

AD-A125 306

SENSITIVITY OF BAYES INFERENCE WITH IGNOREABLE BUT  
DATA-DEPENDENT STOPPING. (U) WISCONSIN UNIV-MADISON  
MATHEMATICS RESEARCH CENTER P R ROSENBAUM ET AL.

1/1

UNCLASSIFIED

DEC 82 MRC-TSR-2455 DAG29-80-C-0041


F/G 12/1

NL



END  
FILMED  
DTIC





MICROCOPY RESOLUTION TEST CHART  
NATIONAL BUREAU OF STANDARDS-1963-A

0

AD A 1 25306

MRC Technical Summary Report #2455

**SENSITIVITY OF BAYES INFERENCE  
WITH IGNORABLE BUT DATA-DEPENDENT  
STOPPING RULES**

Paul R. Rosenbaum and Donald B. Rubin

**Mathematics Research Center  
University of Wisconsin-Madison  
610 Walnut Street  
Madison, Wisconsin 53706**

December 1982

(Received March 16, 1982)

DTIC FILE COPY

DTIC  
ELECTRONIC  
MAR 4 1982  
S  
A

**Approved for public release  
Distribution unlimited**

Sponsored by

U. S. Army Research Office  
P. O. Box 12211  
Research Triangle Park  
North Carolina 27709

National Cancer Institute  
9000 Rockville Pike  
Bethesda, MD 20205

Educational Testing  
Service  
Princeton, NJ 08540

33 02 028 121

UNIVERSITY OF WISCONSIN-MADISON  
MATHEMATICS RESEARCH CENTER

SENSITIVITY OF BAYES INFERENCE WITH  
IGNORABLE BUT DATA-DEPENDENT STOPPING RULES

PAUL R. ROSENBAUM\* AND DONALD B. RUBIN\*\*

Technical Summary Report #2455

December 1982

ABSTRACT

It is sometimes argued that Bayesian inference is unaffected by data dependent stopping rules. Although this is formally true for ignorable rules, there is likely to be heightened sensitivity of inference to prior assumptions when using data dependent rules rather than stopping rules that do not depend on the data. This point is illustrated in a simple example.

AMS (MOS) Subject Classifications: 62P99, 62J05, 62F12, 62F11


Key Words: Sensitivity, Bayesian inference, Stopping rules, Ignorable rules, Missing data.

Work Unit Number 4 (Statistics and Probability)

\*Department of Statistics and Human Oncology, University of Wisconsin-Madison.

\*\*Educational Testing Service, Princeton, NJ 08540

Sponsored in part by the United States Army under Contract No. DAAG29-80-C-0041, by grant P30-CA-14520 from the National Cancer Institute to the Wisconsin Clinical Cancer Center, and by the Educational Testing Service.



Accession For	
UNIVERSITY OF WISCONSIN-MADISON	
Department of Mathematics	
Justification	
By	
Distribution/	
Availability Code	
Dist	Avail and/or Special

AA

## SENSITIVITY OF BAYES INFERENCE WITH IGNORABLE BUT DATA-DEPENDENT STOPPING RULES

Paul R. Rosenbaum\* and Donald B. Rubin\*\*

### 1. Complex Stopping Rules in Applied Inference: Discussion In The Context of An Example.

Complex stopping rules--that is, complex rules for determining how much data to collect -- are often difficult to avoid in practice. One example occurs in the initial phase of experimentation to determine a high yet tolerable dose of a new cancer chemotherapy. The usual procedure is to offer a low dose of the new drug to several patients; if this dose is well tolerated, the next few patients receive a somewhat higher dose. If toxicity is observed, the dose is not increased, but rather several additional patients are given the current dose; based on these additional results, it is decided whether to terminate the experiment or continue increasing the dose. Since the safety of experimental subjects is a primary concern, and since it is not always possible to anticipate the nature of all types of toxicity that might arise, rigidly defined stopping rules are often impractical. Further increasing the difficulties of statistical modeling is the fact that resultant sample sizes are often quite small (i.e., less than 20). Often, inferences are based on statistical procedures that ignore the stopping rule (e.g., Brown and Hu, 1981).

Although we will discuss the effects of stopping rules on Bayesian inferences within the context of phase I trials, our model for these trials is simplified, and in some ways artificial. Consequently, our results are not directly applicable to the practice of these trials, but are mainly suggestive of directions for further work.

Write  $D_i$  for the dose given to patient  $i$  and write  $T_i$  for the vector measure of toxicity observed for patient  $i$ , where  $N$  patients (indexed by  $i = 1, \dots, N$ ) are observed before the study is terminated. Since dose  $D_i$  depends on previous doses,  $D_{i-1} = (D_1, D_2, \dots, D_{i-1})$  and previous toxicities  $T_{i-1} = (T_1, T_2, \dots, T_{i-1})$  and possibly an unknown

\*Department of Statistics and Human Oncology, University of Wisconsin-Madison.

\*\*Educational Testing Service, Princeton, NJ 08540.

Sponsored in part by the United States Army under Contract No. DAAG29-80-C-0041, by grant P30-CA-14520 from the National Cancer Institute to the Wisconsin Clinical Cancer Center, and by the Educational Testing Service.

parameter  $\psi$  of the dose escalation rule, we write for the probability of  $D_i$  given

$(D_{e_{i-1}}, T_{e_{i-1}}, \psi)$

$$(1) \quad \text{pr}(D_i | D_{e_{i-1}}, T_{e_{i-1}}, \psi) .$$

Similarly, since the index,  $N$ , of the last patient entered into the trial depends on the doses and toxicities through the  $N^{\text{th}}$  trial, we write for the data dependent stopping rule

$$(2) \quad \text{pr}(N | D_{e_N}, T_{e_N}, \gamma)$$

where  $\gamma$  is a possibly unknown parameter of the stopping rule. Because it is usually reasonable to assume that dose  $D_i$  determines toxicity  $T_i$ , we complete the specification of the joint distribution of  $N, D_{e_N}, T_{e_N}$  by writing

$$(3) \quad \text{pr}(T_i | D_i, D_{e_{i-1}}, T_{e_{i-1}}, \theta) = p(T_i | D_i, \theta)$$

where  $\theta$  is an unknown parameter that relates toxicity to dose. The parameter  $\theta$  is the unknown parameter of primary interest. Thus, the distribution of the observed data

$(N, D_{e_N}, T_{e_N})$  given the unknown parameters  $(\psi, \gamma, \theta)$  is given by

$$(4) \quad \text{pr}(N | D_{e_N}, T_{e_N}, \gamma) \left\{ \prod_{i=1}^N \text{pr}(T_i | D_i, \theta) \right\} \left\{ \prod_{i=1}^N \text{pr}(D_i | D_{e_{i-1}}, T_{e_{i-1}}, \psi) \right\} .$$

The frequency properties of interval estimates of  $\theta$  generally depend on the factor  $\left\{ \text{pr}(N | D_{e_N}, T_{e_N}, \gamma) \prod_{i=1}^N \text{pr}(D_i | D_{e_{i-1}}, T_{e_{i-1}}, \psi) \right\}$  of the distribution (4) which may be complex, involving high dimensional unknown parameters  $\gamma$  and  $\psi$ . Therefore, small sample frequentist inferences may be difficult if not impossible to obtain, unless carefully constructed and followed stopping rules are used (e.g., Armitage 1975; Pocock 1977; DeMets and Ware 1980; Fleming 1982). There exists the hope, often stated (e.g. Lindley, 1972, p.24), that Bayesian interval estimates produced by ignoring the stopping rule will not lead the investigator astray.

If  $\theta$  and  $(\psi, \gamma)$  are a priori independent, then the dose escalation and stopping rules are ignorable (Rubin, 1976, 1978a, b) in the sense that the marginal posterior distribution of  $\theta$  can be obtained by simply ignoring the factors  $\text{pr}(N | D_{e_N}, T_{e_N}, \gamma)$  and  $\prod_{i=1}^N \text{pr}(D_i | D_{e_{i-1}}, T_{e_{i-1}}, \psi)$ . (To see this, note that the posterior distribution factors into a term involving  $\theta$  and one involving  $(\psi, \gamma)$ ; see also Rubin (1976).) If  $\theta$  and  $(\psi, \gamma)$  are a priori dependent, then in general Bayesian inferences about  $\theta$  explicitly depend on the stopping rule or the dose escalation rule.

The dependence of Bayesian inference on nonignorable experimental designs is well known (Rubin 1978a; Rosenbaum and Rubin 1983). In this paper we study the possible sensitivity of Bayesian inference even to ignorable stopping rules, and for this purpose, we assume that  $\theta$  and  $(\psi, \gamma)$  are independent. We shall see that Bayesian inference can depend on the stopping rule in subtle ways even when it is ignorable. In particular we find that Bayesian inferences from experiments involving a priori fixed stopping rules will frequently be less sensitive to model/prior specifications (i.e. more robust in the sense of Box and Tiao(1962,1964,1973) and Dempster (1977)) than inferences from experiments with ignorable but data-dependent stopping rules.

Thus we see that, as in other contexts (e.g. randomization in experiments (Rubin 1978a; Rosenbaum and Rubin 1983), and nonresponse in sample surveys (Rubin 1978b; Little 1982)), an inability to draw frequency inferences, or at least an inability to do so straightforwardly, is indicative of heightened sensitivity of Bayesian inference to prior specifications.

## 2. A Simple Illustrative Case

In order to focus on the effect of the stopping rule on Bayesian inference we consider a particularly simple case. We begin by assuming that the dose of the drug,  $D_1$ , remains constant at  $D$ , so that for the escalation rule we have

$$(5) \quad \text{pr}(D_1 | D_{e_1-1}, T_{e_1-1}, \psi) = \begin{cases} 1 & \text{if } D_1 = D \\ 0 & \text{otherwise} \end{cases} .$$

Because the dose is constant for all patients, the average toxicity at the fixed dose  $D$  is a natural parameter to estimate. Suppose

$$(6) \quad \text{pr}(T_1 | D_1 = D, \theta) = \frac{1}{\sqrt{2\pi}} \exp\{-\frac{(T_1 - \theta)^2}{2}\} = \phi(T_1 - \theta)$$

so the toxicities are normal with mean  $\theta$  and variance 1, where  $\phi(\cdot)$  is the standard normal probability density.

We complete the model specification by assuming a simple stopping rule: 100 trials will be conducted unless the data suggest that the average toxicity,  $\theta$ , is too high, in which case the study will be terminated. In particular, data will be collected to the

100th trial unless the current t-statistic testing  $\theta = 0$  is larger than  $C$  where  $C$  is a constant, that is, unless  $\bar{T}_N > C/\sqrt{N}$  where  $\bar{T}_N = \sum_{i=1}^N T_i / N$ .

Formally, we have

$$(7) \quad \text{pr}(N|D_{\theta_N}, T_{\theta_N}, \gamma) = \begin{cases} 1 & \text{if } N = 100 \text{ and } \bar{T}_1 < C/\sqrt{1}, \text{ for all } i < 100 \\ 1 & \text{if } N < 100, \bar{T}_N > C/\sqrt{N}, \text{ and } \bar{T}_i < C/\sqrt{i} \text{ for all } i < N \\ 0 & \text{otherwise.} \end{cases}$$

The escalation and stopping rules defined by (5) and (7) are ignorable, and moreover, are free of unknown parameters. Consequently, Bayesian inferences for  $\theta$  with a fixed prior,  $p(\theta)$ , follow from the prior distribution for  $\theta$  and the normal specification (6); that is the posterior distribution of  $\theta$  is proportional to

$$(8) \quad p(\theta) \prod_{i=1}^N \phi(T_i - \theta) = p(\theta) \phi\left\{\frac{(\bar{T}_N - \theta)\sqrt{N}}{\sigma}\right\}.$$

Suppose

$$(9) \quad p(\theta) = \sqrt{\rho} \phi\left\{\frac{(\theta - \mu)\sqrt{\rho}}{\sigma}\right\}$$

so that a priori  $\theta$  is normal with mean  $\mu$  and variance  $1/\rho$ . Then the posterior distribution of  $\theta$  is normal with mean

$$(10) \quad (\rho\mu + N\bar{T}_N)/(\rho + N)$$

and variance

$$(11) \quad (\rho + N)^{-1}.$$

### 3. The Standard Interval for $\theta$ and its Probability Coverage

The standard 95% interval for  $\theta$  under the normal specification (6) is

$$(12) \quad I(\bar{T}_N, N) = \bar{T}_N \pm \frac{2}{\sqrt{N}}.$$

For a priori fixed  $N$ , i.e. for  $C = \infty$  in (7), the interval  $I(\bar{T}_N, N)$  is a 95% confidence interval, covering  $\theta$  in slightly more than 95% of experiments. For ignorable stopping rules, that is, for  $N$  that depends only on the observed data,  $T_{\theta_N}$ , as in this example, the standard interval is the limit of a sequence of 95% highest posterior density intervals as the prior variances  $1/\rho$  tends to  $\infty$ , with any fixed prior mean,  $\mu$ . One justification that has been offered for this interval based on a "flat prior" is that the

interval is in some sense conservative, adding less prior information to the inference than priors which are more peaked. Note that the interval corresponds to the one obtained from Jeffery's (1961) noninformative prior if  $N$  is fixed a priori, but not generally if other stopping rules are used.

We now investigate the sensitivity of interval estimation to variations in prior assumptions about  $\theta$ . In particular, suppose the prior distribution of  $\theta$  is given by (9). Then the posterior distribution of  $\theta$  is specified by (10) and (11).

Hence, the posterior coverage of the standard interval,  $I(\bar{T}_N, N)$ , is

$$\begin{aligned}
 & \text{pr}(\theta \in [\bar{T}_N \pm \frac{2}{\sqrt{N}}] | N, T_{\text{opt}}, D_{\text{opt}}, \mu, \rho) \\
 &= \text{pr}\left\{ \frac{\theta - \frac{N\bar{T}_N + \rho\mu}{N + \rho}}{1/\sqrt{N + \rho}} \in \left[ \frac{\bar{T}_N - \frac{N\bar{T}_N + \rho\mu}{N + \rho}}{1/\sqrt{N + \rho}} \right] \pm \frac{2/\sqrt{N}}{1/\sqrt{N + \rho}} \mid T_{\text{opt}}, D_{\text{opt}}, \mu, \rho \right\} \\
 &= \Phi\left\{ \sqrt{N + \rho} \left( \bar{T}_N - \frac{N\bar{T}_N + \rho\mu}{N + \rho} \right) \pm \frac{2\sqrt{N + \rho}}{\sqrt{N}} \right\} \\
 (13) \quad &= \Phi\left\{ \frac{\rho}{\sqrt{1 + \rho/N}} \frac{c}{\sqrt{N}} \pm 2\sqrt{1 + \rho/N} \right\} = P(N, \bar{T}_N, \mu, \rho),
 \end{aligned}$$

where  $\Phi(\cdot)$  is the standard normal probability measure and  $c = \bar{T}_N - \mu$ . In words,  $P(N, \bar{T}_N, \mu, \rho)$  is the posterior probability, given  $\mu, \rho$ , that  $\theta$  falls in the standard interval,  $I(\bar{T}_N, N)$ .

Inferences about  $\theta$  are relatively insensitive to prior specifications if  $P(N, \bar{T}_N, \mu, \rho)$  is approximately .95 or higher for values of  $(\mu, \rho)$  that are not contradicted by the data. In the next section, we examine the frequency properties of  $P(N, \bar{T}_N, \mu, \rho)$  for various stopping rules, and we find that fixed sample sizes yield less sensitivity than some stochastic stopping rules.

Some preliminary observations are possible based on inspection of expression (13).

1. For every fixed  $\epsilon = T_N - \mu$  and  $N$ ,

$$P(N, \bar{T}_N, \mu, \rho) \rightarrow \Phi\{\pm 2\} = .95 \text{ as } \rho \rightarrow 0;$$

i.e., informally, if the prior (9) is actually diffuse, then the standard interval will be approximately correct regardless of the stopping rule.

2. For every fixed  $\rho$  and  $N$ ,

$$P(N, \bar{T}_N, \mu, \rho) \rightarrow \Phi\{\pm 2\sqrt{1+\rho/N}\} > .95 \text{ as } \epsilon \rightarrow 0;$$

i.e., informally, if the sample mean,  $\bar{T}_N$ , happens to be close to the prior mean,  $\mu$ , then the standard interval will tend to be conservative regardless of the stopping rule.

3. If we stop whenever  $\epsilon/\sqrt{N}$  appears large, then  $P(N, \bar{T}_N, \mu, \rho)$  will be more likely to fall below .95, since the interval will tend to be centered away from zero. This fact is the basis for increased sensitivity under certain data dependent stopping rules.

4. If  $N$  is fixed a priori, so that the stopping rule is not data dependent, then  $\epsilon/\sqrt{N} | \mu, \rho, N \sim N(0, \frac{1}{N}(\frac{1}{N} + \frac{1}{\rho}))$ . Therefore,  $\epsilon/\sqrt{N}$  may be expected to be nearly zero if the prior variance,  $1/\rho$ , of  $\theta$  is small and  $N$  is not too small, and hence, under these conditions, the coverage  $P(N, \bar{T}_N, \mu, \rho)$  of the standard 95% interval will often be close to or above 95%. Observations 1 and 4 together provide a basis for expecting the standard 95% interval to have approximately 95% coverage with fixed, large sample sizes.

#### 4. The Marginal Distribution of Bayesian Coverage Probabilities: A Simulation

This section describes a simulation which shows that the sensitivity of Bayesian inference depends in part on the stopping rule used. In particular, we find that the marginal distribution given  $\mu, \rho$  of the posterior coverage probability  $P(N, \bar{T}_N, \mu, \rho)$  of the usual 95% interval may be less tightly concentrated around .95 if certain data dependent stopping rules are used.

Seven stopping rules are compared, all using the same distribution of  $\theta$ . First,  $\theta$  is sampled from  $N(0, \frac{1}{\rho})$ . Then  $T_1, T_2, \dots, T_{100}$  are independently sampled from  $N(\theta, 1)$ . Four data dependent stopping rules Elev- $\infty$ , Elev-2, Elev-1.5 and Elev-.5 involve stopping at the first  $N < 100$  at which  $\bar{T}_N$  is elevated. In particular, rule Elev-C is given by (7). Of course, Elev- $\infty$  fixes the sample size at  $N=100$ .

To investigate the effect of data dependent stopping rules, it is not sufficient to compare Elev- $\infty$  with Elev-C, for  $C = 2, 1.5, .5$ , since the marginal distribution of sample sizes is not the same for the four stopping rules. Therefore, we also sampled using three other stopping rules, namely Rand-C for  $C = 2, 1.5, .5$ , where Rand-C and Elev-C produce the same marginal distribution of sample sizes  $N$ , but under Rand-C the sample size  $N$  is conditionally independent of  $\bar{T}_N/\sqrt{N}$  given  $\theta$ . The Rand-C rule can be implemented as follows: If one observation from rule Elev-C yielded a sample size of  $N$ , then Rand-C also stopped at sample size  $N$ , but  $\bar{T}_N$  was calculated from  $N$  new independent observations from  $N(\theta, 1)$ .

Table 1 displays estimates of the 10% point of the sampling distribution of Bayesian coverage probabilities  $P(N, \bar{T}_N, 0, \rho)$  of the standard 95% interval,  $\bar{T}_N \pm 2/\sqrt{N}$ . Examination of this table leads to the following observations. First, as one would expect, if the prior distribution of  $\theta$  is diffuse ( $\rho = .01$ ), the coverage of the standard interval is nearly 95% for all seven stopping rules. Second, for both fixed sample sizes (Elev- $\infty$ ) and purely random sample sizes (Rand-C), the estimated 10% point of the distribution of coverage probabilities is at least .83 for all prior precisions in the table. However, this estimated 10% point falls as low as .54 for Elev-2. In other words, it appears that the posterior coverage of the usual interval,  $\bar{T}_N \pm 2/\sqrt{N}$ , may be lower than 55% in 10% of experiments if the Elev-2 stopping rule is used. Inferences based on the standard  $\bar{T}_N \pm 2/\sqrt{N}$  interval are less sensitive to prior/model specifications if data dependent stopping rules such as Elev-2 are avoided.

The comparison of the Elev-C and Rand-C rules in Table 1 has shown that it is not the size of the sample, but rather the reason for stopping that causes the increased sensitivity. Nonetheless, given that an Elev-C rule was used, it is natural to ask: Can we identify, using the terminal sample size  $N$ , the intervals that have poor coverage? Figure 1 addresses this question. Twenty-five independent samples were drawn using Elev-2 and  $\rho=100$ , and the coverage probabilities were plotted against the sample sizes. Fifteen of the twenty-five samples stopped at sample size 100; of these, 14 yielded coverage probabilities of about .99. All ten samples that stopped with less than 100 observations

TABLE 1. Estimated 10% Point of the Distribution of Coverage Probabilities

$P(N, \bar{X}_N, O, \rho)$  for Nine Stopping Rules and Various Values of  $\rho$ .


STOPPING RULE	Prior Precision $\rho$			
	100	10	1	.01
	Prior standard deviation ( $1/\sqrt{\rho}$ )			
	.1	.32	1	10
Elev-∞	.84	.93	.95	.95
Elev-2	.54*	.78	.91	.95
Rand-2	.87*	.92	.94	.95
Elev-1.5	.80**	.84	.91	.95
Rand-1.5	.86**	.91	.95	.95
Elev-.5	.93	.92	.89	.95
Rand-.5	.89	.87	.93	.95

\* Based on 400 pseudoreplications.

\*\* Based on 600 pseudoreplications.

All other values are based on 200 pseudoreplications.

Figure 1



had coverage of less than .9, with the lowest coverages for samples with less than ten observations.

##### 5. Interpretation of Results.

The marginal distribution of Bayesian coverage probabilities that was simulated in the last section has two, not really contradictory, interpretations: a subjectivist interpretation and a frequency interpretation.

In the subjectivist interpretation, the prior distribution on  $\theta$  reflects personal beliefs about  $\theta$  before data are observed. The distribution of coverage probabilities  $P(N, \bar{T}, \mu, \rho)$  is relevant during the design of an experiment if the standard 95% interval (12) will be used in the analysis, since the distribution describes personal predata beliefs about the final postdata coverage of that interval. If for practical reasons the standard interval is to be used in the analysis, some subjectivists may want to design the experiment so that this interval can be expected to be at worst conservative for a range of reasonable prior distributions. The Elev- $\infty$  and Rand-C stopping rules do this quite well, whereas Elev-2 runs a risk of substantial undercoverage.

In the frequency interpretation, the current experiment is viewed (perhaps accurately) as one in a long series of experiments, and the prior distribution of  $\theta$  is the distribution of  $\theta$  values arising in these experiments. In this context, there is a correct Bayesian inference based on the true prior distribution, which is unknown to the experimenter. The experimenter specifies some prior distribution(s)--often flat prior distributions--and hopes that inferences based on the specified prior approximate the correct inferences based on the true but unknown prior. Our simulation illustrates that this hope will be fulfilled with greater frequency if rules Elev- $\infty$  and Rand-C are used in place of Elev-2.

We have examined the heightened sensitivity of Bayesian inferences to misspecification of the prior/model when data-dependent stopping rules are used; however, our investigation has been confined to a highly specialized and somewhat artificial case. Further investigation is required to identify situations which produce more or less sensitivity

than we have observed (e.g., perhaps in §1, stopping at a leverage point when toxicity has been observed at a high dose).

#### REFERENCES

- Armitage, P. (1975) Sequential Medical Trials, 2nd edition. Oxford: Blackwell.
- Box, G.E.P. and Tiao, G.C. (1962). A further look at robustness via Bayes' theorem. Biometrika 49, 419-432.
- Box, G.E.P. and Tiao, G.C. (1964). A note on criterion robustness and inference robustness. Biometrika 51, 169-173.
- Box, G.E.P. and Tiao, G.C. (1973). Bayesian Inference in Statistical Analysis. Reading, Massachusetts: Addison-Wesley.
- Brown, B.W. and Hu, M. (1980). Setting dose levels for the treatment of testicular cancer. In Biostatistics Casebook, Ed. R.G. Miller et. al., pp 123-152. New York: John Wiley.
- DeMets, D.L. and Ware, J.H. (1980) Group sequential methods for clinical trials with a one-sided hypothesis. Biometrika 67, 651-60.
- Dempster, A.P. (1977). Examples relevant to the robustness of applied inferences. In Statistical Decision Theory and Related Topics II, Ed. S.S. Gupta and D.S. Moore, pp. 121-138. New York: Academic Press.
- Fleming, T.R. (1982) One sample multiple testing procedure for phase II clinical trials. Biometrics 38, 143-152.
- Jeffreys, H. (1961). Theory of Probability. Third Edition. London: Oxford University Press.
- Lindley, D. (1972) Bayesian Statistics: A Review. Philadelphia: SIAM.
- Little, R.J.A. (1982). Models for nonresponse in sample surveys. Journal of the American Statistical Association 77, 237-250.
- 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. (1983). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. To appear in the Journal of the Royal Statistical Society, Series B.
- Rubin, D.B. (1976). Inference and missing data. Biometrika 63, 581-592.

Rubin, D.B. (1978a). Bayesian inference for causal effects: The role of randomization. Annals of Statistics 6, 34-58.

Rubin, D.B. (1978b). The phenomenological Bayesian perspective in sample surveys from finite populations: Foundations. Imputation and Editing of Faulty or Missing Survey Data. U. S. Department of Commerce, pp 10-18.

PRR/DBR:jgb

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 2455	2. GOVT ACCESSION NO. AD-A125306	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Sensitivity of Bayes Inference with Ignorable But Data-Dependent Stopping Rules		5. TYPE OF REPORT & PERIOD COVERED Summary Report - no specific reporting period
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Paul R. Rosenbaum and Donald B. Rubin		8. CONTRACT OR GRANT NUMBER(s) DAAG29-80-C-0041 & P30-CA-14520
9. PERFORMING ORGANIZATION NAME AND ADDRESS Mathematics Research Center, University of 610 Walnut Street Madison, Wisconsin 53706		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Work Unit Number 4 - Statistics and Probability
11. CONTROLLING OFFICE NAME AND ADDRESS See 18 below		12. REPORT DATE December 1982
		13. NUMBER OF PAGES 13
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES U.S. Army Research Office      National Cancer Institute      Educational Testing P.O. Box 12211                      9000 Rockville Pike                      Service Research Triangle Park              Bethesda, MD 20205                      Princeton, NJ 08540 North Carolina 27709		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Sensitivity, Bayesian inference, Stopping rules, Ignorable rules, Missing data.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) It is sometimes argued that Bayesian inference is unaffected by data dependent stopping rules. Although this is formally true for ignorable rules, there is likely to be heightened sensitivity of inference to prior assumptions when using data dependent rules rather than stopping rules that do not depend on the data. This point is illustrated in a simple example.		