THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A Bayesian Climate Change Detection and Attribution Assessment

Citation for published version:

Lee, TCK, Zwiers, FW, Hegerl, GC, Zhang, XB & Tsao, M 2005, 'A Bayesian Climate Change Detection and Attribution Assessment', *Journal of Climate*, vol. 18, no. 13, pp. 2429-2440.
<https://doi.org/10.1175/JCLI3402.1>

Digital Object Identifier (DOI):

[10.1175/JCLI3402.1](https://doi.org/10.1175/JCLI3402.1)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Journal of Climate

Publisher Rights Statement:

Published in the Journal of Climate by the American Meteorological Society (2005)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A Bayesian Climate Change Detection and Attribution Assessment

TERRY C. K. LEE

Department of Mathematics and Statistics, University of Victoria, Victoria, British Columbia, Canada

FRANCIS W. ZWIERS

Canadian Centre for Climate Modelling and Analysis, Meteorological Service of Canada, University of Victoria, Victoria, British Columbia, Canada

GABRIELE C. HEGERL

Department of Earth and Ocean Sciences, Duke University, Durham, North Carolina

XUEBIN ZHANG

Climate Monitoring and Data Interpretation Division, Meteorological Service of Canada, Downsview, Ontario, Canada

MIN TSAO

Department of Mathematics and Statistics, University of Victoria, Victoria, British Columbia, Canada

(Manuscript received 8 July 2004, in final form 2 December 2004)

ABSTRACT

A Bayesian analysis of the evidence for human-induced climate change in global surface temperature observations is described. The analysis uses the standard optimal detection approach and explicitly incorporates prior knowledge about uncertainty and the influence of humans on the climate. This knowledge is expressed through prior distributions that are noncommittal on the climate change question. Evidence for detection and attribution is assessed probabilistically using clearly defined criteria. Detection requires that there is high likelihood that a given climate-model-simulated response to historical changes in greenhouse gas concentration and sulphate aerosol loading has been identified in observations. Attribution entails a more complex process that involves both the elimination of other plausible explanations of change and an assessment of the likelihood that the climate-model-simulated response to historical forcing changes is correct. The Bayesian formalism used in this study deals with this latter aspect of attribution in a more satisfactory way than the standard attribution consistency test. Very strong evidence is found to support the detection of an anthropogenic influence on the climate of the twentieth century. However, the evidence from the Bayesian attribution assessment is not as strong, possibly due to the limited length of the available observational record or sources of external forcing on the climate system that have not been accounted for in this study. It is estimated that strong evidence from a Bayesian attribution assessment using a relatively stringent attribution criterion may be available by 2020.

1. Introduction

Statistical analysis plays a major role in climate change detection and attribution studies. The method widely used in such studies is known as “optimal fin-

gerprinting” and utilizes the generalized multiple regression model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (1)$$

where vector $\mathbf{y}$ contains a filtered version of the observed temperature record, matrix $\mathbf{X}$ contains estimated response patterns (signal patterns) to known forcings on the climate system, and vector $\boldsymbol{\varepsilon}$ contains natural variability that is generated by internal climate pro-

Corresponding author address: Dr. Francis W. Zwiers, Canadian Centre for Climate Modelling and Analysis, Meteorological Service of Canada, P.O. Box 1700, STN CSC, Victoria, BC V8W 2Y2, Canada.
E-mail: francis.zwiers@ec.gc.ca

cesses. The filtering process that is applied to the observations typically results in an observed vector $\mathbf{y}$ with only 10–20 elements. Each column in matrix $\mathbf{X}$ contains the same number of elements as $\mathbf{y}$, and a separate column is used for each of the external forcing factors that is taken into account in the analysis. Vector $\boldsymbol{\beta}$ contains one element for each signal taken into account. The signal pattern estimates contained in the columns of $\mathbf{X}$ are typically obtained by using a coupled climate system model (see, e.g., Flato et al. 2000) to simulate the response to known forcing changes in the twentieth century. Vector $\boldsymbol{\varepsilon}$ is assumed to be a realization of a Gaussian random vector consisting of correlated elements.

The forcings that are most often accounted for in optimal fingerprinting studies result from anthropogenic changes in greenhouse gas concentrations and sulphate aerosol loadings. Some studies also account for anthropogenic changes in the distribution of ozone, and variability in natural external forcing factors such as changes in the solar output and in the stratospheric loading of aerosols from explosive volcanic eruptions. Because climate models simulate natural internal variability as well as the response to specified changes in external forcing, the response to forcing is typically estimated by averaging across an ensemble of simulations of the twentieth century. These signal estimates, which must be linearly independent, are inserted into the columns of matrix $\mathbf{X}$. The vector of parameters $\boldsymbol{\beta}$ accounts for the possibility that the amplitudes of the climate model responses to the specified external forcings may not be correct by allowing us to scale the signal patterns to best match the pattern of change that is contained in the observations. A scaling factor different from unity might also indicate that the observations have been affected by one or more additional forcing factors not included in the model that may have similar patterns of response.

Detection and attribution questions are assessed through a combination of deductive reasoning (to determine whether there is evidence that other mechanisms of change not included in the climate model could plausibly explain the observed change) and by testing specific hypotheses on $\boldsymbol{\beta}$. Detailed accounts of the optimal fingerprinting method are given by Hasselmann (1979, 1997), Allen and Tett (1999), and Allen et al. (2004). See also the authoritative review of Mitchell et al. (2001).

The detection of a postulated climate change signal occurs when its amplitude in observations is shown to be significantly different from zero. This is handled in the standard optimal fingerprinting approach by testing the null hypothesis

$$H_D: \boldsymbol{\beta} = \mathbf{0},$$

where $\mathbf{0}$ is a vector of zeros. Rejection of H_D by a one-sided test of significance leads to detection at a specified level of significance.

Attribution refers to the process of establishing a cause and effect relationship between the observed change and the hypothesized forcing agents. The requirements for making an attribution claim (e.g., see Mitchell et al. 2001) are detection, elimination of other plausible causes, and evidence that the observed change is consistent with the estimated response to external forcing; that is, $\boldsymbol{\beta} = \mathbf{1}$ where $\mathbf{1}$ is a vector of units. Classical “optimal fingerprinting” analysis uses a hypothesis test called the *attribution consistency test* (Hasselmann 1997; Allen and Tett 1999) to test the null hypothesis

$$H_A: \boldsymbol{\beta} = \mathbf{1}.$$

Formally, consistency between the observed and climate-model-simulated response to forcing can be claimed when H_A cannot be rejected. While a failure to reject H_A does not constitute direct evidence for the hypothesis, consistency does provide support for attribution, particularly when the assessment of H_A involves the simultaneous testing of several externally forced signals.

The Bayesian approach to assessing evidence (e.g., Hasselmann 1998; Leroy 1998; Berliner et al. 2000) provides an alternative to standard (frequentist) hypothesis testing. The Bayesian approach as presented by Berliner et al. (2000) consists of three main steps:

- (i) “Optimal” (generalized linear) regression to estimate the vector of amplitudes $\boldsymbol{\beta}$ that produce the best fit between the climate model estimates of the response to external forcing and the observed data.
- (ii) Identification of the posterior distribution of the amplitudes. This is achieved by combining prior knowledge about the amplitudes that has been expressed as a prior probability distribution with the optimal estimates obtained in (i).
- (iii) A detection and attribution assessment that is done by examining the posterior probabilities for appropriately defined detection and attribution criteria.

The standard optimal fingerprinting analysis, in contrast, consists of step (i) followed by a step (iii') in which hypotheses H_D and H_A are tested. Differences in the clarity of the inferences that can be made with these two approaches will become apparent in the example that will be described below.

We will see that the choice of inference technique does not affect the conclusion that there is strong evidence that the warming observed during the latter half of the twentieth century is unlikely to have been entirely due to natural internal variability (Houghton et al. 2001). Nevertheless, the Bayesian approach has the advantages that it allows incorporation of prior knowledge into the analysis and, more importantly, that it provides direct probabilistic assessments of the evidence for detection and attribution.

The particular Bayesian approach used in this paper is a modified version of Berliner et al.'s (2000) method. Our modifications include changes in the details of how the observations are filtered, the use of a different climate model to estimate the signals of interest, the use of a different approach for specifying the parameters of the prior distribution on the signal amplitudes, and a somewhat different approach to the assessment of the results. We apply the method to the detection of an anthropogenic signal in observed surface air and sea temperatures, and through this application we highlight issues associated with its implementation. The remainder of this paper is organized as follows. We describe the observations and the Bayesian method with which we study the climate change problem in section 2. Our findings are presented in section 3. Discussion and concluding remarks are given in section 4.

2. Data and methods

a. Data

Our study is set up in a manner similar to Zwiers and Zhang (2003). The observational dataset used is the HadCRUTv dataset (Jones et al. 2001), which is based on monthly values of surface air and sea temperature anomalies relative to 1961–90 and is presented on a 5° latitude by 5° longitude grid.

We use only one externally forced climate signal in this study—that being the combined response to historical changes in greenhouse gas concentrations and the direct effect of sulphate aerosol loading. Following standard convention, we will refer to this signal as the “GS” signal. Our estimates of the GS signal, and of natural climate variability, are constructed from two versions of a coupled global climate model. The models used are the Canadian Center for Climate Modelling and Analysis (CCCma) CGCM1 (Flato et al. 2000; Boer et al. 2000) and CGCM2 (Flato and Boer 2001). Details of the models and an extensive archive of model output are available from the CCCma Web site (<http://www.cccma.ec.gc.ca>). Ensembles of three GS simulations are available from each model, for a total combined ensemble of six GS simulations. The GS signal is

estimated by averaging across this six-member ensemble.

b. Analysis strategy

The data and climate model output that we use are preprocessed to restrict the analysis to large spatial scales and to filter as much natural internal variability as possible from the GS signal estimate. Thus our analysis is conducted on decadal averages of observed annual-mean temperature anomalies relative to the climatology for the twentieth century, calculated from non-missing years in the observed data. These decadal anomalies are aggregated into 30° latitude by 40° longitude grid boxes for both observations and climate model output.

Separate analyses are performed for each of six overlapping five-decade periods (1900–49, 1910–59, 1920–69, 1930–79, 1940–89, and 1950–99). To avoid systematic bias, missing data are not filled in. Model output is flagged as missing whenever the corresponding observations are missing so that the length of the $\mathbf{y}$ vector is the same as the row dimension of the $\mathbf{X}$ matrix. Observed annual means are treated as missing even if one month within the year is missing. Decadal means are treated as missing if fewer than 6 of the 10 years are present. In the absence of any missing data, the observational vector $\mathbf{y}$ would have length $5 \times 6 \times 9 = 270$ for any one five-decade period, where $6 \times 9 = 54$ is the number of 30° latitude by 40° longitude grid boxes that cover the globe. Missing data reduces the dimension length to a number ranging from 220 for the period 1900–49 to 253 for the period 1950–1999.

c. Introduction to Bayesian aspects

Before we proceed with our analysis, readers not familiar with Bayesian statistics may need some clarifications on steps (ii) and (iii). We should begin by noting that the Bayesian aspect of our analysis begins with the reduced information that is obtained from the optimal regression analysis. That is, we use Bayesian tools to make inferences about β beginning with the estimate $\hat{\beta}$ that was obtained by means of optimal regression. A more complete analysis would include the regression step within the Bayesian framework, and perhaps even the signal matrix $\mathbf{X}$. However, this would significantly complicate the application of the Bayesian paradigm (see, e.g., the discussion in Berliner et al. 2000).

Bayesian inferences are based on the *posterior distribution*, which in our case is the conditional distribution of β given the estimated value $\hat{\beta}$ from the “optimal” regression in step (i). To see how such a distribution is obtained, let $\pi(\beta)$ be the density function of the *prior*

distribution of β that expresses our prior knowledge about the true value of the amplitude vector before observing $\mathbf{y}$ and estimating $\hat{\beta}$. After observing $\mathbf{y}$ and estimating $\hat{\beta}$, our knowledge of β is updated by calculating the posterior distribution of β conditional upon $\hat{\beta}$, which is denoted $\pi(\beta|\hat{\beta})$. By Bayes theorem, the posterior distribution is given by

$$\pi(\beta|\hat{\beta}) = \frac{f(\hat{\beta}|\beta)\pi(\beta)}{f(\hat{\beta})} = \frac{f(\hat{\beta}|\beta)\pi(\beta)}{\int f(\hat{\beta}|\beta)\pi(\beta) d\beta}, \quad (2)$$

where $f(\hat{\beta}|\beta)$ is the conditional density of $\hat{\beta}$ given the true amplitude β .

The posterior distribution contains all of the available information about β and thus any statistical inferences concerning β should be drawn from this distribution. Similar to Berliner et al. (2000), our detection criterion requires that β have a large posterior probability of lying outside a predefined neighborhood of $\mathbf{0}$. Because the external signals can be defined so that positive values of β would be expected when the signals are present in the observations, our “neighborhood” of $\mathbf{0}$ includes all negative amplitudes. Thus in our one signal case, the detection region simply reduces to an open-ended interval on the positive β axis. Similarly, in the multisignal case, the resulting multidimensional detection region is a positive quadrant of the β plane. Our criterion for assessing whether there is evidence to support attribution requires that β have large posterior probability of lying in a predefined neighborhood of $\mathbf{1}$. This neighborhood is referred to as the attribution region. In both cases, the posterior probabilities are calculated by integrating the posterior distribution [Eq. (2)] over these predefined regions.

The following subsections provide the details of our analysis procedure.

d. Optimal detection procedure

To perform the “optimal” regression, we implemented the method outlined by Allen and Tett (1999). Thus the estimated scaling factor obtained from the generalized linear regression is

$$\hat{\beta} = (\mathbf{X}^T \hat{\mathbf{C}}_N^{-1} \mathbf{X})^{-1} \mathbf{X}^T \hat{\mathbf{C}}_N^{-1} \mathbf{y} = \mathbf{F}^T \mathbf{y}, \quad (3)$$

where $\hat{\mathbf{C}}_N$ is an estimate of the covariance matrix of ϵ . To ensure the variance of $\hat{\beta}$ is not negatively biased because of possible overfitting (Hegerl et al. 1997), the variance of $\hat{\beta}$ is estimated by

$$\hat{\sigma}^2 = \mathbf{F}^T \tilde{\mathbf{C}}_N \mathbf{F}, \quad (4)$$

where $\tilde{\mathbf{C}}_N$ is a second, statistically independent estimate of the covariance matrix of ϵ .

The covariance matrices must be estimated from control simulation output from climate models because the available instrumental record does not contain sufficient information to do this reliably. Thus as in Zwiers and Zhang (2003), our covariance matrices are estimated from 1600 years of control simulation, composed of 600 years from CGCM1 and 1000 years from CGCM2. The 1600 years of control output is partitioned into two 800-yr subsets consisting of the first 300 years from CGCM1 and the first 500 years from CGCM2 in one case, and the last 300-yr and 500-yr segments in the other.

Covariance matrix $\tilde{\mathbf{C}}_N$ is estimated from the 72 overlapping five-decade control run chunks contained in the first control subset. Note that the chunks do not span the boundary between the 300-yr CGCM1 and 500-yr CGCM2 segments. Each 50-yr chunk is processed in the same manner as the observations to obtain a sequence of five decadal means that are then masked as the observations and aggregated into $30^\circ \times 40^\circ$ grid boxes. The covariance matrix is then estimated from these processed sequences of control run output. Note that this estimation process is repeated for each of the five-decade analysis periods because the details of the masking, which reflect observational coverage, change from one period to the next. Matrix $\tilde{\mathbf{C}}_N$ is estimated similarly from the second subset of control run output.

A difficulty encountered in all detection studies using the optimal fingerprinting technique is that, even with the use of long control simulations, the covariance matrices $\hat{\mathbf{C}}_N$ and $\tilde{\mathbf{C}}_N$ are not invertible. This occurs because the dimension of $\mathbf{y}$ remains large relative to the length of the available control simulations even after preprocessing into decadal and $30^\circ \times 40^\circ$ box means. The exact dimension of $\mathbf{y}$ varies with the particular five-decade period being analyzed, but it is always larger than the number of overlapping five-decade sequences (72) that are available in each of the control simulation subsets. Consequently, the regression is performed in the space spanned by the k leading EOFs of $\tilde{\mathbf{C}}_N$. A residual consistency test (Allen and Tett 1999) is performed to determine the appropriate truncation level k . A simple schematic that outlines the optimal fingerprinting approach can be found in Weaver and Zwiers (2000).

Once the “optimal” regression has been performed and $\hat{\beta}$ computed, we can proceed with steps (ii) and (iii). This requires the selection of a prior distribution for β and the calculation of the resulting posterior distribution.

e. The prior distribution

An important element of the analysis is the choice of prior distribution. This distribution embodies our prior knowledge about the magnitude of the response to the external forcing of the climate system that is being investigated. This knowledge can be either subjective or objective in nature. An example of subjective knowledge that could be incorporated into the prior is a judgment about whether the climate has the potential to be affected by anthropogenic forcing. Such a judgment, which might be based on theoretical considerations, should be independent of observations gathered during the period that is represented by $\mathbf{y}$. An example of how objective knowledge could be incorporated into the prior is that the width of the prior distribution could be chosen to reflect uncertainty in the amplitude of the response of climate models to known external forcings. Such an estimate of uncertainty might be obtained by performing “cross model” studies in which we attempt to detect a signal estimated with one climate model in the output of several other models (e.g., Hegerl et al. 2000). The width of the prior might also reflect uncertainty in the pattern of response, in the sense that a biased (but not totally orthogonal) response pattern may require a scaling factor different from unity to make the best fit with observations. In addition, the width of the prior could be augmented to reflect uncertainty in the forcing itself. Some considerations based on energy balance model (EBM) experiments that informed our specific choices of prior distributions are discussed below.

More generally, uncertainty can be modeled explicitly in the Bayesian framework by building hierarchical models in which uncertainty in the forcing and signal response patterns is explicitly taken into account. Alternatively, one could perform a *robust Bayesian analysis* in which the prior distribution is varied systematically over a range of possibilities that reflects the uncertainty of prior knowledge more completely than a single distribution (Berliner et al. 2000). However, we considered such approaches to be beyond the scope of this paper.

To make the above more concrete, we use an EBM to explore how climate model and forcing uncertainty might translate into uncertainty on β . To determine the uncertainty range, we have employed a suite of EBM simulations that reflect an estimate of climate model and forcing uncertainty. EBM simulations reproduce many of the large-scale temperature responses of general circulation models (e.g., Crowley et al. 1991). EBMs do not simulate internal climate variability (unless forced with noise) and thus they produce climate

change signal estimates that are not affected by internal climate variability. Here we employ a linear two-dimensional (i.e., realistic land–sea distribution) seasonal model that has been used in comparisons to paleoclimate reconstructions (Crowley et al. 2003; Hegerl et al. 2003).

Parameters such as the climate sensitivity and ocean heat uptake must be set explicitly in EBMs. To reflect the range of climate model uncertainty and the most important forcing uncertainty, three parameters in the EBM simulations are varied:

- The *climate sensitivity*, which represents the equilibrium response to CO_2 doubling, was varied in 0.5°C increments from 1.5° to 4.5°C . This reflects the range presently covered by most coupled climate models as reported by Houghton et al. (2001).
- The *diffusivity*, which determines the ocean heat uptake in the EBM, was set to values between 0.95 and $2.5 \text{ cm}^2 \text{ s}^{-1}$, yielding a conservative range embracing an observational estimate from tracer studies of $1.7 \pm 0.2 \text{ cm}^2 \text{ s}^{-1}$ (Li et al. 1984).
- Since EBMs cannot distinguish between direct and indirect tropospheric aerosol forcing, we also employed *direct aerosol forcing* starting in 1850 that reaches values between 0.5 and 4.0 W m^{-2} by 2000 in increments of 0.5 W m^{-2} for the latitude strip of 30° – 90°N . The forcing in other latitude bands was tapered to lower values in successive equal-area strips so as to yield a total global aerosol forcing of approximately 0.3 to 2.5 W m^{-2} . This approximately covers the Houghton et al. (2001) range for the total combined direct and indirect aerosol forcing.

The first two parameters are important determinants of the large-scale response of climate models to external forcing (cf. Forest et al. 2000), while the total value of aerosol forcing is highly uncertain (Houghton et al. 2001) and was therefore also varied. Note that all simulations are forced with changes in greenhouse gases, direct aerosols, and solar and volcanic forcing (see Crowley et al. 2003 for details). For practical reasons, we use the simulated Northern Hemispheric mid- to high-latitude (30° – 90°N) mean annual-mean temperature to assess the variations in response that result from the prescribed variations in parameter settings and aerosol forcing.

We chose one of the simulations, whose parameters characterize the CCCma models (sensitivity 3.5°C , 30° – 90°N aerosol forcing of -2 W m^{-2} , and midrange ocean heat uptake) and determined by simple linear regression the scaling that brings this simulation into closest agreement to each member of the suite of simulations. The scaling factors range between 0.15 and 1.84 for the

FIG. 1. Prior distribution (a) $\pi_1(\beta)$ and (b) $\pi_2(\beta)$.

1950–99 period when all three parameters are varied. The range becomes 0.57 to 1.15 if climate sensitivity is the only free parameter (with diffusivity fixed to $1.9 \text{ cm}^2 \text{ s}^{-1}$ and the $30^\circ\text{--}90^\circ\text{N}$ average aerosol forcing set to -2.0 W m^{-2}). The ranges are wider if the entire twentieth century is used to determine the scaling factors (0.50 to 1.20 if only the sensitivity is varied). The asymmetry of the scaling ranges, with a greater part of the range below unit scaling than above, reflects the fact that CCCma's models have a moderately large climate sensitivity.

Using these results as guidance and acknowledging that not all sources of uncertainty have been taken into account, it would be appropriate for a simple Gaussian prior that gives the greatest prior likelihood to $\beta = 1$ to have variance $\tau^2 \approx 0.25$. This yields a plus or minus two standard deviation prior uncertainty range for β of 0.0 to 2.0 that is inclusive of the ranges derived above, and provides a central plus or minus one standard deviation range for β of 0.5 to 1.5.

To examine the sensitivity of the results to a range of plausible choices of prior, we construct three alternative priors, each of which is noncommittal on the climate change question.

The prior distribution we use most extensively (Fig. 1a) is similar to that used by Berliner et al. (2000). It is a mixture of two Gaussian distributions given by

$$\pi_1(\beta) = p\phi(0, \tau_1^2) + (1 - p)\phi(1, \tau_2^2), \quad (5)$$

where $p = 1/2$ and the component distribution variances τ_1^2 and τ_2^2 are set to 0.25. A value of p other than $1/2$ could be used to reflect the analyst's degree of belief that the signal $\mathbf{X}$ is either present (p near 0) or absent (p near 1) in the observations $\mathbf{y}$ (Berliner et al. 2000; Min et al. 2004). The choice of prior with $p = 1/2$ gives equal weight to the possibilities that the GS signal is either absent or present in the data, and conditional on the latter, reflects our estimates of the effect of climate model and forcing uncertainty on β . The width of the prior conditional on the former possibility is set to the same value for convenience. The resulting mixture distribution gives the greatest amount of weight to β values in the interval $[0, 1]$, but does not give preference to any particular subset of values in this range.

To ensure that our results are not sensitive to the choice of prior, we consider two additional distributions that express prior knowledge and uncertainty, or the lack thereof, in different ways. The first, π_0 (not shown), is identical to π_1 except that the variance of its two components is only 0.1 instead of 0.25. This prior is still noncommittal about whether the GS signal is present in the observations, but it expresses greater confidence that the climate model is able to respond correctly to the specified GS forcing. The other prior distribution we consider is the noninformative prior, $\pi_2(\beta) \propto 1$, an improper distribution that gives equal weight to all values of β between $-\infty$ and ∞ (Fig. 1b). This distribution expresses a complete absence of prior knowledge about β by not favoring any value of β over another.

f. The posterior distribution

As indicated in section 2b, the posterior distribution is derived via Bayes' theorem [Eq. (2)] from the distribution of $\hat{\boldsymbol{\beta}}$ conditional on β , which is also known as the *likelihood function*, and the prior distribution.

Assuming that ε has a multivariate Gaussian distribution, the distribution of $\hat{\boldsymbol{\beta}}$ conditional on β is approximately $f(\hat{\boldsymbol{\beta}}|\beta) = \phi(\hat{\boldsymbol{\beta}}, \hat{\sigma}^2)$, where $\hat{\sigma}^2$ [Eq. (4)] is the estimated variance of $\hat{\boldsymbol{\beta}}$ and $\phi(\cdot, \cdot)$ is a Gaussian probability density function. The distribution is approximate because $\hat{\sigma}^2$ is estimated. However, this estimate is known to be consistent, and thus the quality of this approximation will improve as the size of the samples available for calculating covariance estimates $\hat{\mathbf{C}}_N$ and $\hat{\mathbf{C}}_N$ increases.

Using this approximation in (2) together with prior distribution π_1 given by (5) results in a posterior distribution that is given by

$$\pi_1(\boldsymbol{\beta}|\hat{\boldsymbol{\beta}}) = \bar{p}\phi(\tilde{\boldsymbol{\mu}}_1, \tilde{\tau}_1^2) + (1 - \bar{p})\phi(\tilde{\boldsymbol{\mu}}_2, \tilde{\tau}_2^2). \quad (6)$$

This is again a mixture of two Gaussian distributions (see Berger 1985, 127–128, 206; Berliner et al. 2000), in this case with

$$\begin{aligned}\tilde{\mu}_1 &= \frac{\tau_1^2 \hat{\beta}}{\tau_1^2 + \hat{\sigma}^2}, \quad \tilde{\tau}_1^2 = \frac{\tau_1^2 \hat{\sigma}^2}{\tau_1^2 + \hat{\sigma}^2}, \\ \tilde{\mu}_2 &= \frac{\tau_2^2 \hat{\beta} + \hat{\sigma}^2}{\tau_2^2 + \hat{\sigma}^2}, \quad \tilde{\tau}_2^2 = \frac{\tau_2^2 \hat{\sigma}^2}{\tau_2^2 + \hat{\sigma}^2}, \quad \text{and} \\ \tilde{p} &= \left(1 + \frac{1-p}{p} \sqrt{\frac{\tau_1^2 + \hat{\sigma}^2}{\tau_2^2 + \hat{\sigma}^2}} \right. \\ &\quad \left. \exp \left\{ -0.5 \left[\frac{(\hat{\beta} - 1)^2}{\tau_2^2 + \hat{\sigma}^2} - \frac{\hat{\beta}^2}{\tau_1^2 + \hat{\sigma}^2} \right] \right\} \right)^{-1}.\end{aligned}$$

The posterior that results from prior π_0 , which is a mixture of Gaussian components with variances $\tau_1 = \tau_2 = 0.1$, is identical, except that different values of τ_1 and τ_2 are used in the expressions above. Finally, the posterior distribution that corresponds to the improper (uninformative) prior π_2 can be shown to be

$$\pi_2(\beta|\hat{\beta}) = \phi(\hat{\beta}, \hat{\sigma}^2), \quad (7)$$

where $\hat{\beta}$ and $\hat{\sigma}^2$ are given by Eqs. (3) and (4). Note that in this case the posterior is simply that likelihood function that is often used to make frequentist inferences about β .

g. Detection and attribution assessment

The posterior distributions are used to make detection and attribution assessments by calculating the posterior probabilities that β lies in predefined detection ($\mathcal{D}$) and attribution ($\mathcal{A}$) regions, respectively.

We have chosen to define the detection region as $\mathcal{D} = [0.1, \infty)$. We leave open the question of how we should interpret a scaling factor, of say, $\beta = 0.1$ (the climate model response is 10 times the observed response) or $\beta = 10$ (the model response is only 1/10 the observed response), possibilities that are both contained within $\mathcal{D}$. However, we acknowledge that such an outcome would raise important questions about the climate model and/or the completeness and accuracy of the specification of the external forcing to the climate model. Thus others may want to define $\mathcal{D}$ differently. In any case, given a defensible definition of $\mathcal{D}$, detection is claimed if the posterior probability that $\beta \in \mathcal{D}$, which we denote by $P_{\mathcal{D}|\hat{\beta}} = \Pr(\beta \in \mathcal{D}|\hat{\beta})$, is large.

The exact interpretation of a “large” posterior probability is a matter of judgment, but we can be guided by the use of Bayes’ factors (e.g., see Kass and Raftery 1995, and references therein). The Bayes factor B is the

scaling factor that relates the prior odds of detection to the posterior odds of detection through the equation

$$\frac{P_{\mathcal{D}|\hat{\beta}}}{1 - P_{\mathcal{D}|\hat{\beta}}} = B \frac{\Pr(\beta \in \mathcal{D})}{1 - \Pr(\beta \in \mathcal{D})}. \quad (8)$$

That is, the Bayes factor is the ratio of the posterior odds to the prior odds. Jeffreys (1935) first proposed the Bayes factor as a way of comparing scientific theories. Kass and Raftery (1995), following Jeffreys (1961), suggest that $3 \leq \beta \leq 20$ indicates “positive” evidence for the hypothesis that is being evaluated, that $20 \leq \beta \leq 150$ indicates “strong” evidence, and that $B > 150$ represents “very strong” evidence. The prior odds of detection using definition $\mathcal{D}$ are easily shown to be approximately 9:4 under prior π_1 . Thus a posterior probability of detection $P_{\mathcal{D}|\hat{\beta}} > 0.978$ is required to obtain a Bayes factor that indicates strong evidence for detection, while $P_{\mathcal{D}|\hat{\beta}} > 0.997$ is required to conclude that there is very strong evidence.

As noted previously, statistical evidence that the climate model has responded correctly to historical external forcing is an important aspect of the overall attribution assessment. Thus it is reasonable to define the attribution region $\mathcal{A}$ as a subset of the detection region that contains the point $\beta = 1$. A large posterior probability that $\beta \in \mathcal{A}$, which we denote by $P_{\mathcal{A}|\hat{\beta}} = \Pr(\beta \in \mathcal{A}|\hat{\beta})$, provides evidence in support of attribution.

For purposes of illustration, we consider an attribution region defined as $\mathcal{A} = (0.8, 1.2)$. This particular choice of $\mathcal{A}$ implies that the criteria for making an attribution claim should include a high posterior likelihood that the model response to forcing is within 20% of observed. Further discussion of the choice of $\mathcal{A}$ is given in section 4.

As with detection, we may use the Bayes factors to guide the interpretation of $P_{\mathcal{A}|\hat{\beta}}$. The prior odds of attribution given our particular definition of $\mathcal{A}$ are approximately 0.22:1 under prior π_1 . Under these circumstances a posterior probability $P_{\mathcal{A}|\hat{\beta}} > 0.97$ would represent very strong evidence for consistency between model and observed response to forcing, while $P_{\mathcal{A}|\hat{\beta}} > 0.81$ would still represent strong evidence.

3. Results

The residual consistency test statistic reported in Zwiers and Zhang (2003) suggests that our analysis should be conducted with approximately 15–25 EOFs retained. We therefore use truncations in that range. Figure 2 displays the posterior distribution π_1 obtained for time periods 1900–49 and 1950–99 with 15, 20, and 25 EOFs retained. As can be seen, the posterior is not very sensitive to the number of EOFs retained over this

FIG. 2. Posterior distribution derived using prior $\pi_1(\beta)$ for (a) 1900-49 and (b) 1950-99 when 15, 20, and 25 EOFs are retained.

truncation range. The posterior is somewhat wider for 1900-49 than that for 1950-99 for truncations in the 15-25 EOF range that are shown. This occurs because there is more missing data in the earlier period and because the signal is weaker. The width of the posterior [Eq. (6)] depends upon $\hat{\sigma}^2$, which in turn depends inversely upon the generalized inner product of $\hat{\beta}$ with itself [Eq. (4)]. The latter increases when the signal is stronger and when the signal vector is longer, and thus the posterior narrows under those circumstances.

Most of the posterior probability in the 1950-99 analysis is allocated to β values between 0.4 and 0.8 when 20 EOFs are used, suggesting that the mean model response to the historical forcing changes is somewhat stronger than indicated by the observed temperature changes. This may, in part, be the result of a downward bias in the estimate of β that occurs because some noise from internal climate variability remains in the estimated GS signal even after averaging the six available forced climate simulations. However, fitting the regression model by means of the total least squares algorithm (Allen and Stott 2003) to account for noise in

the estimated signal pattern increases $\hat{\beta}$ only slightly. The somewhat too strong model response may also be due to missing, and probably net negative, forcings (Mitchell et al. 2001) such as the indirect effects of sulphate aerosols and the combined effect of solar and volcanic forcing.

Figure 3 shows the posterior probabilities that β lies in the detection and attribution regions that are calculated from $\pi_1(\beta|\hat{\beta})$ for each time period. The probability of detection $P_{\mathcal{D}|\hat{\beta}}$ is greater than 0.998 for both 1900-49 and 1950-99 when 20 EOFs are retained, indicating that it is highly unlikely (very strong evidence) that the observed temperature changes during these periods are entirely the result of natural internal climate variability. Similar results are obtained for 1910-59 and 1940-89.

FIG. 3. (a) Detection probability $P_{\mathcal{D}|\hat{\beta}}$ for the six time periods 1900-49, 1910-59, ..., 1950-99 as a function of the number of EOFs retained where the detection region is $\mathcal{D} = [0.1, \infty)$. (b) As above, except the posterior attribution probability $P_{\mathcal{A}|\hat{\beta}}$ is displayed. This probability takes only anthropogenic forcing into account and is obtained using the attribution region that is given by $\mathcal{A} = (0.8, 1.2)$. This definition of $\mathcal{A}$ imposes a very stringent attribution criterion. A broader, but still reasonable criterion, such as $\mathcal{A} = (0.5, 1.5)$, results in greater values of $P_{\mathcal{A}|\hat{\beta}}$ in the range 0.7 to 0.9 for the 1950-99 period. See text for a detailed discussion of the interpretation of $P_{\mathcal{A}|\hat{\beta}}$ and the selection of $\mathcal{A}$.

For the other two time periods, the probability of detection ranges between 0.4 and 0.8. Work by others (e.g., Stott et al. 2000, 2001; Hegerl et al. 2003) suggests that the nonmonotonic nature of the warming during the twentieth century that underlies these results can only be adequately explained by considering both the anthropogenic forcing factors investigated here and natural external forcing factors such as solar and volcanic forcing.

The probability of attribution $P_{\mathcal{A}|\hat{\beta}}$ that is computed from $\pi_1(\beta|\hat{\beta})$ with 20 EOFs retained is less than 0.5 for all time periods and ranges between 0.1 and 0.3 for truncations with 15 to 25 EOFs for the 1950–99 period. These probabilities are not large, and thus attribution is not strongly supported when the posterior is interpreted using our specific choice of attribution region $\mathcal{A}$. The choice of $\mathcal{A}$ and the implications for attribution are further discussed in section 4.

We repeated our analysis with the noninformative prior distribution π_2 . The resulting posterior distribution $\pi_2(\beta|\hat{\beta})$ (not shown) is somewhat wider than $\pi_1(\beta|\hat{\beta})$. Correspondingly, the probabilities of detection and attribution calculated from $\pi_2(\beta|\hat{\beta})$ are slightly smaller than those displayed in Fig. 3. Specifically, the probability of detection is about 0.0001 less than that reported above and the probability of attribution is about 0.01 less. Results obtained using prior distribution π_0 are similar to those obtained with π_1 . Overall, results appear to be insensitive to variations in prior distributions that vary the way they express uncertainty about the climate model response to GS forcing but are similarly noncommittal about whether there has been an anthropogenic influence on climate in the twentieth century. Less moderate prior positions, where greater weight is put on the possibility that $\beta = 1$, or depending on perspective, $\beta = 0$, may result in a somewhat more narrowly defined posterior distribution.

4. Conclusions and discussion

We have described a Bayesian approach to detection and attribution that was first proposed by Berliner et al. (2000). The method is based on the standard optimal fingerprinting approach of Hasselmann (1979, 1997). In contrast to Berliner et al., the details of our implementation of the optimal fingerprinting component of the method are similar to those used in recent optimal fingerprinting studies (e.g., Allen and Tett 1999; Allen et al. 2004). In addition, we used a series of sensitivity experiments with an energy balance model to derive “ballpark” estimates of the prior uncertainty of the scaling factor on the climate-model-simulated response

to external forcing. These estimates helped inform our choices of parameters for the prior distributions used in our study.

As an example, we assessed whether the combined effect of greenhouse gases and sulfate aerosols as simulated by the CCCma climate models (the CCCma “GS” signal) is detectable in the observed global surface temperature record in the twentieth century. Furthermore, we also assessed whether the posterior distribution for the scaling factor on the GS signal provides evidence to support attribution. The assessment was made with three prior distributions that are noncommittal on the question of whether there has been an anthropogenic influence on the climate of the twentieth century. We used specific assessment criteria that were defined in terms of predefined sets of scaling factors ($\mathcal{D}$ and $\mathcal{A}$) and predefined posterior probability thresholds. The latter were determined by using Bayes factors to associate posterior probabilities with specific levels of “strength of evidence.”

Consistent with earlier studies, our Bayesian analysis indicates that the anthropogenic forcing signal is detectable in the early and late halves of the twentieth century. Note, however, that the early-century anthropogenic climate change signal may be somewhat confounded with the response to natural forcing changes during, and prior to, this period. Tett et al. (1999) and others have suggested that solar forcing changes may also have contributed to the early warming. However, all studies performed to date suggest that the anthropogenic forcing and response dominate during the latter half of the twentieth century (Mitchell et al. 2001; International Ad Hoc Detection and Attribution Group 2005).

Our results are insensitive to our choice of prior distribution because our selection of priors does not include distributions that express strong, a priori positions on anthropogenic climate change. The resulting posterior distributions have considerable spread, and thus we were not able to contribute strong evidence to an attribution assessment with our specific definition of the attribution region $\mathcal{A}$. For 1950–99, when surface temperature data coverage is the most complete, most of the posterior probability is allocated to β values less than unity, suggesting that the CCCma climate models may have oversimulated the combined response to greenhouse gas and direct sulphate aerosol forcing. This may, in part, be the result of missing forcings such as the indirect effects of sulphate aerosols, the effects of other types of aerosols such as mineral dust and carbonaceous aerosols, the effects of the changing abundance of tropospheric and stratospheric ozone, and the

effects of variations in natural external forcing of the climate system.

Our findings, which are preliminary in nature, do not in any way negate the Houghton et al. (2001) conclusion that "... most of the warming observed over the last 50 years is attributable to human activities." Houghton et al. used a very broad range of evidence to make their attribution assessment, including a number of optimal detection studies using several different coupled climate models. What our findings do tell us is that the limited observational record, together with missing forcings and possible bias in the CCCma model, preclude the possibility of constructing a narrow posterior distribution that is centered directly on $\beta = 1$. Thus strong evidence supporting attribution does not emerge when using an attribution region that imposes a relatively stringent criterion, where the amplitude of the model response must be within 20% of observed. However, if we use a less stringent criterion with, say, $\mathcal{A} = (0.5, 1.5)$ (the model response must be within 50% of observed), we then find that the posterior probability that $\beta \in \mathcal{A}$ increases to levels of approximately 0.7 to 0.9 (depending upon EOF truncation) for the 1950–1999 period. This corresponds to Bayes factors in the range $3.2 < B < 12.4$, giving an indication of "positive" evidence to support an attribution assessment. Relaxing the attribution criterion might be reasonable if taking into consideration external forcings not included in the CCCma GS simulations that may have a pattern of response similar to the GS response.

The absence of strong evidence to support attribution in the presence of "very strong" evidence for detection requires further consideration. As already suggested, there are several factors that may affect the amount of evidence that the posterior distribution on β can provide to support attribution. One of these is the length of the observational record. The posterior distribution should become narrower as the analysis period is made longer and extended into the future.

Another factor that affects the width of the posterior distribution is the signal strength. The latter is expected to continue to increase and therefore should contribute to a reduction in the width of the posterior as the analysis period shifts into the future. To quantify when strong evidence in support of an attribution assessment might emerge, we reevaluated the posterior distribution $\pi_1(\beta|\hat{\beta})$ using the 1950–99 values of the scaling factor $\hat{\beta}$, but estimates of future values of $\hat{\sigma}^2$. The latter is strongly affected by the signal strength. The steps in our calculation were as follows:

- The future estimates of $\hat{\sigma}^2$ [Eq. (4)] were obtained using the 1950–99 observational mask and the corre-

FIG. 4. Estimated future attribution probability $P_{\mathcal{A}|\hat{\beta}}$, where $\mathcal{A} = (\mu_p - 0.2, \mu_p + 0.2)$ and μ_p indicates the mode of the posterior distribution $\pi_1(\beta|\hat{\beta})$. Details of the calculation are given in the text.

sponding 1950–99 covariance matrices, but using a future five-decade signal vector. The latter was masked in the same way as the 1950–99 observations and calculated relative to the same base period. This was repeated for all five-decade periods beginning with 1960–2009 and ending with 2050–2099.

- The estimated future values of $\hat{\sigma}^2$ and the 1950–99 values for $\hat{\beta}$ were substituted into $\pi_1(\beta|\hat{\beta})$. For each "posterior" that was obtained in this way, we calculated the probability that $\beta \in (\mu_p - 0.2, \mu_p + 0.2)$ where μ_p is the mode of the posterior distribution. This, in effect, is the attribution probability $P_{\mathcal{A}|\hat{\beta}}$ that would be obtained for $\mathcal{A} = (0.8, 1.2)$ when $\hat{\beta} = 1$.

The results of this exercise are displayed in Fig. 4 for truncations with 15, 20, and 25 EOFs. The graph has been extended to the left using the posteriors that are appropriate for those historical periods. By centering the attribution region on the mode of the posterior distribution, we have made the assumption that bias in the estimate of β will not be a problem with signals derived from future climate models. With this assumption, the suggestion is that strong or very strong evidence in support of attribution, based on a requirement that the model response to external forcing should be within 20% of observed, should emerge by 2020.

A third factor that affects the posterior distribution is the accuracy of our estimated covariance matrices. Since these matrices are estimated from climate model control runs they may be biased by model errors, although the use of multiple models (Gillett et al. 2002) can help to reduce that bias. The matrices also remain uncertain because of the limited length of the control simulations.

A fourth factor is whether all of the prior information

that is available for inclusion in the analysis via the prior distribution has, in fact, been used. In our example we purposefully used priors that are noncommittal on whether there has been an anthropogenic influence on climate. Use of a prior distribution centered on $\beta = 1$, rather than on $\beta = 0.5$ as in our noncommittal priors, will result in a somewhat sharper posterior distribution and an increase in the posterior probability that $\beta \in \mathcal{A}$. The multiple lines of evidence for attribution from sources other than surface temperature data that are described in Houghton et al. (2001) and subsequent literature (e.g., International Ad Hoc Detection and Attribution Group 2005) justifies the use of a prior distribution that does not sit on the fence. Recent experience with the detection of externally forced signals in the paleoclimate record (e.g., Hegerl et al. 2003) provides further justification.

Using a Gaussian prior $\pi_3 = \phi(1, 0.25)$ results in attribution probabilities $P_{\mathcal{A}|\hat{\beta}} = 0.11$ to 0.37 for 1950–99 when using between 15 and 25 EOFs and when $\mathcal{A} = (0.8, 1.2)$. This is slightly greater than the result obtained with the noncommittal prior. However, the corresponding Bayes factors, which lie in the range $0.3 < B < 1.3$, still do not indicate support for attribution. When we increase the size of the attribution region to $\mathcal{A} = (0.5, 1.5)$ we find that $P_{\mathcal{A}|\hat{\beta}}$ ranges from 0.81 to 0.91 , which is again somewhat larger than obtained with the noncommittal prior. In this case the corresponding Bayes factors range between 2 and 4.7 , indicating that we now have weaker evidence in support of attribution than obtained with the noncommittal prior.

This is an example of the effect that the choice of prior has on the criteria for the assessment of the evidence in support of attribution (and also detection) when evidence is evaluated by means of Bayes factors. The relatively strong prior position on attribution that is expressed with distribution π_3 has substantially larger prior odds of attribution than the noncommittal prior π_1 (approximately $2.2:1$ as opposed to $0.22:1$). Consequently, the posterior odds of attribution must be correspondingly larger to achieve the same assessment of the strength of evidence. For example, to be identified as strong evidence, the posterior odds would have to be approximately $44:1$ when using prior π_3 while posterior odds of approximately $4.4:1$ would be sufficient when using π_1 .

This illustrates an important benefit of using Bayes factors to assess evidence in support of detection or attribution. By comparing posterior odds ratios with corresponding prior odds ratios, the Bayes factor requires inferences to be heavily dependent on evidence that is obtained from the data. If the analyst has taken a strong prior position on detection or attribution, the

Bayes factor requires the data to substantially improve the prior assessment of the odds of detection or attribution before declaring that there is, in fact, strong supporting evidence for that position.

The ability to incorporate prior knowledge into the analysis, particularly with regards to sources of uncertainty, is an important advantage of the Bayesian approach. This knowledge can be both subjective and objective and can include prior information about whether there has been a climatic response to anthropogenic forcing, about climate model uncertainty, and about the nature of the response to natural external forcing changes in, for example, paleo records, as discussed above. Another important advantage of the Bayesian approach is that the analyst must explicitly define the criteria that are used to assess evidence in support of detection and attribution. Admittedly, the choice of prior and the definitions of $\mathcal{D}$ and $\mathcal{A}$ are subjective, but they transparently reflect the perspective of the analyst making the climate change assessment.

In comparison, the frequentist criteria for the assessment of evidence are much less explicit. The frequentist criterion for detection is clear enough (we seek evidence that the observed change global surface temperature is not consistent with H_D), but exactly what the rejection of H_D means in probabilistic terms is not clear. More problematically, the standard criterion for the assessment of evidence in support of attribution is the *failure* to reject H_A (Berliner et al. 2000). This means that the attribution criterion is relaxed, rather than made more stringent, when the significance level is reduced, say, from 10% to 1% . Also, perversely, attribution becomes more difficult with the standard approach as the observed record increases in length for a given signal strength. The analyst does not have direct control over the definition of the criteria for detection or attribution in either case (this happens, at best, by varying the significance level). In the Bayesian setting we must give, and defend, clear definitions of the criteria used to assess the evidence for detection and attribution. Once having defined the criteria, we can calculate the probabilities that they have been met. These posterior probabilities, in turn, can be used in a decision theoretic framework to make rational adaptation and mitigation decisions that minimize expected costs and losses.

Acknowledgments. We gratefully acknowledge that work conducted by Terry Lee was supported by the Canadian Foundation for Climate and Atmospheric Science through the Canadian CLIVAR Research Network. Work by Min Tsao was supported by the Natural Sciences and Engineering Research Council through a Discovery Grant. Francis Zwiers and Xuebin Zhang

both benefited from support provided by the Climate Change Action Fund. This paper was improved by insightful and helpful comments provided by Greg Flato, Nathan Gillett, Slava Kharin, and Andreas Hense.

REFERENCES

- Allen, M. R., and S. F. B. Tett, 1999: Checking for model consistency in optimal fingerprinting. *Climate Dyn.*, **15**, 419–434.
- , and P. A. Stott, 2003: Estimating signal amplitudes in optimal fingerprinting. Part I: Theory. *Climate Dyn.*, **21**, 477–491.
- , and Coauthors, 2004: Quantifying anthropogenic influence on recent near-surface temperature. *Surv. Geophys.*, in press.
- Berger, J. O., 1985: *Statistical Decision Theory and Bayesian Analysis*. 2d ed. Springer-Verlag, 617 pp.
- Berliner, L. M., R. A. Levine, and D. J. Shea, 2000: Bayesian climate change assessment. *J. Climate*, **13**, 3805–3820.
- Boer, G. J., G. Flato, M. C. Reader, and D. Ramsden, 2000: A transient climate change simulation with greenhouse gas and aerosol forcing: Experimental design and comparison with the instrumental record for the twentieth century. *Climate Dyn.*, **16**, 405–425.
- Crowley, T. J., S. K. Baum, and W. T. Hyde, 1991: Climate model comparison of Gondwanan and Laurentide glaciations. *J. Geophys. Res.*, **96**, 9217–9226.
- , —, K.-Y. Kim, G. C. Hegerl, and W. T. Hyde, 2003: Modeling ocean heat content changes during the last millennium. *Geophys. Res. Lett.*, **30**, 1932, doi:10.1029/2003GL017801.
- Flato, G. M., and G. J. Boer, 2001: Warming asymmetry in climate change simulations. *Geophys. Res. Lett.*, **28**, 195–198.
- , —, W. G. Lee, N. A. McFarlane, D. Ramsden, M. C. Reader, and A. J. Weaver, 2000: The Canadian Centre for Climate Modelling and Analysis Global Coupled Model and its climate. *Climate Dyn.*, **16**, 427–450.
- Forest, C. E., M. R. Allen, P. H. Stone, and A. P. Sokolov, 2000: Constraining uncertainties in climate models using climate change detection techniques. *Geophys. Res. Lett.*, **27**, 569–572.
- Gillett, N. P., F. W. Zwiers, A. J. Weaver, G. C. Hegerl, M. R. Allen, and P. A. Stott, 2002: Detecting anthropogenic influence with a multi-model ensemble. *Geophys. Res. Lett.*, **29**, 1970, doi:10.1029/2002GL015836.
- Hasselmann, K., 1979: On the signal-to-noise problem in atmospheric response studies. *Meteorology of Tropical Oceans*, D. B. Shaw, Ed., Royal Meteorological Society, 251–259.
- , 1997: On multifingerprint detection and attribution of anthropogenic 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, U. Cubasch, J. F. B. Mitchell, E. Roeckner, R. Voss, and J. Waszkewitz, 1997: Multifingerprint detection and attribution analysis of greenhouse gas, greenhouse gas-plus-aerosol and solar forced climate change. *Climate Dyn.*, **13**, 613–634.
- , P. Stott, M. Allen, J. F. B. Mitchell, S. F. B. Tett, and U. Cubasch, 2000: Detection and attribution of climate change: Sensitivity of results to climate model differences. *Climate Dyn.*, **16**, 737–754.
- , T. J. Crowley, S. K. Baum, K.-Y. Kim, and W. T. Hyde, 2003: Detection of volcanic, solar and greenhouse gas signals in paleo-reconstructions of Northern Hemispheric temperature. *Geophys. Res. Lett.*, **30**, 1242, doi:10.1029/2002GL016635.
- Houghton, J. T., Y. Ding, D. J. Griggs, M. Noguer, P. J. van der Linden, X. Dai, K. Maskell, and C. A. Johnson, Eds., 2001: *Climate Change 2001: The Scientific Basis*. Cambridge University Press, 881 pp.
- International Ad Hoc Detection and Attribution Group, 2005: Detecting and attributing external influences on the climate system: A review of recent advances. *J. Climate*, **18**, 1291–1314.
- Jeffreys, H., 1935: Some test of significance, treated by the theory of probability. *Proc. Cambridge Philos. Soc.*, **31**, 203–222.
- , 1961: *Theory of Probability*. 3d ed. Clarendon Press, 447 pp.
- Jones, P. D., T. J. Osborn, K. R. Briffa, C. K. Folland, E. B. Horton, L. V. Alexander, D. E. Parker, and N. A. Rayner, 2001: Adjusting for sampling density in grid box land and ocean surface temperature time series. *J. Geophys. Res.*, **106**, 3371–3380.
- Kass, R. E., and A. E. Raftery, 1995: Bayes factors. *J. Amer. Stat. Assoc.*, **90**, 773–795.
- Leroy, S. S., 1998: Detecting climate signals: Some Bayesian aspects. *J. Climate*, **11**, 640–651.
- Li, Y.-H., T.-H. Peng, W. S. Broecker, and H. G. Oestlund, 1984: The average vertical mixing coefficient for the oceanic thermocline. *Tellus*, **36B**, 212–217.
- Min, S.-K., A. Hense, H. Paeth, and W.-T. Kwon, 2004: A Bayesian decision method for climate change signal analysis. *Meteor. Z.*, **13**, 421–436.
- Mitchell, J. F. B., D. J. Karoly, G. C. Hegerl, F. W. Zwiers, M. R. Allen, and J. Marengo, 2001: Detection of climate change and attribution of causes. *Climate Change 2001: The Scientific Basis*, J. T. Houghton et al., Eds., Cambridge University Press, 695–738.
- Stott, P. A., S. F. B. Tett, G. S. Jones, M. R. Allen, J. F. B. Mitchell, and G. J. Jenkins, 2000: External control of twentieth century temperature variations by natural and anthropogenic forcings. *Science*, **15**, 2133–2137.
- , —, —, W. J. Ingram, and J. F. B. Mitchell, 2001: Attribution of twentieth-century temperature change to natural and anthropogenic causes. *Climate Dyn.*, **17**, 1–22.
- 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.
- Weaver, A. J., and F. W. Zwiers, 2000: Uncertainty in climate change. *Nature*, **407**, 571–572.
- Zwiers, F. W., and X. Zhang, 2003: Toward regional-scale climate change detection. *J. Climate*, **16**, 793–797.