



HAL
open science

The multivariate point null testing problem: A bayesian discussion

Miguel A. Gómez-Villegas, Beatriz González-Pérez

► **To cite this version:**

Miguel A. Gómez-Villegas, Beatriz González-Pérez. The multivariate point null testing problem: A bayesian discussion. *Statistics and Probability Letters*, 2010, 78 (17), pp.3070. 10.1016/j.spl.2008.05.010 . hal-00644715

HAL Id: hal-00644715

<https://hal.science/hal-00644715>

Submitted on 25 Nov 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accepted Manuscript

The multivariate point null testing problem: A bayesian discussion

Miguel A. Gómez-Villegas, Beatriz González-Pérez

PII: S0167-7152(08)00269-1
DOI: 10.1016/j.spl.2008.05.010
Reference: STAPRO 5089

To appear in: *Statistics and Probability Letters*

Received date: 15 December 2005

Revised date: 30 October 2007

Accepted date: 1 May 2008

Please cite this article as: Gómez-Villegas, M.A., González-Pérez, B., The multivariate point null testing problem: A bayesian discussion. *Statistics and Probability Letters* (2008), doi:10.1016/j.spl.2008.05.010

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



The Multivariate Point Null Testing Problem: A Bayesian Discussion

Miguel A. Gómez-Villegas and Beatriz González-Pérez

*Departamento de Estadística e Investigación Operativa, Facultad de CC.
Matemáticas, Universidad Complutense de Madrid. Madrid, 28040, Spain*

Abstract

In this paper the problem of testing a multivariate point hypothesis is considered. Of interest is the relationship between the p-value and the posterior probability. A Bayesian test for simple $H_0 : \theta = \theta_0$ versus bilateral $H_0 : \theta \neq \theta_0$, with a mixed prior distribution for the parameter θ , is developed. The methodology consist on fixing a sphere of radius δ around θ_0 and assigning a prior mass, π_0 , to H_0 by integrating the density $\pi(\theta)$ over this sphere and spreading the remainder, $1 - \pi_0$, over H_1 according to $\pi(\theta)$. A theorem that shows when the frequentist and Bayesian procedures can give rise to the same decision is proved. Then, some examples are revisited.

Key words:

point null hypothesis; p-values; posterior probabilities; Lindley's paradox.

Email addresses: ma.gv@mat.ucm.es, beatrizg@mat.ucm.es (Miguel A. Gómez-Villegas and Beatriz González-Pérez).

1 Introduction

Let $f(x_1, \dots, x_n | \theta)$ ($\theta = (\theta_1, \dots, \theta_m)$ unknown) the likelihood of a sample (X_1, \dots, X_n) . For a specified $\theta_0 = (\theta_{01}, \dots, \theta_{0m})$ we want to test

$$H_0 : \theta = \theta_0 \text{ versus } H_1 : \theta \neq \theta_0. \quad (1)$$

As usual, we use the following mixed prior distribution to test (1)

$$\pi^*(\theta) = \pi_0 I_{H_0}(\theta) + (1 - \pi_0) \pi(\theta) I_{H_1}(\theta). \quad (2)$$

For a proper metric d and a value of δ sufficiently small, consider

$$H_{0\delta} : d(\theta_0, \theta) \leq \delta, \text{ versus } H_{1\delta} : d(\theta_0, \theta) > \delta.$$

Denote $B(\theta_0, \delta) = \{ \theta \in \Theta^r, \sum_{i=1}^m (\theta_i - \theta_{0i})^2 \leq \delta^2 \}$. Gómez-Villegas, Maín and Sanz (2007) propose to compute π_0 by means of

$$\pi_0 = \int_{B(\theta_0, \delta)} \pi(\theta) d\theta. \quad (3)$$

By this way, π_0 is the probability assigned to H_0 , through $\pi^*(\theta)$, and $H_{0\delta}$, through $\pi(\theta)$, the Kullback-Leibler discrepancy between $\pi(\theta)$ and $\pi^*(\theta)$ goes to zero as δ goes to zero, and the posterior probability of the null is

$$P_{\pi^*}(\theta_0 | x_1, \dots, x_n) = \left[1 + \frac{1 - \pi_0}{\pi_0} \eta(x_1, \dots, x_n) \right]^{-1}, \quad (4)$$

$$\eta(x_1, \dots, x_n) = \frac{\int_{\Theta} f(x_1, \dots, x_n | \theta) \pi(\theta) d\theta}{f(x_1, \dots, x_n | \theta_0)}. \quad (5)$$

Suppose that H_0 is rejected when $P_{\pi^*}(\theta_0 | x_1, \dots, x_n) \leq 1/2$. On the other hand, if $\Lambda(x_1, \dots, x_n) = \sup_{\theta \in \Theta} f(x_1, \dots, x_n | \theta) / f(x_1, \dots, x_n | \theta_0)$ is the test

statistic, then the associated p-value,

$$p(x_1, \dots, x_n) = P(\Lambda(X_1, \dots, X_n) > \Lambda(x_1, \dots, x_n) | \theta_0), \quad (6)$$

can be approximated by $P(\chi_m^2 > 2 \ln \Lambda(x_1, \dots, x_n))$. In this case, H_0 is rejected when $p(x_1, \dots, x_n) \leq p^*$, for a level of significance p^* . The aim is to find conditions for δ in order that, when p^* is fixed, whatever $(x_1, \dots, x_n) \in \chi$, the same decision is reached with both methods. In section 2, a theorem that shows when both approaches agree is proved. In section 3 some relevant examples are revisited. Conclusions and comments are included in section 4.

2 Comparison between both approaches

In parametric testing of a simple null hypothesis, as opposed to the one-sided (see Casella and Berger , 1987), Bayesian and frequentist procedures can give rise to different decisions (see Lindley , 1957; Berger and Sellke , 1987; Berger and Delampady , 1987, among others). In most of the Bayesian approaches, the infimum of a Bayesian evidence measure over a wide class of priors is computed and is substantially larger than the p-value. Other important references are Oh and Dasgupta (1999), Gómez-Villegas and Sanz (2000), Sellke, Bayarri and Berger (2001), Gómez-Villegas, Maín and Sanz (2002) and Gómez-Villegas and González-Pérez (2005, 2006).

2.1 Preliminaries

In this section we introduce some definitions and results in order to prove a characterization theorem of the agreement between both methods to test (1).

Definition 2.1 Let $\pi(\theta)$ be the prior about θ . $T = T(X_1, \dots, X_n)$ is a *sufficient statistic to test* $H_0 : \theta \in \Theta_0$, versus $H_1 : \theta \in \Theta_1$, with $\Theta_0 \cap \Theta_1 = \phi$ and $\Theta_0 \cup \Theta_1 = \Theta$, if $P_\pi(\Theta_0|x_1, \dots, x_n) = P_\pi(\Theta_0|t)$, when $T(x_1, \dots, x_n) = t$.

Note that the Bayes factor η given in (5) is a *sufficient statistic to test* the multivariate point null given in (1) when the mixed prior (3) is used. In fact, with this prior choice, if T is a *sufficient statistic to test* (1), then there is a function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $g(T) = \eta$. Moreover, the usual definition of *sufficient statistic* is not equal to this new concept of *sufficient statistic to test*. If T is a *sufficient statistic* then T verifies definition 2.1. The reciprocal is not true. We introduce the next concept in order to compare two statistics.

Definition 2.2 Let $T_1 = T_1(X_1, \dots, X_n)$ and $T_2 = T_2(X_1, \dots, X_n)$ be univariate statistics. T_1 is an *increasing tendency statistic* with respect to T_2 in a value $T_1 = t$ if
$$\sup_{T_1(x_1, \dots, x_n) < t} T_2(x_1, \dots, x_n) \leq \inf_{T_1(x_1, \dots, x_n) \geq t} T_2(x_1, \dots, x_n).$$

Proposition 2.1 If $T_1 = h(T_2)$ and $h : \mathbb{R} \rightarrow \mathbb{R}$ is a non-decreasing monotonous function, then T_1 is an *increasing tendency statistic* with respect to T_2 for any value $T_1 = t$. Furthermore, when h is a strictly increasing function, then

$$\sup_{T_1(x_1, \dots, x_n) < t} T_2(x_1, \dots, x_n) = \inf_{T_1(x_1, \dots, x_n) \geq t} T_2(x_1, \dots, x_n) = h^{-1}(t).$$

2.2 Agreement between frequentist and Bayesian approaches

Theorem 2.1 Suppose that we wish to test (1) with the prior distribution given in (2) and π_0 as in (3). If Λ is an *increasing tendency statistic* with respect to η in $\Lambda = \lambda^*$, then for $p^* = P_{\theta_0} \{\Lambda \geq \lambda^*\}$ there is an interval of values of $\pi_0 = \pi_0(\delta)$, $I = I(p^*, n) = (\ell_1, \ell_2)$, where both methods are agree.

Proof Define $A_\eta(\kappa) = \{(x_1, \dots, x_n), \eta(x_1, \dots, x_n) = \kappa\}$. $P_{\pi^*}(\theta_0|\kappa) > 1/2$ is verified when $\pi_0 > \kappa(\kappa + 1)^{-1}$. $\pi^0(\kappa) = \kappa(\kappa + 1)^{-1}$ is a strictly increasing function. Moreover, if $\lambda_1 < \lambda_2$, then $p(\lambda_1) = P_{\theta_0}\{\Lambda \geq \lambda_1\} \geq P_{\theta_0}\{\Lambda \geq \lambda_2\} = p(\lambda_2)$. If Λ is an *increasing tendency statistic* with respect to η in $\Lambda = \lambda^*$ then $\kappa_* = \sup_{(x_1, \dots, x_n), p > p^*} \eta = \sup_{\Lambda < \lambda^*} \eta \leq \inf_{\Lambda \geq \lambda^*} \eta = \inf_{(x_1, \dots, x_n), p \leq p^*} \eta = \kappa^*$ and $\ell_1 = \pi^0(\kappa_*) = \sup_{(x_1, \dots, x_n), p > p^*} \pi^0(\eta) \leq \inf_{(x_1, \dots, x_n), p \leq p^*} \pi^0(\eta) = \pi^0(\kappa^*) = \ell_2$. are verified for $p^* = P_{\theta_0}\{\Lambda \geq \lambda^*\}$. Consider $\pi_0 \in (\ell_1, \ell_2)$. If $(x_1, \dots, x_n) \in A_\eta(\kappa)$, $\kappa < \kappa_* \leq \kappa^*$, then $\pi_0 > \ell_1 > \kappa(\kappa + 1)^{-1}$ and $P_{\theta_0}\{\Lambda \geq \Lambda(x_1, \dots, x_n)\} > p^*$. On the other hand, if $(x_1, \dots, x_n) \in A_\eta(\kappa)$, $\kappa \geq \kappa^* \geq \kappa_*$, then $\pi_0 < \ell_2 \leq \kappa(\kappa + 1)^{-1}$ and $P_{\theta_0}\{\Lambda \geq \Lambda(x_1, \dots, x_n)\} \leq p^*$.

Corollary 2.1 Let $\Lambda = h(\eta)$ with $h : \mathbb{R} \rightarrow \mathbb{R}$ be a non-decreasing monotonous function. If $(x_1, \dots, x_n) \in A_\eta(\kappa)$ is observed, the same decision is reached with the posterior probability $P_{\pi^*}(\theta_0|\kappa) = [1 + (1 - \pi_0(\delta))\kappa/\pi_0(\delta)]^{-1}$ and δ such that $\pi_0(\delta) \in (\ell_1, \ell_2)$,

$$\ell_1 = \ell_1(p^*, n) = \sup_{(x_1, \dots, x_n), p > p^*} \eta(\eta + 1)^{-1}, \quad (7)$$

$$\ell_2 = \ell_2(p^*, n) = \inf_{(x_1, \dots, x_n), p \leq p^*} \eta(\eta + 1)^{-1}, \quad (8)$$

and $p^* = P(\Lambda \geq \lambda^*)$ to quantify the p-value, $p(h(\kappa)) = P_{\theta_0}\{\Lambda \geq h(\kappa)\}$.

Moreover, when h is strictly increasing continuous, then $\eta = h^{-1}(\Lambda)$ and

$$\ell_1 = \ell_2 = \pi_0(\delta, p^*, n) = h^{-1}(\lambda^*) [h^{-1}(\lambda^*) + 1]^{-1}. \quad (9)$$

The proof of Corollary 2.1 is easy by using theorem 2.1 and proposition 2.1. Another immediate outcome is that, when the sample size n and the level of significance p^* are fixed, the verification of $\ell_1 \leq \ell_2$, defined in (7) and (8), is sufficient for the existence of agreement between the p-value and the posterior

probability. Moreover, when $\Lambda = h(\eta)$ and $h : \mathbb{R} \rightarrow \mathbb{R}$ is non-decreasing monotonous, $\ell_1 \leq \ell_2$ is always true whatever p^* , and if h is strictly increasing continuous then $\ell_1 = \ell_2$ and the agreement is obtained for (9).

3 Applications

3.1 Lindley's Paradox

Let (x_1, \dots, x_m) be a sample from a univariate normal of mean θ and known variance σ^2 and the prior mass of $H_0 : \theta = \theta_0$, be c . Distribute the remainder $1 - c$ uniformly over some interval I containing θ_0 . \bar{x} , the arithmetic mean, and a minimal sufficient statistic, is well within I for n sufficiently enough. Then $\bar{x} - \theta_0$ tends to zero as n increases and the posterior probability that $\theta = \theta_0$ is

$$\bar{c} = cK^{-1} \exp[-n(\bar{x} - \theta_0)^2 / (2\sigma^2)], \quad (10)$$

$K = c \exp[-n(\bar{x} - \theta_0)^2 / (2\sigma^2)] + (1 - c) \int_I \exp[-n(\bar{x} - \theta)^2 / (2\sigma^2)] d\theta$. By the assumption about \bar{x} and I , the last integral can be evaluated as $\sigma\sqrt{2\pi/n}$. By using the usual test $\bar{x} = \theta_0 + \lambda_\alpha \sigma\sqrt{n}$ [$\Phi(\lambda_\alpha) = 1 - \alpha/2$, Φ being the normal standard distribution function] is significant at the $\alpha\%$ level and

$$\bar{c} = c \exp(-\lambda_\alpha^2/2) \left[c \exp(-\lambda_\alpha^2/2) + (1 - c) \sigma\sqrt{2\pi/n} \right]^{-1}. \quad (11)$$

From (11) we see that $\bar{c} \rightarrow 1$ when $n \rightarrow \infty$ and whatever c , it can be found $n = n(c, \alpha)$ such that “ \bar{x} is significantly different from θ_0 at the $\alpha\%$ level” and “the posterior probability that $\theta = \theta_0$ is $100(1 - \alpha)\%$ ”. This is the paradox. Both statements are in direct conflict. However, Lindley (1957) re-

mark too that if $A = c \exp(-\lambda_\alpha^2/2)[(1-c)\sqrt{2\pi}]^{-1}$, then $\bar{c} = A(A + \sigma/\sqrt{n})^{-1}$ and $\bar{c} \rightarrow 0$ as $\sigma/\sqrt{n} \rightarrow \infty$. Therefore, in a small experiment, significance at 5% may give good reasons to doubt the null hypothesis. On the other hand when n and α are fixed, there is a value of $c = c(n, \alpha)$, such that both approaches are agree. Observe that the posterior probability (10) may be written $\bar{c} = [1 + (1-c)h^{-1}[T(\bar{x}, \theta_0)]/c]^{-1}$, $h^{-1}(u) = \sigma\sqrt{2\pi/n}\exp(u/2)$, $T(\bar{x}, \theta_0) = n(\bar{x} - \theta_0)^2/\sigma^2$. Lindley's argument to show a paradox consist on finding a value of $n = n(c, \alpha)$ such that both decisions differ radically. However, by applying Corollary 2.2 both agree for $c = h^{-1}(\chi_{1,\alpha}^2) [h^{-1}(\chi_{1,\alpha}^2) + 1]^{-1} = [1 + n^{1/2}(2\pi\sigma^2)^{-1/2}\exp(-\chi_{1,\alpha}^2/2)]^{-1}$. Note that if $\alpha = 0.05$, then $c = 1/2$ when $n/\sigma^2 \approx 300$. In general, for c and α fixed, $n/\sigma^2 = 2\pi c^{-2}(1-c)^2 \exp(\chi_{1,\alpha}^2)$.

3.2 Lower bounds for unimodal and symmetric priors

Let $\mathbf{X} = (X_1, \dots, X_m)' \sim N_m(\theta, \sigma^2 \mathbf{I})$ with σ^2 known, \mathbf{I} the identity matrix $m \times m$ and $\theta = (\theta_1, \dots, \theta_m)'$ unknown. For testing (1) with a sample of size n the usual test statistic is $T(\bar{\mathbf{X}}, \theta_0) = n\sigma^{-2}(\bar{\mathbf{X}} - \theta_0)'(\bar{\mathbf{X}} - \theta_0)$, with $\bar{\mathbf{X}} = (\bar{X}_1, \dots, \bar{X}_m)'$, and the p-value of the observed data, $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_m)'$, is $p(\bar{\mathbf{x}}) \approx P\{\chi_m^2 \geq T(\bar{\mathbf{x}}, \theta_0)\}$. If $\pi^*(\theta)$ is the mixed prior given in (2) with π_0 computed by (3), the infimum of the posterior probability of the point null when $\pi(\theta) \in Q_{US}$, Q_{US} being the unimodal and symmetric priors about θ_0 , is

$$\inf_{\pi \in Q_{US}} P(H_0|\bar{\mathbf{x}}) = [1 + 2^{m/2}\delta^{*-m}\Gamma(m/2 + 1)\exp[T(\bar{\mathbf{x}}, \theta_0)/2]]^{-1},$$

where $\delta^* = \delta\sqrt{n}/\sigma^2$ (see Gómez-Villegas, Maín and Sanz, 2007) Let t^* be such that $P\{\chi_m^2 \geq t^*\} = p^*$. By applying Corollary 2.2 both methods are always agree when $\delta^* = [2^{m/2}\Gamma(m/2 + 1)\exp(t^*/2)]^{1/m}$. Table 1 shows the values of

δ^* obtained for different values of m and p^* and both methods are compared when $t = T(\bar{\mathbf{x}}, \theta_0) = 7$. We compute the p-values, and the infimum of the posterior probabilities over the values of δ^* given in Table 1, when $m = 2$, $m = 5$, $m = 10$ and $m = 20$. For instance, when $m = 2$, the p-value is $p(7) = 0.0302$. Therefore, a frequentist statistician who uses $p^* = 0.05$, rejects H_0 , whereas with $p^* = 0.01$, accepts H_0 . The same result is obtained by a Bayesian with $\delta^* = 6.32$ and $\delta^* = 14.14$. In this case, $\inf_{\pi \in Q_{us}} P(H_0|t = 7, \delta^* = 6.32) = 0.46066$ and $\inf_{\pi \in Q_{us}} P(H_0|t = 7, \delta^* = 14.14) = 0.81027$.

Table 1 goes here

3.3 Lower bounds for scale mixture of Normals

In the same context of subsection 3.2, we want to test (1) with $\theta_0 = (0, \dots, 0)'$ and a sample of size n . The usual test statistic is $T(\bar{\mathbf{X}}) = n\sigma^{-2}\bar{\mathbf{X}}'\bar{\mathbf{X}}$. If $m > 2$ the infimum of the posterior probability of H_0 when $\pi(\theta) \in Q_N$, $Q_N = \{\pi(\theta|v^2) \approx N_m(\mathbf{0}, v^2\mathbf{I}), \pi(v^2) \text{ non-decreasing on } (0, \infty)\}$, is

$$\inf_{\pi \in Q_N} P(H_0|t) = [1 + \delta^{*-2}F_{m-2}(t)/f_m(t)]^{-1},$$

$t = n\bar{\mathbf{x}}'\bar{\mathbf{x}}/\sigma^2$, $\delta^* = \delta\sqrt{n}/\sigma^2$, F_{m-2} is the distribution function of a χ_{m-2}^2 and f_m is the density function of a χ_m^2 (see Gómez-Villegas, Maín and Sanz, 2007). Let t^* such that $P\{\chi_m^2 \geq t^*\} = p^*$. By applying corollary 2.2 both methods always agree when $\delta^*(p^*, n, m) = [F_{m-2}(t^*)/f_m(t^*)]^{1/2}$. Table 2 shows the values of δ^* computed for different values of m and p^* and both methods are compared when $t = T(\bar{\mathbf{x}}) = 20$. The p-values and the infimums on the values of δ^* of Table 3 are computed. The p-value is $p(20) = 0.02925$ when $m = 10$. Therefore, a frequentist statistician who uses $p^* = 0.05$, re-

jects H_0 , whereas with $p^* = 0.01$, accepts H_0 . The same result is reached by a Bayesian by using, respectively, the values $\delta^* = 7.96$ and $\delta^* = 17.01$ given in Table 2. In this case $\inf_{\pi \in Q_{us}} P(H_0 | t = 20, \delta^* = 7.96) = 0.3772$ and $\inf_{\pi \in Q_{us}} P(H_0 | t = 7, \delta^* = 14.14) = 0.7343$.

Table 2 goes here

4 Conclusions and Comments

An important conclusion is that p-values and posterior probabilities can be reconciled in the multivariate point null testing problem. The methodology consist on assigning a prior mass to θ_0 computed by the probability of a sphere of radius δ centered at θ_0 assigned by a density $\pi(\theta)$. The p-value and the posterior probability of the null for the mixed prior $\pi^*(\theta) = \pi_0 I_{H_0}(\theta) + (1 - \pi_0) \pi(\theta) I_{H_1}(\theta)$ are computed. This procedure allows to prove a theorem that shows when and how both approaches agree. The analyzed examples show that such agreement is always possible when $\delta^* = \delta^*(p^*, n, m) = \delta \sqrt{n} / \sigma^2$. This is due to the infimum of the posterior probability over the prior classes used depends on the usual test statistic through an increasing function. When this dependence is not possible, the agreement is in terms of a sufficient condition.

References

- Berger, J. O. and Delampady, M. (1987) Testing Precise Hypotheses, (with Discussion). *Statistical Science*, **2**(3), 317-352.
- Berger, J. O. and Sellke, T. (1987) Testing a point null hypothesis: The irrec-

- onciliability of p-values and evidence, (with discussion). *J. Amer. Statist. Assoc.*, **82**, 112-139.
- Casella, G. and Berger, R. L. (1987) Reconciling bayesian and frequentist evidence in the one-sided testing problem, (with discussion). *J. Amer. Statist. Assoc.*, **82**, 106-135.
- Gómez-Villegas, M. A. and González-Pérez, B. (2005) Bayesian analysis of contingency tables. *Commun. Statist-Theory Meth.*, 34(8), 1743-1754.
- Gómez-Villegas, M. A. and González-Pérez, B. (2006) A condition to obtain the same decision in the homogeneity testing problem from the frequentist and Bayesian point of view. *Commun. Statist-Theory Meth.*, 35, 2211-2222.
- Gómez-Villegas, M. A., Maín, P. and Sanz, L. (2002) A suitable Bayesian approach in testing point null hypothesis: some examples revisited. *Commun. Statist.-Theory Meth.*, **31**(2), 201-217.
- Gómez-Villegas, M. A., Maín, P. and Sanz, L. (2007) Contraste de hipótesis nula puntual multivariante para variables normales correladas. *Actas del XXX Congreso Nacional de Estadística e Investigación Operativa y IV Jornadas de Estadística Pública 2007*.
- Gómez-Villegas, M. A. and Sanz, L. (2000) ε -contaminated priors in testing point null hypothesis: a procedure to determine the prior probability. *Statistics & Probability Letters*, **47**, 53-60.
- Lindley, D. V. (1957) A statistical paradox. *Biometrika*, **44**, 187-192.
- Oh, H. S. and Dasgupta, A. (1999) Comparison of the p-value and posterior probability. *J. of Statist. Planning and Inference*, **76**, 93-107.
- Sellke, T., Bayarri, M. J. and Berger, J. O. (2001) Calibration of p-values for testing point null hypotheses. *The American Statistician*, **55**, 62-71.