

Bayesian Climate Change Assessment

L. MARK BERLINER

The Ohio State University, Columbus, Ohio

RICHARD A. LEVINE

University of California, Davis, Davis, California

DENNIS J. SHEA

National Center for Atmospheric Research, Boulder, Colorado*

(Manuscript received 20 September 1999, in final form 10 January 2000)

ABSTRACT

A Bayesian fingerprinting methodology for assessing anthropogenic impacts on climate was developed. This analysis considers the effect of increased CO₂ on near-surface temperatures. A spatial CO₂ fingerprint based on control and forced model output from the National Center for Atmospheric Research Climate System Model was developed. The Bayesian approach is distinguished by several new facets. First, the prior model for the amplitude of the fingerprint is a mixture of two distributions: one reflects prior uncertainty in the anticipated value of the amplitude under the hypothesis of “no climate change.” The second reflects behavior assuming “climate change forced by CO₂.” Second, within the Bayesian framework, a new formulation of detection and attribution analyses based on practical significance of impacts rather than traditional statistical significance was presented. Third, since Bayesian analyses can be very sensitive to prior inputs, a robust Bayesian approach, which investigates the ranges of posterior inferences as prior inputs are varied, was used. Following presentation of numerical results that enforce the claim of changes in temperature patterns due to anthropogenic CO₂ forcing, the article concludes with a comparative analysis for another CO₂ fingerprint and selected discussion.

1. Introduction

Statistical analyses play a fundamental role in climate change assessment (Barnett et al. 1999). Many of the methods known as “fingerprinting” have strong roots in classical statistical hypothesis testing and regression analysis (e.g., Levine and Berliner 1999). An alternative to the traditional statistical viewpoint is offered by the Bayesian approach. To most easily illustrate the Bayesian approach, we develop a Bayesian fingerprinting analysis for near-surface temperature, restricted to the impact of CO₂ forcing only. Heuristically, the modeling process formalizes the relation:

$$\text{observations} = \mathbf{g} \times a + \text{noise}, \quad (1)$$

where a is an unknown parameter; \mathbf{g} is a spatially vary-

ing CO₂ fingerprint; and “noise” is modeled to account for *both* errors in the data and natural climate variability.

The observational dataset used here is a subset (1961–98) of the Jones annual temperature observations (Jones 1994; Jones et al. 1999). We rely on model output from the National Center for Atmospheric Research (NCAR) Climate System Model (CSM; see Boville and Gent 1998; Meehl et al. 2000) in developing a statistical model. When appropriate, datasets were spatially interpolated to the 5° × 5° Jones data grid. We used a 300-yr CSM control run as a basis for estimating natural climate variability. We used a 120-yr CO₂ forced run in combination with the control run in constructing a CO₂ fingerprint pattern \mathbf{g} . Specifically, we computed differences between temporal averages of the final 100 yr of the forced CSM run and 100 yr of the control run at each grid box. Figure 1 displays this fingerprint. (The double use of the CSM control run in developing a fingerprint and estimating climate variability may be of some concern, though we judged that to be minor for our analyses.) For comparison, we also present analyses based on the CO₂ fingerprint pattern of Santer et al. (1995).

The Bayesian viewpoint is characterized by two main points: uncertainty modeling and learning from data. In

* The National Center for Atmospheric Research is sponsored by the National Science Foundation.

Corresponding author address: Mark Berliner, Department of Statistics, The Ohio State University, 1958 Neil Avenue, Columbus, OH 43210-1247.

E-mail: mb@stat.ohio-state.edu

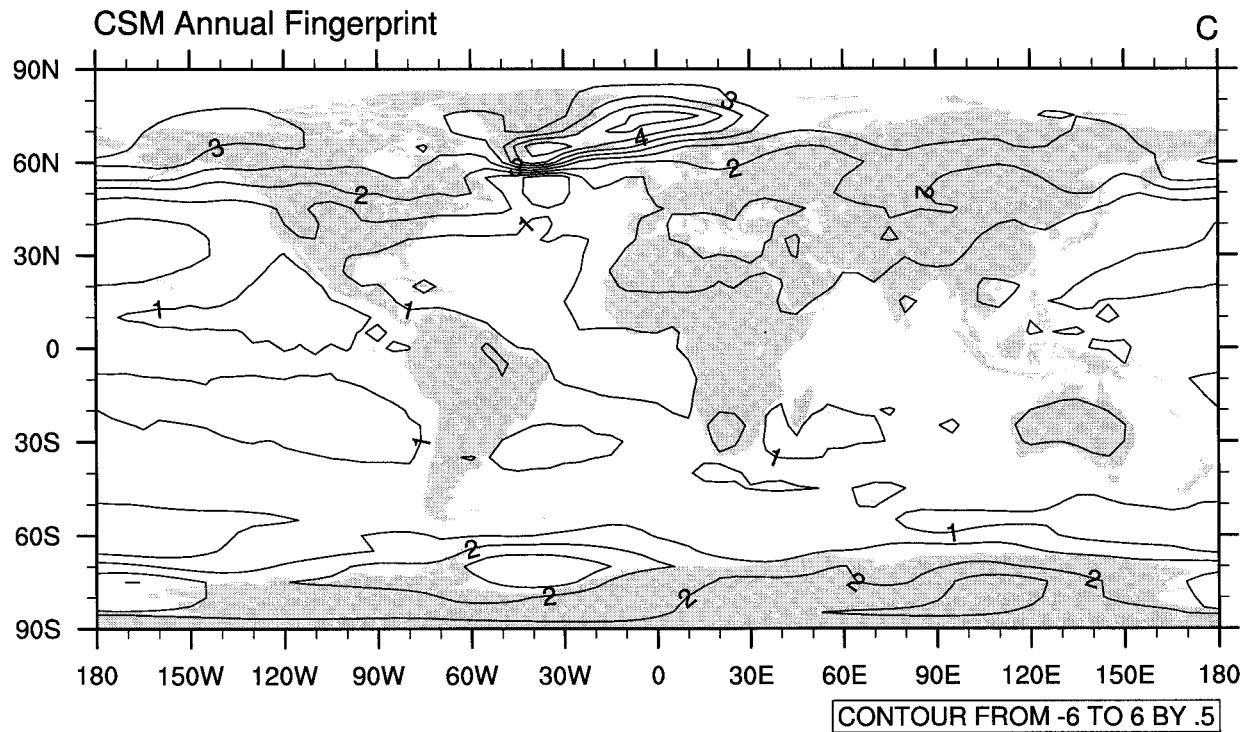


FIG. 1. Fingerprint pattern developed from NCAR CSM control and CO₂ runs on a 5° × 5° grid.

general, the uncertain quantities may be unobserved physical variables, including future values as in forecasting problems, as well as unknown model parameters. In this article, we consider only the latter setting: our analysis will be for the unknown, amplitude a introduced in (1). The formal mechanisms for achieving the above two points are as follows.

- 1) All unknowns are treated as random variables and endowed with a *prior* distribution. For our example, the prior is denoted by $\pi(a)$.
- 2) Learning is formally an application of Bayes's theorem, which yields the *posterior* distribution:

$$\pi(a | \text{data}) \propto f(\text{data} | a)\pi(a), \quad (2)$$

where $f(\text{data} | a)$ is the conditional distribution of the data given a .

This distribution, often known as the *likelihood function*, is typically formulated in conventional statistical analyses as well. An intuitive interpretation of Bayes's theorem is that once the data have been observed, the prior distribution is reweighted by the likelihood function to produce the posterior. General reviews can be found in Berger (1985), Epstein (1985), and Bernardo and Smith (1994).

Specific reasons for adopting the Bayesian approach in climate change analysis include

- inability to adopt the fundamental practices associated with traditional statistical cause–effect analysis, such

as controlled experiments comparing responses to treatments;

- the need to account for various uncertainties in observational data and our knowledge of how the climate system works;
- separation of investigation of the practical significance of potential changes as opposed merely to consideration of statistical significance.

Common approaches for dealing with the first item involve the use of climate models and their responses to various forcings or treatments. We followed suit in this article, using CSM model output, though in general a variety of prior elicitation approaches may be considered (Berger 1985; Bernardo and Smith 1994).

The second item is only partially dealt with here. Our statistical modeling described in section 2 explicitly adjusts for (i) errors in the observational data, (ii) missing data, and (iii) both spatial and temporal dependence structures present in climate variability. There is some evidence that observational uncertainties are negligible in climate change analyses for global, annual near-surface temperatures (Hegerl et al. 2000). Nevertheless, we present an approach for adjusting for data errors here because the method is simple and may be of interest in climate change studies for other physical variables with less reliable data. For most of this article, we focus on uncertainty regarding the amplitude a . First, a prior distribution for a is developed. However, the resulting prior is itself the object of uncertainty, and often a source of

controversy. To address this problem, we develop *robust Bayesian* analyses. The key is that rather than viewing the prior as fixed and perfectly specified, we view our prior information as specifying only a class of plausible priors. The robust Bayesian approach considers the ranges of posterior inferences as the prior varies over the specified class.

Perspective for the third item is best provided via comparison to non-Bayesian approaches. In the simplified framework considered here, traditional fingerprinting involves applications of the triumvirate of classical statistical methods: estimation, confidence intervals, and testing for the amplitude a . Detection of climate change is associated with rejection of the hypothesis that $a = 0$ via classical statistical testing. However, such rejection does not admit the interpretation that the event $a = 0$ has small probability, nor does it imply anything regarding the chance that “ a is near zero.” (Such statements make no sense in the strict classical statistical approach: a is not viewed as a random variable and hence, cannot be the object of probability statements.) Next, assuming a “detection,” standard approaches to attribution are intended to provide statistical assessments of the consistency of a with a value (or values) of the parameter suggested by climate model responses to hypothesized anthropogenic forcings. The idea is that a forced, model-based estimate of a , say a_M , is used as the null hypothesis of another test. Attribution is then associated with the *failure* to reject the hypothesis that $a = a_M$. Again, such a failure to reject does not imply that the event “ a is near a_M ” has high probability. Rather, it means that the observed data are not surprising if we assume that $a = a_M$. Levine and Berliner (1999) note that such an approach is not recognized as being statistically valid even from the classical viewpoint.

The Bayesian’s treatment of parameters such as a as random variables can lead to very different formulations and interpretations of results. As opposed to the indirect inferences associated with classical statistical testing, Bayesian analyses can produce direct probability statements to serve as bases for detection and attribution. In particular, we define detection as occurring when the posterior probability that a departs significantly from 0 is large. “Significance” here refers to practical significance, rather than statistical significance. Specifically, for a specified neighborhood, say \mathcal{D} of $a = 0$, we have a detection if the posterior probability of the event “ a is in \mathcal{D} ” is small. Similarly, for a specified neighborhood, say \mathcal{A} , of $a = a_M$, we (essentially) define attribution as occurring when the posterior probability that “ a is in \mathcal{A} ” is large. (We say “essentially” here, because we include another technical condition for attribution, as clarified in section 3a.)

Forms of Bayesian analyses in climate change research have been discussed in the literature. Some examples include Hasselmann (1998), Leroy (1998), and Tol and de Vos (1998). However, these papers do not provide complete views of the potential of the Bayesian

approach. Neither does this article, though we do provide substantial discussion of the flexibility of modeling and forms of inference associated with the Bayesian approach [also see Chiang et al. (1999, manuscript submitted to *J. Climate*)]. On the other hand, the scope of the model used here is limited. We make no pretense to describing a complete climate change analysis. We use a variety of simplifying assumptions. Some were chosen to enhance readability of the article; some were chosen due to computational limitations. For example, though we use a model for the space–time structure of the noise process in (1), we did not adjust for our uncertainty in that model. Inclusion of prior distributions for features (e.g., natural variability covariance matrix) of the distribution of the noise are mandated. Further, we limit the form of signals to be purely spatial and a to be constant in time. A natural generalization is to allow for time-varying amplitudes and fingerprints (e.g., Hasselmann 1993; North and Stevens 1998; Santer et al. 1995; Tett et al. 1999). In a Bayesian framework, a collection of a , would be endowed with a time series prior developed for climate change analysis. Also, extensions to Bayesian multiattribute fingerprinting are quite accessible (Hasselmann 1997; Hegerl et al. 1997). Attention would focus on a vector, say \mathbf{a} , whose dimension corresponds to the number of fingerprints used in the model. The posterior distribution of \mathbf{a} would present us with dependence structures among the elements of \mathbf{a} , that would be of interest scientifically. We will pursue these avenues elsewhere.

Next, anthropogenic climate change need not be restricted to the study of signals. For example, we might anticipate potential changes in climate variability. We could use richer models and priors to enable consideration of both signal and variability. More generally, the probability distribution of climate may change form. Other potential extensions are mentioned in section 5d.

In section 2 we describe development of a statistical model for the data and a prior distribution for the parameter a particularly tailored for climate change analysis. Section 3 offers development of a variety of posterior Bayesian inferences. These include formalizations of our above notions of Bayesian detection and attribution as well as the robust Bayesian approach. Numerical results are presented in section 4. These results support both detection of change and attribution of that change to anthropogenic CO₂ forcing. In section 5 we present companion Bayesian results for the CO₂ fingerprint described in Santer et al. (1995) and some remarks and discussion.

2. Bayesian fingerprinting: Formulation

We discuss developments of the data model or likelihood function $f(\text{data} | a)$ and prior $\pi(a)$ in the next two sections.

a. Likelihood function: $f(\text{data} | a)$

The observational data used here are the familiar Jones data (Jones 1994), consisting of monthly temperature observations on a $5^\circ \times 5^\circ$ grid; we only use data for the period 1961–98. Our analyses are for yearly averages of these data. The Jones grid leads to $n = 36 \times 72 = 2592$ spatial grid points. Some data are missing. Our analysis uses only actual data, that is, we do not impute any missing observations.

Let \mathbf{T}_t represent the n vector of “true” annual, gridded temperature anomalies in year t , $t = 1, \dots, m$. (We perform analyses for four different choices of m .) We assume that these anomalies represent departures from a baseline compatible with that chosen in preparation of the Jones data; namely, the mean of temperatures over the 30-yr period 1961–90. (In general, definitions of baselines appropriate for climate change studies are very important in the interpretation of results.) Let $\mathbf{T} = (\mathbf{T}'_1, \dots, \mathbf{T}'_m)'$ be the $m \cdot n$ vector collecting temperature anomalies over the entire period where $'$ denotes transpose. Define vectors of observations $\mathbf{Y}_1, \dots, \mathbf{Y}_m$, where \mathbf{Y}_t represents “observed” annual temperatures in year $t = 1, \dots, m$. These vectors may be of various lengths, say n_t , due to missing data. Let $\mathbf{Y} = (\mathbf{Y}'_1, \dots, \mathbf{Y}'_m)'$ be the collection of these vectors.

Our data model includes a formal adjustment for errors present in the observations. We assume that these errors all have mean zero, are independent from year to year, and have normal (Gaussian) distributions. To keep track of missing data, for each year define $n_t \times n$ incidence matrices \mathbf{X}_t , consisting of rows of zeroes, except for a single 1, indicating locations of the observations. This is summarized for each year as

$$\mathbf{Y}_t | \mathbf{T}_t, \mathbf{D}_t \sim N(\mathbf{X}_t \mathbf{T}_t, \mathbf{D}_t), \quad (3)$$

where \mathbf{D}_t is an $n_t \times n_t$ covariance matrix. Alternatively, we can represent the model for all observations as

$$\mathbf{Y} | \mathbf{T}, \mathbf{D} \sim N(\mathbf{X}\mathbf{T}, \mathbf{D}), \quad (4)$$

where \mathbf{X} is a block diagonal matrix with blocks \mathbf{X}_t and \mathbf{D} is a block diagonal matrix with blocks \mathbf{D}_t . [For completeness, note that the dimensions of \mathbf{X} are $\sum_{t=1}^m n_t \times (n \cdot m)$; the dimensions of \mathbf{D} are $\sum_{t=1}^m n_t \times \sum_{t=1}^m n_t$.]

The second stage of our model considers natural climate variability and the spatial fingerprint. In particular we assume that each year $t = 1, \dots, m$

$$\mathbf{T}_t | a, \mathbf{g}, \boldsymbol{\Sigma} \sim N(\mathbf{g}a, \boldsymbol{\Sigma}^s), \quad (5)$$

where $\boldsymbol{\Sigma}^s$ is an $n \times n$ spatial covariance matrix; \mathbf{g} is an n vector denoting the fingerprint. Note that, as in the fingerprint work of Hasselmann (1997) and the references therein, we assume that the fingerprint pattern is constant over time. Under this model, a standard probability calculation implies that *unconditional* on \mathbf{T} ,

$$\mathbf{Y} | a, \mathbf{g}, \mathbf{D}, \boldsymbol{\Sigma} \sim N(\mathbf{X}\mathbf{G}a, \mathbf{D} + \mathbf{X}\boldsymbol{\Sigma}\mathbf{X}'), \quad (6)$$

where \mathbf{G} is an $m \cdot n$ vector of m replicates of \mathbf{g} . The

matrix $\boldsymbol{\Sigma}$ is an $n \cdot m \times n \cdot m$ space–time covariance matrix developed to represent both the spatial dependence reflected in $\boldsymbol{\Sigma}^s$ and temporal dependence. We assume that $\boldsymbol{\Sigma}$ is *space–time separable* as described in appendix A [see North and Wu (1999, manuscript submitted to *J. Climate*) and Cressie and Huang (1999) for potential avenues to relaxing this assumption].

A fully Bayesian analysis would involve formulation of prior probability models for the quantities a , \mathbf{D} , $\boldsymbol{\Sigma}$, and perhaps \mathbf{g} . In view of the high dimensionality of the problem, this is a daunting task. For purposes of illustration, we follow simpler strategies of (i) formulating plausible estimates of \mathbf{D} and $\boldsymbol{\Sigma}$, and (ii) selecting fixed fingerprints. We focus on modeling and inference for the amplitude a here and describe our estimation of \mathbf{D} and $\boldsymbol{\Sigma}$ in appendix A.

Both Bayesian and traditional approaches to statistical inference benefit from application of the so-called *sufficiency principle* (Berger 1985, 126–127; Epstein 1985, 23–25). The idea is to construct functions of the data that summarize all the statistical information contained in the data regarding the unknown parameters. It turns out that for the model (6) a sufficient statistic for a is the *generalized least squares estimator* (GLS), denoted by $\hat{a}(\mathbf{Y})$ (e.g., Graybill 1976; Levine and Berliner 1999). This estimator is given by

$$\hat{a}(\mathbf{Y}) = \frac{\mathbf{H}'\mathbf{V}^{-1}\mathbf{Y}}{\mathbf{H}'\mathbf{V}^{-1}\mathbf{H}}, \quad (7)$$

where

$$\mathbf{H} = \mathbf{X}\mathbf{G} \quad \text{and} \quad \mathbf{V} = \mathbf{D} + \mathbf{X}\boldsymbol{\Sigma}\mathbf{X}'.$$

The quantity $\hat{a}(\mathbf{Y})$ in (7) is often known as the *optimal filter* in the climate science literature, though it should be clarified that it is optimal under a mean squared error criterion, and then only within the class of estimators that are both unbiased and linear in the data vector \mathbf{Y} . It is also familiar as the *most probable* value of a , though this Bayesian-like interpretation is founded on a so-called “uniform prior” assumption; we do not use uniform priors in this article.

A further result from statistics (Graybill 1976) implies that

$$\hat{a}(\mathbf{Y}) | a \sim N(a, \sigma^2), \quad (8)$$

where

$$\sigma^2 = (\mathbf{H}'\mathbf{V}^{-1}\mathbf{H})^{-1}. \quad (9)$$

In summary, we base our Bayesian inferences about a only on (8) and our prior on a . For potential extensions to multiattribute studies, we remark that computations and results for a multivariate GLS sufficient statistic for a vector of amplitudes are readily available (Graybill 1976).

b. Prior: $\pi(a)$

The issues of climate change motivate construction of the prior $\pi(a)$ as a *mixture* of two main components

as follows: with probability p , a has a normal distribution with mean 0 and variance τ^2 , and with probability $1 - p$, a has a normal distribution with mean μ_A and variance τ_A^2 . Symbolically, we write

$$\pi(a) = pn(0, \tau^2) + (1 - p)n(\mu_A, \tau_A^2), \quad (10)$$

where $n(\cdot, \cdot)$ represents the probability density function for a normal distribution; the first argument is the mean and the second is the variance.

In specifying the prior parameters, note that we should think of a and \mathbf{g} as coming in pairs. That is, our prior views concerning the actual values of a depend on the choice of \mathbf{g} . The first component of (10) is to model information about a under the qualitative statement that anthropogenic CO₂ forcing has not impacted global temperatures, at least not in a fashion reflected by the fingerprint \mathbf{g} . The selection of zero as the prior mean of a in this component of the prior seems quite natural. The variance τ^2 is interpreted as the variability, arising naturally, we anticipate in fitting a by projecting temperature fields onto the selected fingerprint \mathbf{g} . The probability p is then a quantification of the degree of belief in the hypothesis of no anthropogenic impacts.

The second component of the prior represents an implied distribution on a assuming significant anthropogenic effects representable as changes in temperature patterns along the fingerprint \mathbf{g} . The quantity μ_A is the hypothesized shift in a , which in turn translates to shifts in temperature anomalies $\mathbf{g}a$. The variance τ_A^2 reflects anticipated variability in a under CO₂ forcing. This parameter reflects, in part, natural climate variability, though there is no requirement that $\tau_A^2 = \tau^2$.

Note that the hypotheses here are modeled as distributions rather than as single points (i.e., either $a = 0$ or $a = a_A$, exactly). This is in contrast to the simplified prior used in Hasselmann (1998), which is essentially (10) with both τ^2 and τ_A^2 set equal to zero, and thus ignores a significant source of uncertainty.

The variation represented by τ^2 would be best estimated through (unavailable) replications of the climate system under a no CO₂ forcing scenario. As a surrogate, we use CSM control run output. A single 300-yr series was partitioned into 100 batches of equal time duration, each representing ‘‘replications’’ of an unforced, surrogate global climate system over short time periods. The batches are overlapping in that they are constructed by taking a running window of 200 yr. For each batch and our CSM fingerprint pattern \mathbf{g} , we computed generalized least squares estimates of a as if the batched data were real observations. That is, in each case, we assumed a model like (6) but with model output playing the role of the data \mathbf{Y} and no measurement error covariance matrix \mathbf{D} . The sample standard deviation (0.007) of the resulting 100 estimates then provides a plausible first guess at the variance τ . However, since the results batching process of the NCAR CSM runs would vary with sample size, number of batches, and

so on, as well as with different CSM runs, we actually use a somewhat larger value $\tau = 0.02$.

Prior parameters under CO₂ forcing were computed analogously to the previous discussion but using batches of the CSM CO₂ forced output. We used the sample mean of the batched results to estimate the prior mean, μ_A . The sample variance provides a guess for τ_A^2 , though we again used a slightly larger value since we are merely sampling from a particular CSM run.

3. Posterior inferences

For a prior of the mixture form (10), the posterior distribution is also a mixture (e.g., Berger 1985, 127–128, p. 206). We have

$$\begin{aligned} \pi(a|\hat{a}) = & p(\hat{a})n[\mu(a|\hat{a}), \tau^2(a|\hat{a})] \\ & + [1 - p(\hat{a})]n[\mu_A(a|\hat{a}), \tau_A^2(a|\hat{a})], \end{aligned} \quad (11)$$

where

$$\mu(a|\hat{a}) = \frac{\tau^2}{\tau^2 + \sigma^2} \hat{a} \quad \text{and} \quad \tau^2(a|\hat{a}) = \frac{\tau^2 \sigma^2}{\tau^2 + \sigma^2}, \quad (12)$$

$$\mu_A(a|\hat{a}) = \frac{\tau_A^2}{\tau_A^2 + \sigma^2} \hat{a} + \frac{\sigma^2}{\tau_A^2 + \sigma^2} \mu_A \quad \text{and}$$

$$\tau_A^2(a|\hat{a}) = \frac{\tau_A^2 \sigma^2}{\tau_A^2 + \sigma^2}. \quad (13)$$

and

$$\begin{aligned} p(\hat{a}) = & \left\langle 1 + \left(\frac{1-p}{p} \right) \sqrt{\frac{\tau^2 + \sigma^2}{\tau_A^2 + \sigma^2}} \right. \\ & \left. \times \exp \left\{ -\frac{1}{2} \left[\frac{(\hat{a} - \mu_A)^2}{\tau_A^2 + \sigma^2} - \frac{\hat{a}^2}{\tau^2 + \sigma^2} \right] \right\} \right\rangle^{-1}. \end{aligned} \quad (14)$$

This distribution is a mixture of the two posteriors corresponding to each of the two components of the prior (10). The mixing weight $p(\hat{a})$ is a function of the observed data. It is important to note that $p(\hat{a})$ is not necessarily interpretable as the posterior probability that the no-change *prior* distribution is correct. Rather, it is a weight associated with the corresponding posterior distribution. Indeed, depending on the data, either one or both of the components of the posterior distribution (11) may be very different from their corresponding priors.

a. Detection and attribution

We next formalize the specialized inspections of the posterior distribution intended to serve as Bayesian procedures for detection and attribution. One motivation is to offer quantitative measures of direct ‘‘practical significance’’ of results as opposed to restriction to indirect ‘‘statistical significance.’’ We *define* ‘‘no significant climate change’’ to be the event that $a \in \mathcal{D}$, where \mathcal{D} is

some neighborhood of zero and ϵ is read as “is an element of.” Similarly, define another set, \mathcal{A} , a neighborhood of the mean μ_A to represent an attribution set reflecting a considered to be plausible in relation to the physics of CO₂ impacts on climate as reflected in the CSM model (in our case). Recall that the physical meanings of a and \mathbf{g} are paired through their product. Hence, the sets \mathcal{D} and \mathcal{A} may vary with \mathbf{g} .

With these definitions, we can use the Bayesian model for inferences regarding detection and attribution. The notions are intuitive:

- Define detection as occurring when the posterior probability that a departs significantly from 0 is large; namely, if

$$\Pr(a \in \mathcal{D} | \hat{a}) \text{ is small.} \quad (15)$$

- Define attribution as occurring when both (i) the posterior probability

$$\Pr(a \in \mathcal{A} | \hat{a}) \text{ is large,} \quad (16)$$

and (ii) $1 - p(\hat{a})$ is large.

The two requirements in the above statement of attribution need not be redundant. Consider the following example (the numbers are completely artificial, chosen for ready illustration). Suppose that $\hat{a} | a \sim N(a, 1)$ [see (8)]. The prior parameters chosen are $p = 0.5$, $\tau^2 = 1$, $\mu_A = 5$, and $\tau_A^2 = 25$. The observed value of \hat{a} is -5 . Computing (14), we obtain $p(\hat{a}) = 0.045$, apparently suggesting that this data offers strong evidence for the anthropogenically forced climate change model. Intuitively, this suggestion seems silly. The observed $\hat{a} = -5$ is actually closer to the center, 0, of the no-change prior distribution, than to the center, 5, of the forced prior distribution. However, in a sense, it is true that the forced model offers a better explanation of the data than does the no-change model. The key is the highly different values of the prior variances. Specifically, \hat{a} is five standard deviations (under the no-change component of the prior) away from 0, but only two standard deviations away from 5 under the forced alternative component. That is, neither alternative offers a good explanation of this data, though in comparing only the two candidates, the forced model is far superior.

Two points should be clarified. First, the above behavior of $p(\hat{a})$ could not occur if $\tau^2 = \tau_A^2$. However, our analyses below do suggest that for our setting we anticipate that $\tau^2 < \tau_A^2$, so this discussion is important. The second, more critical point is that sole reliance on too simple a summary of the posterior can be misleading, though the full posterior is well behaved. Computing the actual posterior (11) for our example yields:

$$\begin{aligned} \pi(a | \hat{a} = -5) &= 0.045n(-2.5, 0.5) \\ &+ 0.995n(-4.615, 0.96). \end{aligned} \quad (17)$$

This posterior is indeed far from either of the two components of the prior; the data have overwhelmed both

of them. If we defined the negligible impact region \mathcal{D} to be the neighborhood $(-0.5, 0.5)$, we have $\Pr(a \in \mathcal{D} | \hat{a} = -5) = 0.00014$. That is, we have a detection (of something); there is strong evidence of significant departure (as expressed in \mathcal{D}) from the no-change hypothesis. On the other hand, if we define attribution regions \mathcal{A} to be even very large neighborhoods of $\mu_A = 5$ [e.g., $(2, 8)$], we have $\Pr(a \in \mathcal{A} | \hat{a} = -5) = 0.0$. That is, we find no evidence for attribution, despite the high posterior probability 0.955 on the forced alternative posterior model.

These examples, suggesting the potential need for care and inspection of the full posterior rather than simple summaries, may cause some readers to wonder if it might be easier to rely on non-Bayesian approaches. We think not, for a variety of reasons outlined in this article. Further, traditional testing is not immune to unintuitive results, but rather also requires care in interpretations. See Berger (1985) and Levine and Berliner (1999, p. 573) for examples.

Two important questions arise. First, who is to quantify what “small” and “large” mean in our definitions of detection and attribution? Second, who is to define the sets \mathcal{D} and \mathcal{A} ? Note that analogs of the first question also arise in non-Bayesian detection and attribution studies. Namely, who is to decide on cutoffs of achieved statistical significance (in the case of detection) and lack of statistical significance (in the case of attribution)? Regarding the specification of \mathcal{D} and \mathcal{A} , some might suggest that such specifications are separate from the science of climate change. Even if that is true, we at least provide descriptions of tools for quantitative assessments of the practical impacts of human influences.

Our proposition of detection and attribution analyses are different from conventional, two-step approaches. This is primarily due to the points of the previous paragraph. The Bayesian framework is not limited to simply computing direct analogs of classical procedures. For example, Hasselmann (1998) wrote “. . . it is not possible in the Bayesian approach to separate formally between detection and attribution.” We disagree, and offered such a separation.

b. Classes of priors

The prior (10) assumes that μ_A , τ^2 , and τ_A^2 are known quantities. In reality, we would be uncertain about these parameters. (This also is true for p ; one could seldom argue that p should be 0.7253, but not 0.73. However, study of the impact of varying p is quite easy.) In the Bayesian paradigm, a *hierarchical* view, in which prior distributions for these quantities are incorporated into the analysis, could be adopted. An alternative is the *robust Bayesian viewpoint*. The idea is that our information seldom is refined enough to identify a single prior distribution. Rather, our prior information specifies a *class* of priors, denoted by Γ . We then study ranges of posterior measures (probabilities, estimates, etc.) as

TABLE 1. For each of the four time periods a–d: GLS estimate \hat{a} of a and associated standard deviation σ ; mean and associated standard deviation of a under the first component (no climate change) of the posterior mixture in (11); posterior mean and associated standard deviation under the second component (CO₂ forcing) of the posterior mixture in (11); posterior weight $p(\hat{a})$, assuming prior weight $p = 0.5$, $\tau = 0.02$, $\mu_A = 0.17$, and $\tau_A = 0.05$.

Time period	1961–98		1970–98		1980–98		1988–98	
\hat{a} , σ	0.059	0.022	0.093	0.025	0.160	0.031	0.220	0.041
$\mu(a \hat{a})$, $\tau(a \hat{a})$	0.027	0.015	0.036	0.016	0.047	0.017	0.042	0.018
$\mu_A(a \hat{a})$, $\tau_A(a \hat{a})$	0.077	0.020	0.108	0.022	0.164	0.026	0.200	0.032
$p(\hat{a})$		0.671		0.062		0.0001		0.00001

the prior varies over Γ . See Berger (1985, chapter 4) for overviews and examples of both the hierarchical and robust Bayesian viewpoints. Those descriptions include more general classes of priors (relaxation of reliance on Gaussian priors, etc.) and analyses.

As in (10), consider priors of the form $\pi(a) = pn(0, \tau^2) + (1 - p)n(\mu_A, \tau_A^2)$. For ease in conveying the idea, we assume that τ^2 is fixed, but we have substantial uncertainty in μ_A and τ_A^2 . Similar calculations without this imposition are quite accessible, but our preliminary analyses for formulating estimates of the prior parameters suggest that this simplification is reasonable.

Define a collection of priors:

$$\Gamma = \{\pi(a) : \mu_l \leq \mu_A \leq \mu_u; \tau_l \leq \tau_A \leq \tau_u\}, \quad (18)$$

for specified bounds $\mu_l < \mu_u$ and $\tau_l < \tau_u$. One example of a robust Bayesian calculation is to find the range of $p(\hat{a})$ as the prior varies. During these calculations, \hat{a} , p , σ^2 , and τ^2 are fixed. Hence, from (14) we see that the maximum of $p(\hat{a})$ is of the form

$$\bar{p}(\hat{a}) = \left[1 + \left(\frac{1-p}{p} \right) M(\hat{a}) \right]^{-1}, \quad (19)$$

and the minimum is of the form

$$\underline{p}(\hat{a}) = \left[1 + \left(\frac{1-p}{p} \right) m(\hat{a}) \right]^{-1}, \quad (20)$$

where the values $M(\hat{a})$ and $m(\hat{a})$ are to be determined. The mathematics for these optimizations are outlined in appendix B. The results have a useful interpretation in terms of *odds ratios*. The quantity $(1-p)/p$ is the prior odds in favor of the anthropogenically forced component of the distribution of a . The values $M(\hat{a})$ and $m(\hat{a})$ are extremal values of the corresponding odds as reflected by the data. The products of these values and the prior odds are extremals of posterior odds ratios. We return to this interpretation when presenting numerical results in the next section.

Computing bounds for $\Pr(a \in \mathcal{D}|\hat{a})$ and $\Pr(a \in \mathcal{A}|\hat{a})$ are more complicated than those for $p(\hat{a})$. However, these bounds are easily found using simple numerical optimization methods.

4. Results

We present posterior inferences using four subsets of the Jones data corresponding to the time periods: (a)

1961–98; (b) 1970–98; (c) 1980–98; (d) 1988–98. The year 1961 was chosen as the initial time point since the baseline from which the Jones data anomalies are computed begins in 1961. These successively shorter time periods are suggested since climate change trends are anticipated to be increasingly more visible during the latter part of the twentieth century. It should be noted that this reasoning suggests that we should use different priors for a depending on the time period analyzed. For simplicity, we do not do so here.

Table 1 presents the generalized least squares estimates \hat{a} of a defined in (7) and the associated standard deviations σ computed from (9) for each of the four time periods considered. Note that, as anticipated above, these estimates grow when computed for time windows increasingly focused on the recent past.

a. Basic analysis

Our prior parameter estimates, resulting from the “batching” of CSM output described earlier, are $\tau = 0.02$, $\mu_A = 0.17$, and $\tau_A = 0.05$. (We actually inflated standard deviations obtained from those model data analyses a bit.) For these prior parameter values and each of the four time periods, Table 1 presents the means and standard deviations [see (12) and (13)] for the two components of the posterior distribution (11). Table 1 also includes the posterior weight $p(\hat{a})$ [see (14)] on the no change posterior model, assuming the prior weight $p = 0.5$. Figure 2 displays the prior and posterior distributions of a for $p = 0.5$. The plots also display the likelihood functions derived from (8). Note that the posterior distributions are bimodal for time periods a and b, but not time periods c and d. The bimodality is a consequence of the mixing of the two components in the posterior distribution (11). In the latter two time periods, the posterior mixture probability $p(\hat{a})$ is essentially zero. Thus the posterior distribution in these time periods is solely the component of the posterior distribution under CO₂ forcing. These results generally suggest that the data most resemble the component of the posterior distributions of a corresponding to anthropogenic forcing; this behavior seems overwhelming for the most recent time periods.

Figure 3 displays the probability $p(\hat{a})$ of the “no change” component of the posterior model, as a function of the prior probability p . These probabilities are

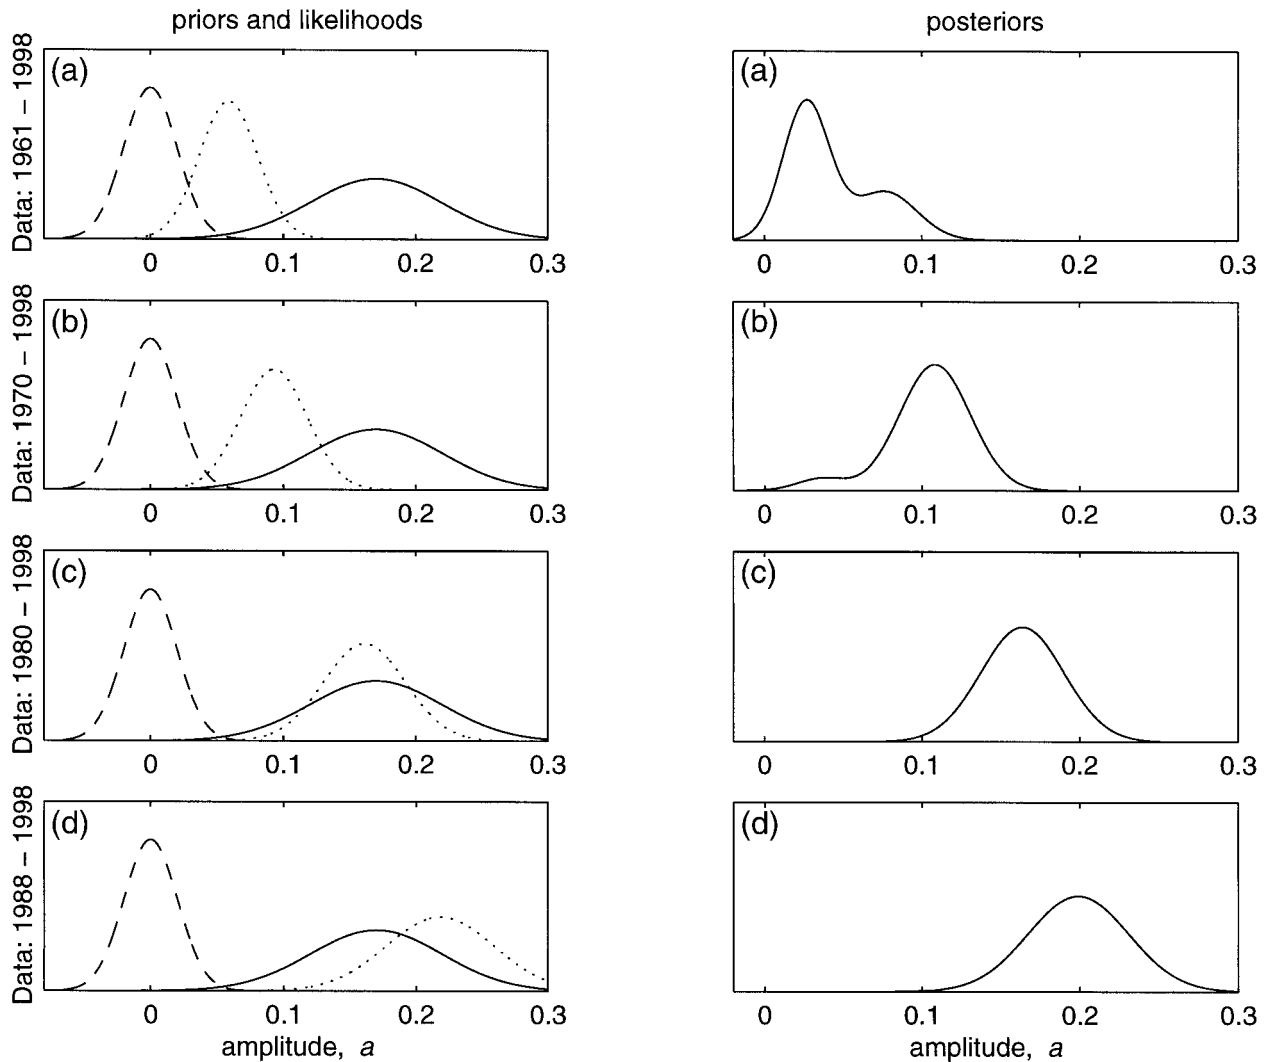


FIG. 2. Likelihood function, prior distribution, and posterior distribution of a using our NCAR CSM fingerprint. For each of the time periods (a)–(d): (left) the likelihood (dotted line) and prior distribution components [anthropogenic CO_2 forcing (solid line); no anthropogenic impacts (dashed line)], and (right) the posterior mixture distribution (11).

of course related. For example, the more confident one is in the hypothesis of “no climate change,” the more difficult it is for data to sway that opinion. The posterior probability $p(\hat{a})$ is of moderate size for time period a. In fact, for $p = 0.5$, $p(\hat{a})$ is about 0.70, providing more weight for the no change component of the posterior distribution. For time period b, the posterior probabilities $p(\hat{a})$ lie well below 0.5 for most of the range of p , thus placing substantial weight on the “anthropogenic forcing” component of the posterior distribution. We have not included corresponding plots for the remaining time periods c and d. The evidence from these periods is so overwhelmingly supportive of the forced model that $p(\hat{a})$ (and its upper bound) is essentially zero unless p is absurdly large (0.99 or so).

b. Detection and attribution results

Figures 4 and 5 present the posterior probabilities of detection and attribution regions, $\text{Pr}(a \in \mathcal{D} | \hat{a})$ and $\text{Pr}(a \in \mathcal{A} | \hat{a})$, respectively, as functions of p , for the time periods a and b. We selected $\mathcal{D} = [0 - 0.05, 0 + 0.05]$ and $\mathcal{A} = [\mu_A - 0.05, \mu_A + 0.05]$. Figure 4 shows moderate probabilities of the detection region across the range of p for time period a suggesting little or no evidence of a detected signal. In fact, when $p = 0.5$, $\text{Pr}(a \in \mathcal{D} | \hat{a})$ is about 0.6. For the same data, Fig. 5 displays small posterior probabilities of the attribution region across the range of p . However, there is some suggestion of a detected change based on time period b in that $\text{Pr}(a \in \mathcal{D} | \hat{a}) < 0.5$ for most of the range of p . However, since the corresponding posterior probability $\text{Pr}(a \in$




FIG. 3. Posterior mixture probability $p(\hat{a})$ versus the prior mixture probability p based on two data subsets: (top) 1961–98 and (bottom) 1970–98. The solid line in each denotes the posterior probability computed from the data. The dotted lines denote the upper and lower bounds on the posterior probabilities from the robust Bayesian analysis.

$\mathcal{A}|\hat{a}$) is on the order of 0.3–0.4 across the range of p , there is little evidence for attribution to the CO₂ forcing represented in our CSM-based fingerprint. These moderate values may be expected after examining Fig. 2. In that figure, the likelihood function lies in between the two components of the prior mixture and the two components of the posterior distribution overlap significantly, particularly for time period a.

For time periods c and d, the data strongly support the CO₂ forcing model. In particular, the posterior probabilities $\Pr(a \in \mathcal{D}|\hat{a})$ are essentially zero (less than 6×10^{-3} for c; less than 10^{-6} for d) for reasonable prior probabilities p . The posterior probabilities $\Pr(a \in \mathcal{A}|\hat{a})$ are constant across the range of p with values 0.92 (period c) and 0.73 (period d). We thus do not include these probabilities for time periods c and d in Figs. 4 and 5. The constant values arise because the posterior weights $1 - p(\hat{a})$ on the forced posterior component is essentially one for all p . (We explain the apparent decrease in this probability based on period d below.)

c. Robust Bayesian results

To illustrate robust Bayesian results for a plausible class of priors, we specified Γ based on the ranges $\mu_A \in [0.15, 0.19]$ and $\tau_A \in [0.05, 0.07]$ (recall that our baseline prior had $\mu_A = 0.17$, $\tau_A = 0.05$).

Table 2 displays the computed values of $M(\hat{a})$ and $m(\hat{a})$ needed to find the extrema of $p(\hat{a})$ described in (19) and (20). These extrema are plotted as functions of p for the first two time periods in Fig. 3. For time period b, the results are quite strong. For example, we have that for $p = 0.5$, $0.041 \leq p(\hat{a}) \leq 0.104$; even for the large value $p = 0.75$, $0.115 \leq p(\hat{a}) \leq 0.258$.




FIG. 4. Posterior probability $\Pr(a \in [-0.05, 0.05]|\hat{a})$ as a function of the prior probability p based on two time periods: (top) 1961–98 and (bottom) 1970–98. The solid lines in each denotes the posterior probabilities for the primary prior. The dotted lines give the upper and lower bounds on these probabilities from the robust Bayesian analysis.

For the two later time periods c and d the very strong evidence against the no change hypothesis remains even over this class of priors. The values of m and M for these time periods are overwhelmingly supportive of the CO₂ forced distribution. The posterior probability of the no change model is essentially zero (to three decimal points) over the entire class for the third data subset, unless one has an absurdly high prior odds on the no change model. For example, one would need a prior odds in favor of no change to about 6300-to-1 to be indifferent on the hypotheses; that is, for the posterior




FIG. 5. Posterior probability $\Pr(a \in [\mu_A - 0.05, \mu_A + 0.05]|\hat{a})$ as a function of the prior probability p based on two time periods: (top) 1961–98 and (bottom) 1970–98. The solid lines in each denotes the posterior probabilities for the data. The dotted lines give the upper and lower bounds on these probabilities from the robust Bayesian analysis.

TABLE 2. Upper and lower bounds on data-based contributions to the odds ratio in favor of the CO₂ forced posterior model for each of the four data subsets.

1961–98		1970–98		1980–98		1988–98	
m	M	m	M	m	M	m	M
1.3339	0.2174	23.1546	8.6383	8776	6283.24	71 150	43 524

range on $p(\hat{a})$ to include 0.5. The results for the last data subset are even more extreme. However, we do not claim that based on the fourth time period, we have proven that starting with an indifferent prior ($p = 0.5$), the odds in favor of the CO₂ forced model are at least 43 524-to-1. Such extreme odds and corresponding probabilities are computed as *tail events* of the Gaussian distributions assumed in this article. While these Gaussian distributions appear plausible, they can imply probabilities of tail events that are subject to round-off error and are not testable or believable. Further, these probabilities are based on all the other assumptions made concerning other parameter specifications. Hence, while we find these results compelling, the actual values obtained should not be considered exact when they are too extreme to be experienced. However, rounding off 43 524-to-1 to 1000-to-1 or even 100-to-1 remains of serious interest.

Figures 4 and 5 present the upper and lower bounds on the detection and attribution probabilities, $\Pr(a \in [-0.05, 0.05] | \hat{a})$ and $\Pr(a \in [\mu_A - 0.05, \mu_A + 0.05] | \hat{a})$, respectively for time periods a and b. Notice that the ranges for the probability of the detection region are tight for each of the two time periods. Though not displayed here, the latter two time periods c and d show similar behavior with probability ranges bounded very close to zero. This suggests that our interpretation of these posterior probabilities is “robust” with respect to priors varying within Γ . Our conclusion is that there is very strong evidence of a detected signal based on the last three time periods.

The attribution region probabilities are also robust over Γ . Based on time period a, the posterior probabilities of \mathcal{A} lie within a range of zero to 0.3, suggesting no ability to attribute a climate change signal to anthropogenic forcing. For time period b, the attribution probabilities may fall slightly above 0.5 depending on prior parameter specification on Γ . The range over which the posterior probability of \mathcal{A} may fall is tight for periods c and d as well, being (0.8, 0.92) and (0.45, 0.95) respectively. These ranges are constant across p and are thus not shown in Fig. 5. Hence our conclusion that the data most resemble the component of the posterior distribution corresponding to anthropogenic forcing is reasonable over this range of prior distributions. The larger range, (0.45, 0.95), on the last time period is a consequence of the large estimated signal, $\hat{a} = 0.22$, actually being greater than the upper bound on μ_A used to define Γ . Since these data do intuitively support strong projection onto our fingerprint,

we still consider this good evidence for attribution. In fact, had both our range for μ_A been more conservative and our definition of \mathcal{A} been larger to include large a , the resulting probabilities of \mathcal{A} would be much greater. (Recall that such increases are actually in accord with our prior beliefs, though we did not vary the prior on a with the time period considered.)

5. Discussion

a. An alternate fingerprint

We performed an analogous analysis using the CO₂ fingerprint (see Fig. 6) developed in Santer et al. (1995). This fingerprint uses the GRANTOUR/CCM1 experimental configuration (Taylor and Penner 1994). Table 3 presents the GLS estimates \hat{a} of a and associated standard deviations σ for the time periods a–d. Note that these results are similar to those based on our CSM fingerprint, except that the estimated amplitudes for the Santer et al. fingerprint are consistently smaller. However, recall our earlier alert that amplitudes and fingerprints should be viewed as arising in pairs. This means that we should redefine our basic prior, class of priors, and detection and attribution regions for use in analyzing the Santer et al. fingerprint. One option is to repeat the batching ideas used earlier. A second option is to view the adjustments needed as a problem of standardization to place each fingerprint on a common scale for a fair comparison. We chose the second option.

Let \mathbf{g} denote our CSM fingerprint, \mathbf{w} denote the Santer fingerprint, and a_g and a_w denote the corresponding amplitudes. We consider renormalization of the fingerprints by their L_2 norms (sum of squared values): $\|\mathbf{g}\|_2 = 90$ and $\|\mathbf{w}\|_2 = 100$. Standardizing each fingerprint by the square root of its norm, we suggest that if we are to compare projections of the data

$$\mathbf{Y} \approx \sqrt{90}a_g(\mathbf{g}\sqrt{1/90}) \approx \sqrt{100}a_w(\mathbf{w}\sqrt{1/100}),$$

then we expect that $a_w \approx 0.95 a_g$.

A simple transformation of the prior developed for a_g takes us to a comparable prior for a_w : namely, $\mu_A = (0.95)(0.17) = 0.162$, $\tau_A = (0.95)(0.05) = 0.0475$, and $\tau = (0.95)(0.02) = 0.019$. Since this argument is quite ad hoc and leaves us less certain about where we anticipate a_w to be, we thought it reasonable to use a larger class of priors for a robust Bayesian analysis: namely, we assumed that Γ_w be generated by the ranges $\mu_A \in [0.13, 0.21]$ and $\tau_A \in [0.05, 0.08]$. (Compare these with $\mu_A \in [0.15, 0.19]$ and $\tau_A \in [0.05, 0.07]$ used for our fingerprint.)

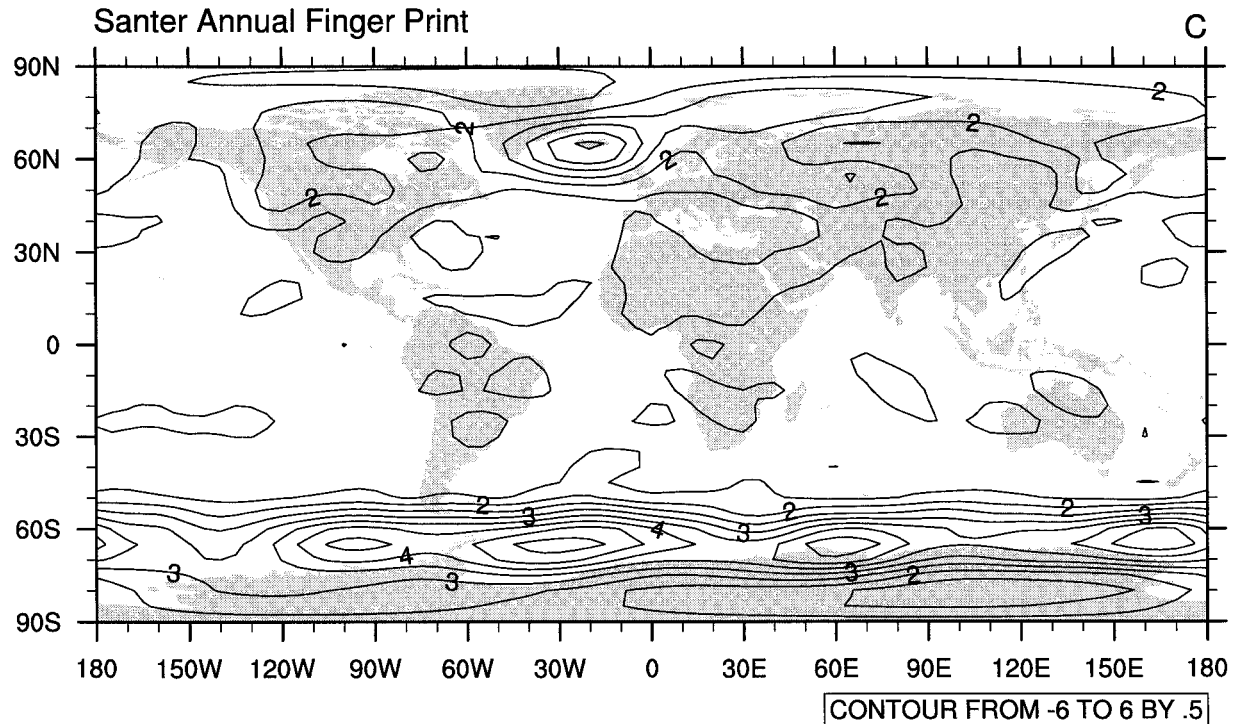


FIG. 6. Fingerprint pattern generated by Santer et al. (1995), after their Fig. 3b, on a $5^\circ \times 5^\circ$ grid.

For brevity, we present results only for detection and attribution analyses. We again associate detection with small probabilities $\Pr(a_w \in [-0.05, 0.05] | \hat{a})$, and attribution with large probabilities $\Pr(a_w \in [\mu_A - 0.05, \mu_A + 0.05] | \hat{a})$ (recall, now the baseline $\mu_A = 0.162$). Figures 7 and 8 are plots of these probabilities and ranges as functions of p for the time periods a and b. For time periods c and d, the probabilities of the detection region are tightly bounded around zero for reasonable values of p and thus not presented in Fig. 7. For these same time periods, the probabilities of the attribution region are constant across the range of p with values 0.65 and 0.95 and constant bounds from the robust analysis (0.05, 0.98) and (0.90, 0.96), respectively.

The results appear somewhat less definitive for the Santer fingerprint than for the CSM fingerprint. In particular, the ranges on probabilities appear larger, but of course the class Γ_w is larger than its counterpart. However, there is one major exception: time period d now yields very strong and robust attribution. Note that the estimate of a_w from this period does fall in the range of means defined by Γ_w . Indeed, our overall conclusion

is in concert with that of our CSM-based fingerprint: namely, there is strong evidence in the recent data for detection and attribution to CO_2 forcing as reflected in either fingerprint.

b. Comparison to traditional results

We provide traditional, non-Bayesian detection and attribution results. By non-Bayesian detection, we mean rejection of the null hypothesis that $a = 0$, when compared to the alternative hypothesis $a > 0$. Such a test rejects when the significance probability or p value for the test is small. By non-Bayesian attribution, we mean failure to reject (or “acceptance” of) the null hypothesis that $a = \mu_A$, when compared to the alternative hypothesis $a \neq \mu_A$. Intuitively this is indicated by a relatively large significance probability. Note that the results given in Table 4 are generally in accordance with simple summaries of our Bayesian results. One exception is the apparent strong, traditional detection results for both fingerprints based on the full data for the time period 1961–98. Our companion Bayesian analyses are more

TABLE 3. GLS estimates of a and associated standard deviations for each of the four data subsets using the Santer et al. (1995) fingerprint.

Fingerprint	1961–98		1970–98		1980–98		1988–98	
	\hat{a}	σ	\hat{a}	σ	\hat{a}	σ	\hat{a}	σ
Santer	0.0394	0.017	0.062	0.019	0.1117	0.024	0.1508	0.031
CSM	0.0589	0.022	0.093	0.025	0.1612	0.031	0.2200	0.041




FIG. 7. Posterior probability $\Pr(a \in [-0.05, 0.05] | \hat{a})$ as a function of the prior probability p based on two time periods: (top) 1961–98 and (bottom) 1970–98. The solid lines in each denotes the posterior probabilities for the data. The dotted lines give the upper and lower bounds on these probabilities from the robust Bayesian analysis.

conservative. This is a result of the fact that the non-Bayesian calculation only considers the plausibility of the hypothesis $a = 0$ exactly, while the our Bayesian formulation studies the plausibility of a lying near zero. Such conservatism through more attention to uncertainty suggests that one may derive comparatively more confidence in our Bayesian detections when they do occur. The situation is somewhat reversed in the context of attribution. The traditional test considers the plausibility of the the value $a = \mu_A$ exactly, while our Bayesian formulation essentially considers events of the form “ a lies near μ_A .” Nevertheless, the use of robust Bayesian calculations in conjunction with inspection of practical significance of results suggest both the inferential value and robustness of the Bayesian viewpoint.

c. Computation

Though the form of GLS estimates used here is simple, the actual calculations for finding \hat{a} and σ^2 are not easy, primarily due to the presence of the inverse of the very high dimensional matrix \mathbf{V} . However, the calculations only require that we compute two quadratic forms [the numerator and denominator in (7)]. Rather, than attempting inversion of \mathbf{V} , we numerically solved the linear equation

$$\mathbf{V}\mathbf{z} = \mathbf{H}. \quad (21)$$

Since $\mathbf{z} = \mathbf{V}^{-1}\mathbf{H}$, the computations are completed by finding two products: $\mathbf{Y}'\mathbf{z}$ and $\mathbf{H}'\mathbf{z}$. We used the *conjugate gradient* algorithm to approximate the solution \mathbf{z} (Golub and van Loan 1996).

d. Extensions

An intriguing and important issue involves the selection of fingerprints. In our analysis we viewed the fingerprint \mathbf{g} as fixed. Of course, we might also quantify our uncertainty about \mathbf{g} (e.g., Allen and Tett 1999). (Before doing so, we must entertain the question of the existence of a \mathbf{g} .) In developing a joint prior on a_g and \mathbf{g} , a two-step thought process may be a value. Any joint prior, $\pi(\mathbf{g}, a_g)$ can be factored as $\pi(\mathbf{g}, a_g) = \pi(a_g | \mathbf{g})\pi(\mathbf{g})$. This means one can build the prior on the fingerprints, and then a prior on amplitudes conditional on fingerprints. For example, we might envision a collection of models or theories, each of which produces particular forms of fingerprints, and then implied distributions on the a_g . One can use the Bayesian strategy to assess the models (see Leroy 1998), as well as to combine information from them. In the statistics liter-




FIG. 8. Posterior probability $\Pr(a \in [\mu_A - 0.05, \mu_A + 0.05] | \hat{a})$ as a function of the prior probability p based on two time periods: (top) 1961–98 and (bottom) 1970–98. The solid lines in each denotes the posterior probabilities for the data. The dotted lines give the upper and lower bounds on these probabilities from the robust Bayesian analysis.

TABLE 4. Significance probabilities for traditional, non-Bayesian detection and attribution test results. Detection is suggested by very small values; attribution may be suggested by moderate or large values.

	1961–98		1970–98		1980–98		1988–98	
	CSM	Santer	CSM	Santer	CSM	Santer	CSM	Santer
Detection	0.004	0.010	0.000	0.001	0.000	0.000	0.000	0.000
Attribution	0.000	0.000	0.000	0.000	0.744	0.037	0.220	0.720

ature, these tracks are known as *model selection* and *model averaging*, respectively.

The specification and interpretations of modeled quantities are intricately related. We have already alluded to the intertwining of fingerprints \mathbf{g} and amplitudes a . In addition, the meaning of climate variability, expressed through Σ must be gauged against the specific that the primary model is essentially a signal-plus-noise model. That is, Σ does not exist in isolation but rather is defined to reflect variation around a particular specification of the mean (taken here to be zero) of anomalies. The difficulty is that the anomalies are in turn defined as departures from some, typically ad hoc, baseline. These apparently subtle issues merit more careful attention to enable unequivocal interpretation of statistically estimated amplitudes.

e. Objectivity versus subjectivity

Bayesian statistics is by nature subjective; people with different priors may arrive at different conclusions. This issue has a long history and many interesting points of discussion (Berger 1985; Bernardo and Smith 1994), especially since traditional statistics methods are often thought to be objective. In complex settings it is difficult to argue that only objectivity really guides analysis; a variety of assumptions are made about the modeling. Indeed, substantial prior information is used in deciding on methods of analysis and data sources to be used, as well as in the final interpretation of results. Rather than “sweeping them under the carpet” (a phrase due to I. J. Good in this context), the Bayesian seeks to make clear what the assumptions are and quantifies them. This enables comparatively simple assessments of model assumptions and permits inspections of the sensitivity of results to these assumptions. Of course, we admit that there is essentially no such thing as a perfect Bayesian, and in that spirit, alluded to the robust Bayesian viewpoint.

However, even if we accept that classical statistical results are objective, we reiterate that they are only indirect answers to the fundamental questions typically of interest. Further, what about climate change science should be objective? The public is not interested in truly objective predictions regarding climate change, and so on. Rather, the most informed, intelligent predictions ought to be the goal. This statement should not be misinterpreted to suggest that analysts should base analyses on their personal hidden agendas, bounties on “proofs” of anthropogenic climate change or the reverse, and so on. Rather than being objective, analyses should be *impartial* (a word used by Laplace in this context), but informed by the best information and scientific opinion available. Again, there may be no such thing as a perfect and impartial Bayesian. Nevertheless, Bayesians suggest that the most effective way to seek impartiality as a community is for analyses to include description and quantification of the prior information used.

Acknowledgments. This research was supported in part by the Geophysical Statistics Project at the National Center for Atmospheric Research in Boulder, Colorado. LMB was also supported by the U. S. Environmental Protection Agency Agreement R827257-01-9. We are grateful to Rol Madden, Doug Nychka, and Chris Wikle for valuable discussions and comments on preliminary versions of this article. Tim Hoar shared valuable advice on large dataset manipulations. We also thank Phil Jones, Tim Osborne, and Mark New for providing “Jones data” and error variance estimates; Lawrence Buja for help in accessing CSM data; and Tom Wigley for providing the “Santer et al. fingerprint.” Last, the Editor, Francis Zwiers, and two referees provided valuable and insightful comments and suggestions for this article.

APPENDIX A

Estimation of Variability

a. Data errors covariance

The observational error covariance matrix \mathbf{D} was obtained from the standard errors of Jones et al. (1997). These standard errors were computed for each decade from 1856 to 1998. Figure A1 presents the average of these decadal standard errors. We assumed the yearly standard errors were equal to each other within each decade. Hence, we set the annual standard errors to be $(10)^{1/2}$ times the decadal standard errors provided in Jones et al. (1997).

b. Natural climate variability

First, we define the Kronecker product of two matrices $\mathbf{B}_{I \times J}$ and $\mathbf{C}_{K \times L}$. Let $B(i, j)$ denote the i, j th element of \mathbf{B} . Define $\mathbf{B} \otimes \mathbf{C}$ to be the $IJK \times IJL$ matrix obtained by expanding the products $B(i, j) \times \mathbf{C}$ over all pairs (i, j) .

Our estimation of Σ is based on a critical assumption. Namely, we assume that Σ is *space–time separable*, in that it can be factored as

$$\Sigma = \Sigma^c \otimes \Sigma^s, \quad (\text{A1})$$

where Σ^c is an $m \times m$ matrix representing temporal dependencies and Σ^s is an $n \times n$ matrix representing spatial dependence structure. This assumption implies that we can estimate the spatial and temporal matrices separately. While we do not believe space–time separability is strongly defensible, we make the assumption for practical reasons. It also offers some adjustment for temporal dependence.

1) ESTIMATION OF Σ^c

We set Σ^c to be the *correlation matrix* for an autoregressive process of order one (Brockwell and Davis 1991). In particular, all diagonal elements are equal to

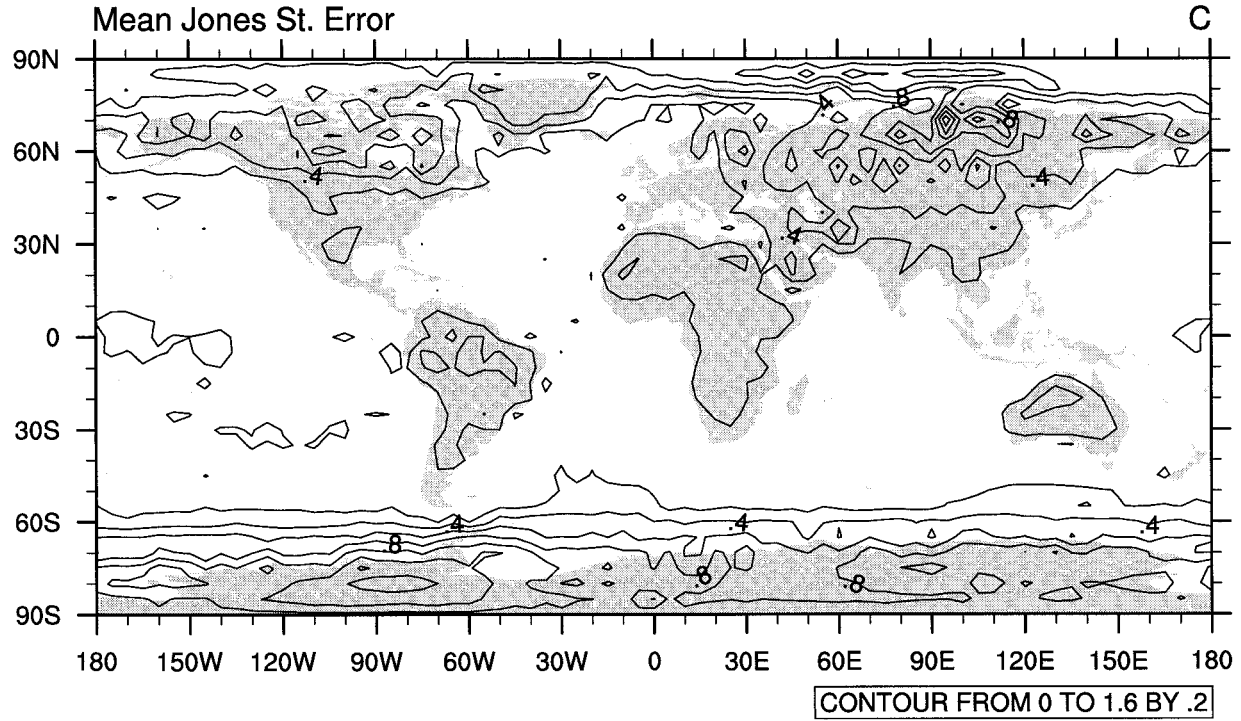


FIG. A1. Standard errors from Jones et al. (1997) for the global temperature data over the period 1856–1998 on a $5^\circ \times 5^\circ$ grid.

one. Off diagonal elements are of the form ρ^d , where d is the time lag in years. We estimated ρ from the CSM control run. The result was $\hat{\rho} = 0.33$.

2) ESTIMATION OF Σ^s

For the space–time separable model with Σ^c specified as a correlation matrix, Σ^s is a spatial covariance of \mathbf{T}_t for every year t . That is, it can be estimated (using surrogate model output) by averaging over years. Data from the 300-yr NCAR CSM control run are denoted by $\mathbf{C} = (\mathbf{C}'_1, \dots, \mathbf{C}'_{300})'$, where \mathbf{C}_i is an n vector of the model output of global temperatures at time $i = 1, \dots, 300$. We model Σ^s as the sum

$$\Sigma^s = \frac{1}{300} \left(\nu \mathbf{I} + \sum_{i=1}^{20} \lambda_i^2 \mathbf{P}_i \mathbf{P}'_i \right), \quad (\text{A2})$$

where the $\lambda_i - \mathbf{P}_i$ are the first 20 singular value EOF pairs from the NCAR CSM control run. Here ν was set to 100. (The division by 300 is an adjustment for sample size.)

APPENDIX B

Robust Bayesian Calculations

We compute the minimum and maximum values of $p(\hat{a})$ over the class Γ defined in (18). We assume that the lower bound, μ_l , on μ_A is a positive value. We also assume that \hat{a} is positive.

We rewrite $p(\hat{a})$ in (14) as

$$p(\hat{a}) = \left\{ 1 + \left(\frac{1-p}{p} \right) \sqrt{\tau^2 + \sigma^2} \exp \left[0.5 \left(\frac{\hat{a}^2}{\sqrt{\tau^2 + \sigma^2}} \right) \right] \times Q(\mu_A, \tau_A^2) \right\}^{-1}, \quad (\text{B1})$$

where

$$Q(\mu_A, \tau_A^2) = \sqrt{\frac{1}{\tau_A^2 + \sigma^2}} \exp \left[-0.5 \left(\frac{(\hat{a} - \mu_A)^2}{\tau_A^2 + \sigma^2} \right) \right]. \quad (\text{B2})$$

Minimization (maximization) of $p(\hat{a})$ involves maximization (minimization) of Q . The special form of Q enables these optimizations in two simple steps.

a. Case 1: $\hat{a} < \mu_l$

To minimize Q we should first set $\mu_A = \hat{a}$, since $Q(\mu_u, \tau_A^2) \leq Q(\mu_A, \tau_A^2)$ for all τ_A in Γ . It remains then to minimize $Q(\mu_u, \tau_A^2)$ with respect to τ_A . The derivative of this function with respect to τ_A^2 is

$$\frac{dQ(\mu_u, \tau_A^2)}{d\tau_A^2} = H \left[\frac{(\mu_u - \hat{a})^2}{\tau_A^2 + \sigma^2} - 1 \right], \quad (\text{B3})$$

where H is strictly positive. Therefore, we must consider three subcases:

Case a: If $(\mu_u - \hat{a})^2 / (\tau_A^2 + \sigma^2) > 1$ on Γ , Q is in-

creasing and the minimizing value of τ_a^2 is the endpoint τ_l^2 .

Case b: If $(\mu_u - \hat{a})^2/(\tau_A^2 + \sigma^2) < 1$ on Γ , Q is decreasing and the minimizing value of τ_a^2 is the endpoint τ_u^2 .

Case c: If $(\mu_u - \hat{a})^2/(\tau_A^2 + \sigma^2) = 1$ somewhere in Γ , Q is minimized at either τ_l^2 or τ_u^2 ; one simply checks these two values numerically.

Next, to maximize Q we first set $\mu_A = \mu_l$, since $Q(\mu_l, \tau_A^2) \geq Q(\mu_A, \tau_A^2)$ for all τ_A in Γ . To maximize $Q(\mu_l, \tau_A^2)$ with respect to τ_A , repeat the differentiation argument as in (B3). This again leads to three subcases:

Case a: If $(\mu_l - \hat{a})^2/(\tau_A^2 + \sigma^2) > 1$ on Γ , Q is increasing and the maximizing value of τ_a^2 is the endpoint τ_u^2 .

Case b: If $(\mu_l - \hat{a})^2/(\tau_A^2 + \sigma^2) < 1$ on Γ , Q is decreasing and the maximizing value of τ_a^2 is the endpoint τ_l^2 .

Case c: If $(\mu_l - \hat{a})^2/(\tau_A^2 + \sigma^2) = 1$ somewhere in Γ , that value is the maximizer; namely, $\tau_A^2 = (\mu_l - \hat{a})^2 - \sigma^2$.

b. Case 2: $\mu_l \leq \hat{a} \leq \mu_u$

To minimize Q we set μ_A to be that endpoint of the interval that is farthest from \hat{a} . Let μ^* denote the result. That is, $\mu^* = \mu_u$, if

$$\hat{a} \leq \frac{\mu_l + \mu_u}{2},$$

otherwise, $\mu^* = \mu_l$. Then we again turn to differentiation to complete the minimization.

Case a: If $(\mu^* - \hat{a})^2/(\tau_A^2 + \sigma^2) > 1$ on Γ , the minimizing value of τ_a^2 is τ_l^2 .

Case b: If $(\mu^* - \hat{a})^2/(\tau_A^2 + \sigma^2) < 1$ on Γ , the minimizing value of τ_a^2 is τ_u^2 .

Case c: If $(\mu^* - \hat{a})^2/(\tau_A^2 + \sigma^2) = 1$ somewhere in Γ , the minimizer is either τ_l^2 or τ_u^2 .

To maximize Q we first set $\mu_A = \hat{a}$, since $Q(\hat{a}, \tau_A^2) \geq Q(\mu_A, \tau_A^2)$ for all τ_A in Γ . Hence, we need only maximize $(\tau_A^2 + \sigma^2)^{-1/2}$, which of course leads to τ_l^2 as the solution.

c. Case 3: $\hat{a} > \mu_u$

This case is similar to Case 1. To minimize Q we set $\mu_A = \mu_l$. Then

Case a: If $(\mu_l - \hat{a})^2/(\tau_A^2 + \sigma^2) > 1$ on Γ , the minimizer is τ_l^2 .

Case b: If $(\mu_l - \hat{a})^2/(\tau_A^2 + \sigma^2) < 1$ on Γ , the minimizer is τ_u^2 .

Case c: If $(\mu_l - \hat{a})^2/(\tau_A^2 + \sigma^2) = 1$ somewhere in Γ , the minimizer is either τ_l^2 or τ_u^2 .

To maximize Q set $\mu_A = \mu_u$. Next,

Case a: If $(\mu_u - \hat{a})^2/(\tau_A^2 + \sigma^2) > 1$ on Γ , the maximizer is τ_u^2 .

Case b: If $(\mu_u - \hat{a})^2/(\tau_A^2 + \sigma^2) < 1$ on Γ , the maximizer is τ_l^2 .

Case c: If $(\mu_u - \hat{a})^2/(\tau_A^2 + \sigma^2) = 1$ somewhere in Γ , the maximizer is $(\mu_u - \hat{a})^2 - \sigma^2$.

REFERENCES

- Allen, M. R., and S. F. B. Tett, 1999: Checking for model consistency in optimal fingerprinting. *Climate Dyn.*, **15**, 419–434.
- Barnett, T. P., and Coauthors, 1999: Detection and attribution of climate change: A status report. *Bull. Amer. Meteor. Soc.*, **80**, 2631–2659.
- Berger, J. O., 1985: *Statistical Decision Theory and Bayesian Analysis*. 2d ed. Springer-Verlag, 617 pp.
- Bernardo, J. M., and A. F. M. Smith, 1994: *Bayesian Theory*. John Wiley and Sons, 586 pp.
- Boville, B. A., and P. Gent, 1998: The NCAR Climate System Model, version one. *J. Climate*, **11**, 1115–1130.
- Brockwell, P. J., and R. A. Davis, 1991: *Times Series: Theory and Methods*. Springer-Verlag, 419 pp.
- Cressie, N., and H.-C. Huang, 1999: Classes of nonseparable, spatio-temporal stationary covariance functions. *J. Amer. Stat. Assoc.*, **94**, 1330–1340.
- Epstein, E. S., 1985: *Statistical Inference and Prediction in Climatology: A Bayesian Approach*. Meteor. Monogr. No. 42, Amer. Meteor. Soc., 199 pp.
- Golub, G. H., and C. F. van Loan, 1996: *Matrix Computations*. 3d ed. Johns Hopkins University Press, 642 pp.
- Graybill, F. A., 1976: *Theory and Application of the Linear Model*. Duxbury Press, 704 pp.
- Hasselmann, K., 1993: Optimal fingerprinting for the detection of time-dependent climate change. *J. Climate*, **6**, 1957–1969.
- , 1997: Multi-pattern fingerprint method for detection and attribution of climate change. *Climate Dyn.*, **13**, 601–611.
- , 1998: Conventional and Bayesian approach to climate-change detection and attribution. *Quart. J. Roy. Meteor. Soc.*, **124**, 2541–2565.
- Hegerl, G. C., K. Hasselmann, C. Cubasch, J. F. B. Mitchell, E. Roeckner, R. Voss, and J. Waszkewitz, 1997: Multi-fingerprint detection and attribution analysis of greenhouse gas, greenhouse gas-plus-aerosol and solar forced climate change. *Climate Dyn.*, **13**, 613–634.
- , P. D. Jones, and T. P. Barnett, 2000: Effect of observational sampling error on the detection of anthropogenic climate change. *J. Climate*, in press.
- Jones, P. D., 1994: Hemispheric surface air temperature variations: A reanalysis and an update to 1993. *J. Climate*, **7**, 1794–1802.
- , T. J. Osborn, and K. R. Briffa, 1997: Estimating sampling errors in large-scale temperature averages. *J. Climate*, **10**, 2548–2568.
- , M. New, D. E. Parker, S. Martin, and I. G. Rigor, 1999: Surface air temperature and its changes over the past 150 years. *Rev. Geophys.*, **37**, 173–199.
- Leroy, S. S., 1998: Detecting climate signals: Some Bayesian aspects. *J. Climate*, **11**, 640–651.
- Levine, R. A., and L. M. Berliner, 1999: Statistical principles for climate change studies. *J. Climate*, **12**, 565–574.
- Meehl, G. A., W. D. Collins, B. A. Boville, J. T. Kiehl, T. M. L. Wigley, and J. M. Arblaster, 2000: Response of the NCAR Climate System Model to increased CO₂ and the role of physical processes. *J. Climate*, **13**, 1879–1898.
- North, G. R., and M. J. Stevens, 1998: Detecting climate signals in the surface temperature record. *J. Climate*, **11**, 563–577.

- Santer, B. D., K. E. Taylor, T. M. L. Wigley, J. E. Penner, P. D. Jones, and U. Cubasch, 1995: Towards the detection and attribution of an anthropogenic effect on climate. *Climate Dyn.*, **12**, 77–100.
- Taylor, K. E., and J. E. Penner, 1994: Response of the climate system to atmospheric aerosols and greenhouse gases. *Nature*, **369**, 734–737.
- Tett, S. F. B., P. A. Stott, M. R. Allen, W. J. Ingram, and J. F. B. Mitchell, 1999: Causes of twentieth century temperature change near the earth's surface. *Nature*, **399**, 569–572.
- Tol, R. S. J., and A. F. de Vos, 1998: A Bayesian statistical analysis of the enhanced greenhouse effect. *Climatic Change*, **38**, 87–112.