

The book's title &  
copyright page  
are at the end.

## CHAPTER 44

# The Effect of Partial-Ordering Utilities on Bayesian Design of Sequential Experiments

**James S. Hodges**

*Division of Biostatistics, School of Public Health, University of Minnesota*

### 1. INTRODUCTION

The archetypal experimental design allocates a fixed number of units to the conditions under study. Sequential experiments, by contrast, can stop early if the results are extreme enough, so their design includes a stopping rule. Sequential experiments are well suited for comparing Bayesian and frequentist methods because the sample space, on which frequentists focus, is so different from the parameter space, on which Bayesians focus.

This chapter discusses an issue in the Bayes-frequentist debate about sequential experiments. Section 2 summarizes that debate, emphasizing a feature of frequentist designs: the stopping boundaries at interim analyses depend on the number of interim analyses. Many Bayesians see this property as inherently frequentist and thus deficient, but Sections 3 through 6 argue that it is not. Sections 3 and 4 catalogue the differences between Bayesian and frequentist sequential designs, a key difference being that Bayesians usually use loss functions imposing a complete ordering on the space of outcomes, while non-Bayesians usually use losses imposing only a partial ordering. Sections 5 and 6 show that the dependence of stopping boundaries on the number of interim analyses arises from this difference, obtaining frequentist designs from Bayesian machinery by weakening the complete-ordering axiom. Thus, the dependence of stopping boundaries on the number of interim analyses is not a consequence

*Bayesian Analysis in Statistics and Econometrics*, Edited by Donald A. Berry, Kathryn M. Chaloner, and John K. Geweke.

ISBN 0-471-11856-7 © 1996 John Wiley & Sons, Inc.

of a deep Bayes-frequentist split, but rather of an axiom about which reasonable people can differ. It is also a reminder that there is no free lunch: complete-ordering losses are more informative than partial-ordering losses, and the advantages of the former come at the price of having to supply the extra information.

There is a large literature about design, monitoring, and analysis of sequential trials. This paper cannot do justice to such innovations as, for example, Freedman and Spiegelhalter (1992) or Carlin et al. (1995). Interested readers can consult these papers or others cited below.

## 2. FREQUENTIST DESIGNS AND THE BAYESIAN CRITIQUE

### 2.1. Frequentist Designs: Interim Analyses Have a Price

Suppose a sequential experiment has at most  $n$  observations. If the sample size is fixed at  $n$  and the probability of falsely rejecting the null hypothesis—Type I error—is also fixed, the usual derivation yields a critical value for rejecting the null. If the experiment can be stopped early, the nominal Type I errors at the interim analyses must be smaller than the overall Type I error, and thus smaller than the Type I error of the single test in the fixed-sample-size design. That is, interim analyses have a price: the more looks, the larger the critical value at each look. The frequentist design problem might then be abstracted as

- Fix the maximum sample size  $n$  and the desired overall Type I error  $\alpha$ .
- Compute stopping boundaries with Type I error  $\alpha$  and (say) one, two, or four analyses.
- Pick the best design according to other criteria, such as expected sample size or power.

The first sequential designs for clinical trials (Pocock, 1977; O'Brien and Fleming, 1979) fixed the number and timing of interim analyses. Lan and DeMets (1983) did not require the interim analyses to be fixed in advance, using instead a function that describes the rate at which the overall Type I error is "spent."

### 2.2. The Bayesian Critique: Do Interim Analyses as Often as Possible

By the likelihood principle, at any stage of an experiment it is permissible to stop, compute the posterior distribution for the parameter of interest, and use it to evaluate designs for the rest of the experiment. Accordingly, Bayesian designs are constructed by

- Specifying a model for observations

- Specifying a loss function incorporating all the elements of value
- Deriving the decision criterion at each stage of the trial by backward induction

Interim analyses need not be foregone and foregoing them in clinical trials may be unethical because it can increase the expected number of patients in the trial. In particular, the frequentist focus on Type I error and the resulting costliness of interim analyses is arbitrary and inappropriate.

Theoretical arguments against the frequentist approach and for the Bayesian approach are given by Berry (1989) and papers cited there. A fully developed example of a Bayesian sequential design is given by Berry et al. (1992).

## 3. HOW DO BAYESIAN AND FREQUENTIST DESIGNS DIFFER?

The process of designing an experiment takes inputs—assumptions and procedures—and produces an output, a design. Section 2 discussed differences between the output of frequentist and Bayesian design processes, the difference of interest being that interim analyses have a price in frequentist design. The object of this chapter is to trace the differences in Bayesian and frequentist designs to differences in inputs, using the terminology of decision theory for the various inputs.

Two inputs to design are not matters of dispute by Bayesians and frequentists: the model for observations given the parameters of interest, and the action space (i.e., stop the trial or continue, and if stopped, chose the null or the alternative).

Another issue, the flexibility of a sequential design, is a matter of dispute. Many Bayesians are not troubled by changing the design during an experiment, relying on the likelihood principle. In frequentist theory, Lan-DeMets designs permit flexibility in the timing of analysis, but otherwise computation of  $p$ -values and confidence intervals rely fundamentally on the design, so it is sacrosanct. (In frequentist reality, designs are changed routinely and certain rigors of the theory are ignored.) This difference between the theories is real and cannot be addressed here; however, this chapter will show that it is independent of the issue of whether interim analyses have a price.

Sections 4 through 6 treat three other differences between the inputs to design used by frequentists and Bayesians. The first of these is the loss function. An axiom of Bayesian decision theory asserts that the decision maker's utility function imposes a complete ordering on the possible outcomes (DeGroot, 1970, p. 87). That is, there is a utility function  $U$  such that for any two outcomes  $X$  and  $Y$ , either  $U(X) > U(Y)$ ,  $U(X) = U(Y)$ , or  $U(X) < U(Y)$ . Loss functions typically used in Bayesian designs embody this axiom by combining all aspects of utility in a single function. (Berry et al., 1992, give a nice example.)

Weaker axioms are possible, for example, that the utility function imposes

only a partial ordering on outcomes. An example is a lexicographic utility (e.g., Abrams, 1980), in which the decision maker's utility has distinct components that she is unwilling to combine into a single function, but which she can rank in order of importance. If two choices are indistinguishable on the most important component of utility—the first echelon of the utility—then the choices are compared on the second most important component—the second echelon—and so on through the echelons of the utility function. The key feature of lexicographic utilities is that some pairs of outcomes can be compared and others cannot. Suppose a lexicographic utility function has two components and that the utility for an outcome is expressed as  $(u_1, u_2)$  if the first component of utility is  $u_1$  and the second component is  $u_2$ . Then two outcomes having utilities  $(x_1, x_2)$  and  $(y_1, y_2)$  are commensurable if and only if one of the following four conditions holds:

- $x_1 = y_1$ .
- $x_2 = y_2$ .
- Both  $x_1 > y_1$  and  $x_2 > y_2$ .
- Both  $x_1 < y_1$  and  $x_2 < y_2$ .

If none of these four conditions holds, then  $(x_1, x_2)$  and  $(y_1, y_2)$  are incommensurable; that is, no preference between them can be stated. The generalization to lexicographic utility functions with more components of utility is straightforward.

In frequentist designs, as abstracted in Section 2.1, the components of loss are Type I error, Type II error, and expected sample size. Type I error is usually the first echelon of the loss function, although the occupant of the second echelon is less standardized.

The second remaining difference between inputs to design used by Bayesians and frequentists is that frequentists build in restrictions on the stopping boundaries and Bayesians do not. These restrictions leave a single "degree of freedom" in the stopping boundary, which is specified by the overall Type I error. Bayesian design imposes no a priori restrictions on the stopping boundary, which instead is purely a consequence of the model, the prior, and the loss function. The final remaining difference is that Bayesians use explicit priors and frequentists do not. This difference is real, primarily because it allows Bayesians to introduce information from sources outside the current experiment.

The rest of this chapter examines these last three differences, focusing mainly on the effect of partial-ordering loss functions. Section 4 discusses the use of partial-ordering losses. Section 5 sets up a sequential experiment permitting consideration of the issues. Section 6 examines Bayesian sequential designs constructed using lexicographic loss functions built from Type I error, Type II error, and expected sample size. For these designs, the number of looks matters regardless of which of the three components of loss is assigned to the first echelon. Thus, it is not the traditional focus on Type I error that implies that the number of looks matters, but rather the partial order imposed by the loss.

It would seem, then, that this Bayes-frequentist difference arises from the axiom that utilities impose a complete ordering on the outcome space. Section 6 also shows that the frequentist restriction on decision rules has no foundational content and that the prior distribution is of no consequence for the result.

#### 4. USING PARTIAL-ORDERING LOSS FUNCTIONS

Using complete-ordering loss functions in sequential design is conceptually simple: select the design that minimizes expected loss by backward induction. Using partial-ordering loss functions is less straightforward because the components of the loss must be handled separately, making backward induction impossible. This section discusses ways of using partial-ordering loss functions and then considers whether such losses are, in fact, merely excuses.

It is helpful to distinguish two groups of partial-ordering loss functions:

- "No-trade" loss functions, for which optimizing on any component of loss yields an absurd test. The abstracted frequentist test in Section 2.1 has this property: optimizing on Type I or II error or on expected sample size yields an absurd test.
- "Trade" loss functions, for which optimizing on at least one component of loss yields a plausible test. An example has as its first echelon a loss function carrying a penalty of 1 for a Type I error and  $\gamma$  for a Type II error, and as its second echelon expected sample size.

"No-trade" loss functions can be used in two ways. One way is to select a cutoff value for each echelon such that any procedure must attain at least that value. Many papers describing frequentist tests read as if this were their procedure: for example, a test must have size 0.05 and power 0.8, implying a particular sample size. Another way to use no-trade losses is to reduce the second echelon of loss until the first echelon reaches an unacceptable value, and use the resulting test.

"Trade" losses allow a different use. If the first echelon allows an explicit trade-off, as in the example above, one can minimize the first-echelon loss function and use the second echelon to select among procedures that minimize the first echelon.

Why not insist on a complete-ordering loss function? Bayesians often cite as a virtue of their approach that it forces a user to be explicit, in this case, about trade-offs among the different echelons of loss. But if someone says that he considers Type I error and expected sample size incommensurable, need we question his integrity? A complete ordering does not come for free: one must add information to the loss in the form of a numerical trade of each kind of loss for the others, throughout the range of their possible values—a difficult and sometimes specious task. It would not be surprising that if one declined to add such information to the loss, one would lose something. The next two sections show that this is the case.

## 5. A CANONICAL SEQUENTIAL EXPERIMENT

Decision theory problems have four parts: probability model, loss function, action space, and decision rules.

### 5.1. The Probability Model

Consider a trial with  $n$  observations; in Section 6,  $n$  will be 2 or 4. Fixed sample-size designs require all  $n$  observations to be taken, while sequential designs can stop earlier. The parameter of interest is  $\delta$ , the difference between a new treatment and the standard, the null hypothesis being  $\delta \leq 0$  (new is no better than standard) and the alternative being  $\delta > 0$  (new is better than standard). The statistical model is  $X_j | \delta \sim N(\delta, 1)$ ,  $j = 1, 2, \dots, n$ , with a  $N(0, \tau^2)$  prior distribution for  $\delta$ . Letting  $\tau^2$  approach infinity permits frequentist procedures to be obtained as limiting cases. This formulation or an equivalent one was used in Pocock (1977), O'Brien and Fleming (1979), and Lan and DeMets (1983), among others.

### 5.2. The Loss Function

The loss function is lexicographic with three components: Type I error, Type II error, and expected sample size. The Bayesian definitions for these are

$$\text{Type I error} = \int_{-\infty}^0 \text{Pr}(\text{reject null} | \delta) f(\delta) d\delta \quad (44.1)$$

$$\begin{aligned} \text{Type II error} &= \int_0^{\infty} \text{Pr}(\text{accept null} | \delta) f(\delta) d\delta \\ &= \int_0^{\infty} \{1 - \text{Pr}(\text{reject null} | \delta)\} f(\delta) d\delta \end{aligned} \quad (44.2)$$

$$\text{Expected sample size} = \int_{-\infty}^{\infty} \left\{ \sum_{j=1}^n j \text{Pr}(\text{trial stops at } j | \delta) \right\} f(\delta) d\delta \quad (44.3)$$

where  $f(\delta)$  is the prior density of  $\delta$ . (The frequentist definitions of these components of loss are conditional on  $\delta$ , as they must be in the absence of a prior distribution.)

### 5.3. The Action Space and Decision Rules

The action space is: at each observation, stop or take the next observation; if you stop, choose the standard or the new treatment. By symmetry, the decision rules can be expressed as  $(t_1, t_2, \dots, t_n)$ , where  $t_j \geq 0$  is defined as follows:

- For  $j = 1, \dots, n - 1$ 
  - Stop after  $X_j$  and accept the null if  $d_j + t_j s_j < 0$ .
  - Stop after  $X_j$  and reject the null if  $d_j - t_j s_j > 0$ .
  - Take the next observation if neither of the above conditions is satisfied.
- After  $X_n$ , reject the null if  $d_n - t_n s_n > 0$  and accept it otherwise.

Here  $d_j$  and  $s_j$  are the posterior mean and standard deviation of  $\delta$ . Fixed-sample-size designs are included by allowing infinite  $t_j$  for  $j < n$ .

## 6. BAYESIAN DESIGN OF SEQUENTIAL EXPERIMENTS USING LEXICOGRAPHIC LOSS FUNCTIONS

Section 6.1 considers the lexicographic loss in which Type I error occupies the first (most important) echelon of the loss function, Section 6.2 puts Type II error in the first echelon, and Section 6.3 puts expected sample size in the first echelon. All give the same qualitative result: interim analyses have a price.

### 6.1. Type I Error as the Most Important Component of Loss

Let  $n = 2$ . The result is familiar, but is included because it sets a pattern for Sections 6.2 and 6.3. First consider the fixed sample size test. For the assumptions given in Section 5, if two observations are taken, the size of the test as a function of  $t_2$  is

$$0.5 - \int_{-\infty}^0 \Phi \left( \frac{t_2}{\sqrt{2} s_2} - \sqrt{2} \delta \right) \pi(\delta) d\delta$$

where  $\Phi$  is the standard normal distribution function. Fixing the size at  $\alpha$  specifies an equation in  $t_2$ ; call its solution  $t_2^f(\alpha)$ . Define  $t_1^f(\alpha)$  analogously to be the value of  $t_1$  such that if a single observation were taken,  $t_1^f(\alpha)$  would be the critical value yielding a test of size  $\alpha$ .

Now consider the sequential test. Fixing the size of the test at  $\alpha$  defines an equation in  $t_1$  and  $t_2$ , so infinitely many sequential designs can be specified with the right size. All of these designs have the familiar property of frequentist sequential designs, by the following argument. If  $t_1 = t_1^f(\alpha)$  and  $t_2 < \infty$ , then the size of the resulting sequential test is clearly greater than  $\alpha$ , so  $t_1$  must be increased above  $t_1^f(\alpha)$  if the test is to remain a sequential test. (That is, the only alternative is to set  $t_2 = \infty$ , in which case it is not a sequential test any more.) If  $t_2 = t_2^f(\alpha)$  and  $t_1 < \infty$ , then the size of the resulting sequential test is

$$\alpha + P(r1 \cap a2 | \delta < 0) - P(a1 \cap r2 | \delta < 0)$$

where

- $r1 = \{(x_1, x_2) | d_1 - s_1 t_1 > 0\}$ , observations such that the first is large enough to stop the trial and reject the null.
- $a1 = \{(x_1, x_2) | d_1 + s_1 t_1 < 0\}$ , observations such that the first is small enough to stop the trial and accept the null.
- $r2 = \{(x_1, x_2) | d_2 - s_2 t_2 > 0\}$ , observations such that the posterior mean after two observations is outside the boundary for rejecting the null.
- $a2 = \{(x_1, x_2) | d_2 - s_2 t_2 \leq 0\}$ , observations such that the posterior mean after two observations is not outside the boundary for rejecting the null, so the null is accepted.

If  $P(r1 \cap a2 | \delta < 0) - P(a1 \cap r2 | \delta < 0)$  is positive, then the size of the sequential test is greater than  $\alpha$  and  $t_2$  must be increased above  $t_2^f(\alpha)$  if the test is to remain a sequential test. By writing  $P(r1 \cap a2 | \delta < 0) - P(a1 \cap r2 | \delta < 0)$  as two triple integrals (in  $\delta$ ,  $x_1$ , and  $x_2$ ), integrating out  $\delta$ , transforming to  $y = (x_1 + x_2)/\sqrt{2}$  and  $z = (x_1 - x_2)/\sqrt{2}$ , and integrating out  $z$ , one is left with

$$K \left\{ \int_{-a_2/\sqrt{2}}^{a_2/\sqrt{2}} \Phi(-\sqrt{2} S_2 y) g(a_1, y) dy + \int_{-\infty}^{-a_2/\sqrt{2}} [2\Phi(-\sqrt{2} S_2 y) - 1] g(a_1, y) dy \right\} \quad (44.4)$$

where  $K$  is a constant,  $g(a_1, y) = \Phi(-\sqrt{2}a_1 + y)$ ,  $\phi(cy)$  is positive,  $\phi$  is the standard normal density,  $a_j = t_j/s_j$ , and  $c = s_2/t_2$ . Both terms are positive, so the design has the familiar feature of frequentist designs: the more looks, the larger the critical values.

The choice between the fixed-sample test and the infinity of sequential tests is then based on the second echelon of loss. If  $n$  is larger than the number of echelons of loss, infinitely many tests will be equally satisfactory. This is why frequentists impose conditions on sequential stopping boundaries: it reduces the problem and permits a relatively easy selection of a test, but it has no foundational content.

## 6.2. Type II Error as the Most Important Component of Loss

Let  $n = 2$  and consider the fixed-sample-size test with Type II error  $\beta$ . By Equation (44.2),

$$0.5 - \beta = \int_0^{\infty} \Pr(r2 | \delta) f(\delta) d\delta$$

where  $r2$  was defined earlier. If Type II error is fixed at  $\beta$ , this specifies an equation in  $t_2$ ; denote the solution as  $t_2^f(\beta)$ . Define  $t_1^f(\beta)$  analogously to be the value of  $t_1$  such that if a single observation were taken,  $t_1^f(\beta)$  would be the critical value yielding Type II error  $\beta$ .

Now consider the sequential test. As in Section 6.1, fixing Type II error at  $\beta$  defines an equation in  $t_1$  and  $t_2$ . If  $t_1 = t_1^f(\beta)$  and  $t_2 < \infty$ , the Type II error of the resulting sequential test is clearly smaller than  $\beta$ , so  $t_1$  must be increased above  $t_1^f(\beta)$  if the test is to remain a sequential test. On the other hand, if  $t_2 = t_2^f(\beta)$  and  $t_1 < \infty$ , then Type II error satisfies

$$0.5 - \beta = \int_0^{\infty} \Pr(r2 | \delta) f(\delta) d\delta + \int_0^{\infty} \Pr(r1 \cap a2 | \delta) f(\delta) d\delta - \int_0^{\infty} \Pr(a1 \cap r2 | \delta) f(\delta) d\delta$$

If the sum of the last two terms is positive, the Type II error of the sequential test is smaller than  $\beta$  and  $t_2$  must be increased above  $t_2^f(\beta)$  if the test is to remain a sequential test. To show this, apply to the last two terms the same sequence of operations that was used to obtain Equation (44.4), yielding

$$M(t_1) \left\{ \int_{-\infty}^{-a_2/\sqrt{2}} [2\Phi(\sqrt{2} S_2 y) - 1] \phi(cy) h(a_1, y) dy + \int_{-a_2/\sqrt{2}}^{a_2/\sqrt{2}} \Phi(\sqrt{2} S_2 y) \phi(cy) h(a_1, y) dy \right\}$$

where  $M$  is positive,  $h(a_1, y) = \Phi(-\sqrt{2}a_1 + y)/\Phi(-\sqrt{2}a_1 - a_2/\sqrt{2})$ , and the other symbols were defined near Equation (44.4). Of the two terms inside the braces the first is negative and the second is positive. For  $t_2 = t_2^f(\beta)$  and  $t_1 \geq t_1^f(\beta)$ ,

- For  $y < -a_2/\sqrt{2}$ ,  $h(a_1, y)$  is in the interval  $(0, 1)$  and monotonically decreasing as a function of  $t_1$ .
- For  $y > -a_2/\sqrt{2}$ ,  $h(a_1, y)$  is greater than 1 and monotonically increasing as a function of  $t_1$ .

But the expression in braces is positive for  $t_1 = t_1^f(\beta)$ ; therefore, it is positive for finite  $t_1 \geq t_1^f(\beta)$ , and the result follows.

Thus, when the first echelon—the most important—of the loss function is Type II error, more interim analyses mean larger critical values.

### 6.3. Expected Sample Size as the Most Important Component of Loss

To show the result for expected sample size (ESS), the maximum sample size  $n$  must be greater than 2. If  $n = 2$  and the probability of stopping at the first observation is positive, all sequential tests beat the fixed-size test on ESS. So let  $n = 4$ , and consider two sequential tests: one allowing stopping after the second and fourth observations, and the other allowing stopping after all four observations. (The fixed-sample-size test is dominated by the sequential tests.) Define  $t_j^m$  to be the value of  $t_j$  in the sequential test with  $m$  analyses; that is,  $t_4^2$  is the value of  $t_4$  for the test that permits stopping after the second and fourth observations.

Fix the ESS at  $S$ . This and Equation (4.3) define a single equation in  $t_2^2$  and  $t_4^2$  and a single equation in  $(t_1^4, t_2^4, t_3^4, t_4^4)$ . Suppose  $t_4^4 = t_4^2$ . If  $t_2^4 < t_2^2$ , the test with four analyses will have  $ESS < S$ , so  $t_2^4 \geq t_2^2$ . If  $t_2^4 = t_2^2$  and either of  $t_1^4$  or  $t_3^4$  is finite, the test with four analyses will have  $ESS < S$ . In other words, compared to the test with two analyses, the test with four analyses must give up something at either observation 2 or observation 4: more look implies larger critical values.

### 6.4. The Effect of the Prior Distribution

In Sections 6.1 through 6.3, the prior for  $\delta$  was  $N(0, \tau^2)$ . What happens with more general priors? When the first echelon of loss is expected sample size, the argument in Section 6.3 works without change. When the first echelon of loss is Type I error, it is easy to show that  $P(r1 \cap a2|\delta) - P(a1 \cap r2|\delta)$  is nonnegative for all  $\delta$ , so the result holds for general priors on  $\delta$ . I have not been able to extend the result in Section 6.2 to general priors for  $\delta$ , but I conjecture that it is true.

## 7. DISCUSSION

This chapter's purpose was to explore the difference between Bayesian and frequentist solutions to a problem arising in sequential experiments. It showed that Bayesian machinery generates frequentist designs for sequential experiments if the loss function is a lexicographic loss involving Types I and II error and expected sample size—in other words, that part of the clash between frequentists and Bayesians can be traced to a single axiom of Bayesian decision theory. This axiom is about utility, not probability, so one may accept Savage's axioms of subjective probability and still obtain "frequentist" designs in sequential experiments.

This result reaffirms that there is no free lunch: to gain the benefits of a complete-ordering loss, one must pay with information inserted into the loss function. Others have also argued that the apparent free lunch of Bayesian sequential design has some costs. Rosenbaum and Rubin (1984) and Rubin

(1984) show that the use of data-dependent stopping rules and informative priors can cause posterior intervals to be out of calibration. That result is, of course, only of interest to those who think that frequencies matter sometimes. The result here does not require acceptance of anything foreign to a personalistic view of probability.

## REFERENCES

- Abrams, R. (1980). *Foundations of Political Analysis: An Introduction to the Theory of Collective Choice*. Columbia Univ. Press, New York.
- Berry, D. A. (1989). Monitoring accumulating data in a clinical trial. *Biometrics* **45**, 1197-1211.
- Berry, D. A., Wolff, M. C., and Sack, D. (1992). Public health decision making: A sequential vaccine trial. In *Bayesian Statistics 4* (J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, eds.), pp. 79-96. Oxford Univ. Press, Oxford.
- Carlin, B. P., Chaloner, K., Church, T., Louis, T. A., and Matts, J. P. (1995). Bayesian approaches to monitoring clinical trials with an application to toxoplasmic encephalitis prophylaxis. In *Bayesian Case Studies 2* (C. Gatsonis, J. S. Hodges, R. E. Kass, and N. D. Singpurwalla, eds.). Springer-Verlag, New York (in preparation).
- DeGroot, M. H. (1970). *Optimal Statistical Decisions*. McGraw-Hill, New York.
- Freedman, L. S., and Spiegelhalter, D. J. (1992). Application of Bayesian statistics to decision making during a clinical trial. *Stat. Med.* **11**, 23-35.
- Lan, K. K. G., and DeMets, D. L. (1983). Discrete sequential boundaries for clinical trials. *Biometrika* **70**, 659-663.
- O'Brien, P. C., and Fleming, T. R. (1979). A multiple testing procedure for clinical trials. *Biometrics* **35**, 549-556.
- Pocock, S. J. (1977). Group sequential methods in the design and analysis of clinical trials. *Biometrika* **64**, 191-199.
- Rosenbaum, P. R., and Rubin, D. B. (1984). Sensitivity of Bayes inference with data-dependent stopping rules. *Am. Stat.* **38**, 106-109.
- Rubin, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *Ann. Stat.* **12**, 1151-1172.

# Bayesian Analysis in Statistics and Econometrics

Essays in Honor of Arnold Zellner

Edited by

DONALD A. BERRY  
Duke University

KATHRYN M. CHALONER  
University of Minnesota

JOHN K. GEWEKE  
University of Minnesota

*a profoundly  
obscure book*

ICS

UEL S. WILKS

*sher, J. Stuart Hunter,  
an F. M. Smith,*

end of this volume



A Wiley-Interscience Publication  
JOHN WILEY & SONS, INC.

New York • Chichester • Brisbane • Toronto • Singapore

# Contents

## Contributors

## Preface

## PART I FORECASTING AND ASSESSMENT

1. **Aggregating Forecasts: An Bayesian Methods**  
*Robert T. Clemen, Steven K*
2. **Recursive Identification, E Nonstationary Economic T GNP International Data**  
*A. García-Ferrer, J. del Ho*
3. **Algorithms for Conditiona**  
*Angelo Gilio*
4. **A Comparison of Techniq Points in Regional Emplo**  
*James P. LeSage*

## PART II INFERENCE, ESTI

5. **Bayesian Modeling for Ca Information**  
*N. Sedransk and Ruby C. H*

This text is printed on acid-free paper.

Copyright © 1996 by John Wiley & Sons, Inc.

All rights reserved. Published simultaneously in Canada.

Reproduction or translation of any part of this work beyond that permitted by Section 107 or 108 of the 1976 United States Copyright Act without the permission of the copyright owner is unlawful. Requests for permission or further information should be addressed to the Permissions Department, John Wiley & Sons, Inc., 605 Third Avenue, New York, NY 10158-0012.

### *Library of Congress Cataloging in Publication Data:*

Bayesian analysis in statistics and econometrics: essays in honor of Arnold Zellner / edited by Donald A. Berry, Kathryn M. Chaloner, and John K. Geweke.

p. cm. — (Wiley series in probability and statistics. Applied probability and statistics)

“A Wiley-Interscience publication.”

Includes bibliographical references and index.

ISBN 0-471-11856-7 (cloth)

1. Econometrics. 2. Bayesian statistical decision theory.  
3. Zellner, Arnold. I. Berry, Donald A. II. Chaloner, Kathryn.  
III. Geweke, John K. IV. Zellner, Arnold. V. Series.  
HB139.B3936 1995  
330'.01'5195—dc20 95-22393

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1