

Title: Redefine Statistical Significance

Authors: Daniel J. Benjamin^{1*}, James O. Berger², Magnus Johannesson^{3*}, Brian Nosek⁴, E.J. Wagenmakers⁵, Richard Berk⁶, Kenneth A. Bollen⁷, Björn Brembs⁸, Lawrence Brown⁹, Colin Camerer¹⁰, David Cesarini^{11,12}, Chris Chambers¹³, Merlise Clyde², Tom Cook^{14,15}, Paul de Boeck¹⁶, Zoltan Dienes¹⁷, Anna Dreber³, Kenny Easwaran¹⁸, Charles Efferson¹⁹, Ernst Fehr²⁰, Fiona Fidler²¹, Andy P. Field¹⁷, Malcolm Forster²², Ed George⁹, Richard Gonzales²³, Steven Goodman²⁴, Edwin Green²⁵, Don Green²⁶, Anthony Greenwald²⁷, Jarrod Hadfield²⁸, Larry Hedges²⁹, Leonhard Held³⁰, Teck Hua Ho³¹, Herbert Hoijtink³², James Holland Jones^{38,39}, Dan Hruschka³³, Kosuke Imai³⁴, Guido Imbens³⁵, John Ioannidis³⁶, Minjeong Jeon³⁷, Michael Kirchler⁴⁰, David Laibson⁴¹, John List⁴², Roderick Little⁴³, Skip Lupia⁴⁴, Edouard Machery⁴⁵, Scott Maxwell⁴⁶, Michael McCarthy⁴⁷, Don Moore⁴⁸, Stephen Morgan⁴⁹, Marcus Munaf⁵⁰, Shinichi Nakagawa⁵¹, Brendan Nyhan⁵², Tim Parker⁵³, Luis Pericchi⁵⁴, Marco Perugini⁵⁵, Jeff Rouder⁵⁶, Judith Rousseau⁵⁷, Victoria Savalei⁵⁸, Felix Schoenbrodt⁵⁹, Thomas Sellke⁶⁰, Richard Shiffrin⁶¹, Betsy Sinclair⁶², Dustin Tingley⁶³, Trish Van Zandt⁶⁴, Simina Vazire⁶⁵, Duncan Watts⁶⁶, Christopher Winship⁶⁷, Robert Wolpert², Yu Xie⁶⁸, Cristobal Young⁶⁹, Jonathan Zinman⁷⁰, Valen E. Johnson^{71*}

Affiliations:

¹Center for Economic and Social Research, University of Southern California, Los Angeles, CA 90089-3332, USA.

²Department of Statistical Science, Duke University, Durham, NC 27708-0251, USA.

³Department of Economics, Stockholm School of Economics, SE-113 83 Stockholm, Sweden.

⁴University of Virginia, Charlottesville, VA 22908, USA.

⁵University of Amsterdam, Department of Psychology, 1018 VZ Amsterdam, The Netherlands.

⁶University of Pennsylvania, School of Arts and Sciences and Department of Criminology, Philadelphia, PA 19104-6286, USA.

⁷University of North Carolina Chapel Hill, Department of Psychology and Neuroscience Department of Sociology, Chapel Hill, NC 27599-3270, USA.

⁸Institute of Zoology - Neurogenetics, Universität Regensburg, Universitätsstrasse 31 93040 Regensburg, Germany.

⁹The Wharton School Department of Statistics, University of Pennsylvania, Philadelphia, PA 19104, USA.

¹⁰Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, USA.

- ¹¹Department of Economics, New York University, New York, NY 10012, USA.
- ¹²The Research Institute of Industrial Economics (IFN), SE- 102 15 Stockholm, Sweden.
- ¹³Cardiff University Brain Research Imaging Centre (CUBRIC), CF24 4HQ UK.
- ¹⁴Northwestern University, Evanston, IL 60208, USA.
- ¹⁵Mathematica Policy Research, Washington, DC, 20002-4221, USA.
- ¹⁶Department of Psychology, Quantitative Program, Ohio State University, Columbus, OH 43210, USA.
- ¹⁷School of Psychology, University of Sussex, Brighton BN1 9QH UK.
- ¹⁸Department of Philosophy, Texas A&M University, College Station, TX 77843-4237, USA.
- ¹⁹Department of Psychology, Royal Holloway University of London, Egham Surrey TW20 0EX.
- ²⁰Department of Economics, University of Zurich, 8006 Zurich, Switzerland.
- ²¹School of BioSciences and School of Historical & Philosophical Studies, University of Melbourne, Vic 3010, Australia.
- ²²Department of Philosophy, University of Wisconsin - Madison, Madison, WI 53706, USA.
- ²³Department of Psychology, University of Michigan, Ann Arbor, MI 48109-1043, USA.
- ²⁴Stanford University, General Medical Disciplines, Stanford, CA 94305, USA.
- ²⁵Department of Ecology, Evolution and Natural Resources SEBS, Rutgers University, New Brunswick, NJ 08901-8551, USA.
- ²⁶Department of Political Science, Columbia University in the City of New York, New York, NY 10027, USA.
- ²⁷Department of Psychology, University of Washington, Seattle, WA 98195-1525, USA.
- ²⁸Institute of Evolutionary Biology School of Biological Sciences, The University of Edinburgh, Edinburgh EH9 3JT, UK.
- ²⁹Weinberg College of Arts & Sciences Department of Statistics, Northwestern University, Evanston, IL 60208, USA.
- ³⁰Department of Biostatistics, Epidemiology, Biostatistics & Prevention Institute, University of Zurich, 8001 Zurich, Switzerland.
- ³¹National University of Singapore, Singapore 119077.
- ³²Department of Methods and Statistics, Universiteit Utrecht Padualaan, 3584 CH Utrecht, The Netherlands.

- ³³School of Human Evolution and Social Change, Arizona State University, Tempe, AZ 85287-2402, USA.
- ³⁴Department of Politics and Center for Statistics and Machine Learning, Princeton University, Princeton NJ 08544, USA.
- ³⁵Stanford University, Stanford, CA 94305-5015, USA.
- ³⁶Departments of Medicine, of Health Research and Policy, of Biomedical Data Science, and of Statistics and Meta-Research Innovation Center at Stanford (METRICS), Stanford University, Stanford, CA 94305, USA.
- ³⁷ Advanced Quantitative Methods, Social Research Methodology, Department of Education, Graduate School of Education & Information Studies, University of California, Los Angeles, CA 90095-1521, USA.
- ³⁸Department of Life Sciences, Imperial College London, Ascot SL5 7PY, UK.
- ³⁹Woods Institute for the Environment, Stanford, CA 94305- 4216, USA.
- ⁴⁰Department of Banking and Finance, University of Innsbruck and University of Gothenburg, A-6020 Innsbruck, Austria.
- ⁴¹Department of Economics, Harvard University, Cambridge, MA 02138, USA.
- ⁴²Department of Economics, University of Chicago, Chicago, IL 60637, USA.
- ⁴³Department of Statistics, University of Michigan, Ann Arbor, MI 48109-2029, USA.
- ⁴⁴Department of Political Science Center for Political Studies (CPS), Institute for Social Research, University of Michigan, Ann Arbor, MI 48109-1248, USA.
- ⁴⁵Department of History and Philosophy of Science, University of Pittsburgh, Pittsburgh PA 15260, USA.
- ⁴⁶Department of Psychology, University of Notre Dame, Notre Dame, IN 46556, USA.
- ⁴⁷School of BioSciences, University of Melbourne, Vic 3010, Australia.
- ⁴⁸Haas School of Business, University of California at Berkeley, Berkeley, CA 94720-1900A, USA.
- ⁴⁹Johns Hopkins University, Baltimore, MD 21218, USA.
- ⁵⁰MRC Integrative Epidemiology Unit at the University of Bristol and UK Centre for Tobacco and Alcohol Studies, School of Experimental Psychology, University of Bristol, Bristol BS8 1TU, UK.
- ⁵¹Research, Evolution & Ecology Research Centre and School of Biological, Earth and Environmental Sciences, University of New South Wales, Sydney, NSW 2052, Australia.
- ⁵²Department of Government, Dartmouth College, Hanover, NH 03755, USA.
- ⁵³Department of Biology, Whitman College, Walla Walla, WA 99362, USA.

- ⁵⁴Department of Mathematics, University of Puerto Rico, San Juan, PR 00936-8377.
- ⁵⁵Department of Psychology, University of Milan - Bicocca, 20126 Milan, Italy.
- ⁵⁶Cognition and Neuroscience Area, Department of Psychological Sciences, University of Missouri Columbia, MO 65211, USA.
- ⁵⁷Université Paris Dauphine, 75016 Paris, France.
- ⁵⁸Department of Psychology, The University of British Columbia, Vancouver, BC Canada V6T 1Z4.
- ⁵⁹Department Psychology (Psychological Methods and Assessment) Ludwig-Maximilians-University Munich, Leopoldstraße 13 80802 Munich, Germany.
- ⁶⁰Department of Statistics, Purdue University, West Lafayette, IN 47907-2067, USA.
- ⁶¹Indiana University, IU Bloomington, College of Arts and Sciences, Psychological and Brain Sciences, Bloomington, IN 47405-7007, USA.
- ⁶²Department of Political Science, Washington University in St. Louis, St. Louis, MO 63130-4899, USA.
- ⁶³Government Department, Harvard University, Cambridge, MA 02138, USA.
- ⁶⁴Department of Psychology, Ohio State University, Columbus, OH 43210, USA.
- ⁶⁵University of California, Davis Campus, Davis, CA 95616, USA.
- ⁶⁶Microsoft Research, Columbia University in the City of New York, New York, NY 10027, USA.
- ⁶⁷Department of Sociology, Harvard University, Cambridge, MA 02138, USA.
- ⁶⁸Department of Sociology, Princeton University, Princeton NJ 08544, USA.
- ⁶⁹Department of Sociology, Stanford University, Stanford, CA 94305-2047, USA.
- ⁷⁰Department of Economics, Dartmouth College, Hanover, NH 03755-3514, USA.
- ⁷¹Department of Statistics, Texas A&M University, College Station, TX 77843, USA.
- *Correspondence to: Daniel J. Benjamin, daniel.benjamin@gmail.com; Magnus Johannesson, magnus.johannesson@hhs.se; Valen E. Johnson, vejohanson@exchange.tamu.edu.

One Sentence Summary: We propose to change the default *P*-value threshold for statistical significance from 0.05 to 0.005.

Main Text:

The lack of reproducibility of scientific studies has caused growing concern over the credibility of claims of new discoveries based on “statistically significant” findings. There has been much progress toward documenting and addressing several causes of this lack of reproducibility (e.g., multiple testing, P-hacking, publication bias, and under-powered studies). However, we believe that a leading cause of non-reproducibility has not yet been adequately addressed: Statistical standards of evidence for claiming discoveries in many fields of science are simply too low. Associating “statistically significant” findings with $P < 0.05$ results in a high rate of false positives *even in the absence of other experimental, procedural and reporting problems*.

For fields where the threshold for defining statistical significance is $P < 0.05$, we propose a change to $P < 0.005$. This simple step would immediately improve the reproducibility of scientific research in many fields. Results that would currently be called “significant” but do not meet the new threshold should instead be called “suggestive.” While statisticians have known the relative weakness of using $P \approx 0.05$ as a threshold for discovery and the proposal to lower it to 0.005 is not new (1, 2), a critical mass of researchers now endorse this change.

We restrict our recommendation to claims of discovery of new effects. We do not address the appropriate threshold for confirmatory or contradictory replications of existing claims. We also do not advocate changes to discovery thresholds in fields that have already adopted more stringent standards (e.g., genomics and high-energy physics research; see Potential Objections below).

We also restrict our recommendation to studies that conduct null hypothesis significance tests. We have diverse views about how best to improve reproducibility, and many of us believe that other ways of summarizing the data, such as Bayes factors or other posterior summaries based on clearly articulated model assumptions, are preferable to P -values. However, changing the P -value threshold is simple and might quickly achieve broad acceptance.

Strength of evidence from P -values

In testing a point null hypothesis H_0 against an alternative hypothesis H_1 based on data x_{obs} , the (two-tailed) P -value is defined as the probability, calculated under the null hypothesis, that a test statistic is as extreme or more extreme than its observed value. The null hypothesis is typically rejected—and the finding is declared “statistically significant”—if the P -value falls below the (current) Type I error threshold $\alpha = 0.05$.

From a Bayesian perspective, a more direct measure of the strength of evidence for H_1 relative to H_0 is the ratio of their probabilities. By Bayes’ rule, this ratio may be written as:

$$\frac{\Pr(H_1|x_{\text{obs}})}{\Pr(H_0|x_{\text{obs}})} = \frac{f(x_{\text{obs}}|H_1)}{f(x_{\text{obs}}|H_0)} \times \frac{\Pr(H_1)}{\Pr(H_0)} \equiv BF \times (\text{prior odds}), \quad (1)$$

where BF is the Bayes factor that represents the evidence from the data, and the prior odds can be informed by researchers' beliefs, scientific consensus, and validated evidence from similar research questions in the same field. The effects of multiple hypothesis testing, P-hacking, and publication bias are to reduce the prior odds of H_1 relative to H_0 . Prediction markets (3) and analyses of replication results (4) both suggest that for psychology experiments, the prior odds of H_1 relative to H_0 may be only about 1:10. A similar number has been suggested in cancer clinical trials, and the number is likely to be much lower in preclinical biomedical research (5).

There is no unique mapping between the P -value and the Bayes factor since the Bayes factor depends on H_1 . However, the connection between the two quantities can be evaluated for particular test statistics under certain classes of plausible alternatives (Figure 1).




Fig. 1. Relationship between the P -value and the Bayes Factor. The Bayes factor (BF) is defined as $\frac{f(x_{\text{obs}}|H_1)}{f(x_{\text{obs}}|H_0)}$. The figure assumes that observations are drawn i.i.d. according to $x \sim N(\mu, \sigma^2)$, where the mean μ is unknown and the variance σ^2 is known. The P -value is from a two-sided z test (or equivalently a one-sided χ_1^2 test) of the null hypothesis $H_0: \mu = 0$.

“Power”: BF obtained by defining H_1 as putting $\frac{1}{2}$ probability on $\mu = \pm m$ for the value of m that gives 75% power for the test of size $\alpha = 0.05$. This H_1 represents an effect size typical of that which is implicitly assumed by researchers during experimental design.

“Likelihood Ratio Bound”: BF obtained by defining H_1 as putting $\frac{1}{2}$ probability on $\mu = \pm \hat{x}$, where \hat{x} is approximately equal to the mean of the observations. These BFs are upper

bounds among the class of all H_1 's that are symmetric around the null, but they are improper because the data are used to define H_1 . “UMPBT”: BF obtained by defining H_1 according to the uniformly most powerful Bayesian test (5) that places $\frac{1}{2}$ probability on $\mu = \pm w$, where w is the alternative hypothesis that corresponds to a one-sided test of size 0.0025. This curve is indistinguishable from the “Power” curve that would be obtained if the power used in its definition was 80% rather than 75%. “Local- H_1 Bound”: $\text{BF} = \frac{1}{-ep \ln p}$, where p is the P -value, is a large-sample upper bound on BF from among all unimodal alternative hypotheses that have a mode at the null and satisfy certain regularity conditions (5). For more details, see the Supplementary Online Materials (SOM).

Figure 1 shows that a two-sided P -value of 0.05 corresponds to Bayes factors in favor of H_1 that range from about 2.5 to 3.4 under reasonable assumptions about H_1 . This is weak evidence from at least three perspectives. First, conventional Bayes factor categorizations (6) characterize this range as “weak” or “very weak.” Second, we suspect many scientists would guess that $P \approx 0.05$ implies stronger support for H_1 than a Bayes factor of 2.5 to 3.4. Third, using equation (1) and prior odds of 1:10, a P -value of 0.05 corresponds to *at least 3:1 odds* (i.e., the reciprocal of the product $\frac{1}{10} \times 3.4$) *in favor of the null hypothesis!*

Why 0.005?

The choice of any particular threshold is arbitrary and involves a trade-off between Type I and II errors. We propose 0.005 for two reasons. First, a two-sided P -value of 0.005 corresponds to Bayes factors between approximately 14 and 26 in favor of H_1 . This range represents “substantial” to “strong” evidence according to conventional Bayes factor classifications (6).

Second, in many fields the $P < 0.005$ standard would reduce the false positive rate to levels we judge to be reasonable. If we let ϕ denote the proportion of null hypotheses that are true, $(1 - \beta)$ the power of tests in rejecting false null hypotheses, and α the Type I error/significance threshold, then as the population of tested hypotheses becomes large, the false positive rate (i.e., the proportion of true null effects among the total number of statistically significant findings) can be approximated by

$$\text{false positive rate} \approx \frac{\alpha\phi}{\alpha\phi + (1 - \beta)(1 - \phi)}. \quad (2)$$

For different levels of the prior odds that there is a true effect, $\frac{1-\phi}{\phi}$, and for significance thresholds $\alpha = 0.05$ and $\alpha = 0.005$, Figure 2 shows the false positive rate as a function of power $1 - \beta$.




Fig. 2. Relationship between the P -value threshold, power, and the false positive rate. Calculated according to Equation (2), with prior odds defined as $\frac{1-\phi}{\phi} = \frac{\Pr(H_1)}{\Pr(H_0)}$. For more details, see the Supplementary Online Materials (SOM).

In many studies, statistical power is low (e.g., ref. 7). Figure 2 demonstrates that low statistical power and $\alpha = 0.05$ combine to produce high false positive rates.

For many, the calculations illustrated by Figure 2 may be unsettling. For example, the false positive rate is greater than 33% with prior odds of 1:10 and a P -value threshold of 0.05, *regardless of the level of statistical power*. Reducing the threshold to 0.005 would reduce this minimum false positive rate to 5%. Similar reductions in false positive rates would occur over a wide range of statistical powers.

Empirical evidence from recent replication projects in psychology and experimental economics provide insights into the prior odds in favor of H_1 . In both projects, the rate of replication (i.e., significance at $P < 0.05$ in the replication in a consistent direction) was roughly double for initial studies with $P < 0.005$ relative to initial studies with $0.005 < P < 0.05$: 50% versus 24% for psychology (8), and 85% versus 44% for experimental economics (9). Although based on relatively small samples of studies (93 in psychology, 16 in experimental economics, after excluding initial studies with $P > 0.05$), these numbers are suggestive of the potential gains in reproducibility that would accrue from the new threshold of $P < 0.005$ in these fields. In biomedical research, 96% of a sample of recent papers claim statistically significant results with the $P < 0.05$ threshold (10). However, replication rates were very low (5) for these studies, suggesting a potential for gains by adopting this new standard in these fields as well.

Potential Objections

We now address the most compelling arguments against adopting this higher standard of evidence.

The false negative rate would become unacceptably high. Evidence that does not reach the new significance threshold should be treated as suggestive, and where possible further evidence should be accumulated; indeed, the combined results from several studies may be compelling even if any particular study is not. Failing to reject the null hypothesis does *not* mean accepting the null hypothesis. Moreover, the false negative rate will not increase if sample sizes are increased so that statistical power is held constant.

For a wide range of common statistical tests, transitioning from a P -value threshold of $\alpha = 0.05$ to $\alpha = 0.005$ while maintaining 80% power would require an increase in sample sizes of about 70%. Such an increase means that fewer studies can be conducted using current experimental designs and budgets. But Figure 2 shows the benefit: false positive rates would typically fall by factors greater than two. Hence, considerable resources would be saved by not performing future studies based on false premises. Increasing sample sizes is also desirable because studies with small sample sizes tend to yield inflated effect size estimates (11), and publication and other biases may be more likely in an environment of small studies (12). We believe that efficiency gains would far outweigh losses.

The proposal does not address multiple hypothesis testing, P-hacking, publication bias, low power, or other biases (e.g., confounding, selective reporting, measurement error), which are arguably the bigger problems. We agree. Reducing the P -value threshold complements—does not substitute for—solutions to these other problems, which include good study design, ex ante power calculations, pre-registration of planned analyses, replications, and transparent reporting of all statistical analyses conducted.

The appropriate threshold for statistical significance should be different for different research communities. We agree that the significance threshold selected for claiming a new discovery should depend on the prior odds that the null hypothesis is true, the number of hypotheses tested, the study design, the relative cost of Type I versus Type II errors, and other factors that vary by research topic. For exploratory research with very low prior odds (well outside the range in Figure 2), even lower significance thresholds than 0.005 are needed. Recognition of this issue led the genetics research community to move to a “genome-wide significance threshold” of 5×10^{-8} over a decade ago. And in high-energy physics, the tradition has long been to define significance by a “5-sigma” rule (roughly a P -value threshold of 3×10^{-7}). We are essentially suggesting a move from a 2-sigma rule to a 3-sigma rule.

Our recommendation applies to disciplines with prior odds broadly in the range depicted in Figure 2, where use of $P < 0.05$ as a default is widespread. Within those disciplines, it is helpful for consumers of research to have a consistent benchmark. We feel the default should be shifted.

Changing the significance threshold is a distraction from the real solution, which is to replace null hypothesis significance testing (and bright-line thresholds) with more

focus on effect sizes and confidence intervals, treating the P -value as a continuous measure, and/or a Bayesian method. Many of us agree that there are better approaches to statistical analyses than null hypothesis significance testing, but as yet there is no consensus regarding the appropriate choice of replacement (13). Even after the significance threshold is changed, many of us will continue to advocate for alternatives to null hypothesis significance testing.

Concluding remarks

Ronald Fisher understood that the choice of 0.05 was arbitrary when he introduced it (14). Since then, theory and empirical evidence have demonstrated that a lower threshold is needed. A much larger pool of scientists are now asking a much larger number of questions, often with much lower prior odds of success.

For research communities that continue to rely on null hypothesis significance testing, reducing the P -value threshold for claims of new discoveries to 0.005 is an actionable step that will immediately improve reproducibility. We emphasize that this proposal is about standards of evidence, not standards for policy action nor standards for publication. Results that do not reach the threshold for statistical significance (whatever it is) can still be important and merit publication in top journals if they address important research questions with rigorous methods. This proposal should not be used to reject publications of novel findings with $0.005 < P < 0.05$ properly labeled as suggestive evidence. We must reward quality and transparency of research as we impose these more stringent standards, and we should monitor how researchers' behaviors are affected by this change. Otherwise, science runs the risk that the more demanding threshold for statistical significance will be met to the detriment of quality and transparency.

Journals can help transition to the new statistical significance threshold. Authors and readers can themselves take the initiative by describing and interpreting results more appropriately in light of the new definition of "statistical significance." The new significance threshold will help researchers and readers to understand and communicate evidence more accurately.

References and Notes:

1. A. G. Greenwald *et al.*, Effect sizes and p values: What should be reported and what should be replicated? *Psychophysiology* **33**, 175-183 (1996).
2. V. E. Johnson, Revised standards for statistical evidence. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 19313-19317 (2013).
3. A. Dreber *et al.*, Using prediction markets to estimate the reproducibility of scientific research. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 15343-15347 (2015).
4. V. E. Johnson *et al.*, On the reproducibility of psychological science. *J. Am. Stat. Assoc.* **112**, 1-10 (2016).
5. G. C. Begley, J. P. A. Ioannidis, Reproducibility in science: Improving the standard for basic and preclinical research. *Circ. Res.* **116**, 116-126 (2015).
6. R. E. Kass, A. E. Raftery, Bayes Factors. *J. Am. Stat. Assoc.* **90**, 773-795 (1995).

7. D. Szucs, J. P. A. Ioannidis, Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biol.* **15**, (2017).
8. Open Science Collaboration, Estimating the reproducibility of psychological science. *Science* **349**, (2015).
9. C. Camerer *et al.*, Evaluating replicability of laboratory experiments in economics. *Science* **351**, 1433-1436 (2016).
10. D. Chavalarias *et al.*, Evolution of reporting p values in the biomedical literature, 1990-2015. *JAMA* **315**, 1141-1148 (2016).
11. A. Gelman, J. Carlin, Beyond power calculations: Assessing Type S (Sign) and Type M (Magnitude) errors. *Perspect. Psychol. Sci.* **9**, 641-651 (2014).
12. D. Fanelli, R. Costas, J. P. A. Ioannidis, Meta-assessment of bias in science. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 3714-3719 (2017).
13. R. L. Wasserstein, N. A. Lazar, The ASA's statement on p-values: Context, process, and purpose. *Am. Stat.* **70** (and online comments), 129-133 (2016).
14. R. A. Fisher, *Statistical Methods for Research Workers* (Oliver & Boyd, Edinburgh, 1925).
15. T. Sellke, M. J. Bayarri, J. O. Berger, Calibration of p-values for testing precise null hypotheses. *Am. Stat.* **55**, 62-71 (2001).

Acknowledgements: We thank Rebecca Royer and Anh Tuan Nguyen Viet for excellent research assistance.