A NOTE ON JEFFREYS-LINDLEY PARADOX

by

Christian P. Robert
Université Paris VI and Purdue University

Technical Report #91-52C

Department of Statistics
Purdue University

September 1991

A Note on Jeffreys-Lindley Paradox

CHRISTIAN P. ROBERT[1]

Université Paris VI and Purdue University

ABSTRACT. The Jeffreys-Lindley paradox, namely the fact that a point null hypothesis will always be accepted when the prior variance goes to infinity, has often been argued to imply prohibiting the use of improper priors in hypothesis testing. We reevaluate this paradox by considering the role of the prior hypotheses probabilities and obtain a noninformative answer which is equivalent decisionwise to the classical *p-value*.

1. Introduction

In hypothesis testing, it is well-known that Bayesian and frequentist answers may differ drastically. For instance, Berger and Sellke (1987) and Berger and Delampady (1987) have shown that the smallest posterior probability of a point null hypothesis is usually much larger than the corresponding frequentist answer, i.e. the *p-value*. Lindley (1957) shows that the disagreement may be dramatic, in the following sense. Let $x \sim \mathcal{N}(\theta, 1)$ and the null hypothesis to test is $H_0 : \theta = 0$. If one uses conjugate priors, $\theta \sim \mathcal{N}(0, \sigma^2)$, the posterior probability of H_0,

$$(1.1) \qquad \left[1 + \frac{1-\pi_o}{\pi_0} \frac{e^{-x^2/2(\sigma^2+1)}}{e^{-x^2/2}} \frac{1}{\sqrt{\sigma^2+1}}\right]^{-1},$$

goes to 1 as σ^2 goes to infinity, whatever π_0 and x are.

This result is statistically paradoxical because, first, large σ^2 somehow correspond to a noninformative setup and, therefore, noninformative answers seem to be impossible to provide for this problem. Secondly, it is usually the case in estimation settings that the limit of conjugate estimators is equivalent to a "classical" frequentist answer and this property does not seem to occur for hypothesis

Key words and phrases. Hypothesis testing, vague prior, noninformative answer, p-value.
AMS Subject Classification (1985): 62A15, 62F03, 62C10.

[1]This research was supported by NSF Grant DMS-8717799, while the author was visiting Purdue University. I wish to thank Jim Berger and Ed George for helpful discussions.

testing. Obviously, the fact that (1.1) goes to 1 is not a mathematical paradox since the prior sequence is giving less and less mass to any neighborhood of 0 as σ^2 goes to infinity.

Many authors have commented on this paradox, either to criticize the Bayesian approach (Shafer, 1982) or to dismiss the use of improper priors for testing (Jeffreys, 1961 and DeGroot, 1982). Berger (1991) considers that it shows that a noninformative answer is not possible in this context and moreover, that it is in accordance with "Occam's razor" rule, i.e. that between two equally likely explanations, we should always choose the simplest one if no additional argument supports the other one (Berger and Jeffreys, 1991).

However, recent decision-theoretic considerations of the testing problem in Hwang *et al.* (1991) have shown that improper priors were definitely necessary. For instance, in the Jeffreys-Lindley setup, the p-value $p(x) = 2(1 - \Phi(|x|))$ is inadmissible under squared-error loss,

$$(\mathbb{I}_0(\theta) - p(x))^2,$$

where $\mathbb{I}$ denotes the indicator function, but cannot be dominated by a proper Bayes estimator, i.e. a true posterior probability. Furthermore, generalized Bayes answers, i.e. solutions of the form

$$\frac{\pi_0 \varphi(x)}{\pi_0 \varphi(x) + (1 - \pi_0) \int_{-\infty}^{+\infty} \varphi(x - \theta) \pi_1(\theta) d\theta} \tag{1.2}$$

where φ is the standard normal density, π_1 is a σ-finite measure and π_0 a prior weight, are also admissible under squared-error loss and form a *minimal complete class*.

As pointed out in DeGroot (1982), the trouble with using improper priors is that if one replaces the σ-finite measure $\pi_1(\theta)$ by $c\pi_1(\theta)$, the constant c can be chosen to give any desired answer. We will show in the next section that there exists a way to obtain the "proper" constant c for the Jeffreys prior by considering again a sequence of conjugate priors. The resulting noninformative answer is then no more uniformly equal to 1 and, furthermore, provides an estimator which is surprisingly close to the classical p-value (for most decision purposes).

2. Reweighting the alternatives

The fundamental (and simple) argument underlying our reevaluation of the Jeffreys-Lindley paradox is that the prior probability π_0 of the null hypothesis H_0 should depend on the prior variance under the alternative hypothesis H_1, σ^2. Such a dependence may seem absurd but consider that, from a Bayesian point of view, we are actually testing $H_0 : \theta = 0$ versus $H_1 : \theta \sim \mathcal{N}(0, \sigma^2)$. Therefore, the prior probability of H_1 (and therefore of H_0) may vary with σ^2. Indeed, while taking $\pi_0 = 1/2$ may appear as the fairest (or the most objective) choice, it does not take into account the fact that the alternative prior π_1 considers a larger set of possible values of θ as σ^2 increases, i.e. that the "effective support" of π_1 (say, the 99% HPD region) is getting larger as σ^2 goes to infinity. Larger σ^2 do not exclude smaller values of θ but, in the contrary increase the range of values of θ compatible with H_1. In this sense, an increasing sequence of σ^2 leads to a sequence of *imbedded* alternative hypotheses. Therefore, the prior probability of H_1 should *increase* with σ^2. Such a dependency is also justified if we look at it the other way: a restriction of the range of possible values for θ under H_1 can result from some observations which are incompatible with the previous range of π_1 and which, therefore, partially argue against H_1. It is thus coherent to lower the prior probability of H_1 when the range of π_1 is decreasing. It is because π_0 is kept constant that the Jeffreys-Lindley paradox occurs; we have to prevent the alternative prior mass from going to $\pm\infty$ too quickly. Casella and Berger (1987) noticed that $\pi_0 = 1/2$ was "too large" but they did not pursue the reasoning up to a prior dependent π_0.

A natural requirement on the sequence of priors is that they should give sufficient weight to the range of values of θ which actually caused H_0 to be tested, i.e. the θ's in a neighborhood of 0 which generate x's which could also originate from a $\mathcal{N}(0,1)$ distribution. Since, for σ large enough and a arbitrary, we have

$$\begin{aligned}\pi([-a,0[\cup]0,a]) &= (1-\pi_0)[\Phi(a/\sigma)-\Phi(-a/\sigma)]\\ &\simeq (1-\pi_0)\frac{2a}{\sigma}\varphi(0),\end{aligned}$$

it seems reasonable to impose the following restriction on π_0,

$$\frac{1-\pi_0(\sigma)}{\sigma} = c,$$

where c is a constant to be determined.

However, this constraint is too strong to hold as σ goes to infinity, since the prior probability of any fixed interval must go to 0. A more realistic requirement is therefore to choose $\pi_0(\sigma)$ in such a way that the ratio of the prior probability of the null hypothesis to the prior probability of the "reasonable" range, $[-a,a]\backslash\{0\}$, remains constant as σ goes to infinity, i.e.

$$(1-\pi_0(\sigma))[\Phi(a/\sigma)-\Phi(-a/\sigma)] \propto \pi_0(\sigma). \tag{2.1}$$

For σ large enough, this condition leads to the following equation

$$\frac{1-\pi_0(\sigma)}{\sigma} \propto \pi_0(\sigma). \tag{2.2}$$

In order to determine completely the dependency of π_0 on σ^2, i.e. the proportionality factor in the above relation, we consider that 0 should have the same weight under both alternatives, namely that the densities are equal at 0,

$$\pi_0 = (1-\pi_0)\frac{1}{\sqrt{2\pi}\sigma}. \tag{2.3}$$

This implies that 0 is "indifferent" under both alternatives, whatever σ is. Note that, under this constraint, $\pi_0(\sigma)$ goes to 0 when σ^2 goes to infinity. Such a behavior was also observed by Bernardo (1980) when implementing the reference prior approach in this setting. The posterior probability associated with (2.1) is then

$$\left[1+\sqrt{\frac{\sigma^2}{\sigma^2+1}}\sqrt{2\pi}\; e^{\sigma^2 x^2/2(\sigma^2+1)}\right]^{-1}, \tag{2.4}$$

which converges to

$$(1+\sqrt{2\pi}e^{x^2/2})^{-1} \tag{2.5}$$

when σ^2 goes to infinity. Note that (2.2) converges to 1 when σ^2 goes to 0, as it should since H_0 is then true a priori, while $\pi_0 = 1/2$ leads to 1/2 in (1.1).

A most interesting feature of (2.3) is that it also corresponds to the generalized Bayes answer associated with the Jeffreys prior $\pi_1(\theta) = 1$ and $\pi_0 = 1/2$ since (1.2) leads to

$$\frac{e^{-x^2/2}/\sqrt{2\pi}}{e^{-x^2/2}/\sqrt{2\pi}+1}$$

in this case. Therefore, when the prior probability of H_0 depends on the prior variance σ^2, the Jeffreys estimator is the limit of the conjugate answers, as it is in the point estimation case. Moreover, this result indicates that $c = 1$ is the proper constant in this case.

3. The resulting noninformative answer

The dependency of π_0 on σ^2 thus avoids the undesirable convergence to 1 and provides an estimator which can be considered as a noninformative answer, Bayesian counterpart to the p-value. However, the validity of our derivation may be questioned, since the limiting prior resulting from (2.1) has also some undesirable features. Actually, for every $\varepsilon > 0$, one has

$$\pi([-\varepsilon,\varepsilon]) = \pi_0(\sigma) + (1-\pi_0(\sigma))[\Phi(\varepsilon/\sigma) - \Phi(-\varepsilon/\sigma)],$$

where Φ is the standard normal cdf. Given (2.1), we get

$$\begin{aligned}\pi([-\varepsilon,\varepsilon]) &= \left[\frac{1}{\sqrt{2\pi}\sigma} + \Phi(\varepsilon/\sigma) - \Phi(-\varepsilon/\sigma)\right]\left(1+\frac{1}{\sqrt{2\pi}\sigma}\right)^{-1} \\ &= \frac{1}{1+\sqrt{2\pi}\sigma}\left(1+\sqrt{2\pi}\sigma[\Phi(\varepsilon/\sigma)-\Phi(-\varepsilon/\sigma)]\right),\end{aligned}$$

which converges to 0 as σ^2 goes to infinity. Therefore, the limiting prior gives no positive probability to any neighborhood of 0 and this behavior seems to be quite unreasonable. But this is usually the case with improper priors: they cannot be handled in the same way than subjective priors and, as pointed out in DeGroot (1982), *they should not be regarded as representing ignorance.* This feature of improper priors is present in most statistical problems and, therefore, should not prevent us to consider (2.3) as a possible noninformative answer.

Let us turn now to the behavior of the estimator (2.3). First, it is strictly smaller than the lower bound on the Bayesian estimators obtained by Berger and Sellke (1987),

$$(1 + e^{x^2/2})^{-1}. \tag{3.1}$$

Again, it may seem paradoxical that the noninformative answer does not belong to the range of the Bayesian answers but, contrary to point estimation, testing settings allow for discontinuities between proper and improper priors. Moreover, the bound (3.1) was obtained for $\pi_0 = 1/2$, while π_0 depends on σ^2 in our case. The difference between (2.3) and (3.1) also shows that, although (3.1) appears as the *least favorable Bayesian answer*, it still corresponds to an informative setting and, therefore, that the use of an *informative* (i.e. proper) *prior* makes a significant difference in the answer to a testing problem. This feature definitely separates testing from usual estimation problems but does not necessarily imply that improper priors should not be used.

Table 1 provides some numerical values of the noninformative estimator (2.3) for some values of x. In addition to the above mentioned discrepancy with the least favorable answer, an interesting feature of Table 3.1 is the closeness of (2.3) and the p-value, $p(x)$. Namely, when the p-value is between 0.10 and 0.01, (2.3) produces essentially the same numerical values. In other words, for the range of x's for which the exact value of $p(x)$ really matters, the noninformative approach leads to the same *decision* than the p-value. (Actually, H_0 will be usually accepted for an answer larger than 0.10 and rejected for an answer smaller than 0.01.) Therefore, *decisionwise*, the two approaches are somehow equivalent.

TABLE 1. Comparison of answers for the normal point null test.

x	0	1.68	1.96	2.58
Least favorable Bayesian answer	0.5	0.196	0.128	0.035
Noninformative answer	0.285	0.089	0.055	0.014
p-value	1	0.093	0.05	0.01

Obviously, this equivalence does not "rehabilitate" the p-value since the numerous undesirable features pointed out in the previously mentioned papers still exist and the noninformative answer is not necessarily a "good" answer. In the contrary, we could argue that the similarity we exhibited in this paper rather points out the need for additional (prior) information. Moreover, the coincidence between (2.3) and $p(x)$ only occurs on a small (although crucial) range of values of x and (2.3) is admissible under squared error loss, while $p(x)$ is not (see Hwang et al., 1991). However, it may also explain why the p-value has survived for such a long period despite its multiple drawbacks. The coincidence of the classical answer with a noninformative answer actually holds in other settings, as shown in Caron and Robert (1991) (who also consider an alternative noninformative approach leading to the same conclusions).

References

1. Berger, J. O. (1991) Personal communication.
2. Berger, J. O. and Delampady, M. (1987). Testing precise hypotheses. *Statist. Sci.* **2**, 317–352.
3. Berger, J. O. and Jeffreys, W. (1991). The application of robust Bayesian analysis to hypothesis testing and Occam's razor. Tech. Report #91-04, Stat. Dept., Purdue University.
4. Berger, J. O. and Sellke, T. (1987). Testing a point null hypothesis: the irreconciability of P-values and evidence (with discussion). *J. Amer. Statist. Assoc.* **82**, 112–139.
5. Bernardo, J. (1980). A Bayesian analysis of classical hypothesis testing. In *Bayesian Statistics*, J. Bernardo, M. DeGroot, D. Lindley and A. Smith (Eds.). University Press, Valencia.
6. Caron, N. and Robert, C. (1991). Noninformative Bayesian testing and neutral Bayes factors. Tech. report #148, L.S.T.A., Université Paris 6.
7. Casella, G. and Berger, R. (1987). Reconciling Bayesian and frequentist evidence in the one-sided testing problem (with discussion). *J. Amer. Statist. Assoc.* **82**, 106–111.
8. DeGroot, M. (1982). Comment on Shafer (1982). *J. Amer. Statist. Assoc.* **77**, 336–338.
9. Hwang, J., Casella, G., Robert, C., Wells, M. and Farrell, R. (1991). Estimation of accuracy in testing. *Ann. Statist.* (to appear).
10. Jeffreys, H. (1961) *Theory of Probability* (3rd edition). Oxford University Press.
11. Lindley, D. (1957). A statistical paradox. *Biometrika* **44**, 187–192.
12. Shafer, G. (1982). Lindley's Paradox (with discussion). *J. Amer. Statist. Assoc.* **77**, 325–334.