or theory. Bias across articles may merely reflect a desire among authors and journals to publish about topics that tend to reject the null hypothesis with

G. Francis (✉)
Department of Psychological Sciences, Purdue University, 703 Third Street, West Lafayette, IN 47906, USA
e-mail: gfrancis@purdue.edu

typical experimental designs and such a bias does not necessarily cast doubt on the findings or theories within any specific article. When judging the quality of scientific work, a finding of bias *within* an article is more important than bias across articles, because the presence of bias within an article undermines that article's theoretical conclusions. Recent investigations (Bakker, van Dijk, & Wicherts, 2012; Francis, 2012a, 2012b, 2012c, 2012d, 2012e, 2013a, 2013b; Renkewitz, Fuchs, & Fiedler, 2011; Schimmack, 2012) have used an objective bias analysis to indicate that some articles (or closely related sets of articles) in the field of psychological science appear to be biased. However, such analyses of individual studies do not indicate whether the appearance of bias within an article is rare or common in psychology.

Partly to estimate the within-article bias rate, I have applied the bias analysis to articles published over the last several years of the journal *Psychological Science*, which is the flagship journal of the Association for Psychological Science, has enormous reach to scientists and journalists, and presents itself as an outlet for only the very best research in the field. Although perhaps the journal is not representative of the field of psychological science in general, it would be valuable to know what proportion of findings (and which specific findings) appear to be biased in a journal that seeks to publish the field's best work. This article summarizes the analyses of the investigated articles from the journal; the article selection criteria and a full description of the analyses (and accompanying computer code) are provided in the Electronic Supplementary Material.

In lay usage, the term "bias" means unfair prejudice, but that is not the intended meaning in this article. Here, the term "bias" is used in a statistical sense—namely, to indicate that the frequency of producing a significant result, or an effect's magnitude, is systematically overestimated. A prejudicial bias by authors may produce a statistical bias, but it is not necessary, because statistical bias can be introduced despite good

intentions from researchers. Moreover, it is not necessary to know the exact cause or source of statistical bias for a reader to be skeptical about the published empirical findings and their theoretical conclusions.

Analyzing the probability of experimental success

The bias analysis is based on the "test for excess significance" (TES) proposed by Ioannidis and Trikalinos (2007). The test contrasts estimates of experimental power with the reported frequency of significant findings. Suppose that an article describes five experiments that are presented as empirical support for a theory. The experiments could be direct replications, conceptual replications, converging evidence, or explorations of different parts of the theory. Further suppose that every experiment produces a successful outcome (e.g., rejects the null hypothesis). Such a positive outcome might commonly be considered vindication of the theory, but this view is unwarranted without considering the statistical power of the experiments.

Due to random sampling, even when an effect is nonzero, some experiments will generate samples that produce test statistics and p values that do not satisfy the criterion for statistical significance. *Power* is the probability of rejecting the null hypothesis for a specified nonzero effect. Suppose that the power for each of five reported experiments is .6. Following Ioannidis and Trikalinos (2007), the probability that all five of such experiments would reject the null hypothesis is the product of the power values, $P_{TES} = .6^5 = .078$, which indicates that the observed pattern of consistent rejections should be rare. A common criterion in these kinds of investigations is a probability of .1 (Begg & Mazumdar, 1994; Francis, 2012e; Ioannidis & Trikalinos, 2007). A power probability below this criterion suggests that the experiment set should be considered biased, and the results and conclusions be treated with skepticism.

The TES is sensitive to the distribution of reported statistics in a way that is similar to other tests for publication bias (Egger, Davey Smith, Schneider, & Minder, 1997) and *p*-hacking (Simonsohn, Nelson, & Simmons, in press). However, these alternative tests require a set of findings to have fixed effect sizes or independent p values, which is often not the case for a series of experiments in psychology. The TES's focus on experimental power allows it to consider inhomogeneous effect sizes and statistical dependencies between experimental outcomes, provided that the appropriate statistical information about the dependencies is published.

The TES can be generalized to consider the probability of experimental success, including a pattern of significant and nonsignificant results. For the remainder of this article, the term "TES" will stand for the test for excess *success*, which has the same acronym and follows the same principles as the

test for excess significance that was proposed by Ioannidis and Trikalinos (2007). If experimental success is defined as statistical significance, then the tests are identical. Table 1 describes the properties of a hypothetical article with five experiments that were interpreted as providing unanimous support for a theoretical conclusion. The middle two columns of Table 1 summarize the statistics and hypothesis tests that were used to support the theoretical claims. Although the statistics and analyses reported in Table 1 are artificial, they reflect the kinds of statistics, hypotheses, and analyses that are common in *Psychological Science* articles. For all of the estimated success probabilities, it was assumed that the original hypothesis tests were appropriate for the reported data (e.g., the data were randomly sampled from normal distributions with a common variance).

The first row describes the properties of Experiment 1, which was a between-subjects design having three different groups. A successful outcome was for every pair of means to be significantly different. The reported statistics satisfy this definition of success (all ps less than .05), using an analysis of variance (ANOVA) and contrasts. The last column of Table 1 provides the post hoc estimated probability of success for this type of experiment with these samples sizes. The probabilities are calculated by supposing that the reported statistics accurately reflect the population parameters. One hundred thousand simulated experiments with the same sample sizes were then generated by drawing random samples from the simulated populations, running the hypothesis tests, and computing the proportion of tests that satisfy the desired outcome (rejecting the null hypothesis). The ANOVA and the contrast between Conditions 1 and 3 are very likely to reject the null. However, the contrasts between Conditions 1 and 2 and between Conditions 2 and 3 have only modest power values. Moreover, the joint probability of all four hypothesis tests being successful is only .435. This low probability reflects the multiple constraints that are imposed by the definition of the experiment's success.

Experiment 2 was a between-subjects design, with Condition 1 corresponding to an experimental condition and Conditions 2 and 3 acting as controls. A successful outcome was that the experimental condition would differ from each of the controls, which would not differ from each other. For the given statistics, all of the tests are successful. The last column indicates the probability of success (estimated with one hundred thousand simulated experiments). For the comparison of the control conditions, success is defined as producing a nonsignificant result. Although the probabilities are relatively high for two of the tests, the joint success probability is only slightly above one half. Again, the multiple constraints on the pattern of results, including a predicted nonsignificant result, reduce the probability of the experiment's success.

Experiment 3 was a between-subjects design that measured two variables, *X* and *Y,* for an experimental group and a control

Table 1 Statistical properties, hypotheses, and estimated probabilities of success for a hypothetical set of five experiments

	Statistics	Hypotheses	Probability of Success
Exp. 1	$n_1 = 15, n_2 = 16, n_3 = 19$ $\overline{X}_1 = 17.2, \overline{X}_2 = 15.3, \overline{X}_3 = 13.4$ $s = 2.15$	ANOVA	.996
		$\mu_1 \neq \mu_2$	.674
		$\mu_1 \neq \mu_3$	.999
		$\mu_2 \neq \mu_3$	.723
		Joint	*.435*
Exp. 2	$n_1 = 23, n_2 = 22, n_3 = 20$ $\overline{X}_1 = 5.5, \overline{X}_2 = 4.6, \overline{X}_3 = 4.9$ $s = 0.85$	$\mu_1 \neq \mu_2$	.934
		$\mu_1 \neq \mu_3$	.617
		$\mu_2 = \mu_3$	.801
		Joint	*.532*
Exp. 3	$n_1 = 25, n_2 = 25$ $\overline{X}_1 = 5.5, \overline{X}_2 = 4.3, s_X = 1.5$ $\overline{Y}_1 = 11, \overline{Y}_2 = 8, s_Y = 4.25$	$\mu_{X1} \neq \mu_{X2}$	.779
		$\mu_{Y1} \neq \mu_{Y2}$	.673
		Joint	*.673*
Exp. 4	$n_1 = 35, n_2 = 36$ $\overline{X}_1 = 24.1, \overline{X}_2 = 20.1, s_X = 6.9$ $\overline{Y}_1 = 33.1, \overline{Y}_2 = 29.3, s_Y = 6.8$ $r_{XY} = .23$	$\mu_{X1} \neq \mu_{X2}$	.672
		$\mu_{Y1} \neq \mu_{Y2}$	.639
		Joint	*.460*
Exp. 5	$n_{1A} = 56, n_{2A} = 53, n_{1B} = 52, n_{2B} = 51$ $\overline{X}_{1A} = 15.0, \overline{X}_{1B} = 20.5$ $\overline{X}_{2A} = 14.2, \overline{X}_{2B} = 15.5$ $s = 6.75$	ANOVA interaction	.615
		$\mu_{1A} \neq \mu_{1B}$	.989
		$\mu_{2A} = \mu_{2B}$	.836
		Joint	*.587*
P_{TES}			.042

Note: A success probability in italics indicates the estimated joint probability of success for the full set of hypotheses

group. The predicted outcome was that both variables would produce a significant difference between the groups. The power value varies across the two tests, and with the given statistics it is not possible to directly estimate the joint power. If the two measured variables were independent, the joint power could be estimated by multiplying the powers of the separate variables. However, because the X and Y scores are correlated (they come from a common set of participants), this product will likely underestimate the true power. A conservative approach, which likely overestimates the true power, is to take the lower of the two power values as an upper limit on the joint power.

Experiment 4 had a design similar to that of Experiment 3, in that it measured two variables for both an experimental and a control group, and in that the predicted outcome was for both variables to reject the null. Unlike Experiment 3, the reported statistics include the correlation between the two measures, and knowing the correlation enables the joint power to consider the probability of both analyses producing significant outcomes. The estimated joint power, derived from simulated experiments, is smaller than either of the individual powers, but larger than their product.

Experiment 5 had a 2 × 2 between-subjects design, and the expected outcome was a significant interaction coupled with a significant difference between scores in Condition 1, but a nonsignificant difference between scores in Condition 2. The probability of success is mostly dominated by the probability

of the significant interaction, with the other tests only modestly reducing the success probability.

The probability of five experiments like these producing uniformly successful outcomes is the product of the joint success probabilities, which is .042. This probability is so low that readers of such an article should be skeptical that the experiments were run properly, analyzed properly, and interpreted correctly relative to the theoretical ideas.

The logic of the TES analysis is similar to traditional hypothesis testing, with the null hypothesis being that the experiments were run properly and without bias. With that premise, one estimates the probability of producing successful experimental outcomes that are equivalent to, or more extreme than, the observed outcomes. If this probability is low, then the analysis suggests that the null is not viable: The set of experiments appears to be biased or flawed in some way. As Francis (2013b) noted, the analysis checks for the consistency of outcomes across a set of experiments. If the experiments are run properly, analyzed properly, and published fully, then the rate of success should be consistent with the estimated probabilities of experimental success. It is also appropriate to think of P_{TES} as the estimated probability that a direct replication of a set of experiments with the same sample sizes would produce results at least as successful as those that were published.

The TES analysis deviates from traditional hypothesis testing in that the .1 criterion for P_{TES} does not define the

frequency of rejecting the null hypothesis when it is really true. Precise control of the Type I error rate for a judgment of bias requires knowing the true probability of experimental success, but almost every practical application of the TES analysis must estimate experimental success using the reported data. Francis (2013b) showed that for two-sample t tests, using the $P_{TES} \leq .1$ criterion usually produces a Type I error rate of around .01, so the test is conservative.

Applying the TES analysis to articles in *Psychological Science*

I downloaded all 951 articles published in *Psychological Science* during the years 2009–2012. Articles were considered for the TES analysis only if the success probabilities could be estimated for at least four experiments, because the analysis needs at least that many experiments in order to have much chance of detecting any type of bias (see the simulations in Francis, 2012e, 2013b). The count of experiments included subdivisions (e.g., Studies 1a and 1b) and occasionally included an experiment that was described outside of a formal header (e.g., an experiment summarized in the conclusion section as a follow-up). A total of 44 articles had four or more experiments that could provide success probability estimates. The supplemental material further discusses the article selection and describes the success probability calculations for each article. The supplemental material also includes source code for the simulation-based probability calculations.

Table 2 lists the P_{TES} value for each of the 44 analyzed articles. The most striking property is that 36 out of the 44 articles have a P_{TES} value smaller than the .1 criterion. Despite the conservative nature of the TES analysis, bias appears to be present for 82 % of the articles in *Psychological Science* with four or more experiments having designs and reported statistics that can enable success probability calculations. This high rate is troubling, because biased articles do not provide appropriate scientific arguments for their derived theories.

The TES analysis cannot identify the source of bias for an article that appears to have too much success, but publication bias and questionable research practices (John, Loewenstein, & Prelec, 2012; Simmons, Nelson, & Simonsohn, 2011) are plausible explanations. For an apparently biased experiment set, it is possible that the reported experiments are part of a larger set containing unsuccessful studies that are relevant to the theoretical ideas but are not reported. It is also possible that the experiments utilized inappropriate data collection or analysis methods to produce p values below the .05 criterion. Such manipulations tend to overestimate two variables that scientists care about: effect size and replication rate (Francis, 2012b, 2012c, 2013b, 2013c; Lane & Dunlap, 1978). As a result of these biases, readers cannot estimate the true effect sizes with the provided data (they might be zero) or predict the

outcome of future studies. It is important to note that the presence of bias should not be taken as an indication that the theoretical claims are necessarily wrong. The proper interpretation is that the presented justification for the theoretical claims is invalid. In some cases, it may be possible to interpret some subset of the reported findings as being valid (Francis, 2013c), but such a reanalysis requires subject matter expertise, and possibly a new theoretical interpretation of the findings.

A skeptical attitude toward the validity of a biased experiment set mirrors how researchers typically think about bias within the context of a single experiment. Scientists would not accept a conclusion based on data from an experiment in which a researcher excluded participants who did not show a desired result or insisted that each participant perform a task until showing a desired result. Similar kinds of biases exist at the level of experiment sets, and they can create experiment sets in which the reported rate of experimental success is incongruent with computations of the experimental success probability. The TES analysis detects this discrepancy.

To what extent the rate of apparent bias in Table 2 generalizes beyond the immediate sample depends on whether one interprets the sample as being representative of a given population. The majority of articles in *Psychological Science* include fewer than four experiments, and such articles may have different rates of bias (lower or higher). Moreover, the findings may not generalize beyond *Psychological Science*. The journal deliberately publishes articles of broad interest with innovative and groundbreaking findings, and it could be that such an emphasis tends to attract articles that appear to be biased. Other journals with different publishing goals may be less prone to include articles that appear to be biased. Regardless of whether the findings in Table 2 generalize, the high rate of apparent bias is a serious concern for *Psychological Science*, and it raises concerns about the quality of findings in other journals.

Unless the TES analysis is flawed (for debates, see Francis, 2013b, 2013c; Johnson, 2013; Johnson & Yuan, 2007; Morey, 2013; Simonsohn, 2012, 2013), two broad explanations could account for the high rate of apparent bias among the articles in *Psychological Science*: malfeasance or ignorance. The former interpretation supposes that researchers deliberately introduce bias into their findings, which is essentially fraud. Although a few unscrupulous researchers may generate flawed investigations with a full understanding that what they are doing misleads the field, the TES analysis does not necessarily lead to such a conclusion. An alternative explanation is that these authors were unaware that some of the methods used to gather data, analyze the data, or theorize would introduce a bias. Such a charge may seem nearly as disparaging as an accusation of fraud, but there are no pleasant choices here (Gelman, 2013). Ignorance is a plausible explanation for the findings in Table 2, because it is easy to introduce bias even when a researcher attempts to run proper studies (Francis, 2012d, 2013b, 2013c; Gelman & Loken, 2013; John et al., 2012; Simmons et al.,

Table 2 Results of the TES analysis for each of 44 articles in *Psychological Science*

Year	Authors	Short Title	P_{TES}
2012	Anderson, Kraus, Galinsky, & Keltner	Sociometric Status and Subjective Well-Being	.167
2012	Bauer, Wilkie, Kim, & Bodenhausen	Cuing Consumerism	.062
2012	Birtel & Crisp	Treating Prejudice	.133
2012	Converse & Fishbach	Instrumentality Boosts Appreciation	.110
2012	Converse, Risen, & Carter	Karmic Investment	.043
2012	Keysar, Hayakawa, & An	Foreign-Language Effect	.091
2012	Leung et al.	Embodied Metaphors and Creative "Acts"	.076
2012	Rounding, Lee, Jacobson, & Ji	Religion and Self-Control	.036
2012	Savani & Rattan	Choice and Inequality	.064
2012	van Boxtel & Koch	Visual Rivalry Without Spatial Conflict	.071
2011	Evans, Horowitz, & Wolfe	Weighting of Evidence in Rapid Scene Perception	.426
2011	Inesi, Botti, Dubois, Rucker, & Galinsky	Power and Choice	.026
2011	Nordgren, Morris McDonnell, & Loewenstein	What Constitutes Torture?	.090
2011	Savani, Stephens, & Markus	Interpersonal and Societal Consequences of Choice	.063
2011	Todd, Hanko, Galinsky, & Mussweiler	Difference Mind-Set and Perspective Taking	.043
2011	Tuk, Trampe, & Warlop	Inhibitory Spillover	.092
2010	Balcetis & Dunning	Wishful Seeing	.076
2010	Bowles & Gelfand	Status and Workplace Deviance	.057
2010	Damisch, Stoberock, & Mussweiler	How Superstition Improves Performance	.057
2010	de Hevia & Spelke	Number-Spacing Mapping in Human Infants	.070
2010	Ersner-Hershfield, Galinsky, Kray, & King	Counterfactual Reflection	.073
2010	Gao, McCarthy, & Scholl	The Wolfpack Effect	.115
2010	Lammers, Stapel, & Galinsky	Power and Hypocrisy	.024
2010	Li, Wei, & Soman	Physical Enclosure and Psychological Closure	.079
2010	Maddux et al.	Culture and the Endowment Effect	.014
2010	McGraw & Warren	Benign Violations	.081
2010	Sackett, Meyvis, Nelson, Converse, & Sackett	When Time Flies	.033
2010	Savani, Markus, Naidu, Kumar, & Berlia	What Counts as a Choice?	.058
2010	Senay, Albarracín, & Noguchi	Interrogative Self-Talk and Intention	.090
2010	West, Anderson, Bedwell, & Pratt	Red Diffuse Light Suppresses Fear Prioritization	.157
2009	Alter & Oppenheimer	Fluency and Self-Disclosure	.071
2009	Ashton-James, Maddux, Galinsky, & Chartrand	Affect and Culture	.035
2009	Fast & Chen	Power, Incompetence, and Aggression	.072
2009	Fast, Gruenfeld, Sivanathan, & Galinsky	Power and Illusory Control	.069
2009	Garcia & Tor	The *N*-Effect	.089
2009	González & McLennan	Hemispheric Differences in Sound Recognition	.139
2009	Hahn, Close, & Graf	Transformation Direction	.348
2009	Hart & Albarracín	Describing Actions	.035
2009	Janssen & Caramazza	Phonology and Grammatical Encoding	.083
2009	Jostmann, Lakens, & Schubert	Weight and Importance	.090
2009	Labroo, Lambotte, & Zhang	The Name-Ease Effect and Importance Judgments	.008
2009	Nordgren, van Harreveld, & van der Pligt	Restraint Bias	.0998
2009	Wakslak & Trope	Construal Level and Subjective Probability	.061
2009	Zhou, Vohs, & Baumeister	Symbolic Power of Money	.041

2011). The TES analysis may cast doubt on the validity of a report, but it should not, by itself, be used to denigrate the ethics of any author whose work is listed in Table 2.

There is another possible explanation for the appearance of bias for any specific article: chance. It sometimes happens that an experiment set that is properly sampled, analyzed, and

reported may produce results that are uncommonly successful relative to the estimated effects. This possibility is an additional reason why a conclusion of apparent bias from a TES analysis should not, by itself, be used to denigrate an author's ethics. However, such caution does not mean that scientists should ignore the TES conclusion; skepticism about the validity of the reported experiments is appropriate, even though there is a chance of a false alarm. To behave otherwise would be to dismiss the principles of hypothesis testing, which formed the foundation for all of the articles in Table 2.

Conclusions

Table 2 provides the first estimate of the frequency of biased articles for an important psychology journal. The finding that 82 % of the analyzed articles in *Psychological Science* appear to be biased indicates that the problems within the field are severe. Many readers, editors, and reviewers accepted the articles listed in Table 2 as containing innovative and groundbreaking findings in the field, but the TES analyses suggest that scientists should actually be skeptical about the validity of many of those articles. The frequency of such flawed articles implies fundamental misunderstandings about how to generate and identify good scientific arguments from multiple experiments, which suggests a need for radical changes in statistical analysis, experimental design, and theory development.

The findings reported here validate the general recognition that something is wrong with common scientific practice in psychology. This concern has prompted many proposals for reform, including increased emphasis on replication (Koole & Lakens, 2012; Roediger, 2012), enhanced disclosure of experimental methods (Eich, 2014; LeBel et al., 2013), improved access to experimental data (Nosek, Spies, & Motyl, 2012), focusing on confidence intervals and meta-analysis rather than hypothesis testing (Cumming, 2014), adoption of Bayesian data analysis methods (Kruschke, 2010; Rouder, Speckman, Sun, Morey, & Iverson, 2009; Wagenmakers, 2007), preregistration of hypotheses and methods (Chambers, 2013; Wagenmakers, Wetzels, Borsboom, van der Maas, & Kievit, 2012; Wolfe, 2013), and badges for articles that follow some of these reforms (Eich, 2014). These changes to data analysis and publishing will hopefully reduce the frequency of articles with excess success, and it would be valuable for future TES analyses to compare the frequencies of articles with excess success before and after a journal institutes such reforms.

Although these reforms aim for a better future, it is also important to understand and interpret the problems of the past. It is not enough to simply recognize that a high proportion of past studies appear to be biased. To be able to plan future studies and develop theories, scientists need to know which past studies contain flawed findings or theories. If a study appears to have an excess of success, then scientists know not

to trust the results, and they can run new experiments to check on the theoretical ideas. By noting past problems and motivating verification checks, TES analyses complement reform efforts and improve both scientific practice and the accumulation of scientific knowledge.

References

Alter, A. L., & Oppenheimer, D. M. (2009). Suppressing secrecy through metacognitive ease: Cognitive fluency encourages self-disclosure. *Psychological Science, 20,* 1414–1420.

Anderson, C., Kraus, M. W., Galinsky, A. D., & Keltner, D. (2012). The local-ladder effect: Social status and subjective well-being. *Psychological Science, 23,* 764–771.

Ashton-James, C. E., Maddux, W. W., Galinsky, A. D., & Chartrand, T. L. (2009). Who I am depends on how I feel: The role of affect in the expression of culture. *Psychological Science, 20,* 340–346.

Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science, 7,* 543–554. doi:10.1177/1745691612459060

Balcetis, E., & Dunning, D. (2010). Wishful seeing: More desired objects are seen as closer. *Psychological Science, 21,* 147–152.

Bauer, M. A., Wilkie, J. E. B., Kim, J. K., & Bodenhausen, G. V. (2012). Cuing consumerism: Situational materialism undermines personal and social well-being. *Psychological Science, 23,* 517–523.

Begg, C. B., & Mazumdar, M. (1994). Operating characteristics of a rank correlation test for publication bias. *Biometrics, 50,* 1088–1101.

Birtel, M. D., & Crisp, R. J. (2012). "Treating" prejudice: An exposure-therapy approach to reducing negative reactions toward stigmatized groups. *Psychological Science, 23,* 1379–1386.

Bowles, H. R., & Gelfand, M. (2010). Status and the evaluation of workplace deviance. *Psychological Science, 21,* 49–54.

Chambers, C. D. (2013). Registered reports: A new publishing initiative at *Cortex. Cortex, 49,* 609–610.

Converse, B. A., & Fishbach, A. (2012). Instrumentality boosts appreciation: Helpers are more appreciated while they are useful. *Psychological Science, 23,* 560–566.

Converse, B. A., Risen, J. L., & Carter, T. J. (2012). Investing in karma: When wanting promotes helping. *Psychological Science, 23,* 923–930.

Cumming, G. (2014). The new statistics: Why and how. *Psychological Science, 25,* 7–29. doi:10.1177/0956797613504966

Damisch, L., Stoberock, B., & Mussweiler, T. (2010). Keep your fingers crossed! How superstition improves performance. *Psychological Science, 21,* 1014–1020.

de Hevia, M. D., & Spelke, E. S. (2010). Number-space mapping in human infants. *Psychological Science, 21,* 653–660.

Egger, M., Davey Smith, G., Schneider, M., & Minder, C. (1997). Bias in a meta-analysis detected by a simple, graphical test. *British Medical Journal, 315,* 639–634.

Eich, E. (2014). Business not as usual [Editorial]. *Psychological Science, 25,* 3–6. doi:10.1177/0956797613512465

Ersner-Hershfield, H., Galinsky, A. D., Kray, L. J., & King, B. G. (2010). Company, country, connections: Counterfactual origins increase organizational commitment, patriotism, and social investment. *Psychological Science, 21,* 1479–1486.

Evans, K. K., Horowitz, T. S., & Wolfe, J. M. (2011). When categories collide: Accumulation of information about multiple categories in rapid scene perception. *Psychological Science, 22,* 739–746. doi:10.1177/0956797611407930

Fanelli, D. (2010). "Positive" results increase down the hierarchy of the sciences. *PLoS ONE, 5,* e10068. doi:10.1371/journal.pone.0010068

Fast, N. J., & Chen, S. (2009). When the boss feels inadequate: Power, incompetence, and aggression. *Psychological Science, 20,* 1406–1413.

Fast, N. J., Gruenfeld, D. H., Sivanathan, N., & Galinsky, A. D. (2009). Illusory control: A generative force behind power's far-reaching effects. *Psychological Science, 20,* 502–508.

Francis, G. (2012a). Evidence that publication bias contaminated studies relating social class and unethical behavior. *Proceedings of the National Academy of Sciences, 109,* E1587.

Francis, G. (2012b). The psychology of replication and replication in psychology. *Perspectives on Psychological Science, 7,* 585–594. doi:10.1177/1745691612459520

Francis, G. (2012c). Publication bias and the failure of replication in experimental psychology. *Psychonomic Bulletin & Review, 19,* 975–991. doi:10.3758/s13423-012-0322-y

Francis, G. (2012d). The same old New Look: Publication bias in a study of wishful seeing. *i-Perception, 3,* 176–178.

Francis, G. (2012e). Too good to be true: Publication bias in two prominent studies from experimental psychology. *Psychonomic Bulletin & Review, 19,* 151–156. doi:10.3758/s13423-012-0227-9

Francis, G. (2013a). Publication bias in "Red, rank, and romance in women viewing men" by Elliot et al. (2010). *Journal of Experimental Psychology: General, 142,* 292–296. doi:10.1037/a0027923

Francis, G. (2013b). Replication, statistical consistency, and publication bias. *Journal of Mathematical Psychology, 57,* 153–169. doi:10.1016/j.jmp.2013.02.003

Francis, G. (2013c). We should focus on the biases that matter: A reply to commentaries. *Journal of Mathematical Psychology, 57,* 190–195. doi:10.1016/j.jmp.2013.06.001

Gao, T., McCarthy, G., & Scholl, B. J. (2010). The wolfpack effect: Perception of animacy irresistibly influences interactive behavior. *Psychological Science, 21,* 1845–1853.

Garcia, S. M., & Tor, A. (2009). The N-effect: More competitors, less competition. *Psychological Science, 20,* 871–877.

Gelman, A. (2013). Is it possible to be an ethicist without being mean to people? *Chance, 26,* 51–53.

Gelman, A., & Loken, E. (2013). The garden of forking paths: Why multiple comparisons can be a problem, even when there is no "fishing expedition" or "p-hacking" and the research hypothesis was posited ahead of time. In Technical report, Department of Statistics, Columbia University. Retrieved January 30, 2014, from www.stat.columbia.edu/~gelman/research/unpublished/p_hacking.pdf

González, J., & McLennan, C. T. (2009). Hemispheric differences in the recognition of environmental sounds. *Psychological Science, 20,* 887–894.

Hahn, U., Close, J., & Graf, M. (2009). Transformation direction influences shape-similarity judgments. *Psychological Science, 20,* 447–454.

Hart, W., & Albarracín, D. (2009). What I was doing versus what I did: Verb aspect influences memory and future actions. *Psychological Science, 20,* 238–244.

Inesi, M. E., Botti, S., Dubois, D., Rucker, D. D., & Galinsky, A. D. (2011). Power and choice: Their dynamic interplay in quenching the thirst for personal control. *Psychological Science, 22,* 1042–1048.

Ioannidis, J. P. A., & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical Trials, 4,* 245–253.

Janssen, N., & Caramazza, A. (2009). Grammatical and phonological influences on word order. *Psychological Science, 20,* 1262–1268.

John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth-telling. *Psychological Science, 23,* 524–532. doi:10.1177/0956797611430953

Johnson, V. E. (2013). On biases in assessing replicability, statistical consistency and publication bias. *Journal of Mathematical Psychology, 57,* 177–179. doi:10.1016/j.jmp.2013.04.003

Johnson, V., & Yuan, Y. (2007). Comments on "An exploratory test for an excess of significant findings" by JPA Ioannidis and TA Trikalinos. *Clinical Trials, 4,* 254–255. doi:10.1177/1740774507079437. disc. 256–257.

Jostmann, N. B., Lakens, D., & Schubert, T. W. (2009). Weight as an embodiment of importance. *Psychological Science, 20,* 1169–1174.

Keysar, B., Hayakawa, S. L., & An, S. G. (2012). The foreign-language effect: Thinking in a foreign tongue reduces decision biases. *Psychological Science, 23,* 661–668.

Koole, S. L., & Lakens, D. (2012). Rewarding replications: A sure and simple way to improve psychological science. *Perspectives on Psychological Science, 7,* 608–614. doi:10.1177/1745691612462586

Kruschke, J. K. (2010). Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science, 1,* 658–676. doi:10.1002/wcs.72

Labroo, A. A., Lambotte, S., & Zhang, Y. (2009). The "name-ease" effect and its dual impact on importance judgments. *Psychological Science, 20,* 1516–1522.

Lammers, J., Stapel, D. A., & Galinsky, A. D. (2010). Power increases hypocrisy: Moralizing in reasoning, immorality in behavior. *Psychological Science, 21,* 737–744.

Lane, D. M., & Dunlap, W. P. (1978). Estimating effect size: Bias resulting from the significance criterion in editorial decisions. *British Journal of Mathematical and Statistical Psychology, 31,* 107–112.

LeBel, E. P., Borsboom, D., Giner-Sorolla, R., Hasselman, F., Peters, K. R., Ratliff, K. A., & Smith, C. T. (2013). PsychDisclosure.org: Grassroots support for reforming reporting standards in psychology. *Perspectives on Psychological Science, 8,* 424–432.

Leung, A. K., Kim, S., Polman, E., Ong, L. S., Qiu, L., Goncalo, J. A., & Sanchez-Burks, J. (2012). Embodied metaphors and creative "acts". *Psychological Science, 23,* 502–509.

Li, X., Wei, L., & Soman, D. (2010). Sealing the emotions genie: The effects of physical enclosure on psychological closure. *Psychological Science, 21,* 1047–1050.

Maddux, W. W., Yang, H., Falk, C., Adam, H, Adair, W., Endo, Y., … Heine, S. J. (2010). For whom is parting with possessions more painful? Cultural differences in the endowment effect. *Psychological Science, 21,* 1910–1917.

McGraw, A. P., & Warren, C. (2010). Benign violations: Making immoral behavior funny. *Psychological Science, 21,* 1141–1149. doi:10.1177/0956797610376073

Morey, R. D. (2013). The consistency test does not—and cannot—deliver what is advertised: A comment on Francis (2013). *Journal of Mathematical Psychology, 57,* 180–183. doi:10.1016/j.jmp.2013.03.004

Nordgren, L. F., Morris McDonnell, M.-H., & Loewenstein, G. (2011). What constitutes torture? Psychological impediments to an objective evaluation of enhanced interrogation tactics. *Psychological Science, 22,* 689–694.

Nordgren, L. F., van Harreveld, F., & van der Pligt, J. (2009). The restraint bias: How the illusion of self-restraint promotes impulsive behavior. *Psychological Science, 20,* 1523–1528.

Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific Utopia II: Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science, 7,* 615–631. doi:10.1177/1745691612459058

Renkewitz, F., Fuchs, H. M., & Fiedler, S. (2011). Is there evidence of publication biases in JDM research? *Judgment and Decision Making, 6,* 870–881.

Roediger, H. L., III. (2012, February). Psychology's woes and a partial cure: The value of replication. *APS Observer, 25*(2). Retrieved from www.psychologicalscience.org/index.php/publications/observer/2012/february-11-2012-observer-publications/psychology's-woes-and-a-partial-cure-the-value-of-replication.html

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null

hypothesis. *Psychonomic Bulletin & Review, 16,* 225–237. doi:10.3758/PBR.16.2.225

Rounding, K., Lee, A., Jacobson, J. A., & Ji, L.-J. (2012). Religion replenishes self-control. *Psychological Science, 23,* 635–642.

Sackett, A. M., Meyvis, T., Nelson, L. D., Converse, B. A., & Sackett, A. L. (2010). You're having fun when time flies: The hedonic consequences of subjective time progression. *Psychological Science, 21,* 111–117.

Savani, K., Markus, H. R., Naidu, N. V. R., Kumar, S., & Berlia, N. (2010). What counts as a choice? U.S. Americans are more likely than Indians to construe actions as choices. *Psychological Science, 21,* 391–398.

Savani, K., & Rattan, A. (2012). A choice mind-set increases the acceptance and maintenance of wealth inequality. *Psychological Science, 23,* 796–804.

Savani, K., Stephens, N. M., & Markus, H. R. (2011). The unanticipated interpersonal and societal consequences of choice: Victim blaming and reduced support for the public good. *Psychological Science, 22,* 795–802.

Schimmack, U. (2012). The ironic effect of significant results on the credibility of multiple-study articles. *Psychological Methods, 17,* 551–566. doi:10.1037/a0029487

Senay, I., Albarracín, D., & Noguchi, K. (2010). Motivating goal-directed behavior through introspective self-talk: The role of the interrogative form of simple future tense. *Psychological Science, 21,* 499–504.

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,* 1359–1366. doi:10.1177/0956797611417632

Simonsohn, U. (2012). It does not follow: Evaluating the one-off publication bias critiques by Francis (2012a, 2012b, 2012c, 2012d, 2012e, in press). *Perspectives on Psychological Science, 7,* 597–599. doi:10.1177/1745691612463399

Simonsohn, U. (2013). It really just does not follow, comments on Francis (2013). *Journal of Mathematical Psychology, 57,* 174–176. doi:10.1016/j.jmp.2013.03.006

Simonsohn, U., Nelson, L. D., & Simmons, J. P. (in press). *P*-curve: A key to the file-drawer. *Journal of Experimental Psychology: General.*

Sterling, T. D. (1959). Publication decisions and the possible effects on inferences drawn from test of significance—or vice versa. *Journal of the American Statistical Association, 54,* 30–34.

Sterling, T. D., Rosenbaum, W. L., & Weinkam, J. J. (1995). Publication decisions revisited: The effect of the outcome of statistical tests on the decision to publish and vice versa. *American Statistician, 49,* 108–112.

Todd, A. R., Hanko, K., Galinsky, A. D., & Mussweiler, T. (2011). When focusing on differences leads to similar perspectives. *Psychological Science, 22,* 134–141.

Tuk, M. A., Trampe, D., & Warlop, L. (2011). Inhibitory spillover: Increased urination urgency facilitates impulse control in unrelated domains. *Psychological Science, 22,* 627–633.

van Boxtel, J. J. A., & Koch, C. (2012). Visual rivalry without spatial conflict. *Psychological Science, 23,* 410–418.

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review, 14,* 779–804. doi:10.3758/BF03194105

Wagenmakers, E.-J., Wetzels, R., Borsboom, D., van der Maas, H. L. J., & Kievit, R. A. (2012). An agenda for purely confirmatory research. *Perspectives on Psychological Science, 7,* 632–638. doi:10.1177/1745691612463078

Wakslak, C., & Trope, Y. (2009). The effect of construal level on subjective probability estimates. *Psychological Science, 20,* 52–58.

West, G. L., Anderson, A. K., Bedwell, J. S., & Pratt, J. (2010). Red diffuse light suppresses the accelerated perception of fear. *Psychological Science, 21,* 992–999. doi:10.1177/0956797610371966

Wolfe, J. M. (2013). Registered reports and replications in *Attention, Perception, & Psychophysics* [Editorial]. *Attention, Perception, & Psychophysics, 75,* 781–783. doi:10.3758/s13414-013-0502-5

Zhou, X., Vohs, K. D., & Baumeister, R. F. (2009). The symbolic power of money: Reminders of money alter social distress and physical pain. *Psychological Science, 20,* 700–706. doi:10.1111/j.1467-9280.2009.02353.x