

Probabilistic Estimates of Transient Climate Sensitivity Subject to Uncertainty in Forcing and Natural Variability

LAUREN E. PADILLA

Mechanical and Aerospace Engineering, Princeton University, Princeton, New Jersey

GEOFFREY K. VALLIS

GFDL, Princeton University, Princeton, New Jersey

CLARENCE W. ROWLEY

Mechanical and Aerospace Engineering, Princeton University, Princeton, New Jersey

(Manuscript received 9 August 2010, in final form 27 March 2011)

ABSTRACT

In this paper, the authors address the impact of uncertainty on estimates of transient climate sensitivity (TCS) of the globally averaged surface temperature, including both uncertainty in past forcing and internal variability in the climate record. This study provides a range of probabilistic estimates of the TCS that combine these two sources of uncertainty for various underlying assumptions about the nature of the uncertainty. The authors also provide estimates of how quickly the uncertainty in the TCS may be expected to diminish in the future as additional observations become available. These estimates are made using a nonlinear Kalman filter coupled to a stochastic, global energy balance model, using the filter and observations to constrain the model parameters. This study verifies that model and filter are able to emulate the evolution of a comprehensive, state-of-the-art atmosphere–ocean general circulation model and to accurately predict the TCS of the model, and then apply the methodology to observed temperature and forcing records of the twentieth century.

For uncertainty assumptions best supported by global surface temperature data up to the present time, this paper finds a most likely present-day estimate of the transient climate sensitivity to be 1.6 K, with 90% confidence the response will fall between 1.3 and 2.6 K, and it is estimated that this interval may be 45% smaller by the year 2030. The authors calculate that emissions levels equivalent to forcing of less than 475 ppmv CO₂ concentration are needed to ensure that the transient temperature response will not exceed 2 K with 95% confidence. This is an assessment for the short-to-medium term and not a recommendation for long-term stabilization forcing; the equilibrium temperature response to this level of CO₂ may be much greater. The flat temperature trend of the last decade has a detectable but small influence on TCS. This study describes how the results vary if different uncertainty assumptions are made and shows they are robust to variations in the initial prior probability assumptions.

1. Introduction

The steady-state response of the global-mean, near-surface temperature to an increase in greenhouse gas concentrations (e.g., a doubling of CO₂ levels) is given, definitionally, by the equilibrium climate sensitivity (ECS), and this is evidently an unambiguous and convenient measure of the sensitivity of the climate system

to external forcing. However, given the long time scales involved in bringing the ocean to equilibrium, the ECS may only be realized on a time scale of many centuries or more, and so its relevance to policy makers, and indeed to present society, has been debated. Of more relevance to the short and medium term—that is, time scales of a few years to about a century—is the transient climate response (TCR; Hegerl et al. 2007), which is the global and annual mean surface temperature response after about 70 yr given a 1% CO₂ doubling rate. (Sometimes an average may be taken from 60 to 80 yr or similar to ameliorate natural variability.) Although the detailed

Corresponding author address: G. K. Vallis, AOS Program, Princeton University, Princeton, NJ 08544.
E-mail: gkv@princeton.edu

response of the atmosphere to a doubling in CO_2 will likely depend on the rate at which CO_2 is added to the atmosphere, recent work with comprehensive models suggests that surface temperatures respond quite quickly to a change in radiative forcing, reaching a quasi-equilibrium on the time scale of a few years (in part determined by the mixed layer depth) prior to a much slower evolution to the true equilibrium (e.g., Held et al. 2010). In the quasi-equilibrium state, the rate of change of surface temperature is a small fraction of its initial increase, and the response following a doubling of CO_2 may be denoted the transient climate sensitivity (TCS). The TCS may be expected to be very similar to the TCR, but its definition does not depend so strictly on there being a particular rate of increase of greenhouse gases. As long as the CO_2 doubles over a time period short enough for deep ocean temperature to remain far from equilibrium (less than 100 yr, for example), the response to that doubling will likely be nearly independent of the emissions path. The ECS, in contrast, will take centuries to be fully realized. Given the time-scale separation between the transient and equilibrium responses, the TCS is a useful parameter characterizing the climate system, and it is this quantity that is the focus of this paper.

In addition to its relevance, the TCS may be easier to determine from observations than the ECS, in part because there are fewer free parameters to constrain. When estimating the TCS, we sum the atmospheric feedback strength and the rate of ocean heat uptake [also an uncertain quantity (Hegerl et al. 2007; Forest et al. 2002)], rather than constraining each factor separately. The overall response uncertainty, however, may still be dominated more by uncertainty in atmospheric feedbacks than the uptake of heat by the ocean (Knutti and Tomassini 2008; Baker and Roe 2009).

Various observationally based estimates have been made of both ECS and TCS (or TCR) using a variety of statistical techniques and a range of model complexity; Knutti et al. (2008) provide a useful review. Giorgi and Mearns (2002), Tebaldi et al. (2005), and Greene et al. (2006), for example, employ ensembles of comprehensive climate models, such as are described in the Intergovernmental Panel on Climate Change (IPCC) reports (e.g., Randall et al. 2007). These models try to represent the physical processes of the climate system, including processes determining aerosol forcing, in as explicit a way as possible. Although the physical parameterizations are tuned to simulate climate consistent with that observed, the ECS and TCS are not directly tuned by fitting to past climates; rather, they are obtained by integration of the model into the future. Still, model agreement with twentieth-century climates seems to

depend in part on the trade-off between aerosol level and climate sensitivity, and so some implicit tuning of climate sensitivity may occur (Kiehl 2007; Knutti 2008; Huybers 2010). Results from a collection of models may be combined to give a distribution of model sensitivities, but the distributions are effectively distributions of opportunity, rather than being properly controlled, and may be compromised by the repeated use of observations for model development and verification. Relatedly, in optimal fingerprinting (e.g., Hasselmann 1997; Allen et al. 2000; Stott and Kettleborough 2002; Stott et al. 2006, and others), a single comprehensive model's response patterns to specific forcing agents are scaled to achieve the best agreement with past observations. In evident contrast to these calculations and at the other end of the model-complexity spectrum are methods based on linear regression of past forcing and observed climate such as Gregory and Forster (2008) and Murphy (2010).

Another way, and the way that we shall proceed, is to construct a simple but physically based model and then to try to constrain the parameters that determine the model's transient climate sensitivity by a direct comparison with observations. In terms of model complexity, our methodology is closer to simple regression calculations than to use of GCMs, but it differs notably in that we seek to obtain time-dependent, probabilistic information. Specifically, we will constrain a simple energy balance model by observations of the twentieth-century surface temperature record, using a particular nonlinear form of the Kalman filter as a way of estimating parameters. This approach allows us to explicitly examine the way in which probability distributions depend on the underlying assumptions and length of the observed record. Set against this, compared to the general circulation models, is the less comprehensive nature and the lack of detail of the predictions made.

Our paper is structured as follows. In the next section we describe the simple energy balance model (EBM) we will use to constrain TCS. In section 3 we describe how we use a Kalman filter in conjunction with observations of the temperature record and estimates of the radiative forcing over the twentieth century to constrain the model, and in particular to estimate the model parameter that determines its TCS. We then show, in section 4, that the methodology can be applied to data from a comprehensive climate model and that, given only the globally averaged surface temperature and global perturbation radiative forcing of the model over the twentieth century, we are able to reproduce that temperature record with our EBM and to predict the transient climate sensitivity of the comprehensive model. In section 5 we apply the same methodology to the observations of

twentieth-century temperature and estimates of radiative forcing, providing probabilistic estimates of TCS for the real climate system. We also discuss whether and how our estimates are sensitive to the various assumptions we make about the nature of the uncertainties in the temperature and forcing record and the level of natural variability. We follow this by a general discussion and concluding remarks.

2. The energy balance model

To estimate the TCS we will use a zero-dimensional energy balance model similar to that used, for example, by Raper et al. (2002) and Held et al. (2010), and that we verify in section 4 is able to accurately emulate the evolution and predict the TCS of the Geophysical Fluid Dynamics Laboratory Climate Model version 2.1 (GFDL CM2.1) climate model. A feature of the model is that the parameter that governs its transient climate sensitivity may be wholly determined from observations.

To motivate the model, first consider a minimal model of the climate system that might be appropriate for determining both ECS and TCS. Such a model contains two independent variables representing perturbation surface temperature (T) and deep ocean perturbation temperature (T_o), namely

$$\begin{aligned} C \frac{dT}{dt} &= F - \gamma T - \beta(T - T_o) \quad \text{and} \\ C_o \frac{dT_o}{dt} &= \beta(T - T_o), \end{aligned} \quad (2.1a,b)$$

where γ and β are positive parameters; C and C_o are heat capacities of the mixed layer and deep ocean, respectively, with $C_o \gg C$; and F is the net perturbation to the climate forcing (including both natural and anthropogenic factors). In final equilibrium, $T = T_o$, and the temperature response (the ECS) to a specified forcing, say $F_{2\text{CO}_2}$, is given by $T_{\text{ECS}} = F_{2\text{CO}_2}/\gamma$. On decadal time scales the response of the deep ocean is small and $T_o \approx 0$. The system then obeys

$$C \frac{dT}{dt} = F - \lambda T, \quad (2.2)$$

with $\lambda = \gamma + \beta$. The separate values of γ and β , and thus T_{ECS} , are poorly constrained by observations so we focus on λ and the transient problem. We now add an additional term S to parameterize natural variability, so that our model equation becomes

$$C \frac{dT}{dt} = F - \lambda T + S. \quad (2.3)$$

The term S parameterizes internally forced temperature variability and satisfies the Ornstein–Uhlenbeck process (Majda et al. 2001; Vallis et al. 2004),

$$\frac{dS}{dt} = -\frac{S}{\tau} + \frac{\sigma_S \sqrt{2}}{\sqrt{\tau}} w_S, \quad (2.4)$$

where τ determines the temporal correlation of the variability, σ_S is the standard deviation of the variability, and w_S is a white noise process.

We emphasize that the parameter λ determines the transient climate sensitivity, not the equilibrium climate sensitivity, because it includes the rate of heat uptake by the deep ocean as well as the outgoing infrared radiation. Our λ is the same as the quantity ρ , termed the climate resistance by Gregory and Forster (2008), although our methods of finding it differ—we account for the time delay due to mixed layer heat capacity and therefore may make use of volcanic effects in constraining λ . The combination of λ and β is also similar to the sum of positive and negative feedbacks discussed in Baker and Roe (2009).

The parameter C represents the heat capacity of the system on decadal time scales, and we take its value to be that corresponding to a mixed layer of 60 m deep. Our results are fairly insensitive to this, and indeed a value of $C = 0$ does not give significantly different results (section 5f). Because the response is relatively rapid we may use (2.2) to define the transient climate sensitivity,

$$T_{\text{TCS}} \equiv \frac{F_{2\text{CO}_2}}{\lambda}, \quad (2.5)$$

where $F_{2\text{CO}_2}$ is the forcing corresponding to doubled CO_2 , and this is approximately equal to the quasi-equilibrium response to a forcing change in the time-dependent system (2.1). For instance, with plausible climate values—a deep ocean heat capacity corresponding to depth of 5000 m, an ECS of 3.5 K, and $\lambda = 2 \text{ W m}^{-2} \text{ K}^{-1}$ —the surface temperature reaches quasi-equilibrium, after instantaneous double forcing, in just 20 yr (5 times the transient time constant, $C/\lambda \approx 4$ yr). The surface temperature, $T_{\text{TCS}} = 1.86 \text{ K}$, is about 50% of the way to ECS. Surface temperature in this case will not reach even 75% of equilibrium for another 800 yr. Given the rapid initial response of the surface temperature (Held et al. 2010), the TCS will approximately equal the TCR. For the same parameters, the TCR is 1.77 K. The TCS is slightly greater because it is computed at a time when the rate of temperature change is closer to zero. The differences are illustrated in more detail in appendix A.

3. Assimilation of observations with a nonlinear Kalman filter

To estimate the parameter λ from past observations of temperature and forcing we use an adaptation of the Kalman filter applicable to nonlinear systems called the sigma-point Kalman filter (SPKF) based on Julier (2002). [The term λT in (2.3) is formally nonlinear because both λ and T are regarded as state variables. Physically λ is a constant parameter, but the Kalman filter adjusts its value to find the best fit.] Although there are many methods by which to find probability distributions for unknown parameters, we use the nonlinear filter because it is a simple method to implement and provides well-founded probability estimates. The recursive method has the additional advantage that in computing the posterior distribution given a time series of observations from t_1 to t_N , the posterior at every intermediate time t_i is automatically calculated. This feature thus enables one to study the evolution of uncertainty over time with the addition of observed temperatures. The filter also accounts for model dynamics and time delays. A simple, static regression of the temperature against the forcing would, given sufficient data, give similar values for the TCS but less probabilistic information, with less ability to determine the effects of forcing uncertainty and natural variability separately.

The SPKF resembles the classical Kalman filter for linear systems and Gaussian random variables in that it is an approximate recursive Bayesian method. Given prior and observed distributions of a model's state, the posterior state is the linear combination of the prior and observed states that minimizes the posterior error covariance. In each iteration, the forecast of the posterior at t_i becomes the prior at t_{i+1} . For nonlinear systems, the minimization and so the posterior mean and covariance cannot be solved exactly. Instead, they are computed from the statistics of an ensemble of state estimates. Member states are selected as ± 1 standard deviation perturbations about the mean, thus they are called sigma points.

The SPKF may be thought of as a particular type of the ensemble Kalman filters (Evensen 1994, 2007) frequently used in data assimilation and sometimes applied to parameter estimation (Annan et al. 2005a,b). For small systems, the SPKF requires far fewer ensemble members and has equal or better accuracy than the standard implementations of the ensemble Kalman filter. Also, the sigma points are deterministically recomputed, enhancing accuracy and aiding in avoiding collapse of the ensemble. Nonetheless, had a conventional ensemble Kalman filter been used it would likely have given similar results.

Since λ is regarded as a variable with respect to the filter, the state to be estimated is the 3-dimensional vector $[T, \lambda, S]^T$. There are several sources of uncertainty in our system: noisy inputs— w_F representing forcing uncertainty and w_S representing noise driving the natural variability of the EBM—as well as measurement error, with v representing error in the temperature record. Each of these sources is perturbed, along with the state variables, so the ensemble for our model consists of perturbations in 6 dimensions yielding a total of 13 sigma points, 12 symmetric perturbations, plus 1 to include the mean. When considering alternative models for forcing uncertainty in sections 4a and 5, the state dimension increases by one with the introduction of an aerosol forcing–scale factor α , and the noise dimension decreases by one with the elimination of w_F .

Each sigma point is forecast according to the nonlinear model (3). The forecast mean state and error covariance are computed from the statistics of the forecast points. (This is in contrast to the extended Kalman filter, which loses some accuracy because it relies on linearized system equations and a single state to propagate the error covariance.) The forecast sigma points are mapped onto the space of observed variables by the measurement equation for our system, $\hat{T}_{\text{obs}}^{(i)} = T^{(i)} + v^{(i)}$, where v is the zero-mean normal measurement error with standard deviation σ_v . The mean and covariance of the forecast sigma points are updated (or corrected) with the weighted difference between a real observation of the global average surface air temperature and the mean of modeled observations, $T_{\text{obs}} - \hat{T}_{\text{obs}}$, according to the standard, scaled-unscented equations (Julier 2002).

TCS probability

After completing iterations of the filter, we extract from the mean state and covariance estimates the mean $\bar{\lambda}$ and variance σ_λ^2 to form the marginal Gaussian probability density for the transient sensitivity parameter $\text{Pr}(\lambda)$. We map $\text{Pr}(\lambda)$ to transient temperature response densities $\text{Pr}(T_i)$ through the following relationship:

$$\begin{aligned} \text{Pr}(T_i) &= \text{Pr}(\lambda) \cdot \frac{d\lambda}{dT_i} \\ &= \frac{1}{\sqrt{2\pi}\sigma_\lambda} \exp\left[-\frac{1}{2\sigma_\lambda^2} \left(\frac{F_t}{T_i} - \bar{\lambda}\right)^2\right] \cdot \frac{F_t}{T_i^2}, \end{aligned} \quad (3.6)$$

where the derivative $d\lambda/dT_i$ comes from reorganizing (2.5) as $\lambda = F_t/T_i$, where F_t and T_i are the forcing and response corresponding to an arbitrary perturbation in greenhouse gases, including but not limited to CO_2 doubling. A similar probabilistic mapping was used by Roe and Baker (2007). The main feature of the mapping

is that Gaussian λ densities map to skewed transient response probabilities and the smaller the mean value of λ , the larger the tail probabilities toward large climate change. Note that rewriting λ in terms of a new parameter, $\lambda = 1/\mu$, and letting μ have a Gaussian uncertainty leads directly to a Gaussian $\text{Pr}(T_t)$. While it may seem more straightforward for λ to be Gaussian (or uniform—see Frame et al. 2005), the most appropriate formulation remains an open question; in any case, we find that with more observations the importance of the skewness of $\text{Pr}(T_t)$ to our TCS estimates diminishes.

4. Application to a general circulation model

In this section we show that the use of the method (i.e., the energy-balance model in conjunction with the Kalman filter) is able to emulate the evolution over the twentieth century of a comprehensive climate model (GFDL CM2.1; Delworth et al. 2006), and furthermore that the method can predict the TCS of the GCM using only its forcing and temperature record of the twentieth century. Since only a single realization of the real temperature record exists, we examine the extent to which single realizations of the GCM can be used to constrain transient climate sensitivity, rather than the average over an ensemble of integrations.

To do this we consider separately as constraining data the five CM2.1 Assessment Report 4 (AR4) twentieth-century integrations and their mean shown in Fig. 1. (The individual runs are smoothed with a 3-yr moving average to reduce some of the unforced variability in the time series since this will be accounted for in the magnitude of σ_S .) We model the perturbation forcing as the sum of a mean forcing record and white noise in each year, $F = \bar{F} + w_F$. The mean forcing \bar{F} for these experiments is that computed by Held et al. (2010), and the uncertainty in the forcing w_F has standard deviation $\sigma_F = 1 \text{ W m}^{-2} \text{ K}^{-1}$.

The estimate of the transient climate sensitivity parameter (λ) and its standard deviation σ_λ , as determined by the nonlinear Kalman filter, are shown in Fig. 1. By the year 2000, the mean estimate of $\lambda = 2.6 \text{ W m}^{-2} \text{ K}^{-1}$ corresponds to a most likely value for $T_{\text{TCS}} = 1.4 \text{ K}$, which agrees well with the known TCS of CM2.1 of 1.5 K. The value of λ from the individual runs is less constrained but the estimates remain within a standard deviation of the mean with the exception of run 4 (cyan in Fig. 1), which remains within the 90% confidence interval. The uncertainty range throughout the 100-yr time period is a little greater for the individual runs than the mean because they have greater natural variability. For run 4 the estimate deviates the farthest from the true λ as variability on longer time scales in the temperature

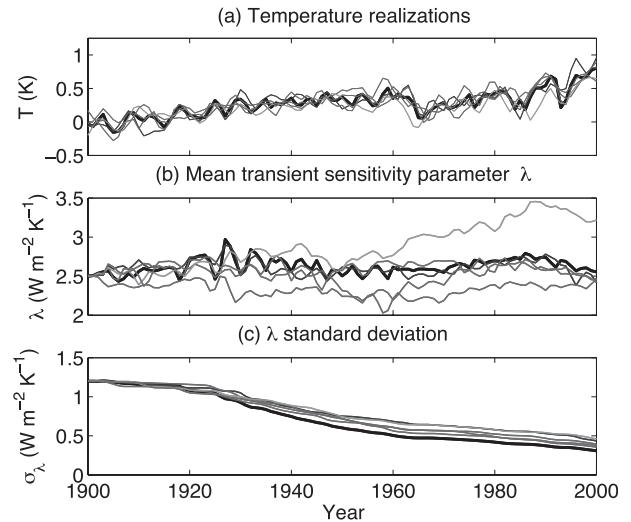


FIG. 1. (a) The five CM2.1 AR4 twentieth-century runs (thin lines) and their mean (thick line). (b) The mean estimate of transient sensitivity parameter λ , corresponding to each of the five temperature realizations (thin lines) and the ensemble mean (thick line). (c) The evolution of the standard deviations of λ for each run as in (b).

obscures forced features that help determine sensitivity. In the experiments run with the real temperature record, we partially offset this difficulty by removing the ENSO signal.

Fixing the value of λ at the mean estimate, $2.6 \text{ W m}^{-2} \text{ K}^{-1}$, the EBM reproduces the temperature response to forcing perturbations over the twentieth-century forcing, as shown in Fig. 2. Each realization of the EBM is initialized with a different random seed for w_S . The three simple model realizations are virtually indistinguishable from the individual five AR4 CM2.1 integrations (shown in Fig. 1a). We also plot a realization of the EBM in which the stochastic component is set to zero. This is an estimate of the forced response without any internal variability, and it agrees well with the main features of the CM2.1 response such as to volcanic eruptions in the years 1902, 1963, 1982, and 1991 as well as the overall trend.

a. Effect of assumptions on forcing uncertainty

In this section we further explore the effects of our assumptions regarding the uncertainty in the radiative forcing, the uncertainty that is often regarded as the biggest single impediment to calculating the equilibrium climate sensitivity from the past record. In the calculations corresponding to Figs. 1 and 2, this was modeled as a white noise, which is a good assumption for the forcing uncertainty in CM2.1, as is shown later in Fig. 4a. However, this is not a good assumption for the uncertainty in the real forcing, so in this section we introduce a more

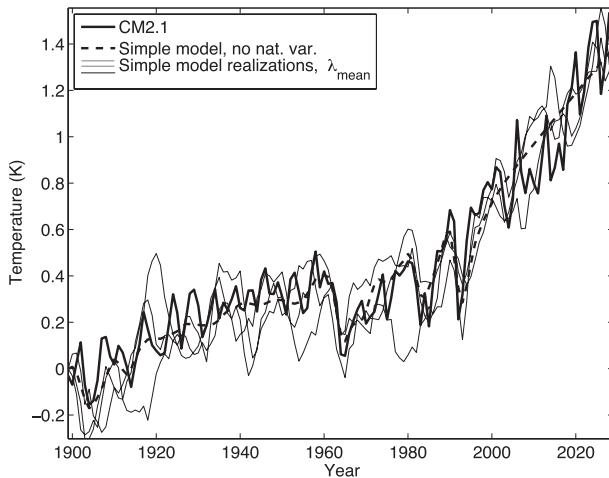


FIG. 2. Three realizations of the simple stochastic model response [Eq. (2.3), thin lines] compared to the mean of GFDL CM2.1 twentieth-century temperature response appended with the first 30 yr of the $1\% \text{ yr}^{-1}$ to CO_2 doubling run (thick solid line). Also shown is the deterministic response of the simple model where the random variable S is fixed at zero (thick dotted line).

realistic model for forcing uncertainty and present results for CM2.1 data, as a precursor to doing the same for observed data.

The IPCC attributes the greatest source of forcing uncertainty to anthropogenic aerosols, reporting a 90% confidence range of -0.5 to -2.2 W m^{-2} in 2005. Although other sources of uncertainty are not insignificant, we restrict uncertainty in our new forcing model to anthropogenic aerosols, and from here forward, aerosols (without a qualifying adjective) means those of anthropogenic origin. We separate the total historical forcing into aerosol and all other components:

$$F(t) = F_{\text{other}}(t) + \alpha F_{\text{aero}}^{\text{est}}(t). \quad (4.7)$$

We suppose aerosols are known only within a multiplicative-scale factor α , which is a unity-mean, normally distributed random variable. Scaling the magnitude of aerosol forcing is an approach that has been adopted previously by Harvey and Kaufmann (2002), Forest et al. (2006), and others. The variance of α and the variance of F are related by

$$\sigma_F^2 = (F_{\text{aero}}^{\text{est}})^2 \sigma_\alpha^2, \quad (4.8)$$

which is consistent with the idea that the greater the magnitude of the aerosol forcing, the greater is the uncertainty about it. For us, $F_{\text{other}}(t)$ is known exactly. Since the individual components of CM2.1 forcing are not available, we set $F_{\text{other}}(t)$ to the sum of greenhouse gas, solar, and volcanic contributions estimated

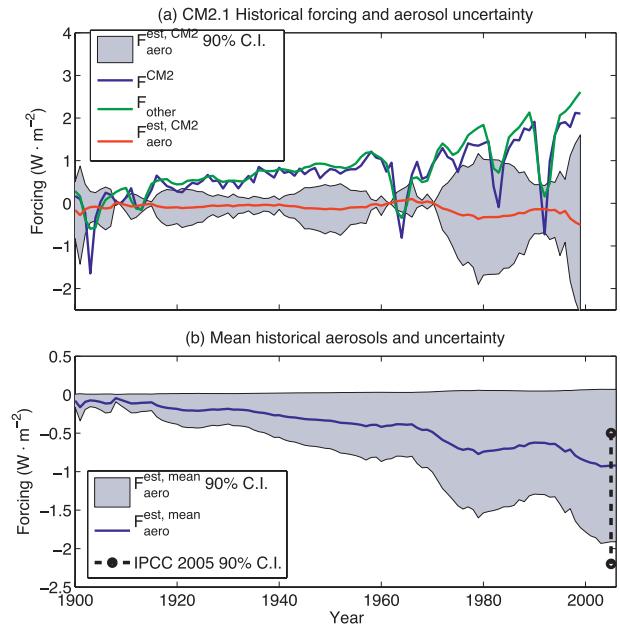


FIG. 3. (a) GFDL CM2.1 historical forcing estimate (blue) and components: *other* (green) and *aerosol* (red) with 90% confidence interval (gray). (b) Mean estimate of historical aerosol forcing based on three sources (GFDL, GISS, GF08) and 90% confidence interval (gray). This is the prior aerosol distribution assumed in the experiments with observed data in section 5.

by Gregory and Forster (2008). The nominal aerosol estimate $F_{\text{aero}}^{\text{est}}$ varies throughout this work depending on the experiment. In this application to a GCM, we are concerned with CM2.1 forcing and temperature, so we estimate CM2.1 aerosol forcing as the smoothed difference $F_{\text{aero}}^{\text{est}} = F^{\text{CM2}} - F_{\text{other}}$. We plot net CM2.1 forcing along with components *aerosol* and *other* in Fig. 3a. Also shown is the 90% confidence interval about the aerosols. Notice that since the variance of the forcing scales with the magnitude of the aerosols, when $F_{\text{aero}}^{\text{est}}$ is near zero, uncertainty is quite small. Here, the prior variance of the scale factor σ_α^2 is chosen such that the forcing variance σ_F^2 in 2005 is consistent with the IPCC confidence interval. By fitting a Gaussian and rounding up, we approximate the IPCC variance as $\sigma_F^2(2005) = 0.36 \text{ W}^2 \text{ m}^{-4}$. Then by Eq. (4.8), the variance of α at the start of assimilation is $\sigma_\alpha^2 = 1.4$.

We now allow the nonlinear Kalman filter to simultaneously constrain the parameters α and λ with the CM2.1 temperature data. In Fig. 4, we compare the results of a calculation with the new scaled forcing model to the results from the previous calculation with additive white forcing uncertainty. The spread of the λ density increases, as expected, since errors in the longer-term forcing trend are now taken into account. The mean estimate of λ remains very similar, with most likely TCS at

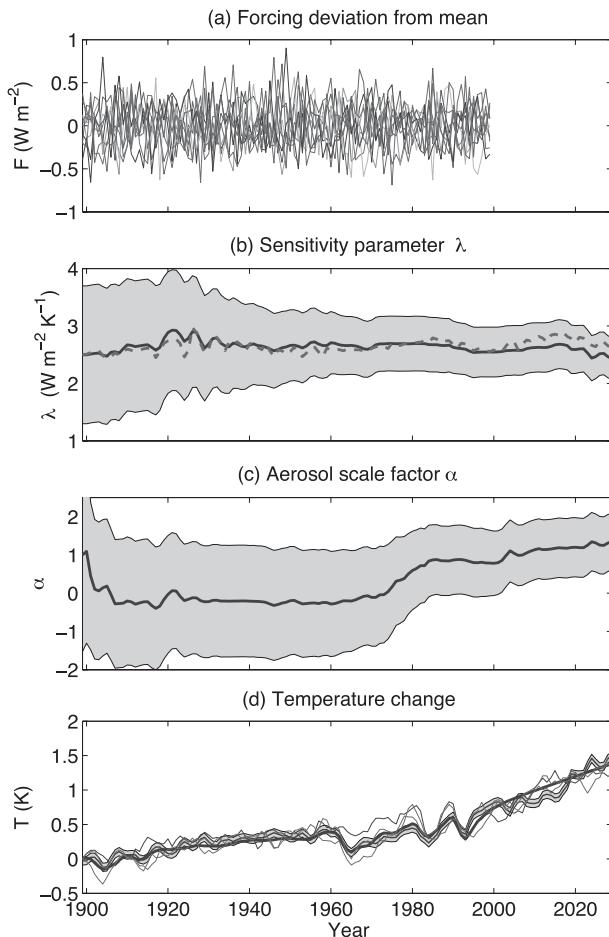


FIG. 4. (a) The difference from the mean of each of 10 forcing estimates derived for the CM2.1 forcing of Held et al. (2010). (b) The estimate of λ assuming scaled forcing uncertainty (solid) and the original estimate from white noise assumptions (dashed) for comparison. The gray band shows a range bounded by one standard deviation. (c) The estimate of α with $\pm 1\sigma$ in gray. (d) Comparison of CM2.1 twentieth-century temperature change and $1\% \text{ yr}^{-1}$ response after 2000 (gray band) to simple model temperature realizations with λ fixed at $2.6 \text{ W m}^{-2} \text{ K}^{-1}$ and $\alpha = 1$ (thick line).

1.4 K closely matching the TCS of CM2.1, which is really 1.5 K. The mean scale factor estimate remains between about 0.5 and 1 ending the assimilation at about 1. The posterior uncertainty about α narrows slightly, indicating that temperature observations do marginally constrain α . In allowing the filter to estimate part of the forcing trend, there are now two mechanisms by which the model may be corrected to emulate the increasing observed temperature record: decreased aerosol forcing, achieved when $\alpha < 1$, and decreased λ (more sensitive system). This makes it difficult to uniquely determine either parameter from the temperature time series alone; thus, we must consider entire probability distributions and avoid the temptation to focus solely on

mean estimates. Nevertheless, and as shown in Fig. 4c, realizations of the EBM over the twentieth century with these estimates of sensitivity parameter and forcing scale factor closely resemble realizations of the GCM.

b. Conclusions from the exercise with the GCM

The main conclusion to be drawn from the above exercise is that the methodology of using the EBM in conjunction with a sigma-point Kalman filter, when applied to the twentieth-century record of globally averaged surface temperature and forcing taken from a comprehensive climate model, is able to estimate, within reasonable error bounds, the transient climate sensitivity of the comprehensive model of about 1.5 K for a doubling of CO_2 . Using this value of TCS along with estimates of natural variability, the EBM is able to produce plausible trajectories of twentieth-century warming that are visually indistinguishable from trajectories of the GCM. There is, however, some sensitivity to the nature of the assumed uncertainty in the forcing. Nevertheless, these results give us confidence to proceed with applying the method to real data.

5. Application to the observed record

a. Observations and forcing

We now apply the same methods to estimate a probability density for λ from the observations of temperature and forcing over the twentieth century, estimates of natural variability, and the errors in these records. The observed temperature, Fig. 5a, is taken from Thompson et al. (2009), which itself is derived from the Hadley Centre and Climate Research Unit gridded near-surface temperature dataset (HadCRUT) data (Brohan et al. 2006). We have annually averaged the residual after subtracting the ENSO signal to remove some unforced variability. In the figure we also show some temperatures after 2008; these have been extrapolated with synthetic data generated by the model with no natural variability and $\lambda = 2.0 \text{ W m}^{-2} \text{ K}^{-1}$, which is approximately the mean estimate we obtain in 2008, and a forcing corresponding to a 1% increase per year in CO_2 with no change in aerosol forcing. That is, we essentially create a climate realization after 2008 using the simple model that we can then analyze with the Kalman filter. (The results are not especially sensitive to the slope of the extrapolated data on the time scales considered and yearly variability in the record is not a major factor in the results, as we show in section 5d.) The standard deviation of measurement errors in the observed temperature record is taken to be $\sigma_v = 0.06 \text{ K}$, as estimated by Brohan et al. (2006). To account for natural variability we assume for our base case that the natural variability

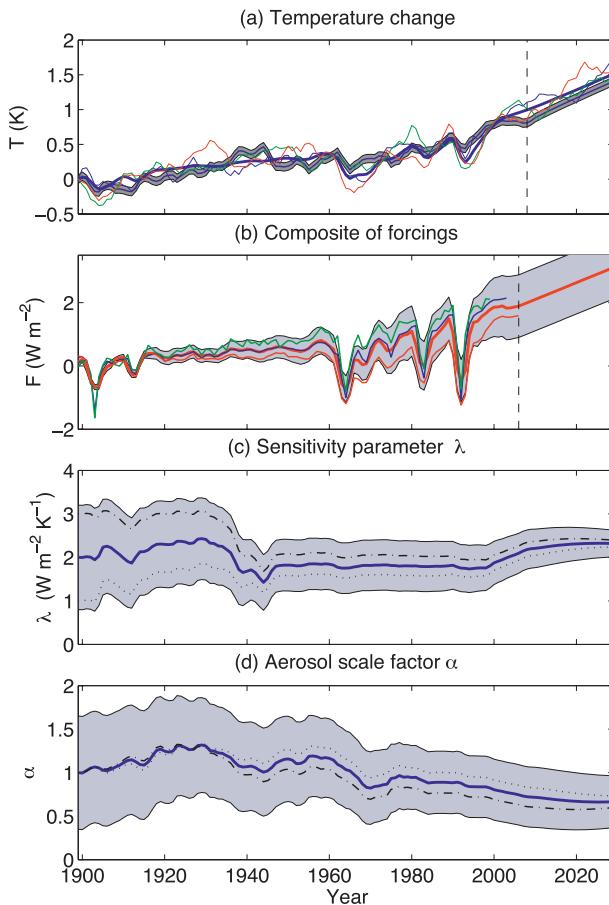


FIG. 5. (a) Observed global temperature after removal of ENSO signal, as well as temperature linearly extrapolated into the first decades of the twenty-first century, with ± 0.06 -K measurement error shaded gray. Also, three realizations and the mean EBM response with $\lambda \approx 2.1 \text{ W m}^{-2} \text{ K}^{-1}$. (b) Twentieth-century combined GFDL, GF08, and GISS forcing, with the 5%–95% confidence interval (shaded) and the three independent forcing estimates superimposed for comparison. (c) Mean estimate of transient sensitivity parameter λ for three prior means, $\lambda(1900) = 2.0$ (solid), 1 (dotted), and $3 \text{ W m}^{-2} \text{ K}^{-1}$ (dashed-dotted) and $\pm 1\sigma$ (gray shaded). (d) Mean estimate of scale factor α for the three prior mean λ values and $\pm 1\sigma$ uncertainty (gray shaded). The vertical dashed line marks the end of real data and beginning of use of synthetic data in the future.

[S in Eq. (2.3)] has a standard deviation of $\sigma_s = 0.07 \text{ W m}^{-2}$, which gives $\sigma_T \approx 0.13 \text{ K}$, this being the standard deviation of the observed detrended twentieth-century temperature record. We also calculate results for a range of values of σ_s and thus σ_T .

The forcing is taken from three sources, which we denote as the GISS [from Hansen et al. (2007), Goddard Institute for Space Studies], GFDL [from Held et al. (2010), Geophysical Fluid Dynamics Laboratory], and GF08 [from Gregory and Forster (2008)]. These forcings, shown in Fig. 5b, are obtained by slightly different techniques and represent slightly different levels of

atmospheric adjustment. For our base case we take F to be the mean of the three, although due to the differing lengths of the records after 2000, the mean forcing is the average of GISS and GF08, and from 2005 to 2006 it is GF08 alone. After 2006 we extrapolate the forcing by assuming $1\% \text{ yr}^{-1}$ increases in CO_2 with no change to other components. Differences in the time series become apparent in the middle of the twentieth century and continue to grow throughout the time period due mainly to unknown aerosol forcing. Therefore, we model this growing uncertainty as an unknown scale factor as we did for GCM data in section 4a by Eq. (4.7). Had we just used one of the forcing records, instead of the mean, there would have been small quantitative differences that are within the uncertainty bounds that we also calculate. Further discussion of the consistency of results with respect to forcing assumptions is provided in section 5e. Here, F is similarly separated into F_{other} , the sum of Gregory and Forster (2008) solar, volcanic, and greenhouse gas contributions, and $F_{\text{aero}}^{\text{est}}$, the smoothed difference between F and F_{other} . We plot $F_{\text{aero}}^{\text{est}}$ in Fig. 3b and infer the time-dependent uncertainty about $F_{\text{aero}}^{\text{est}}$, also shown, from the IPCC aerosol uncertainty range in 2005, assuming that the uncertainty is proportional to the aerosol level itself. Approximating the IPCC 90% confidence interval of -0.5 to -2.2 W m^{-2} as Gaussian as we did previously for the GCM experiments, by Eq. (4.8) the prior variance of the scale factor at initialization of the filter should be $\sigma_\alpha^2 = 0.42$. This yields the prior 90% confidence interval of the net forcing shown as the shaded region in Fig. 5. Recall from section 4a that the variance of the forcing is proportional to both the variance of the scale factor and the magnitude of the aerosol forcing. Since we let the Kalman filter determine the most likely trajectory of the scale factor and its uncertainty, the posterior confidence interval of the forcing will narrow slightly as uncertainty about the scale factor narrows.

b. Results for transient climate sensitivity

With the forcing and temperature data described above, beginning in 1900, we employ the nonlinear filter to sequentially update the probability density for λ and α as more observations are included as constraints. The time-varying mean and standard deviation of the λ density and α density are shown in Figs. 5c and 5d. Year 2000 estimates of mean λ are slightly more sensitive compared to 2008. This is a result of the flat to decreasing temperature trend in the last decade while forcing continued to increase. Beyond 2008, there is (as expected) little change in the mean estimate because we have fixed the sensitivity of the synthetic data close to the 2008 sensitivity. Also shown in Fig. 5c are the results of shifting the prior λ mean to high and low values

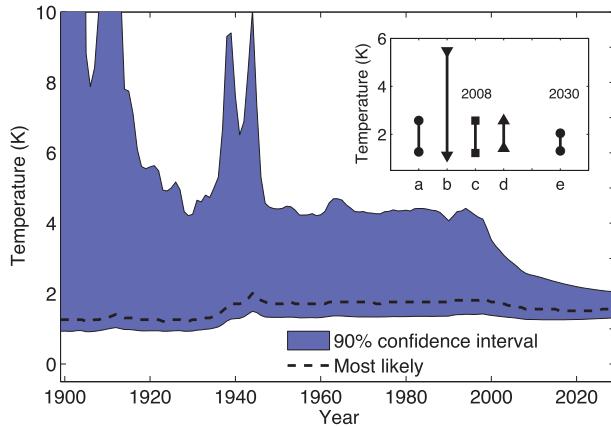


FIG. 6. Evolution of TCS probability as more data from the past are used as a constraint. The most likely TCS is indicated by the dashed line; the shaded region indicates the 90% confidence interval. Inset are the 90% confidence intervals for the uncertainty scenarios outlined in Table 1: (a) most probable, (b) high past natural variability, (c) forcing uncertainty is high, and (d) forcing uncertainty is low, all in 2008, and (e) most probable in 2030.

of 1 and $3 \text{ W m}^{-2} \text{ K}^{-1}$. The means converge steadily over time. The uncertainty in the distribution declines throughout the period of observation as a result of more data points unveiling the temperature trend. Halving the prior λ uncertainty had a minimal effect on the mean and uncertainty after 2000. Increasing prior uncertainty caused greater variance in λ before 1940 with quick convergence to the distributions shown in Figs. 5c and 5d. See appendix B for figures detailing the results of varying the prior.

Mapping the time evolution of the Gaussian λ density of Fig. 5 to the transient climate sensitivity via Eq. (3.6) yields the skewed TCS probability density whose 90% confidence interval as a function of time is shown as the shaded region in Fig. 6. The peak in the distribution is plotted as the dashed line. The prior distribution in year 1900 is noticeably skewed, with the 95th percentile including temperatures in excess of 10 K. As more data become available, the posterior TCS distribution continues to narrow until around 1940, when a temperature perturbation causes the sensitivity estimate to decrease. Even though the overall spread of the λ density narrows, the spread in TCS actually increases. This is a feature of the nonlinear relationship between λ and TCS in Eq. (2.5). For small λ , large improvements in our understanding of λ translate into only modest improvements in the confidence bounds of the TCS. Similarly, the decline in uncertainty after 2000 may be attributed to the increase in λ . In later years, the distribution of λ has narrowed sufficiently that the skewness of the TCS distribution is no longer a prominent feature. In 2008, the 90% TCS confidence interval is 1.3–2.6 K, and that range is reduced by 45% by 2030. (To avoid the effects

TABLE 1. Summary of natural variability and scale factor uncertainty for each of four different experiments. The resulting TCS confidence intervals (CI) given each set of assumptions are also included.

Experiment	σ_T (K)	σ_α	90% CI	
			2008 (K)	2030 (K)
Most plausible	0.13	0.65	1.3–2.6	1.3–2.0
Large F uncertainty	0.13	0.98	1.2–2.6	1.3–2.0
Small F uncertainty	0.13	0.33	1.4–2.6	1.4–2.1
Large natural variability	0.5	0.65	1.1–5.5	1.2–3.5

of high aerosol uncertainty, in section 5g we describe the effects of only using data from 1970 on; in fact, these lead to similar estimates of λ and TCS.)

For comparison with the results described above, which were obtained with assumptions about uncertainty detailed in section 5a that we consider most plausible, we also consider three limiting cases for past uncertainty: forcing uncertainty 50% larger, forcing uncertainty 50% smaller—both with our plausible estimate of unforced variability—and plausible forcing uncertainty with larger natural variability in the temperature record. These uncertainty scenarios are summarized in Table 1 along with the corresponding 90% confidence intervals after assimilation of surface temperature data up to 2008 and 2030. The confidence intervals are also shown in the inset plot of Fig. 6. The combined range of these intervals is an indication of the effects of uncertainty in our uncertainty estimates. As expected, the larger the forcing uncertainty and natural variability, the broader the spread becomes in the estimated λ .

To further describe the relationship among forcing uncertainty, natural variability, and transient climate sensitivity, we plot contours of the mean estimate and the standard deviation of the λ probability density in Fig. 7 for $(\sigma_\alpha, \sigma_T)$ pairs held at constant values for the 108 year period. The dashed lines indicate our most plausible estimate for $(\sigma_\alpha, \sigma_T)$. Notice in Fig. 7b that setting just one of these uncertainties to zero does not lead to zero uncertainty in λ . The flatness of the contours below about $\sigma_T = 0.13 \text{ K}$ indicates that we are primarily limited by natural variability. Large changes in the uncertainty about α have a very minor effect on σ_λ . The same observation can be made from the similarity in the confidence intervals a, c, and d of the Fig. 6 inset. Figure 7a shows that the mean transient sensitivity parameter is not very sensitive to the level of forcing uncertainty.

We also examine how uncertainty estimates change should only data from 1970 projected to 2030 constrain the parameters for the many $(\sigma_\alpha, \sigma_T)$ combinations. Figure 8a shows that λ is largely unaffected by variations in this shorter time period. Figure 8b shows that the standard deviation of λ decreases as the natural variability

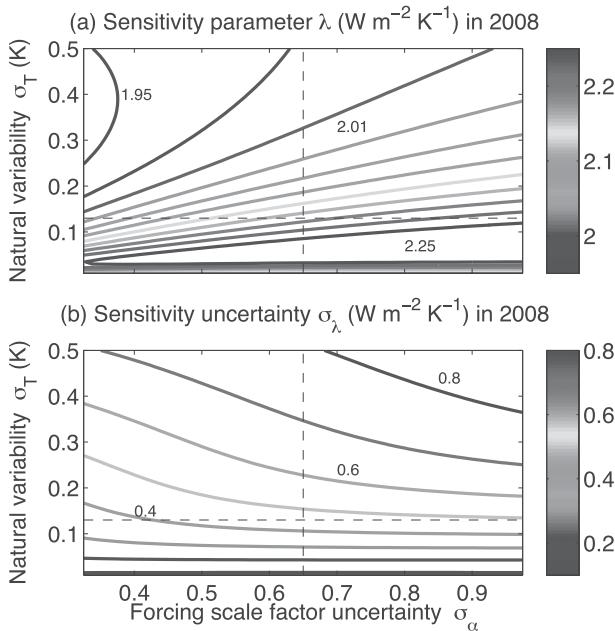


FIG. 7. Contours of mean (a) transient sensitivity parameter λ and (b) uncertainty σ_λ in 2008, as a function of varying levels of forcing uncertainty and natural variability held constant over the 108-yr filter period, 1900–2008. Dashed lines indicate our most plausible estimate of forcing uncertainty and natural variability in the past record. Contour intervals are $0.03 \text{ W m}^{-2} \text{ K}^{-1}$ in (a) and $0.1 \text{ W m}^{-2} \text{ K}^{-1}$ in (b).

and forcing uncertainty decrease, although λ uncertainty is much more sensitive to changes in natural variability over this shorter time period.

c. Implications for forcing levels

For the same scenarios described in Table 1 of the previous section, rather than mapping λ to TCS by fixing $F_{2\text{CO}_2}$ in Eq. (2.5), we ask what level of future forcing would be admissible if we required 95% confidence that temperature rise remained below various thresholds. Motivation for this analysis comes in part from work suggesting that dangerous climate change can be avoided if temperatures do not exceed a 2-K increase over preindustrial climate (Schellnhuber et al. 2006; Pachauri and Reisinger 2007), although we emphasize our results pertain to transient change, not equilibrium response, which may be much larger. The limits are plotted as the temperature threshold increases continuously in Fig. 9 in units of parts per million CO_2 by volume (ppmv). The net forcing, which may be attributed to any mix of greenhouse gases or aerosols, was converted from W m^{-2} to ppmv CO_2 equivalent using a standard empirical formula (Myhre et al. 1998).

Figure 9 shows that for 95% confidence transient temperature change remains below 2°C , given the 2008 λ distribution based on plausible uncertainty assumptions,

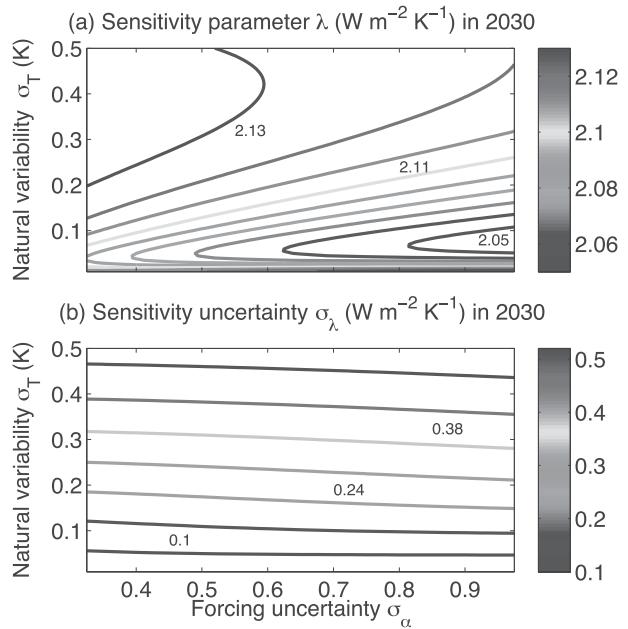


FIG. 8. Contours of the mean (a) transient sensitivity parameter λ and (b) uncertainty σ_λ in the year 2030 for various values of (σ_F, σ_T) held constant between 2070 and 2030. Contour intervals are $0.01 \text{ W m}^{-2} \text{ K}^{-1}$ in (a) and $0.07 \text{ W m}^{-2} \text{ K}^{-1}$ in (b).

short-term future forcing would be equivalent to no more than 2.9 W m^{-2} or 475 ppmv CO_2 if the forcing agent were CO_2 alone. In the natural variability–dominated case, the lowest of the curves, the temperature probability density becomes significantly flatter, leading to more stringent limits on emissions at every threshold compared to the central estimate; for a 2-K threshold, the limit is then 355 ppmv CO_2 equivalent (1.3 W m^{-2}). The current forcing level, already at 375 ppmv equivalent CO_2 (1.6 W m^{-2}), which includes the net effect of all anthropogenic agents (Pachauri and Reisinger 2007), has exceeded this limit. In 2030, with more constraining data and provided the temperature increase between now and then is no more than expected with our currently calculated sensitivity parameter, the 95th percentile of the temperature distribution decreases sufficiently to increase permitted forcing levels as high as 540 ppmv CO_2 equivalent. Relaxing the confidence level requirement would, of course, also increase the target forcing levels. For example, with only 50% confidence in remaining below the 2-K threshold, emissions could be as high as 630 ppmv CO_2 equivalent, and the equivalent CO_2 level that allows 2 K to be the most likely temperature increase is approximately 670 ppmv (4.7 W m^{-2}), although because of the thick tail in the distribution this would leave a greater than 65% probability that 2 K is exceeded.

Care must be taken in interpreting the above numbers because of the effects of aerosols and because of the

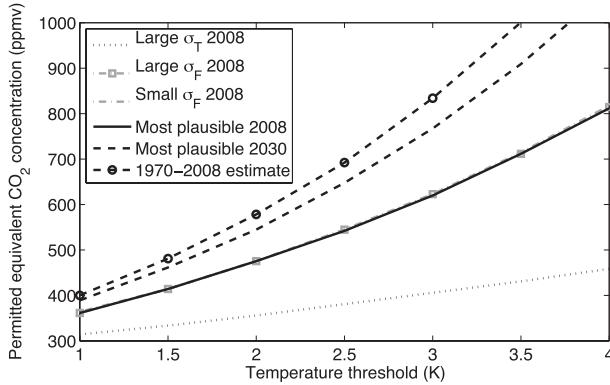


FIG. 9. Maximum forcing allowed to keep TCS below given thresholds with 95% confidence corresponding to the uncertainty scenarios in Table 1: natural variability dominated uncertainty in λ (dotted), most plausible estimates of uncertainty in λ (2008: solid; 2030: dashed), and high (dashed square) and low (dashed–dotted) forcing uncertainty dominated λ . Note, these last 3 curves lay nearly on top of each other.

time scales involved. If we were to attempt to limit future forcing to 475-ppmv CO_2 equivalent by reducing fossil fuel usage, then aerosol levels would likely also fall considerably, potentially enhancing the warming effects. For example, if aerosol effects are approximately -1 W m^{-2} , decreasing their effect to -0.5 W m^{-2} would mean reduction in target CO_2 concentration by 50 ppmv

to maintain a 475 ppmv CO_2 equivalent. Furthermore, the equivalent CO_2 concentrations presented here should be interpreted as upper bounds on forcing in the short-to-medium term (i.e., decadal timeframe), on which the transient climate sensitivity is relevant. If the forcing is such that transient surface temperature change approaches 2 K, then for sustained forcing at this level the equilibrium response almost certainly will exceed 2 K, and may exceed it by a considerable amount, because of the thermal inertia in earth’s oceans. Also, the simple model and short observation time series we use limit our analysis to the effect of feedbacks evident to date. This precludes any evaluation of the effect of delayed or nonlinear feedbacks, perhaps initiated once certain thresholds have been crossed. Thus, the numbers above should not be used as stabilization targets for the long-term (century scale) climate response.

d. Sensitivity to natural fluctuations

Given that we have only one realization of the temperature record, and with forcing estimates that are partly dependent on AOGCM results, we now show that our results are not overly sensitive to fluctuations in each trajectory. In Fig. 10, we repeat experiments, replacing the actual temperature and forcing data with alternate time series shown in Figs. 10a and 10b.

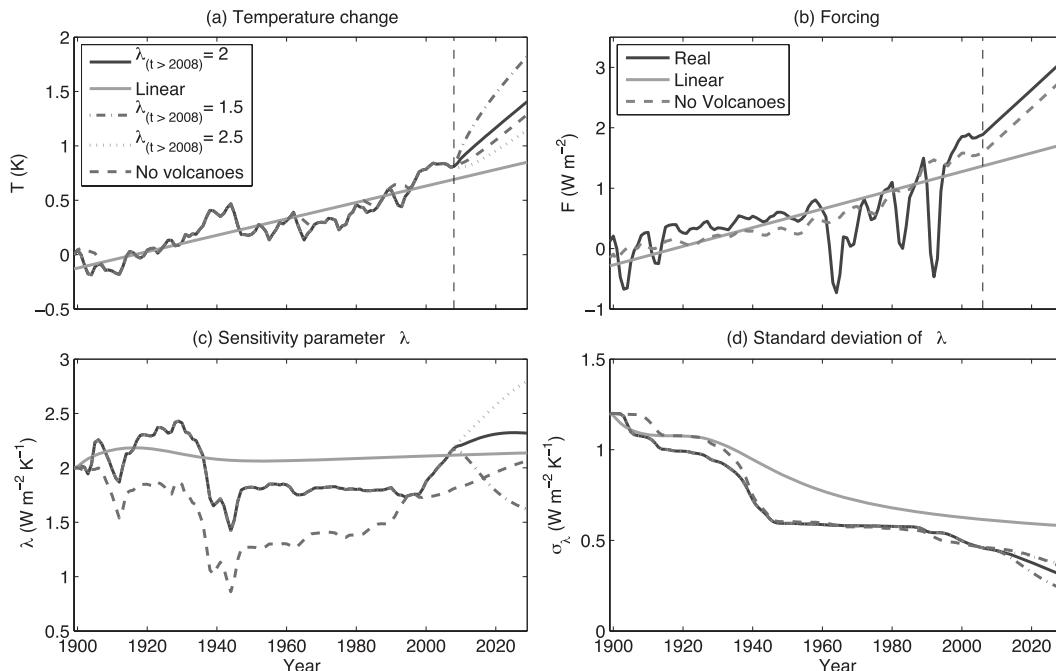


FIG. 10. Alternative (a) temperature and (b) forcing time series for experiments in which the data has been linearly fit, the temperatures after 2008 (denoted with a vertical dashed line) vary in slope, and volcanoes have been interpolated out, respectively. Also shown are the evolution of the (c) mean and (d) standard deviation of the probability density for λ when constrained by the alternative data.

First, we apply straight line fits to the data from 1900 to 2008 as in the green lines of Fig. 10, effectively removing all unforced variability as well as some forced features such as due to volcanoes. For the linear data, the results are not significantly different. In another experiment, we remove volcanic features through linear interpolation between the data in years bounding an eruption (dashed lines). Estimates of λ without volcanoes are slightly lower, indicating a higher T_{TCS} , or more sensitive climate—(see Fig. 12). In general, the observed response to volcanoes is, arguably, somewhat weaker than might be expected, presuming the forcing to be correct. It is possible that, because volcanoes cool, the effective oceanic heat capacity is larger in the period after an eruption, but further study of this seems warranted; however, overall, estimates with and without volcanoes are similar.

To study how uncertainty will change in the future as we collect more observations, we derived future temperatures from the simple model with a nominal choice of $\lambda = 2.0 \text{ W m}^{-2}$, which was near the mean estimate in 2008. We now look at the effects of varying the assumed future data modeled with $\lambda = 1.5$ and $\lambda = 2.5 \text{ W m}^{-2} \text{ K}^{-1}$, which just modifies the slope of the temperature record after 2008. The λ estimates diverge immediately from their 2008 value toward the prescribed values of 1.5 and $2.5 \text{ W m}^{-2} \text{ K}^{-1}$ but do not deviate more than the standard deviation of λ in the 20-yr period to 2030. Also, the reduction in uncertainty about λ is very similar for both records, so our estimate of what we can expect to learn by 2030 is valid even though we may be assuming incorrect future temperatures.

e. Sensitivity to forcing uncertainty

In this section we revisit the issue of the forcing. As is evident from Fig. 7, the uncertainty in forcing plays a major role in limiting the skill in our observational estimates of TCS. We have modeled the forcing uncertainty as an unknown scaling of the contribution of anthropogenic aerosols because we note that the three forcings, GISS, GFDL, and GF08, that we have used diverge from each other with time (Fig. 11b) rather than oscillating round each other, as in the GCM forcing error of Fig. 4a.

Given this, we may see if our estimates of the sensitivity parameter are particularly sensitive to the specific forcing time series. Thus, for example, suppose that we only had available one of the three forcings we have used—would the computed parameter λ differ noticeably? The results of doing just this are shown in Fig. 11a. The estimates of λ remain within about one standard deviation of the mean estimate as indicated by the gray shaded region regardless of which forcing trajectory is

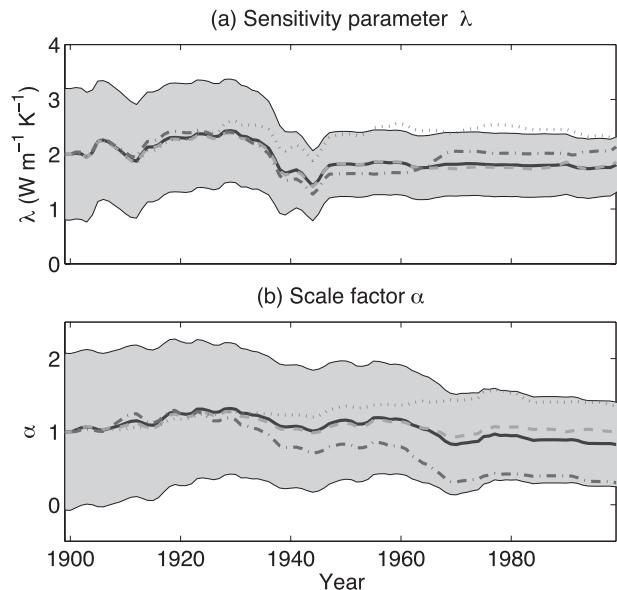


FIG. 11. (a) The evolution of λ when the forcing is taken to be GISS alone (dashed), GF08 alone (dashed-dotted), GFDL alone (dotted), and the three forcing mean (solid). The gray band is one standard deviation about the mean estimate. (b) As in (a), but for α . Here, the gray band is the 90% confidence range.

assumed; GISS and GF08 forcing yield λ well within the envelope and GFDL forcing leads to λ slightly outside. A white noise uncertainty model would underestimate the uncertainty in λ and cause the estimates using the three forcings individually to fall well outside the one standard deviation envelope. The estimates of the scale factor, which vary more widely due to the quite different forcing time series, exceed the standard deviation but remain within the 90% confidence range about the mean.

f. Sensitivity to mixed layer depth

All calculations up to this point have been made assuming an effective heat capacity C in the energy balance model (2.3) corresponding to a mixed layer 60 m deep. We now explore the effects of varying the mixed layer depth H on estimates of TCS and the sensitivity parameter λ . As shown in Fig. 12a, we find that the mean λ , constrained by surface temperatures through 2008, is largely unaffected by depths ranging from $H = 40$ to $H = 200$ m.

However, for $H < 40$ m, the mean estimate of λ increases (so the temperature response becomes less sensitive to forcing). This occurs because low-heat-capacity models have a large temperature response to volcanic forcing whereas the observed temperature responds very little. The Kalman filter corrects the model's sensitivity so that the response to volcanoes is not overly exaggerated, but it is clear from Fig. 12b that with $H = 1$ m and the mean λ and α from 2008, the simulated

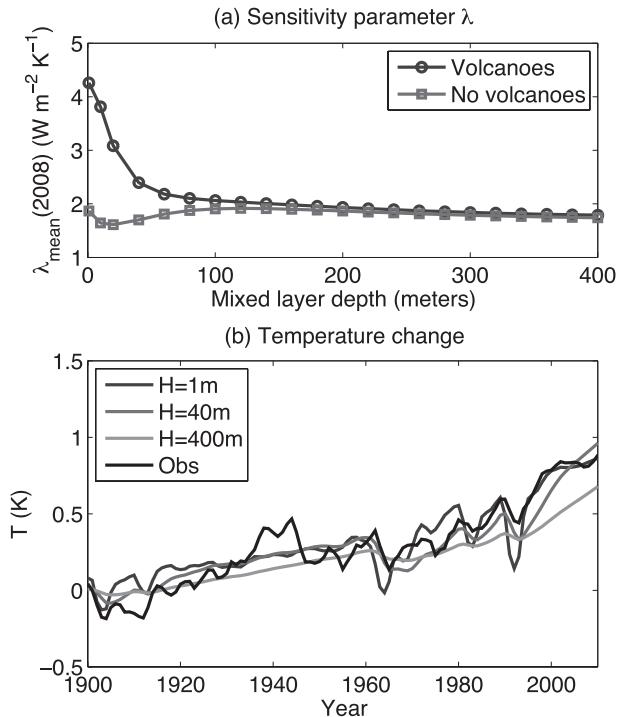


FIG. 12. (a) Estimates of mean transient climate sensitivity parameter by 2008 as a function of mixed layer depth in the simple model, with (circle) and without (square) volcanic response included in the constraining temperature data. (b) Surface temperature response of the simple model [Eq. (2.2)], with λ and α fixed at 2008 means, for mixed layer depths of 1, 40, and 400 m compared to historical temperatures (black).

temperatures still overshoot observation in volcano years. When we removed the volcanic component from the forcing time series, λ remains quite constant for all mixed layer depths. On the other hand, for large depths, $H > 200$ m, the model is quite insensitive to forcing fluctuations, so mean λ is estimated as a bit more sensitive in order to better simulate the historical record. For example, see the simulation with $H = 400$ m in Fig. 12b: the volcanic response is too smooth and the fit becomes increasingly poor beyond 1970. Therefore, our choice of $H = 60$ m is appropriate for use with forcing including volcanoes.

The uncertainty surrounding estimates of the parameters λ and α increases fairly steadily as the mixed layer depth increases. This is expected as the greater the thermal inertia of the model, the longer the time series of temperature observations needed to increase the signal to noise ratio and thus reduce uncertainty in the nonlinear Kalman filter. Overestimating the mixed layer depth, therefore, may lead to slow-to-converge parameter estimates that ignore much of the observed signal. On the other hand, underestimating the depth could result in overconfident probability densities. In Fig. 13,

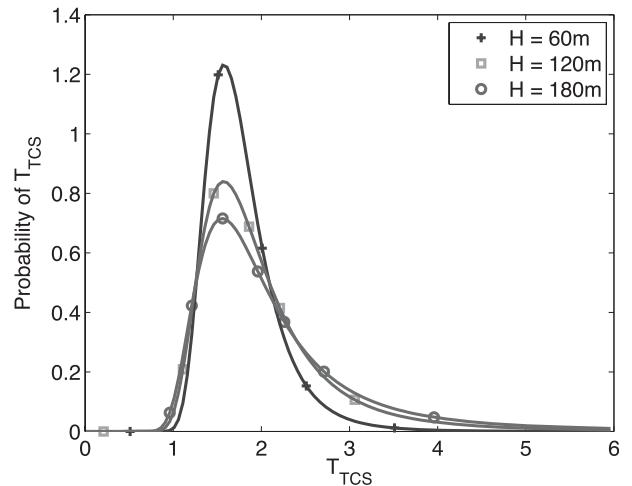


FIG. 13. Probability density distributions for the transient climate sensitivity in 2008 with mixed layer depths of 60, 120, and 180 m.

we compare estimates of the TCS probability density for $H = 60$, 120, and 180 m. Increasing uncertainty with depth is evident in the lengthening of the tails of the distribution.

g. Results using recent data only

The uncertainty in aerosol forcing is generally believed to be greatest in the middle part of the last century and less in the last third of the century and in the early part of the 21st century. Therefore, one may be interested in estimates of TCS using only the data from 1970 on, and this route was taken by Gregory and Forster (2008). The disadvantage is that the shorter time period means that the uncertainty due to natural variability can be expected to be larger. In Fig. 14 we show results using only data from 1970 and using otherwise the same uncertainty assumptions as used previously. The change in the aerosols and thus the three forcing estimates (GISS, GFDL, and GF08) in this period are now very similar to each other, and so our estimate of the standard deviation of the forcing scale factor is rather conservative (i.e., probably too large).

Rather encouragingly, estimates of λ and so of TCS calculated in this manner are similar to the estimates constrained by data all the way back to 1900, although the relatively steady temperatures after 2000 dominate the shorter time window and lead to slightly less sensitive estimates of TCS, with a 90% confidence interval of 1.1–1.9 K by 2008. The spread of this interval is slightly narrower, even with the smaller dataset, than the estimates from the longer record, and this may appear somewhat paradoxical. One reason for the reduced spread may be that the simple EBM better simulates the more recent period and so the model learns faster. A second reason is

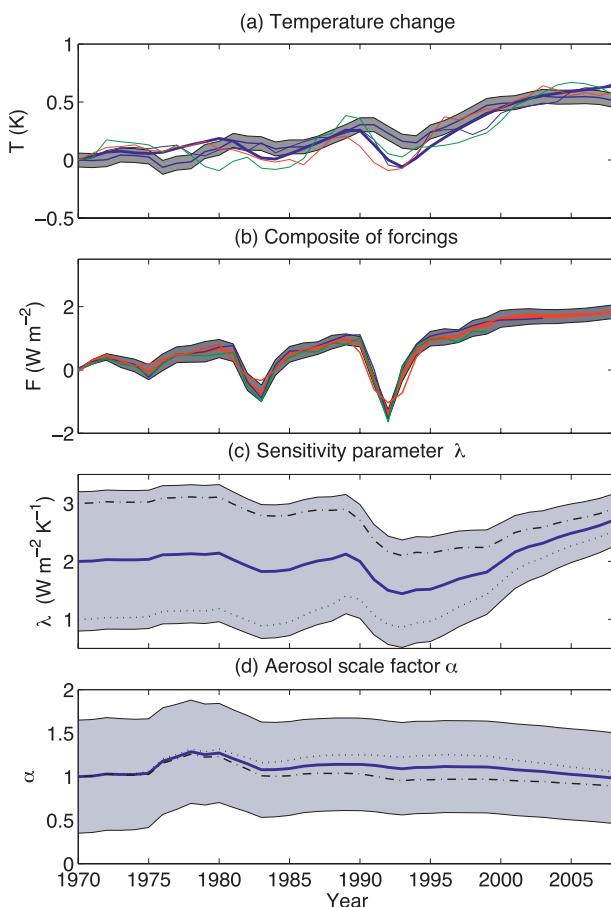


FIG. 14. As in Fig. 5, but using data only from 1970.

that the generally larger λ leads to a TCS distribution with less skew and a reduced likelihood of a very large temperature response: given two lambda distributions with the same spread but centered about different mean values, the distribution with the less sensitive mean will have a smaller TCS uncertainty range. The results using the shorter record are also a little more sensitive to the priors that are chosen, and so the results arising from the use of temperature record over the entire century probably reflect better the true uncertainty; however, as noted, the differences are small.

6. Discussion and conclusions

We conclude with a few general and summary remarks. The use of the nonlinear Kalman filter in conjunction with a semiempirical model allows us to estimate the distribution of transient climate sensitivity, how the distribution explicitly depends on forcing uncertainty and natural variability, and how the distribution may change in the future as more data becomes available. The nature of the forcing uncertainty, whether scaled or not, as well

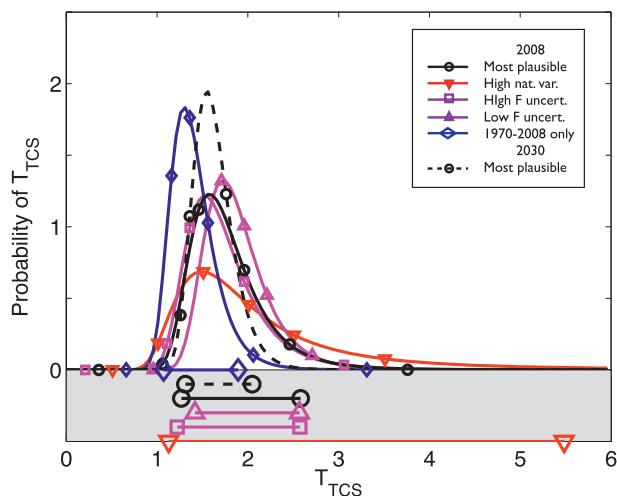


FIG. 15. The range of PDFs and 90% confidence intervals for the transient climate sensitivity for the various sets of assumptions made. With data up to 2008 the estimates are based on the most plausible forcing uncertainty and natural variability, forcing uncertainty underestimated, forcing uncertainty overestimated, natural variability overestimated, and most plausible with only 1970–2008 data. With data up to 2030 the most plausible forcing uncertainty and natural variability are shown.

as its magnitude, also affects our resulting probability distributions.

Although our estimates are certainly sensitive to these uncertainties and to natural variability, they may be sufficiently narrow as to still be useful. For uncertainties ranging from very large forcing uncertainty to very small forcing uncertainty, our confidence intervals for TCS range from 1.2–2.6 to 1.4–2.6 K. With a much larger portion of the observed temperature change attributed to natural variability, our TCS interval increases to 1.1–5.5 K. Our probabilistic estimate of the range of TCS that we believe to be best justified by data, namely 1.3–2.6 K with a most probable estimate of 1.6 K, is broadly consistent with the TCR range of IPCC AR4 climate models whose median and mean are 1.6 and 1.8 K, with 90% confidence interval of 1.2–2.4 K (Randall et al. 2007; Meehl et al. 2007). Figure 15 summarizes our range of probabilistic estimates given data from 1900 to 2008 and 2030. The collection of probability densities and corresponding confidence intervals indicate both the state of TCS uncertainty today and potential for improved understanding 20 years in the future.

To obtain much tighter estimates using methodology similar to ours would require significantly reduced uncertainties in the forcing. Without that, attributing the observed temperature rise definitively to climate feedbacks or erroneous forcing is rather difficult. It may be possible to bypass this difficulty by looking at the temperature increase of the two hemispheres separately or

by looking at still more regional changes, as in Harvey and Kaufmann (2002). The idea would be to increase the amount of constraining data while introducing as few new underdetermined parameters as possible.

Given our most plausible uncertainty assumptions, if the medium-term future temperature increase is to be kept below 2 K with 95% certainty, the data and calculations suggest that equivalent CO₂ levels should be kept below about 475 ppmv. Our calculations suggest that the uncertainty in TCS may be reduced by approximately 45% by the year 2030, and that if the temperature before then does not increase more than currently expected, the target emissions level may rise to 540-ppmv CO₂ equivalent.

The results we obtain are largely independent of those from comprehensive climate models, although we verify that our methodology works in part by comparison with a model (as well as the ability to reproduce the observed twentieth century record), and the forcings shown in Fig. 5 do to some degree involve model calculations. Finally, we emphasize that our results provide upper bounds on emissions targets that constrain temperature increases only in the short and medium term (decadal to century time scale). Because of the heat uptake by the deep ocean, it is likely that if the greenhouse gas levels were to remain indefinitely at some level then the temperature increase at true equilibrium would be much larger than the transient response we calculate for that greenhouse gas level, although this increase may be far in the future. But if anthropogenic emissions were to actually cease and greenhouse gas levels were then to slowly diminish, the long-term response might be more similar to the transient response at the peak greenhouse gas level, although the uncertainty in this regard is large.

Acknowledgments. We thank Jonathan Gregory for providing the forcing used in Gregory and Forster (2008) in digital form, and Isaac Held and Mike Winton for many conversations on this matter. We also thank three anonymous reviewers for their detailed and constructive comments. This work was supported by DOE Grant DE-SC0005189 and the FAA under the Joint University Program.

APPENDIX A

Comparing TCS and TCR

In Fig. A1, we illustrate the (small) differences between TCR and TCS in our model for three different emissions pathways to double CO₂ concentration: an instantaneous doubling, increases at a rate of 1% yr⁻¹, and increases at a rate of 0.7% yr⁻¹, which ensures doubling is achieved in 100 yr. We plot the surface and

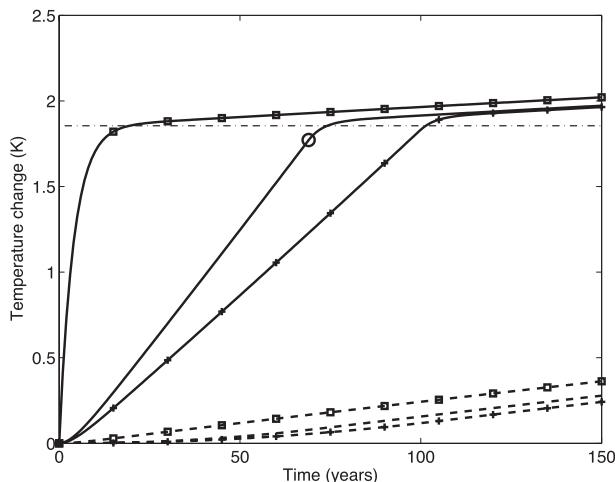


FIG. A1. Surface (solid) and deep ocean (dashed) responses to three different emissions pathways to double CO₂: instantaneous (squares), 1% increases per year compounded continuously (black, with TCR marked as open circle), and 0.7% increases per year such that doubling is achieved in exactly 100 yr (crosses). TCS as computed by Eq. (2.5) is the horizontal line (dashed-dotted).

deep ocean temperature responses to each of the forcing scenarios using the simple two-time-constant system [Eq. (2.1)]. Key parameter values of the model are as follows: ECS = 3.5 K, $\lambda = 2 \text{ W m}^{-2} \text{ K}^{-1}$, deep ocean depth of 5000 m, and mixed layer depth of 60 m. In the first 150 yr of the responses shown, the deep ocean temperature changes remain quite small, increasing to 10% of the ECS or less. The TCS computed by Eq. (2.5) of 1.86 K is marked as the horizontal line. For each of the three emissions pathways, the surface temperature change at the end of the fast transient is almost identical because of the slow deep ocean temperature response. This would be the case for any other emissions pathway, as long as CO₂ doubling is achieved well before the ocean reaches equilibrium. The TCR in this model is 1.77 K and is indicated in Fig. A1 at the time of CO₂ doubling in the 1% yr⁻¹ scenario. The TCR is about 4% smaller than TCS because TCS is evaluated at the end of the fast transient while there may be some committed yet unrealized warming at the time TCR is measured. If, in reality, there is much less separation between the time scale of the mixed-layer response and that of the whole ocean then the TCS and TCR will differ more.

APPENDIX B

Varying Prior Uncertainty

Figures B1 and B2 show the effect the prior uncertainty of λ has on the posterior distributions for the parameters λ and α by the end of a century's worth of

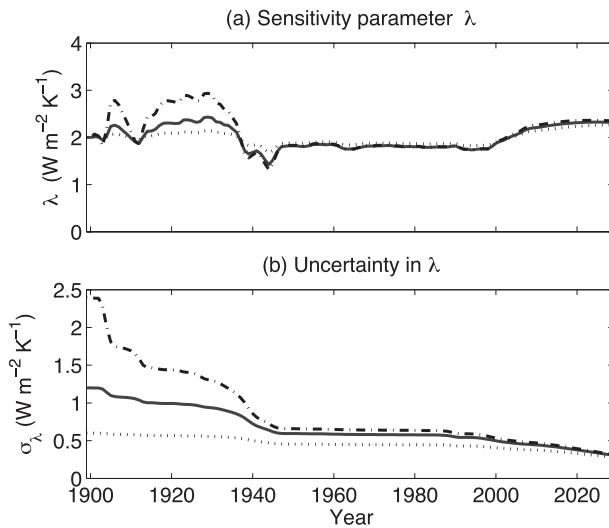


FIG. B1. Estimates of the (a) mean and (b) standard deviation of the transient climate sensitivity parameter λ for three different prior assumptions on uncertainty in λ : $\sigma_\lambda(1900) = 0.6$ (dotted), 1.2 (solid), and 2.4 (dashed-dotted) $\text{W m}^{-2} \text{K}^{-1}$.

data assimilation. The prior in any Bayesian analysis is often difficult to assess and may require some subjective analysis. Our aim is to show that any subjectivity in prior λ distribution does not compromise the objectivity of the data and posterior. We vary the spread of the Gaussian prior to get a sense of how strongly it influences the posterior results (the Kalman filter requires all distributions to be Gaussian so a completely uninformative prior is not an option).

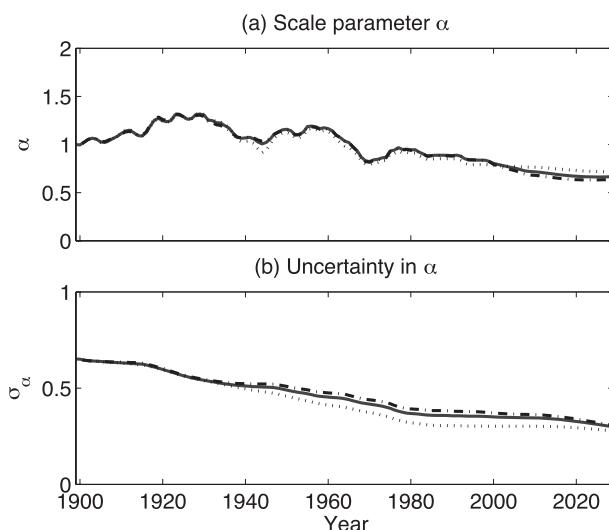


FIG. B2. Estimates of the (a) mean and (b) standard deviation of the forcing scale factor α for the same prior assumptions on uncertainty in λ as in Fig. 16.

As evidenced by Fig. B1, about 50 yr of data contain enough information to effectively forget the prior; the posteriors have converged to the same distribution. The narrowest prior used, $\sigma_\lambda(1900) = 0.6 \text{ W m}^{-2} \text{K}^{-1}$, is probably too small since its posterior standard deviation remains slightly smaller than the others until 2030 (Fig. B1b) and the data hardly reduce its value. In Fig. B2, it is clear that the distribution of the parameter α is not affected by prior assumptions in λ . Note that varying the prior mean λ is addressed in the main text. These experiments give us confidence in our choice of prior uncertainty, $\sigma_\lambda = 1.2 \text{ W m}^{-2} \text{K}^{-1}$, and the objectivity of the final posterior distributions discussed in the main text.

REFERENCES

- Allen, M. R., P. Stott, J. F. B. Mitchell, J. F. B. Schnur, and T. L. Delworth, 2000: Quantifying the uncertainty in forecasts of anthropogenic climate change. *Nature*, **407**, 617–620.
- Annan, J. D., J. C. Hargreaves, R. Ohgaito, A. Abe-Ouchi, and S. Emori, 2005a: Efficiently constraining climate sensitivity with ensembles of paleoclimate simulations. *SOLA*, **1**, 181–184.
- , D. J. Lunt, J. C. Hargreaves, and P. J. Valdes, 2005b: Parameter estimation in an atmospheric GCM using the ensemble Kalman filter. *Nonlinear Processes Geophys.*, **12**, 363–371.
- Baker, M., and G. Roe, 2009: The shape of things to come: Why is climate change so predictable? *J. Climate*, **22**, 4574–4589.
- Brohan, P., J. J. Kennedy, I. Harris, S. F. B. Tett, and P. D. Jones, 2006: Uncertainty estimates in regional and global observed temperature changes: A new dataset from 1850. *J. Geophys. Res.*, **111**, D12106, doi:10.1029/2005JD006548.
- Delworth, T. L., and Coauthors, 2006: GFDL's CM2 global coupled climate models. Part I: Formulations and simulation characteristics. *J. Climate*, **19**, 643–674.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**, 10 143–10 162.
- , 2007: *Data Assimilation: The Ensemble Kalman Filter*. Springer, 280 pp.
- 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.
- Frame, D. J., 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.
- Giorgi, F., and L. O. Mearns, 2002: Calculation of average, uncertainty range, and reliability of regional climate changes from AOGCM simulations via the reliability ensemble averaging method. *J. Climate*, **15**, 1141–1158.
- Greene, A. M., L. Goddard, and U. Lall, 2006: Probabilistic multimodel regional temperature change projections. *J. Climate*, **19**, 4326–4343.
- Gregory, J. M., and P. M. Forster, 2008: Transient climate response estimated from radiative forcing and observed

- temperature change. *J. Geophys. Res.*, **113**, D23105, doi:10.1029/2008JD010405.
- Hansen, J., and Coauthors, 2007: Dangerous human-made interference with climate: A GISS model study. *Atmos. Chem. Phys.*, **7**, 2287–2312, doi:10.5194/acp-7-2287-2007.
- Harvey, L. D. D., and R. Kaufmann, 2002: Simultaneously constraining climate sensitivity and aerosol radiative forcing. *J. Climate*, **15**, 2837–2861.
- Hasselmann, K., 1997: Multi-pattern fingerprint method for detection and attribution of climate change. *Climate Dyn.*, **13**, 601–611.
- Hegerl, G. C., and Coauthors, 2007: Understanding and attributing climate change. *Climate Change 2007: The Physical Science Basis*, S. Solomon et al., Eds., Cambridge University Press, 663–745.
- Held, I., M. Winton, K. Takahashi, T. L. Delworth, F. Zeng, and G. K. Vallis, 2010: Probing the fast and slow components of global warming by returning abruptly to preindustrial forcing. *J. Climate*, **23**, 2418–2427.
- Huybers, P., 2010: Compensation between model feedbacks and curtailment of climate sensitivity. *J. Climate*, **23**, 3009–3018.
- Julier, S., 2002: The scaled unscented transformation. *Proc. American Control Conf.*, Anchorage, AK, Institute of Electrical and Electronics Engineers, Vol. 6, 4555–4559.
- Kiehl, J. T., 2007: Twentieth-century climate model response and climate sensitivity. *Geophys. Res. Lett.*, **34**, L22710, doi:10.1029/2007GL031383.
- Knutti, R., 2008: Why are climate models reproducing the observed global surface warming so well? *Geophys. Res. Lett.*, **35**, L18704, doi:10.1029/2008GL034932.
- , and L. Tomassini, 2008: Constraints on the transient climate response from observed global temperature and ocean heat uptake. *Geophys. Res. Lett.*, **35**, L09701, doi:10.1029/2007GL032904.
- , and Coauthors, 2008: A review of uncertainties in global temperature projections over the twenty-first century. *J. Climate*, **21**, 2651–2663.
- Majda, A., I. Timofeyev, and E. Eijnden, 2001: A mathematical framework for stochastic climate models. *Commun. Pure Appl. Math.*, **54**, 891–974.
- Meehl, G., and Coauthors, 2007: Global climate projections. *Climate Change 2007: The Physical Science Basis*, S. Solomon et al., Eds., Cambridge University Press, 747–845.
- Murphy, D. M., 2010: Constraining climate sensitivity with linear fits to outgoing radiation. *Geophys. Res. Lett.*, **37**, L09704, doi:10.1029/2010GL042911.
- Myhre, G., E. Highwood, K. Shine, and F. Stordal, 1998: New estimates of radiative forcing due to well mixed greenhouse gases. *Geophys. Res. Lett.*, **25**, 2715–2718.
- Pachauri, R. K., and A. Reisinger, Eds., 2007: *Climate Change 2007: Synthesis Report*. Cambridge University Press, 104 pp.
- Randall, D. A., and Coauthors, 2007: Climate models and their evaluation. *Climate Change 2007: The Physical Science Basis*, S. Solomon et al., Eds., Cambridge University Press, 589–662.
- Raper, S. C., J. Gregory, and R. Stouffer, 2002: The role of climate sensitivity and ocean heat uptake on AOGCM transient temperature response. *J. Climate*, **15**, 124–130.
- Roe, G. H., and M. B. Baker, 2007: Why is climate sensitivity so unpredictable? *Science*, **318**, 629–632.
- Schellnhuber, H., W. Cramer, N. Nakicenovic, T. Wigley, and G. Yohe, Eds., 2006: *Avoiding Dangerous Climate Change*. Cambridge University Press, 392 pp.
- 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.
- , J. F. B. Mitchell, M. R. Allen, T. L. Delworth, J. M. Gregory, G. A. Meehl, and B. D. Santer, 2006: Observational constraints on past attributable warming and predictions of future global warming. *J. Climate*, **19**, 3055–3069.
- 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.
- Thompson, D. W. J., J. M. Wallace, P. D. Jones, and J. J. Kennedy, 2009: Identifying signatures of natural climate variability in time series of global-mean surface temperature: Methodology and insights. *J. Climate*, **22**, 6120–6141.
- Vallis, G. K., E. P. Gerber, P. J. Kushner, and B. A. Cash, 2004: A mechanism and simple dynamical model of the North Atlantic Oscillation and annular modes. *J. Atmos. Sci.*, **61**, 264–280.