

NBER WORKING PAPER SERIES

STATISTICAL NON-SIGNIFICANCE IN EMPIRICAL ECONOMICS

Alberto Abadie

Working Paper 24403

<http://www.nber.org/papers/w24403>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2018

I thank Isaiah Andrews, Joshua Angrist, Amy Finkelstein, Guido Imbens, Ben Olken, and especially Gary Chamberlain and Max Kasy for comments and discussions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Alberto Abadie. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Statistical Non-Significance in Empirical Economics
Alberto Abadie
NBER Working Paper No. 24403
March 2018
JEL No. C01,C12

ABSTRACT

Significance tests are probably the most common form of inference in empirical economics, and significance is often interpreted as providing greater informational content than non-significance. In this article we show, however, that rejection of a point null often carries very little information, while failure to reject may be highly informative. This is particularly true in empirical contexts that are typical and even prevalent in economics, where data sets are large (and becoming larger) and where there are rarely reasons to put substantial prior probability on a point null. Our results challenge the usual practice of conferring point null rejections a higher level of scientific significance than non-rejections. In consequence, we advocate a visible reporting and discussion of non-significant results in empirical practice.

Alberto Abadie
Department of Economics, E52-546
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
abadie@mit.edu

1. Introduction

Non-significant empirical results (usually in the form of t -statistics smaller than 1.96) relative to some null hypotheses of interest (usually zero coefficients) are notoriously hard to publish in professional/scientific journals (see, e.g., Andrews and Kasy, 2017; Ziliak and McCloskey, 2008). This state of affairs is in part maintained by the widespread notion that non-significant results are non-informative. After all, lack of statistical significance derives from the absence of extreme or surprising outcomes under the null hypothesis. In this article, we argue that this view of statistical inference is misguided. In particular, we show that non-significant results are informative, and argue that they are more informative than significant results in scenarios that are common, even prevalent, in empirical practice in economics.

To discuss the informational content of different statistical procedures, we formally adopt a limited information Bayes perspective. In this setting, agents representing journal readership or the scientific community have priors, \mathcal{P} , over some parameters of interests, $\theta \in \Theta$. That is, a member p of \mathcal{P} is a probability density function (with respect to some appropriate measure) on Θ . While agents are Bayesian, we will consider a setting where journals report frequentist results, in particular, statistical significance. Agents construct limited information Bayes posteriors based on the reported results of significance tests. We will deem a statistical result informative when it has the potential to substantially change the prior of the agents over a large range of values for θ .

Notice that, like Ioannidis (2005) and others, we restrict our attention to the effect of statistical significance on beliefs. We adopt this framework not because we believe it is (always) representative of empirical practice (in fact, journals typically report additional statistics, beyond statistical significance), but because isolating the informational content of statistical significance has immediate implications for how we should interpret its occurrence or lack of it. Correct interpretation of statistical significance is important because, while many other statistics are reported in practice, the scientific discussion of empirical results is often framed in terms of statistical significance of some parameters of interest and non-significant results may be under-reported as

discussed above.

2. A Simple Example

In this section, we consider a simple example with Normal priors and data that captures the essence of our argument. In section 3 we will consider the case where the priors and the distribution of the data are not restricted to be in a particular parametric family. Assume an agent has a prior $\theta \sim N(\mu, \sigma^2)$ on θ , with $\sigma^2 > 0$. A researcher observes n independent measurement of θ with Normal errors mutually independent and independent of θ , and with variance normalized to one. That is, x_1, \dots, x_n are independent $N(\theta, 1)$. Let

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^n x_i \sim N(\theta, 1/n).$$

θ is deemed significant if $\sqrt{n}|\hat{\theta}| > c$, for some $c > 0$. In empirical practice, c is often equal to 1.96, the 0.975-quantile of the Standard Normal distribution. Suppose a journal reports on statistical significance. We will calculate the limited information posteriors of the agents conditional on significance and lack thereof. These posteriors are the distributions of θ conditional on $\sqrt{n}|\hat{\theta}| > c$ and $\sqrt{n}|\hat{\theta}| \leq c$. First, notice that

$$\begin{aligned} \Pr(\sqrt{n}|\hat{\theta}| > c|\theta) &= \Pr(\hat{\theta} > c/\sqrt{n}|\theta) + \Pr(-\hat{\theta} > c/\sqrt{n}|\theta) \\ &= \Phi(\sqrt{n}\theta - c) + \Phi(-\sqrt{n}\theta - c). \end{aligned}$$

Therefore,¹

$$\Pr(\sqrt{n}|\hat{\theta}| > c) = \Phi\left(\frac{\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right) + \Phi\left(\frac{-\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right). \quad (1)$$

¹This calculation uses the following fact of integration

$$\int \Phi\left(\frac{\lambda - \theta}{\xi}\right) \frac{1}{\sigma} \phi\left(\frac{\theta - \mu}{\sigma}\right) d\theta = \Phi\left(\frac{\lambda - \mu}{\sqrt{\sigma^2 + \xi^2}}\right)$$

for arbitrary real λ and μ and positive σ and ξ . Alternatively, the result can be easily derived after noticing that the distribution of $\hat{\theta}$ integrated over the prior is Normal with mean μ and variance $\sigma^2 + 1/n$.

The limited information posteriors given significance and non-significance are:

$$p(\theta|\sqrt{n}|\hat{\theta}| > c) = \frac{\frac{1}{\sigma}\phi\left(\frac{\theta - \mu}{\sigma}\right)\left(\Phi(\sqrt{n}\theta - c) + \Phi(-\sqrt{n}\theta - c)\right)}{\Phi\left(\frac{\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right) + \Phi\left(\frac{-\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right)},$$

and

$$p(\theta|\sqrt{n}|\hat{\theta}| \leq c) = \frac{\frac{1}{\sigma}\phi\left(\frac{\theta - \mu}{\sigma}\right)\left(1 - \Phi(\sqrt{n}\theta - c) - \Phi(-\sqrt{n}\theta - c)\right)}{1 - \Phi\left(\frac{\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right) - \Phi\left(\frac{-\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right)}.$$

The two posteriors, along with the Normal prior, are plotted in Figure 1 for $\mu = 1$, $\sigma = 1$, $c = 1.96$, and $n = 10$. This figure illustrates the informational value of a significance test. Rejection of the null carries probability mass around zero in the limited information posterior, while failure to reject concentrates probability mass around zero. Notice that failure to reject carries substantial information, even in the rather under-powered setting generated by the values of μ , σ , c , and n adopted for Figure 1, which imply $\Pr(\sqrt{n}|\hat{\theta}| > c) = 0.7028$.

Figure 2 shows how prior and posteriors after significance compare as a function of the sample size. When n is small, significance affects the posterior over a large range of values. When n is large, significance provides only local to zero information. That is, significance is not informative in large samples. This is explained by the fact that the probability of rejection in equation (1) converges to one as the sample size increases. By the law of total probability, it follows that conditional on non-significance probability mass concentrates around zero as n increases. That is, the occurrence of an event that is very unlikely given the prior has a large effect on beliefs.

The full information posterior is

$$p(\theta|x_1, \dots, x_n) = \frac{1}{\sigma_n}\phi\left(\frac{\theta - \mu_n}{\sigma_n}\right),$$

where

$$\mu_n = \frac{\mu + n\sigma^2\hat{\theta}}{1 + n\sigma^2},$$




Figure 1: Posterior Distributions After a Significance Test

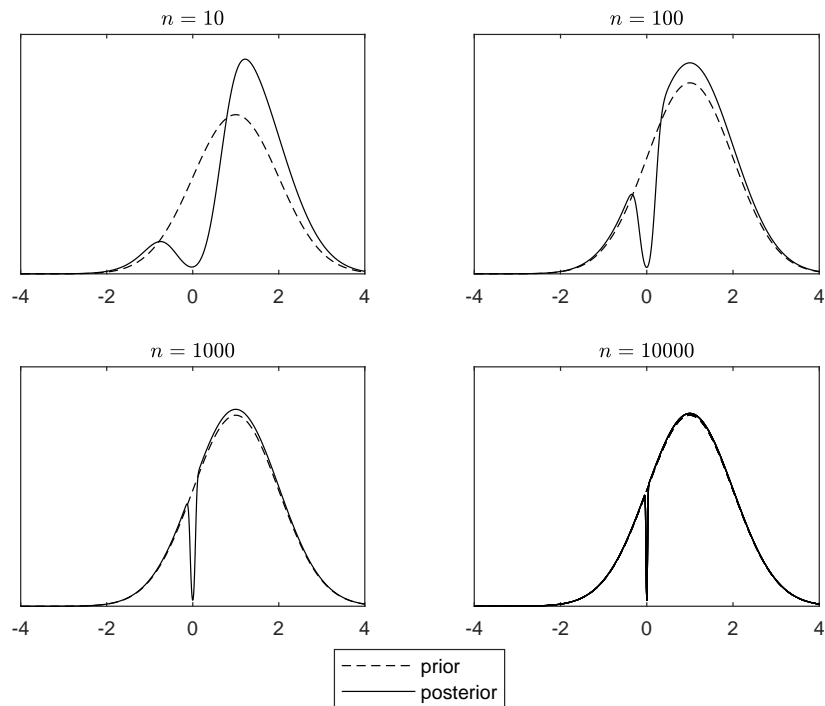


Figure 2: Prior and Posterior with Significance for Different Sample Sizes

and

$$\sigma_n^2 = \frac{\sigma^2}{1 + n\sigma^2}.$$

So, in this very particular context, knowledge of the t -ratio ($\sqrt{n}\widehat{\theta}$) is sufficient to go back to the full information posterior. The same is true for the combined information given by the P -value, $2\Phi(-\sqrt{n}|\widehat{\theta}|)$, and the sign of $\widehat{\theta}$.

These results have immediate counterparts in in large samples settings with asymptotically Normal distributions. They can also be generalized to non-parametric settings, as we demonstrate in the next section.

3. General Case

3.1. Finite Sample Results

Results like that in Figure 1 are rather general and do not depend on Normal priors or data. Consider a test statistic, \widehat{T}_n , such that rejection of the null is given by $\widehat{T}_n > c$. Let $p(\cdot)$ be a prior on θ , and $p(\cdot|\widehat{T}_n > c)$ and $p(\cdot|\widehat{T}_n \leq c)$ be the limited information posteriors under significance and non-significance, respectively. Regardless of the shape of the prior and/or the distribution of the data, by the law of total probability we obtain

$$\left| 1 - \frac{p(\theta|T_n \leq c)}{p(\theta)} \right| = \left(\frac{\Pr(\widehat{T}_n > c)}{\Pr(\widehat{T}_n \leq c)} \right) \left| 1 - \frac{p(\theta|T_n > c)}{p(\theta)} \right| \quad (2)$$

for $\Pr(\widehat{T}_n \leq c) > 0$ and θ such that $p(\theta) > 0$. The absolute value expressions on both sides of Equation (2) provide a measure of informativeness of significance (right) and non-significance (left). For a rather underpowered setting with $\Pr(\widehat{T}_n > c) = 1/2$, Equation (2) implies that non-significant is exactly as informative as significance. Moreover, the relative informativeness of non-significance increases with the statistical power of the test. Next section provides large sample calculations.

3.2. Large Sample Analysis

To extend the large sample results of the previous section beyond Normal priors and data, we will consider a test statistic, \widehat{T}_n , such that

$$\Pr(\widehat{T}_n > c | \theta = 0) \rightarrow \alpha,$$

and

$$\Pr(\widehat{T}_n > c | \theta, \theta \neq 0) \rightarrow 1.$$

That is, we consider significance tests that are consistent under fixed alternatives and have asymptotic size equal to α .

3.2.1. Continuous Prior

We will first assume a prior that is absolutely continuous with respect to the Lebesgue measure, with a version of the density that is positive and continuous at zero. By dominated convergence, we obtain:

$$\Pr(\widehat{T}_n > c) \rightarrow 1.$$

We first derive the posterior densities under significance,

$$p(0 | \widehat{T}_n > c) = \frac{\Pr(\widehat{T}_n > c | \theta = 0)}{\Pr(\widehat{T}_n > c)} p(0) \rightarrow \alpha p(0),$$

and

$$p(\theta | \widehat{T}_n > c) = \frac{\Pr(\widehat{T}_n > c | \theta)}{\Pr(\widehat{T}_n > c)} p(\theta) \rightarrow p(\theta),$$

for $\theta \neq 0$. So, again, significance only changes beliefs locally around zero. The posterior densities after non-significance are

$$p(0 | \widehat{T}_n \leq c) = \frac{\Pr(\widehat{T}_n \leq c | \theta = 0)}{\Pr(\widehat{T}_n \leq c)} p(0) \rightarrow \infty,$$

and

$$p(\theta|\widehat{T}_n \leq c) = \frac{\Pr(\widehat{T}_n \leq c|\theta)}{\Pr(\widehat{T}_n \leq c)} p(\theta)$$

for $\theta \neq 0$. Typically, for $\theta \neq 0$ (using large deviation results)

$$-\frac{1}{n} \log \left(\Pr(\widehat{T}_n \leq c|\theta) \right) \rightarrow d_\theta,$$

with $0 < d_\theta < \infty$. Therefore, $\Pr(\widehat{T}_n \leq c|\theta)$ converges to zero exponentially for $\theta \neq 0$. Let

$$\beta_n(\theta) = \Pr(\widehat{T}_n \leq c|\theta)$$

be the probability of Type II error (one minus the power). Assume that

$$\int \liminf_{n \rightarrow \infty} \beta_n(z/\sqrt{n}) dz > 0.$$

This rules out perfect local asymptotic power. Then, by change of variable $z = n^{1/2}\theta$ and Fatou's lemma, we obtain²

$$\begin{aligned} \liminf_{n \rightarrow \infty} n^{1/2} \Pr(\widehat{T}_n \leq c) &= \liminf_{n \rightarrow \infty} n^{1/2} \int \beta_n(\theta) p(\theta) d\theta \\ &= \liminf_{n \rightarrow \infty} \int \beta_n(z/\sqrt{n}) p(z/\sqrt{n}) dz \\ &\geq \int \liminf_{n \rightarrow \infty} (\beta_n(z/\sqrt{n}) p(z/\sqrt{n})) dz \\ &= \int \liminf_{n \rightarrow \infty} \beta_n(z/\sqrt{n}) \lim_{n \rightarrow \infty} p(z/\sqrt{n}) dz \\ &= p(0) \int \liminf_{n \rightarrow \infty} \beta_n(z/\sqrt{n}) dz > 0. \end{aligned}$$

It follows that

$$p(\theta|\widehat{T}_n \leq c) \rightarrow 0,$$

for $\theta \neq 0$.

That is, like in the Normal case of section 2, conditional on non-significance the posterior converges to a degenerate distribution at zero.

²For the second to last equality, notice that if $a_n \geq 0$ and $b_n \rightarrow b > 0$ as $n \rightarrow \infty$, then

$$\liminf_{n \rightarrow \infty} (a_n b_n) = \liminf_{n \rightarrow \infty} a_n \lim_{n \rightarrow \infty} b_n.$$




Figure 3: Limit of $p(\theta|\widehat{T}_n > c)/p(\theta)$ as a function of q ($\theta \neq 0$, $\alpha = 0.05$)

3.3. Prior with Probability Mass at Zero

We now consider the case when the prior has probability mass q at zero, with $0 < q < 1$. Then

$$\Pr(\widehat{T}_n > c) \rightarrow q\alpha + (1 - q) \in (\alpha, 1).$$

Now, the posterior after significance is,

$$p(0|\widehat{T}_n > c) = \frac{\Pr(\widehat{T}_n > c|\theta = 0)}{\Pr(\widehat{T}_n > c)}p(0) \rightarrow \left(\frac{\alpha}{q\alpha + (1 - q)}\right)q \leq q,$$

and

$$p(\theta|\widehat{T}_n > c) = \frac{\Pr(\widehat{T}_n > c|\theta)}{\Pr(\widehat{T}_n > c)}p(\theta) \rightarrow \left(\frac{1}{q\alpha + (1 - q)}\right)p(\theta) \geq p(\theta),$$

for $\theta \neq 0$. Now, in contrast to the continuous prior case, significance changes beliefs away from zero in large samples. In particular, if we start with a prior that assigns a large probability to $\theta = 0$, then significance greatly affects beliefs for values of θ different from zero. Notice, however,

that for moderate values of q the effect of significance on beliefs may be negligible. Figure 3 show the limit of $p(\theta|\widehat{T}_n > c)/p(\theta)$ as a function of q , for $\theta \neq 0$ and $\alpha = 0.05$. This limit is close to one for modest values of q . In order for significance to at least double the probability of $\theta \neq 0$ we need $q \geq 1/(2(1 - \alpha)) = 0.5263$. Notice that reducing the value of α does not substantially change the value of the limit of $p(\theta|\widehat{T}_n > c)/p(\theta)$, except for very large values of q . For example, with $\alpha = 0.005$ (as advocated in Benjamin et al., 2017), for significance to at least double the probability of $\theta \neq 0$ we need $q \geq 1/(2(1 - \alpha)) = 0.5025$. In fact, regardless of the size of the test, q needs to be bigger than 0.5 in order for significance to double the probability density function of beliefs at non-zero values of θ .

The posterior after non-significance is,

$$p(0|\widehat{T}_n \leq c) = \frac{\Pr(\widehat{T}_n \leq c|\theta = 0)}{\Pr(\widehat{T}_n \leq c)}p(0) \rightarrow \frac{1 - \alpha}{q(1 - \alpha)}q = 1,$$

and for $\theta \neq 0$,

$$p(\theta|\widehat{T}_n \leq c) = \frac{\Pr(\widehat{T}_n \leq c|\theta)}{\Pr(\widehat{T}_n \leq c)}p(\theta) \rightarrow 0.$$

Again, non-significance seems to have a stronger effect on beliefs than significance.

Some remarks about priors with probability mass at a point null are in order. First, it is difficult to think of relevant settings in empirical economics where reasonable prior beliefs assign probability mass to point nulls. For example, beliefs on the causal effect of a policy intervention may sometimes concentrate probability smoothly around zero, but more rarely in such a way that a large probability mass at zero is a good description of a reasonable prior. Moreover, priors with probability mass at a point null generate a drastic discrepancy, know as Lindley's paradox, between frequentist and Bayesian testing procedures (see, e.g., Berger, 1985). Lindley's paradox arises in settings with a fixed value of \widehat{T}_n and a large n . In those settings, frequentists would reject the null hypothesis when $\widehat{T}_n > c$. Bayesians, however, would typically find that the posterior probability of the point null far exceeds the posterior probability of the alternative. Lindley's paradox can be explained by the fact that, as n increases, the distribution of the test statistic under the alternative diverges. Therefore, a fixed value of the test statistic as n increases can

only be explained by the null hypothesis. Notice that conditioning on the event $\{\widehat{T}_n \leq c\}$ (as opposed to conditioning on the value of \widehat{T}_n) is not subject to Lindley's paradox and it may be the natural choice to evaluate a testing procedure for which significance depends on the value of $\{\widehat{T}_n \leq c\}$ only.

4. Testing an Interval Null

In view of the lack of informativeness of non-significance in large samples (under a point null), one could instead try to reinterpret significance tests as tests of the implicit null “ θ is close to zero”.

To accommodate this possibility, we will now concentrate in the problem of testing the null that the parameter θ is in some interval around zero. Under the null hypothesis, $\theta \in [-\delta, \delta]$, where δ is some positive number. Under the alternative hypothesis, $\theta \notin [-\delta, \delta]$. Consider the Normal model of section 2. To obtain a test of size α we control the supremum of the probability of Type I error:

$$\Pr(\sqrt{n}|\widehat{\theta}| > c \mid |\theta| = \delta) = \Phi(\sqrt{n}\delta - c) + \Phi(-\sqrt{n}\delta - c).$$

Therefore, we choose c such that $\Phi(\sqrt{n}\delta - c) + \Phi(-\sqrt{n}\delta - c) = \alpha$. While there is no closed-form solution for c , its value can be calculated numerically for any given value of $\sqrt{n}\delta$, and a very accurate approximation for large $\sqrt{n}\delta$ is given by

$$c = \Phi^{-1}(1 - \alpha) + \sqrt{n}\delta.$$

That is, controlling size in this setting implies that the critical value has to increase with the sample size at a root- n rate. In turn, this implies that the probability of rejection, $\Pr(\sqrt{n}|\widehat{\theta}| > c \mid \theta) = \Phi(\sqrt{n}\theta - c) + \Phi(-\sqrt{n}\theta - c)$ converges to one if $\theta \notin [-\delta, \delta]$, and converges to zero if $\theta \in (-\delta, \delta)$. As a result, the large sample posterior distributions with and without significance are truncated versions of the prior, with the prior truncated at $(-\delta, \delta)$ under significance, and at $(-\infty, -\delta) \cup (\delta, \infty)$ under no significance. If δ is large both significance and non-significance are informative. If, however, δ is small, we go back to the setting where significance carries only

local-to-zero information. Figure 4 reports posterior distributions for $\delta = \{0.5, 1, 2\}$, $\alpha = 0.05$ and $n = 10000$.

5. Conditioning on the sign of the estimated coefficient

In previous sections we have shown that statistical significance may carry very little information in large samples. As a result, the values of other statistics should be taken into account along with significance when the null is rejected in a significance test. As discussed above, in a Normal (or asymptotically Normal) setting it does not take much to go back to full information (e.g., P -value and the sign of $\hat{\theta}$). Here we consider the question of whether minimally augmenting the information on significance with the sign of $\hat{\theta}$ results in informativeness when the null is rejected. This exercise is motivated by the possibility that the sign of the estimated coefficient is implicitly taken into account in many discussions of results from significance tests.

For concreteness, we will concentrate on the case of a positive coefficient estimate, $\hat{\theta} > 0$. That is, the limited information posterior under significance and positive $\hat{\theta}$ conditions on the event $\sqrt{n}\hat{\theta} > c$. The case with negative $\hat{\theta}$ is analogous. Using similar calculations as in section 1, we obtain:

$$p(\theta | \sqrt{n}\hat{\theta} > c) = \frac{\frac{1}{\sigma}\phi\left(\frac{\theta - \mu}{\sigma}\right)\Phi(\sqrt{n}\theta - c)}{\Phi\left(\frac{\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right)},$$

and

$$p(\theta | 0 < \sqrt{n}\hat{\theta} \leq c) = \frac{\frac{1}{\sigma}\phi\left(\frac{\theta - \mu}{\sigma}\right)\left(1 - \Phi(\sqrt{n}\theta - c) - \Phi(-\sqrt{n}\theta)\right)}{1 - \Phi\left(\frac{\sqrt{n}\mu - c}{\sqrt{1 + n\sigma^2}}\right) - \Phi\left(\frac{-\sqrt{n}\mu}{\sqrt{1 + n\sigma^2}}\right)}.$$

Figure 5 reproduces the setting of Figure 1 but for the case when the posterior is conditional on sign of the estimate in addition to significance. Like in Figure 1, failure to reject carries substantial

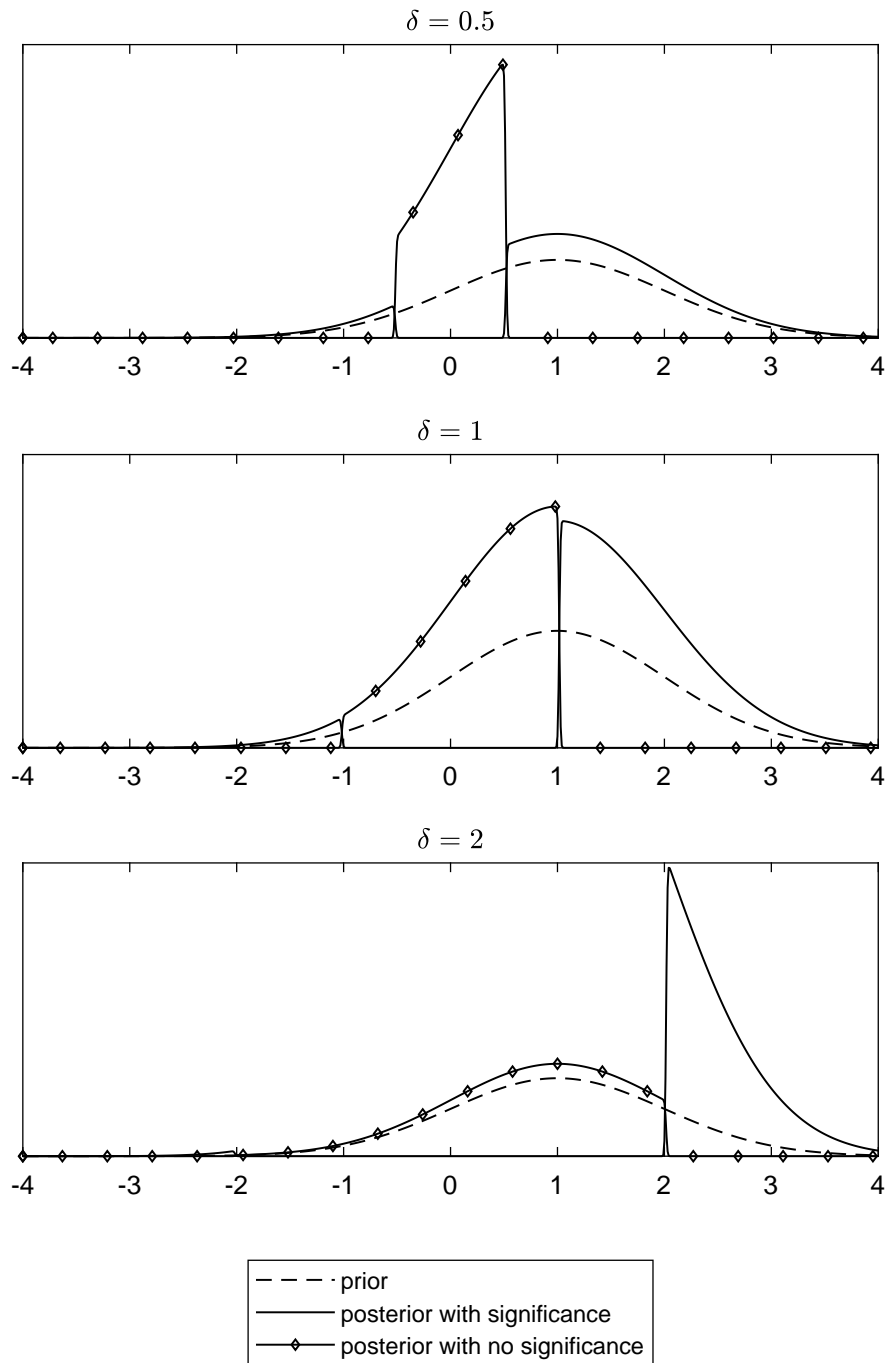


Figure 4: Posterior After a Test of the Null $\theta \in [-\delta, \delta]$ ($n = 10000$, $\alpha = 0.05$)




Figure 5: Posterior Distributions Conditional of Significance and Coefficient Sign

information. In fact, both outcomes of the significance test carry additional information, with respect to the setting in Figure 1, which of course is explained by the additional information in the sign of $\hat{\theta}$.

Notice that, in this case, under significance, the ratio between the posterior and the prior converges to

$$\lim_{n \rightarrow \infty} \frac{p(\theta | \sqrt{n}\hat{\theta} > c)}{p(\theta)} = \begin{cases} 0 & \text{if } \theta < 0, \\ \Phi(-c)/\Phi(\mu/\sigma) & \text{if } \theta = 0, \\ 1/\Phi(\mu/\sigma) & \text{if } \theta > 0. \end{cases}$$

Without significance, the ratio between the posterior and the prior converges to

$$\lim_{n \rightarrow \infty} \frac{p(\theta | 0 < \sqrt{n}\hat{\theta} \leq c)}{p(\theta)} = \begin{cases} 0 & \text{if } \theta \neq 0, \\ \infty & \text{if } \theta = 0. \end{cases}$$

That is, as $n \rightarrow \infty$ non-significance is highly informative. Under significance, the posterior of θ converges to the prior truncated at zero. As a result, in this case the informational content of significance depends on the value of $\Pr(\theta > 0) = \Phi(\mu/\sigma)$. If this quantity is small, significance

with a positive sign is highly informative. Unsurprisingly, when μ/σ is large (that is, in cases where there is little uncertainty about the sign of the parameter of interest), a positive sign of $\hat{\theta}$ does not add much to the informational content of the test. Moreover, the limit of $p(\theta|\sqrt{n}\hat{\theta} > c)$ cannot be more than double the value of $p(\theta)$ as long as μ is non-negative. This is relevant to many instances in economics where there are strong belief about the sign of the estimated coefficients (e.g., the slope of the demand function, or the effect of schooling on wages) and specifications reporting “wrong” signs for the coefficients of interest are rarely reported or published.³

6. Conclusions

Significance testing on a point null is the most extended form of inference in empirical economics. In this article, we have shown that rejection of a point null often carries very little information, while failure to reject is highly informative. This is especially true in empirical contexts that are typical in economics, where data sets are large (and, if anything, are becoming larger) and where there are rarely reasons to put substantial prior probability on a point null. Our results challenge the usual practice of conferring point null rejections a higher level of scientific significance than non-rejections. In consequence, we advocate a visible reporting and discussion of non-significant results in empirical practice (e.g., as in Angrist et al., 2017; Cantoni, 2018; Krueger and Malečková, 2003).

References

- Andrews, I. and M. Kasy (2017). Identification of and correction for publication bias. Working paper.
- Angrist, J. D., V. Lavy, J. Leder-Luis, and A. Shany (2017). Maimonides rule redux. NBER Working Paper 23486.
- Benjamin, D. J., J. O. Berger, M. Johannesson, B. A. Nosek, E.-J. Wagenmakers, R. Berk, K. A.

³The notion that empirical estimates might display “wrong” signs is widespread to the point that econometric articles and textbooks discuss this phenomenon and, in some cases, potential remedies. See, e.g., Wooldridge (2016) and Kennedy (2005).

- Bollen, B. Brembs, L. Brown, C. Camerer, et al. (2017). Redefine statistical significance. *Nature Human Behavior*, 6–10.
- Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*. New York: Springer-Verlag.
- Cantoni, E. (2018). Got ID? The zero effects of voter ID laws on county-level turnout, vote shares, and uncounted ballots, 1992-2014. Working paper.
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLOS Medicine* 2(8), 696–701.
- Kennedy, P. E. (2005). Oh no! I got the wrong sign! What should I do? *The Journal of Economic Education* 36(1), 77–92.
- Krueger, A. B. and J. Malečková (2003, December). Education, poverty and terrorism: Is there a causal connection? *Journal of Economic Perspectives* 17(4), 119–144.
- Leamer, E. E. (1978). *Specification searches: Ad hoc inference with nonexperimental data*. New York: Wiley.
- Pratt, J. W. (1965). Bayesian interpretation of standard inference statements. *Journal of the Royal Statistical Society. Series B (Methodological)* 27(2), 169–203.
- Robert, C. P. (2014). On the Jeffreys-Lindley paradox. *Philosophy of Science* 81(2), 216–232.
- Wasserstein, R. L. and N. A. Lazar (2016). The ASA’s statement on p-values: Context, process, and purpose. *The American Statistician* 70(2), 129–133.
- Wooldridge, J. M. (2016). *Introductory econometrics : A modern approach*. Boston: Cengage Learning.
- Ziliak, S. and D. McCloskey (2008). *The Cult of Statistical Significance: How the Standard Error Costs Us Jobs, Justice, and Lives*. Economics, Cognition, And Society. University of Michigan Press.