Robust Bayesian Uncertainty Analysis of Climate System Properties Using Markov Chain Monte Carlo Methods

LORENZO TOMASSINI AND PETER REICHERT

Swiss Federal Institute of Aquatic Science and Technology, Dübendorf, Switzerland

Reto Knutti

Climate and Global Dynamics, National Center for Atmospheric Research,* Boulder, Colorado

THOMAS F. STOCKER

Climate and Environmental Physics, Physics Institute, University of Bern, Bern, Switzerland

MARK E. BORSUK

Department of Biological Sciences, Dartmouth College, Hanover, New Hampshire

(Manuscript received 23 January 2006, in final form 13 July 2006)

ABSTRACT

A Bayesian uncertainty analysis of 12 parameters of the Bern2.5D climate model is presented. This includes an extensive sensitivity study with respect to the major statistical assumptions. Special attention is given to the parameter representing climate sensitivity. Using the framework of robust Bayesian analysis, the authors first define a nonparametric set of prior distributions for climate sensitivity *S* and then update the entire set according to Bayes' theorem. The upper and lower probability that *S* lies above 4.5° C is calculated over the resulting set of posterior distributions. Furthermore, posterior distributions under different assumptions on the likelihood function are computed. The main characteristics of the marginal posterior distributions of climate sensitivity are quite robust with regard to statistical models of climate variability and observational error. However, the influence of prior assumptions on the tails of distributions is substantial considering the important political implications. Moreover, the authors find that ocean heat change data have a considerable potential to constrain climate sensitivity.

1. Introduction

Uncertainty and risk will be a cross-cutting theme in the forthcoming Intergovernmental Panel on Climate Change (IPCC) Fourth Assessment Report (Manning et al. 2004), and there has been considerable effort recently to quantify uncertainty in climate system properties.

Although attempts have been made (Allen et al. 2000; Stott and Kettleborough 2002; Murphy et al. 2004; Stainforth et al. 2005), there are limitations in using

© 2007 American Meteorological Society

atmosphere–ocean general circulation models for this purpose, since these models are very expensive to run. Therefore in many studies, climate models of intermediate complexity are employed. Their short runtimes make the application of rigorous statistical methods possible, while they are still able to reproduce the mean climate characteristics on a global scale. Furthermore, due to the simplified nature of such models, modelindependent climate system properties, such as the climate sensitivity, can be treated as explicit parameters, which in turn allows for exploration of a wide range of model behavior.

To determine values for climate system properties and estimate their uncertainty, it is reasonable to use comparisons between model predictions and climate observations. The methods of Bayesian statistics seem to be best suited to combine prior knowledge about the climate system with observational evidence (Knutti et al. 2002, 2003; Forest et al. 2002; Tebaldi et al. 2005),

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

Corresponding author address: Lorenzo Tomassini, Swiss Federal Institute of Aquatic Science and Technology, Ueberlandstrasse 133, CH-8600 Duebendorf, Switzerland. E-mail: lorenzo.tomassini@eawag.ch

DOI: 10.1175/JCLI4064.1

although other approaches have also been adopted [such as bootstrapping techniques, see Andronova and Schlesinger (2001)]. However, it has been proven that, in a certain sense, Bayes' theorem is an optimal information processing rule (Zellner 1988).

There are further advantages of the Bayesian approach. Every statistical analysis is based on specific assumptions that can only partly be corroborated by empirical data. This is, amongst other things, the reason for the different estimates of probability distributions for climate sensitivity in Forest et al. (2002, 2006) and Knutti et al. (2002). Bayesian statistics is a consistent framework in which the adopted statistical assumptions are transparent and explicitly formulated. This allows extensive sensitivity and robustness studies to be conducted, a particular strength of this approach (Box and Tiao 1962).

Robust Bayesian analysis (Berger 1984, 1994) provides a suitable framework for such sensitivity and robustness studies. In robust Bayesian analysis, instead of single prior distributions or single likelihood functions, sets of prior distributions and sets of likelihood functions are considered. Then upper and lower expectations of quantities of interest are computed (Wasserman and Kadane 1992). The range of these expectations provides a measure of robustness of the statistical analysis [see Berliner et al. (2000) for an application of robust Bayesian analysis to the climate change detection problem].

In situations where the computation of upper and lower expectations over a set of distributions is too computationally expensive, metrics on the space of distributions, such as the relative entropy distance, can be used to measure the distances between posterior distributions under different assumptions on the prior distributions and likelihood functions. These distances again provide a sensitivity or robustness measure (see Dey and Birmiwal 1994 for an overview).

In decision science, uncertainty in probabilistic estimates is often referred to as ambiguity. It has been shown in empirical studies that ambiguity affects decision making (Camerer and Weber 1992). Traditional decision science fails to capture ambiguity. Instead, under ambiguity, additional decision criteria, such as the precautionary principle, have to be considered (Borsuk and Tomassini 2005). One possible method to represent ambiguity is by sets of probability distributions as applied in robust Bayesian statistics [cf. Camerer and Weber (1992) for a comprehensive review and Kriegler and Held (2005) for a non-Bayesian application to climate change assessment].

We turn to the description of the content of the paper. In section 2 we briefly describe the climate model used, the parameters that were subject to the uncertainty analysis, the basic statistical methods that we applied, and the datasets that were used.

In section 3 the results of our analysis are presented. Throughout the paper, we focus mainly on the marginal distribution for climate sensitivity, since climate sensitivity is a property of the climate system, which does not depend on the type of model or parameterizations used.

In section 3a we show the result of the uncertainty analysis under our baseline statistical assumptions.

In section 3b we define a set of prior distributions for climate sensitivity *S* and calculate upper and lower posterior probabilities for the event that *S* lies above 4.5° C, which is the upper limit of the canonical [1.5, 4.5] IPCC range. This allows the sensitivity of the uncertainty analysis with respect to prior assumptions to be assessed.

In section 3c we investigate likelihood robustness. The sensitivity of the outcome of the uncertainty analysis to the scaling of the observational error and the estimate of natural climate variability is examined. The rationale behind this is that the observational errors that are published together with the data do not usually comprise all possible sources of uncertainty (Gregory et al. 2004). Also, our estimate of natural variability is uncertain (Collins et al. 2001; Gent and Danabasoglu 2004). For this reason it is important to assess the influence of the observational error and the estimate of climate variability on the result of the uncertainty analysis. In this section, we also study robustness with respect to the normality assumption in the likelihood function and quantify the effect of learning from the ocean heat content change data.

Discussions and conclusions follow in section 4.

2. Model and methods

a. Climate model

We use the Bern2.5D climate model, an earth system model of intermediate complexity. It consists of a zonally averaged dynamic ocean model (Stocker and Wright 1991; Wright and Stocker 1991) resolving the Atlantic, Pacific, Indian, and Southern Oceans, coupled to a zonally and vertically averaged energy and moisture-balance model of the atmosphere (Stocker et al. 1992; Schmittner and Stocker 1999).

The additional radiative forcing at the top of the atmosphere is specified as

$$\Delta F_{\rm toa}(t) = \Delta F_{\rm dir}(t) + \mu \Delta T_{\rm atm}(t), \qquad (1)$$

where ΔF_{dir} is the direct radiative forcing reconstructed over the industrial period. Feedback processes that increase the climate sensitivity *S* are represented by the feedback term $\mu \Delta T_{\text{atm}}(t)$, where ΔT_{atm} is the timedependent atmospheric temperature increase and μ is a constant, which is prescribed to lead to different climate sensitivities for the same radiative forcing.

The climate sensitivity *S* is defined as the equilibrium global mean near-surface warming for a doubling of preindustrial atmospheric carbon dioxide, equivalent to a radiative forcing of about 3.71 W m^{-2} (Myhre et al. 1998).

We used a model version with a lateral ocean mixing scheme and constant vertical ocean diffusivity K_v (Wright and Stocker 1991). Knutti et al. (2000) investigated the effect of different ocean mixing parameterizations. For the present study, the relevance of other mixing schemes is small and was subsequently neglected (Knutti et al. 2002).

The net heat uptake by the ocean F is given by Stocker et al. (1992) as

$$F = (1 - \kappa)Q^{\text{short}} - \sigma e_O T^4 + \sigma e_A T_A^4 - D(T - T_A) - E,$$
(2)

where *E* is the temperature-dependent evaporation, T_A is the atmospheric temperature, *T* is the ocean surface temperature, Q^{short} is the net incoming shortwave radiation, κ is a constant atmospheric absorptivity, e_A and e_O are atmospheric and oceanic emissivities, σ is the Boltzmann constant, c_E is the bulk coefficient of evaporation, and *D* is a constant transfer coefficient for sensible heat (Haney 1971).

The anthropogenic radiative forcing from changes in well-mixed greenhouse gases (CO₂, CH₄, N₂O, SF₆, and 28 halocarbons including those controlled by the Montreal Protocol), stratospheric O₃, the direct forcing of black and organic carbon, stratospheric H₂O due to CH₄ changes, and the direct and indirect effects of aerosols are individually prescribed from reconstructions for the years 1765–2000 (Joos et al. 2001). The radiative forcing by volcanoes (Ammann et al. 2003) and variations in solar irradiance (Crowley 2000) are prescribed for the historical period.

b. Basic statistical method, observations, and the likelihood function

Here we present a Bayesian uncertainty analysis for parameters of the Bern2.5D climate model. That is, for a vector of parameters $\theta = (\theta_1, \ldots, \theta_p)$ and observations $y = (y_1, \ldots, y_m)$, we calculate the multivariate posterior probability density of the parameters given the data according to Bayes' theorem:

$$p(\theta|y) \propto \ell(y, \theta) p(\theta),$$
 (3)

where $\ell(y, \theta)$ is the likelihood function (the conditional probability density of the observations *y* given the pa-

rameters θ) and $p(\theta)$ is the prior probability density for the vector of parameters.

For the computation of the posterior distribution, we apply a Metropolis–Hastings Markov chain algorithm with componentwise transitions (Gamerman 1997, section 6.4.1) using the computer program UNCSIM (Reichert 2005). To this end the vector of parameters is divided into two blocks (see section 2c for details). Componentwise transitions can accelerate the convergence of the chain because the rejection rate in each step of the algorithm may decrease. The sample size is 150 000 for all distributions presented in this paper. The corresponding probability density functions were estimated with a kernel density estimator (Silverman 1986).

The observations that we use consist of global annual mean surface temperature data (Jones and Moberg 2003) from the years 1861 to 2003 and annual mean change in World Ocean heat content down to 700-m depth (Levitus et al. 2005) from the years 1955 to 2003. Both datasets are publicly available. Thus the length m of vector y is 192.

We can write the likelihood function as

$$\ell(y,\theta) = p_{\varepsilon}[y - y_M(\theta)], \qquad (4)$$

where y_M is the output of the deterministic climate model, and p_{ε} is the probability density function of a random variable ε with $E[\varepsilon] = 0$.

The likelihood function implies an assumption on the distribution of the measurements around the computer model simulation, assuming that the computer model reproduces the trend in the data correctly. This distribution comprises the observational error as well as a statistical model for the natural processes that are not included in the dynamics of the climate model. These processes are called climate variability and are assumed to be stationary in the long run.

In section 3 we assess the sensitivity of the uncertainty analysis to the assumptions contained in the baseline likelihood function. Also, in section 3a we compare the distribution of residuals (the differences between data and the computer model) with the statistical error model formulated as the likelihood function in the baseline case.

We assume that the observation errors are independent from each other. The corresponding standard deviations for surface temperature and change in ocean heat content are taken from Jones and Moberg (2003) and Levitus et al. (2005), respectively. We are left to specify the variance–covariance structure of the climate variability part of the likelihood function.

We cannot estimate natural variability from data

Parameter	Distribution	Mean	Std dev
Climate sensitivity S	Uniform	5.5ª K	2.6ª K
Vertical ocean diffusivity K_{ν}	Uniform	$8.25^{\rm b} \ 10^{-5} \ {\rm m}^2 \ {\rm s}^{-1}$	$4.0^{\rm b} \ 10^{-5} \ {\rm m}^2 \ {\rm s}^{-1}$
Transfer coefficient D	Normal	$10 \text{ W m}^{-2} \text{ K}^{-1}$	$2.5 \text{ W} \text{ m}^{-2} \text{ K}^{-1}$
Greenhouse gas forcing scale s_{GHG}	Normal	1	0.05
Stratospheric O_3 forcing scale $s_{\text{strat}O_3}$	Normal	1	0.335
Tropospheric O_3 forcing scale s_{tropO_3}	Normal	1	0.215
Direct aerosol forcing scale s_{dir}	Lognormal	1	0.375
Indirect aerosol forcing scale s _{indir}	Uniform	1.5 ^c	0.58°
Organic and black carbon forcing scale s_{carbon}	Lognormal	1.163	0.69
Stratospheric water vapor forcing scale $s_{\text{stratH}_{2}O}$	Lognormal	1.163	0.69
Volcanic forcing scale s_{volc}	Lognormal	0.9	0.379
Solar forcing scale s_{solar}	Normal	1	0.335

TABLE 1. Parameters and prior distributions.

^a Implies a range of [1, 10].

^b Implies a range of [1.375, 15.125].

^c Implies a range of [0.5, 2.5].

alone because of the difficulty in separating natural variability from the underlying trend and because the data time series are relatively short. As is common practice in climate change attribution and detection studies (cf., e.g., Stott et al. 2001), we therefore consider a control run of a complex climate model, in our case the Third Hadley Centre Coupled Ocean–Atmosphere General Circulation Model (HadCM3) [see Collins et al. (2001) for a detailed discussion of the internal variability of HadCM3], as a representation of climate variability. This control run contains processes such as short-term weather fluctuations and ENSO-related variability that are not included in the parameterizations of the Bern2.5D model.

The control run has a length of 341 yr and was first detrended by a local polynomial fit. We then estimated the autocovariance in the HadCM3 control run time series of global annual mean surface temperature and annual mean change of World Ocean heat content down to 700 m. For the cross correlation, we first fitted autoregressive processes to each time series [the Akaike information criterion (AIC) selected a process of order 3 for global annual mean temperature and a process of order 19 for annual mean change of World Ocean heat content down to 700 m] and then estimated the cross correlation of the residuals. From this, the cross correlation of the original time series can be calculated. We found this cross correlation to be insignificant. The cross covariance in ε between global average annual mean surface temperature and annual mean change of World Ocean heat content was therefore set to zero.

Under the assumption that ε is Gaussian, the variance–covariance matrix Σ of ε is the sum of the diagonal matrix **D** that contains the variances of the obser-

vational error and the matrix Ξ that is made up of the estimated variance-covariance structure of climate variability (see Berliner et al. 2000 for the precise argument).

c. Parameters and prior distributions

We present a comprehensive Bayesian uncertainty analysis of climate system properties that includes 12 parameters of the Bern2.5D climate model. Although in the following we will focus on the marginal posterior distributions of climate sensitivity under different statistical assumptions, our study produces a multivariate posterior distribution of all the considered parameters.

The set of considered parameters consists of the climate sensitivity S, the vertical ocean diffusivity K_{ν} , the transfer coefficient for sensible heat D [see Eq. (2)], and nine forcing scale parameters.

Although latent heat dominates the heat exchange between ocean and atmosphere, in the final analysis we included only the transfer coefficient D of sensible heat and not the bulk coefficient of evaporation c_E , since in our climate model the influence of c_E on the ocean heat uptake is limited. Because in the model the atmosphere has no capacity to store water, the moisture evaporated in one ocean cell of the climate model falls as precipitation in the same time step in another ocean cell, and therefore the exchange of energy between ocean and atmosphere globally is independent of the value of c_E .

The forcing parameters consist of multiplicative dimensionless factors by which the historical forcing reconstructions for the various atmospheric components were individually scaled (see Table 1). This includes a scale parameter s_{indir} for the indirect aerosol forcing (Knutti et al. 2002). Greenhouse gases CO₂, CH₄, N₂O, SF₆, and halocarbons are combined in one group. The standard deviations of the prior distributions for the forcing scale parameters are derived from the assumption that the uncertainties given by IPCC in the Third Assessment Report (Houghton et al. 2001) represent a range of ± 1 standard deviation. A Gaussian prior distribution is assumed where the uncertainties are given in percent, and a lognormal distribution is used where the uncertainty is given as a factor (Knutti et al. 2003).

For climate sensitivity and the vertical ocean diffusivity, a uniform prior distribution is applied in our baseline assumptions (but see section 3b for different prior assumptions on climate sensitivity S). The uniform prior for climate sensitivity is restricted to the interval [1, 10] because the climate sensitivity S of the model is controlled by the parameter μ [see Eq. (1)]. The relation between S and μ is nonlinear and approximated numerically. This approximation is only accurate on the interval [1, 10]. Similarly, the uniform prior for K_{ν} is restricted to the interval [1.375, 15.125]. Previous studies (and the results of the present paper) suggest that these are reasonable ranges (Knutti et al. 2003).

For the transfer coefficient of sensible heat D an objectively estimated best-guess value is available (Haney 1971) and a Gaussian prior was used [for uncertainty estimates see also Ganachaud and Wunsch (2000)].

In Table 1 all the parameters with the baseline prior distributions are listed. For simplicity we assume a priori that all parameters are independent. Note that all prior distributions used are proper.

For the computation of the Markov chain Monte Carlo samples from the various posterior distributions, the 12-dimensional vector of parameters was divided into two blocks (cf. section 2b). One block contains all the forcing scale parameters except for the indirect aerosol forcing scale, and the second block contains climate sensitivity S, the indirect aerosol forcing scale, vertical ocean diffusivity K_v , and the transfer coefficient for sensible heat D. The two blocks are then sequentially updated in each iteration of the Markov chain Monte Carlo algorithm.

d. Classes of priors

The use of informative prior distributions in a Bayesian uncertainty analysis for climate sensitivity is controversial. One might argue that the knowledge used to construct such priors draws at least partly on the same observational evidence as that used in the Bayesian update procedure. That is why uniform priors have been used in some previous studies.

Nevertheless, the use of a uniform prior for climate sensitivity is itself questionable for another reason (Frame et al. 2005). In the literature, instead of climate sensitivity *S* one often encounters the feedback parameter λ defined as $\lambda = F^{\text{double}/S}$ where F^{double} is the forcing that corresponds to a doubling of CO₂ concentrations with respect to preindustrial times. If there is no prior knowledge on the climate sensitivity, then clearly there is no prior knowledge on the parameter λ . The use of a uniform prior on λ nevertheless implies a prior for climate sensitivity that is far from being uniform. This suggests that there are difficulties with uniform priors as a model of prior ignorance.

Jeffreys (1946) proposed to use "Jeffreys' prior," which is proportional to the square root of the Fisher information and therefore invariant under invertible parameter transformations, in cases of prior ignorance (see, e.g., Gelman et al. 1995, section 2.8 for a more extensive discussion). But in circumstances where the likelihood function is not given explicitly as a function of the parameters and has to be approximated by a Monte Carlo sample using a multitude of computer model simulations, the application of Jeffreys' prior is not a feasible option.

According to robust Bayesian practice, we take a somewhat different stance and introduce a nonparametric set of prior distributions for climate sensitivity that includes the uniform distribution as well as other informative priors. This entire set of prior distributions is then updated according to Bayes' theorem, which results in a set of posterior distributions.

There are several advantages to this approach. First, it allows the introduction of measures of the sensitivity of the uncertainty analysis to prior assumptions. Second, sets of distributions allow for better models of prior ignorance (Pericchi and Walley 1991).

There are different classes of distributions commonly used in robust Bayesian analysis, for example, mixture class, contamination class, density-bounded class, and density ratio class [see Berger (1994) for an overview and Borsuk and Tomassini (2005) for a discussion of different classes in the context of climate change assessment].

In the present paper we make use of the density ratio class of distributions (De Robertis and Hartigan 1981). This is defined by the set of distributions of the form

$$B(l,u) = \left\{ p \in C(\mathbb{R}) | p = \frac{q}{\int_{\mathbb{R}} q \, d\mu}, \, l(x) \le q(x) \le u(x) \forall x \in \mathbb{R} \right\},\tag{5}$$



FIG. 1. Example of a set of distributions of the density ratio class with upper bound u and lower bound l.

where l and u are specific functions, called the *lower* bound and the *upper bound*, respectively, of B(l, u).

In other words, for any probability density function p in B(l, u) there is a continuous function q, which is bounded by l and u, such that p is equal to q times a normalizing constant (which guarantees that p is normalized to 1).

Figure 1 shows an example of a density ratio class of distributions for a parameter θ , with upper bound *u* and lower bound *l*.

The probability density functions contained in this set of distributions are by definition all probability densities for which the *shape* can be placed within the upper and lower bounds, regardless of the normalization.

e. Calculation of the set of posterior distributions

We opted for the density ratio class because it enjoys the important property of update invariance (De Robertis and Hartigan 1981). This means that updating all distributions contained in B(l, u) according to Bayes' theorem results again in a set of distributions of density ratio class $B(\tilde{l}, \tilde{u})$, with new bounds \tilde{l} and \tilde{u} . This allows for incremental learning by updating the set of prior distributions sequentially (when more data become available), for a concise graphical representation of the posterior class, and for simple and inexpensive computations of upper and lower posterior probabilities (and, more generally, expectations).

The calculation of the posterior bounds \tilde{l} and \tilde{u} is based on the observation that the ratio $\rho(x) = u(x)/l(x)$ between the upper and lower bounds is invariant under updating, that is, $\tilde{\rho}(x) = (\tilde{u}(x)/\tilde{l}(x)) = \rho(x)$. One can therefore choose a certain reference prior $p_r(x) = cq(x)$, $l(x) \le q(x) \le u(x)$, (where *c* is the appropriate normalization constant) in the prior class, calculate the prior ratios $\rho_{\rm lr} = l(x)/q(x)$ and $\rho_{\rm ur} = u(x)/q(x)$, and produce a Markov chain Monte Carlo sample of the posterior \tilde{p}_r . The posterior lower bound is then simply $\tilde{l}(x) = \tilde{p}_r(x)\rho_{\rm lr}(x)$ and similar for the posterior upper bound $\tilde{u}(x) = \tilde{p}_r(x)\rho_{\rm ur}(x)$. For a more general discussion (but using the same principle) see Geweke and Petrella (1998).

Note that in order to calculate the upper and lower bound of the posterior class only one single standard Markov chain Monte Carlo sample has to be produced. Robust Bayesian analysis using the density ratio class does not therefore substantially increase the computational burden compared to standard Bayesian calculations.

3. Results

a. Results under baseline assumptions

We present the marginal posterior distribution for all parameters included in the uncertainty analysis with our baseline prior distributions (as described in section 2c).

In the baseline case, we assume a Gaussian likelihood function. The variance–covariance matrix of the likelihood function is estimated as described in section 2b.

For the baseline case, we scaled the observational standard deviation of the ocean heat content data by a factor of 1.5. We believe this is justified because in the calculation of the observational error not all types of uncertainties are considered (Levitus et al. 2005). The observations for heat content are sparse in some regions of the earth, especially in the Southern Ocean, and there is considerable uncertainty in the choice of an interpolation scheme between data points [see Gregory et al. (2004) for a thorough discussion of this issue].

Also, for the baseline case, we scaled our estimate of the autocovariance in the climate variability for ocean heat content by a factor of 1.25 to account for the fact that there is some variance in the control runs of HadCM3 for ocean heat content and that complex climate models tend to underestimate the climate variability of ocean heat content [Collins et al. (2001); Gent and Danabasoglu (2004) for a detailed discussion of this point with respect to the Community Climate Model version 2].

For the surface temperature data, no such scaling was introduced since these observations are more reliable



FIG. 2. Solid lines: Posterior distributions of all parameters under baseline statistical assumptions. Dashed lines: Assumed prior distributions (for details on definitions of the prior distributions and the units of the parameters see Table 1). The vertical axes denote the density of the distributions. (a) Greenhouse gas forcing; (b) stratospheric O_3 forcing; (c) tropospheric O_3 forcing; (d) direct aerosol forcing; (e) indirect aerosol forcing; (f) organic and black carbon forcing; (g) stratospheric H_2O forcing; (h) volcanic forcing; (i) solar forcing; (j) climate sensitivity *S*; (k) vertical ocean diffusivity K_v ; and (l) transfer coefficient *D* of sensible heat.

(Jones et al. 2001; Jones and Moberg 2003) and climate variability is believed to be well reproduced by HadCM3 (Collins et al. 2001).

In section 3c we will investigate the sensitivity of our results to scalings of the observational error as well as the estimate of climate variability.

In Fig. 2 the marginal prior and posterior probability distributions for all parameters under baseline statistical assumptions are summarized. Some posterior distributions seem to be multimodal. The multimodality of the marginal posterior for the transfer coefficient of sensible heat D is probably caused by lack of full con-



FIG. 3. Trace plot of the Markov chain Monte Carlo sample of size 150 000 for (top) climate sensitivity S, (middle) indirect aerosol forcing scale, and (bottom) vertical ocean diffusivity K_v . The trace of K_v hints at a multimodal marginal posterior distribution for K_v .

vergence. The coefficient D is expected to be insufficiently identifiable by the data used in this analysis, and the posterior distribution is supposed to converge to the prior distribution in this case. However, the influence of D on the uncertainty analysis for the remaining parameters appears to be minor anyway. The multimodality of K_v on the other hand seems to be genuine. Figure 3 shows a plot of the trace of the 150 000 iterations of the Markov chain for the three most important parameters: climate sensitivity S, vertical ocean diffusivity K_v , and the indirect aerosol forcing scale. The trace plot of K_v shows the rather typical pattern caused by a multimodal posterior distribution.

The multimodality of the posterior distribution for

 K_v may be due to the limited number of degrees of freedom of the relatively simple model and the small number of observations. It cannot be concluded that the multimodality of the distribution of K_v is a universal, model-independent phenomenon. To our knowledge, there is no physical, model-independent evidence that K_v should have a multimodal distribution.

Also the possibility that the Markov chain has not fully converged can never be completely excluded. Convergence is intrinsically slow if the target distribution is multimodal. However, the trace plots of the forcing scale parameters (including the ones not shown in Fig. 3) indicate good convergence of the respective marginal distributions.



FIG. 4. Plots of pairs of parameters from the Markov chain Monte Carlo posterior sample for climate sensitivity S, vertical ocean diffusivity K_v , and the indirect aerosol forcing scale. Only every 10th sample point is shown, a total of 15 000 points in each panel.

Besides the marginal posterior distributions, the correlations between parameters are of interest. Figure 4 depicts plots of pairs of parameters from the Markov chain Monte Carlo sample for climate sensitivity S, the indirect aerosol forcing scale, and the vertical ocean diffusivity K_{ν} . The pairs plots suggest that the full multidimensional posterior distribution has a very complicated structure and is not characterized by clear and distinct correlations between the parameters.

Figure 5 shows the model simulation with parameter values derived from the Bayesian parameter estimation. The parameter values correspond to the highest posterior probability density value of the full 12-dimensional posterior probability distribution. The difference between model output and observations is supposed to be accounted for by the statistical model of the observational error and climate variability. The shaded area represents ± 2 standard deviations of the statistical error model formulated by means of the baseline likeli-



FIG. 5. (top) Model-simulated anomalies in global mean surface temperature relative to 1961–90 (solid line) for the parameter set with the highest posterior probability density value (climate sensitivity $S = 2.49^{\circ}$ C, vertical ocean diffusivity $K_{\nu} = 6.46 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$). Observed annual mean values of global surface temperature are given as dots (Jones and Moberg 2003). (bottom) Model simulated anomalies in global mean ocean heat change down to 700 m relative to 1955–95 (solid line) for the parameter set with the highest posterior probability density value (same values as for the top panel). Observed annual mean ocean heat change down to 700 m is given as dots (Levitus et al. 2005). The shaded area represents ±2 standard deviations of the statistical error model formulated by means of the baseline likelihood function (the assumed correlation cannot be depicted in this figure).

hood function (the assumed correlation cannot be visualized in this figure).

Figure 6 presents quantile–quantile plots of the residuals (differences between the data and the best fit solution shown in Fig. 5) versus samples of the distribution of the statistical error model specified by means of the likelihood function for surface temperature and ocean heat uptake. If the data are well represented by the likelihood model, the points will fall on the 1:1 line. Whereas for surface temperature the statistical error model is overall adequate, there seem to be systematic deviations for ocean heat uptake in the region of high values. However, it has to be considered that the data time series is short compared to the range of the correlation in the HadCM3 control run for ocean heat uptake that serves as a model of cli-



FIG. 6. Quantile–quantile plots of the residuals (differences between the data and the best fit solution shown in Fig. 5) vs samples of the distribution of the statistical error model specified by means of the likelihood function for (left) surface temperature and (right) ocean heat uptake. If the data are well represented by the model, the points will fall on the 1:1 line.

mate variability (cf. section 2b). The difference between the distribution of the residuals versus the statistical error model might diminish if the data time series for ocean heat uptake spanned over a time period of 100 yr or more (and not just 49 yr as in our analysis, section 2b).

There is a discussion in the literature (e.g., Foukal et al. 2004) that the solar forcing has been overestimated in the past. Our results (Fig. 2i) suggest that the scaling factor for the solar forcing is overestimated by about 25%, that is, variations in solar forcing might be slightly overestimated in the reconstruction used here. This is consistent with recent findings that question the existence of long-term changes in solar irradiance and suggest that they might be substantially smaller than estimated in early reconstructions. However, it should be noted that the energy balance model used here neglects any possible indirect effects of solar variability on temperature. These include variations in the solar ultraviolet portion of the spectrum that affect stratospheric ozone concentrations, which in turn change incoming solar and outgoing infrared variations. These produce thermal gradients that drive motion and alter circulation patterns and therefore can influence global temperature by processes other than just the total amount of energy arriving at the top of the atmosphere.

The marginal posterior distribution for climate sensitivity under the described baseline assumptions is the focus of our analysis (Fig. 2j). The probability that the climate sensitivity *S* lies above 4.5° C is 0.16 in this case. The sensitivity of this result to prior assumptions will be investigated in section 3b.

b. Prior robustness

We defined a set of prior distributions of density ratio class (see section 2d) for climate sensitivity *S* with lower bound *l* and upper bound *u* as shown in Fig. 7a. These bounds were chosen in such a way that B(l, u)contains the uniform distribution as well as several expert priors on climate sensitivity of Morgan and Keith (1995). Of course it would be desirable to elicit the bounds *l* and *u* from experts directly, but appropriate elicitation techniques have yet to be developed.

The upper bound is a rescaled density function of the uniform distribution restricted to the interval [1, 10]. The lower bound is a rescaled density of the lognormal distribution with mean m = 3 and standard deviation std = 1.2 (again restricted to [1, 10]). Expert priors for climate sensitivity often have the form of a lognormal distribution, and the value of 3°C is a reasonable best estimate for climate sensitivity. The prior distributions for all the other parameters were kept fixed.

If we update all prior distributions that are contained in B(l, u) according to Bayes' theorem using the baseline likelihood function, we get a set of posterior distributions $B(\tilde{l}, \tilde{u})$ that is again characterized by upper and lower bounds \tilde{l} and \tilde{u} , as shown in Fig. 7b (see section 2e for more details; in the present case we chose the uniform prior on the interval [1, 10] as our reference prior p_r).

As a measure of sensitivity of our analysis to the



FIG. 7. (a) Set of prior distributions for climate sensitivity characterized by the lower bound l and the upper bound u; (b) set of posterior distributions of climate sensitivity after updating all distributions in the prior set. The set of posterior distributions is again characterized by lower and upper bounds.

prior assumption on climate sensitivity, we calculated the upper and lower posterior probabilities of the event that *S* lies above 4.5° C. This yielded an upper probability of 0.24 and a lower probability of 0.01, which can be compared to the (precise) probability of 0.16 in the baseline case (see section 3a).

We emphasize that the exact numbers of these upper and lower probabilities are not of primary importance. Rather they serve as a measure of sensitivity to assess the influence of prior assumptions on the uncertainty analysis. Also recall that the likelihood function was kept fixed in this calculation.

From Fig. 7b one can see that prior assumptions considerably influence the upper tail of the posterior distribution. Large climate sensitivities cannot be excluded by means of the data alone. However, the use of a uniform prior for the feedback parameter λ would imply a strongly informative prior for climate sensitivity *S*, which would result, after Bayesian updating, in a posterior distribution for climate sensitivity whose support is almost exclusively bounded to the [1.5, 4.5] range [see Allen et al. (2005) for a further discussion of this point].

c. Likelihood robustness

1) Observational error of surface temperature data

The scaling of the observational error of the surface temperature data (Jones and Moberg 2003) may influence the uncertainty analysis. Figure 8a (solid line) shows the marginal posterior distribution of climate sensitivity after a scaling of the observational error of the surface temperature data by a factor of 1.5.

The difference between the two distributions is within sampling uncertainty and therefore considered to be insignificant.

2) SURFACE TEMPERATURE VARIABILITY

The effect of a scaling of the estimated natural variability in surface temperature on the uncertainty analysis is addressed next. Recall that this variability was estimated using a HadCM3 control integration.

Figure 8b (solid line) depicts the marginal posterior distribution of climate sensitivity after a scaling of the estimated variance–covariance matrix of natural variability by a factor of 2. A narrowing of the distribution can be observed.

This might seem somewhat surprising at first, as the likelihood function is widened by the scaling. Nevertheless the increase of the natural variability in surface temperature gives less weight to the strong surface temperature increase in the last 30 yr compared to the whole temperature time series, which in turn reduces the likelihood of very large values for climate sensitivities. This effect corresponds to the fact that a strong climate reaction caused by anthropogenic CO_2 emissions is more difficult to detect if natural variability in surface temperature is assumed to be large.

3) Observational error of ocean heat content change

In this section we investigate the effect of a scaling of the observational error of the ocean heat content data (Levitus et al. 2005) on the uncertainty analysis.



FIG. 8. (a)–(f) Posterior distributions of climate sensitivity under different likelihood functions. (a) Solid line: observational error of surface temperature data scaled by a factor of 1.5; dashed line: baseline case. (b) Solid line: surface temperature variability scaled by a factor of 2; dashed line: baseline case. (c) Dotted line: observational error of the ocean heat content change data scaled by a factor of 1.25; dashed line: baseline case (observational error of the ocean heat content change data scaled by a factor of 1.5); solid line: observational error of the ocean heat content change data scaled by a factor of 1.5); solid line: observational error of the ocean heat content change data scaled by a factor of 1.5); solid line: observational error of the ocean heat content change; dashed line: baseline case (estimate of natural variability of ocean heat content change; dashed line: baseline case (estimate of natural variability of ocean heat content change scaled by a factor of 1.5); solid line: posterior distribution of climate sensitivity derived by means of a likelihood function that consists of the density of a multivariate t distribution with 3 degrees of freedom; dashed line: baseline case. (f) Solid line: posterior distribution of climate sensitivity derived without considering the ocean heat content change data of Levitus et al. (2005); dashed line: baseline case *with* the ocean heat content change data included.

Figure 8c shows the marginal posterior distribution of climate sensitivity after a scaling of the observational error of the ocean heat content data by a factor of $s_{data}^{ocean} = 1.25$ (dotted line). The dashed line is the baseline case. Recall that for the baseline case the observa-

tional error of the ocean heat content data was scaled by a factor of 1.5. Finally, the solid line represents the marginal posterior distribution of climate sensitivity with $s_{data}^{ocean} = 1.8$.

The widening of the distribution that can be observed

when comparing the dotted and the solid lines indicates that the data on the change in ocean heat content help to constrain the distribution of climate sensitivity. See section 3d for a further discussion of this point.

4) OCEAN HEAT CONTENT VARIABILITY

Figure 8d depicts the marginal posterior distribution of climate sensitivity without any scaling of the natural variability of the change in ocean heat content (dotted line). The solid line corresponds to $s_{var}^{ocean} = 1.5$. The dashed line again is the baseline case, for which the natural variability of change in ocean heat content, estimated from a HadCM3 control run, was scaled by a factor of $s_{var}^{ocean} = 1.25$.

An increase of the assumed natural variability in the change in ocean heat content results in a somewhat more pronounced tail of the posterior distribution of climate sensitivity. The general width of the distribution though is not significantly affected in the range of the scaling that was considered.

5) DISTRIBUTIONAL ASSUMPTION

In this section, we present the marginal posterior distribution of climate sensitivity with the same variance– covariance structure in the likelihood function as in the baseline case but with a likelihood function assumed to be the density of a multivariate t distribution with 3 degrees of freedom instead of a multivariate normal distribution. The small number of degrees of freedom was chosen in order to let the likelihood function become distinctly heavy tailed. The result is shown in Fig. 8e

We do not suggest that a t-distributed likelihood function would be more appropriate than a normal likelihood. Annual means of globally averaged climate observations, as drawn upon in this study, do not show a distinct extreme event pattern that would justify the use of a particularly heavy-tailed distribution. The results of this section should be interpreted in the sense of a sensitivity study with respect to assumptions on the likelihood function. Nevertheless the use of the probability density of a multivariate t distribution as likelihood function is not unsubstantiated. In the present study we assume that the variance-covariance structure that enters the likelihood function is known and can be determined independently from the data used (in our case it was estimated using a HadCM3 control run, see section 2b). If however the variance-covariance structure was considered uncertain, this would lead to a likelihood function for the standardized residuals that is the density of a multivariate t distribution.

d. Ocean heat content change data

In this section we quantify the effect of learning from the ocean heat change data (Levitus et al. 2005). Figure 8f shows the posterior probability density of climate sensitivity under baseline statistical assumptions but *without* the ocean heat change data included. Only the surface temperature record is considered.

A comparison against the posterior distribution *with* the ocean heat change data included reveals that the ocean heat content change data considerably help to constrain climate sensitivity in the present study.

4. Discussion

The posterior probability distribution for climate sensitivity that was derived under our baseline statistical assumptions is in line with the results of other studies (Andronova and Schlesinger 2001; Forest et al. 2002; Murphy et al. 2004). Compared to previous work with the same Bern2.5D climate model (Knutti et al. 2002, 2003), we succeeded in further constraining climate sensitivity. This was accomplished mainly by formulating a more realistic statistical error model for the observations.

We took advantage of past global surface temperature and change in ocean heat content records that consist of yearly measurements. This led to more constrained posterior distributions compared to studies where differences of zonal means between specific time periods were used (Forest et al. 2006). Clearly, the inclusion of data related to the seasonal cycle in the uncertainty analysis would improve the results and most probably further narrow the distributions (Knutti et al. 2006).

Given a set of plausible priors, the range of posterior probabilities of the event that climate sensitivity lies above 4.5°C is considerable given the important political implications and highlights the significance of a careful choice of a prior distribution. However, it should be noted in this context that tails of distributions are intrinsically difficult to estimate.

The use of a nonparametric set of prior distributions of the density ratio class is computationally inexpensive and suitable for assessing the sensitivity of the uncertainty analysis to prior assumptions. It is, in our opinion, to be favored over ad hoc choices of prior distributions.

An important result of our study is the finding that the estimate of natural variability in surface temperature has a quite significant effect on the uncertainty analysis: large natural variability in surface temperature leads to a narrower distribution for climate sensitivity. This should be considered in other, similar studies. However, natural variability in surface temperature estimated by control runs of complex general circulation models is generally believed to be realistic and therefore well represented in our baseline case.

Our sensitivity analysis with respect to the influence of ocean heat uptake observations on the uncertainty analysis suggests that there is indeed a significant potential in ocean diagnostics to constrain climate sensitivity (Barnett et al. 2001). However, this potential is hampered by the fact that natural variability in change of World Ocean heat content is difficult to estimate. Also a comprehensive analysis of the observational error and the reliability of the dataset is lacking (Gregory et al. 2004). Our study accentuates the need for a greater effort in monitoring ocean temperature and understanding its natural variability.

Certain main characteristics are shared by all estimated distributions of climate sensitivity under the various statistical assumptions, such as the general shape and the region of highest likelihood.

Our work makes a first step toward a more appropriate representation of uncertainty in climate system properties by conducting an extensive sensitivity analysis with respect to statistical assumptions. Robust Bayesian analysis proves to be a useful framework for this purpose. Using sets of probability distributions instead of single distributions whenever computationally feasible not only allows for a better communication of uncertainties to policy makers but also provides a sound basis for further developments in decision and risk theory.

Acknowledgments. The authors thank H.-R. Künsch for advice on statistical issues and C. Ammann for providing the volcanic forcing data. The authors acknowledge support by the Swiss National Science Foundation. R.K. acknowledges support by NCAR. Two anonymous reviewers provided helpful comments to improve this paper.

REFERENCES

- Allen, M. R., P. A. Stott, J. F. B. Mitchell, R. Schnur, and T. L. Delworth, 2000: Quantifying the uncertainty in forecasts of anthropogenic climate change. *Nature*, 407, 617–620.
- —, and Coauthors, 2005: Observational constraints on climate sensitivity. Avoiding Dangerous Climate Change, J. S. Schellnhuber et al., Eds., Cambridge University Press, 281– 289.
- Ammann, C. M., G. A. Meehl, W. M. Washington, and Ch. S. Zender, 2003: A monthly and latitudinally varying volcanic forcing dataset in simulations of 20th century climate. *Geophys. Res. Lett.*, **30**, 1657, doi:10.1029/2003GL016875.
- Andronova, N., and M. E. Schlesinger, 2001: Objective estimation

of the probability distribution for climate sensitivity. J. Geophys. Res., **106**, 22 605–22 612.

- Barnett, T. P., D. W. Pierce, and R. Schnur, 2001: Detection of anthropogenic climate change in the world's oceans. *Science*, 292, 270–274.
- Berger, J., 1984: The robust Bayesian viewpoint. *Robustness of Bayesian Analyses*, J. B. Kadane, Ed., North-Holland, 63–144.
- —, 1994: An overview of robust Bayesian analysis. Test, 3, 5–124.
- Berliner, L. M., R. A. Levine, and D. J. Shea, 2000: Bayesian climate change assessment. J. Climate, 13, 3805–3820.
- Borsuk, M., and L. Tomassini, 2005: Uncertainty, imprecision, and the precautionary principle in climate change assessment. *Water Sci. Technol.*, 52, 213–225.
- Box, G. E. P., and G. C. Tiao, 1962: A further look at robustness via Bayes's theorem. *Biometrika*, 49, 419–432.
- Camerer, C., and M. Weber, 1992: Recent developments in modeling preferences: Uncertainty and ambiguity. J. Risk Uncertain., 5, 325–370.
- Collins, M., S. F. B. Tett, and C. Cooper, 2001: The internal climate variability of HadCM3, a version of the Hadley Centre coupled model without flux adjustments. *Climate Dyn.*, **17**, 61–81.
- Crowley, T. J., 2000: Causes of climate change over the past 1000 years. Science, 289, 270–277.
- De Robertis, L., and J. A. Hartigan, 1981: Bayesian inference using intervals of measures. *Ann. Stat.*, **9**, 235–244.
- Dey, D. K., and L. R. Birmiwal, 1994: Robust Bayesian analysis using divergence measures. *Stat. Prob. Lett.*, 20, 287–294.
- Forest, C. E., P. H. Stone, A. P. Sokolov, M. R. Allen, and M. D. Webster, 2002: Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science*, **295**, 113–117.
- —, —, and —, 2006: Estimated PDFs of climate system properties including natural and anthropogenic forcings. *Geophys. Res. Lett.*, **33**, L01705, doi:10.1029/2005GL023977.
- Foukal, P., G. North, and T. Wigley, 2004: A stellar view on solar variations and climate. *Science*, **306**, 68–69.
- Frame, D. J., B. B. B. Booth, J. A. Kettleborough, D. A. Stainforth, J. M. Gregory, M. Collins, and M. R. Allen, 2005: Constraining climate forecasts: The role of prior assumptions. *Geophys. Res. Lett.*, **32**, L09702, doi:10.1029/2004GL022241.
- Gamerman, D., 1997: Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference. Chapman and Hall, 245 pp.
- Ganachaud, A., and C. Wunsch, 2000: Improved estimates of global ocean circulation, heat transport and mixing from hydrographic data. *Nature*, **408**, 453–457.
- Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin, 1995: Bayesian Data Analysis. Chapman and Hall, 526 pp.
- Gent, P. R., and G. Danabasoglu, 2004: Heat uptake and the thermohaline circulation in the community climate system model, version 2. J. Climate, 17, 4058–4069.
- Geweke, J., and L. Petrella, 1998: Density ratio class robustness in econometrics. *J. Bus. Econ. Stat.*, **16**, 469–478.
- Gregory, J. M., H. T. Banks, P. A. Stott, J. A. Lowe, and M. D. Palmer, 2004: Simulated and observed decadal variability in ocean heat content. *Geophys. Res. Lett.*, **31**, L15312, doi:10.1029/2004GL020258.
- Haney, R. L., 1971: Surface thermal boundary condition for ocean circulation models. J. Phys. Oceanogr., 1, 241–248.
- 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, 944 pp.

- Jeffreys, H., 1946: An invariant form for the prior probability in estimation problems. *Proc. Roy. Soc. London*, **186**, 453–461.
- Jones, P. D., and A. Moberg, 2003: Hemispheric and large-scale surface air temperature variations: An extensive revision and an update to 2001. J. Climate, 16, 206–223.
- —, T. J. Osborne, 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.
- Joos, F., I. C. Prentice, S. Sitch, R. Meyer, G. Hooss, G.-K. Plattner, S. Gerber, and K. Hasselmann, 2001: Global warming feedbacks on terrestrial carbon uptake under the IPCC emission scenarios. *Global Biogeochem. Cycles*, **15**, 891–907.
- Knutti, R., T. F. Stocker, and D. G. Wright, 2000: The effects of subgrid-scale parameterizations in a zonally averaged ocean model. J. Phys. Oceanogr., 30, 2738–2752.
- —, —, F. Joos, and G.-K. Plattner, 2002: Constraints on radiative forcing and future climate change from observations and climate model ensembles. *Nature*, **416**, 719–723.
- —, —, and —, 2003: Probabilistic climate change projections using neural networks. *Climate Dyn.*, 21, 257–272.
- —, G. A. Meehl, M. R. Allen, and D. A. Stainforth, 2006: Constraining climate sensitivity from the seasonal cycle in surface temperature. J. Climate, 19, 4224–4233.
- Kriegler, E., and H. Held, 2005: Utilizing belief functions for the estimation of future climate change. *Int. J. Approx. Reasoning*, **39**, 185–209.
- Levitus, S., J. Antonov, and T. Boyer, 2005: Warming of the world ocean, 1955–2003. *Geophys. Res. Lett.*, **32**, L02604, doi:10.1029/2004GL021592.
- Manning, M., M. Petit, D. Easterling, J. Murphy, A. Patwardhan, H.-H. Rogner, R. Swart, and G. Yohe, Eds., 2004: *IPCC Workshop on Describing Scientific Uncertainties in Climate Change to Support Analysis of Risk and of Options*. National University of Ireland, 138 pp.
- Morgan, M. G., and D. W. Keith, 1995: Subjective judgments by climate experts. *Environ. Sci. Technol.*, 29, 468–476.
- Murphy, J. M., D. M. H. Sexton, D. N. Barnett, G. S. Jones, M. J. Webb, M. Collins, and D. A. Stainforth, 2004: Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature*, **430**, 768–772.

- Myhre, G., E. J. Highwood, K. P. Shine, and F. Stordal, 1998: New estimates of radiative forcing due to well mixed greenhouse gases. *Geophys. Res. Lett.*, 25, 2715–2718.
- Pericchi, L. R., and P. Walley, 1991: Robust Bayesian credible intervals and prior ignorance. Int. Stat. Rev., 58, 1–23.
- Reichert, P., 2005: UNCSIM—A computer programme for statistical inference and sensitivity, identifiability, and uncertainty analysis. Proc. 2005 European Simulation and Modelling Conference (ESM 2005), Porto, Portugal, City Hall of Porto, 51–55.
- Schmittner, A., and T. F. Stocker, 1999: The stability of the thermohaline circulation in global warming experiments. J. Climate, 12, 1117–1133.
- Silverman, B. W., 1986: *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 175 pp.
- Stainforth, D. A., and Coauthors, 2005: Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature*, 433, 403–406.
- Stocker, T. F., and D. G. Wright, 1991: A zonally averaged ocean model for the thermohaline circulation. Part II: Interocean circulation in the Pacific–Atlantic basin system. J. Phys. Oceanogr., 21, 1725–1739.
- —, —, and L. A. Mysak, 1992: A zonally averaged, coupled ocean-atmosphere model for paleoclimate studies. J. Climate, 5, 773–797.
- Stott, P. A., and J. A. Kettleborough, 2002: Origins and estimates of uncertainty in predictions of twenty-first century temperature rise. *Nature*, **416**, 723–726.
- —, S. F. B. Tett, G. S. Jones, M. R. Allen, W. J. Ingram, and J. F. B. Mitchell, 2001: Attribution of twentieth century temperature change to natural and anthropogenic causes. *Climate Dyn.*, **17**, 1–21.
- Tebaldi, C., R. L. Smith, D. Nychka, and L. O. Mearns, 2005: Quantifying uncertainty in projections of regional climate change: A Bayesian approach to the analysis of multimodel ensembles. J. Climate, 18, 1524–1540.
- Wasserman, L., and J. B. Kadane, 1992: Computing bounds on expectations. J. Amer. Stat. Assoc., 87, 516–522.
- Wright, D. G., and T. F. Stocker, 1991: A zonally averaged ocean model for the thermohaline circulation. Part I: Model development and flow dynamics. J. Phys. Oceanogr., 21, 1713– 1724.
- Zellner, A., 1988: Optimal information processing and Bayes's theorem (with invited discussion). *Amer. Stat.*, **42**, 278–284.