Compensation between model feedbacks and curtailment of climate sensitivity

Peter Huybers *

* *Corresponding author address:* Harvard University, 20 Oxford St., Cambridge, MA 02138.

E-mail: phuybers@fas.harvard.edu

ABSTRACT

The spread in climate sensitivity obtained from twelve general circulation model runs used in the Fourth Assessment of the Intergovernmental Panel on Climate Change indicates a 95% confidence interval of 2.1°C to 5.5°C, but this reflects compensation between model feedbacks. In particular, cloud feedback strength negatively covaries with the albedo feedback as well as the combined water vapor plus lapse rate feedback. The degree of negative covariance between feedbacks is unlikely to result from chance alone. It may, however, result from the method by which the feedbacks were estimated, physical relationships represented in the models that have not yet been fully diagnosed, or from conditioning the models upon some unknown combination of observations and expectations. In so much as the latter case holds, attempts to test or constrain model results through comparison with observations runs the risk of being circular. Furthermore, this compensation between model feedbacks—when taken together with indications that variations in radiative forcing and the rate of ocean heat uptake play a similar compensatory role in models—suggests systematic curtailment of the inter-model spread in climate sensitivity. If the compensation between feedbacks is removed, the 95% confidence interval for climate sensitivity expands to 1.9°C to 8.0°C. Neither of the qouted 95% intervals adequately reflect the understanding of climate sensitivity, but their differences illustrate that model interdependencies must be understood before model spread can be correctly interpreted.

1. Introduction

Collections of global climate model runs are the backbone of efforts to predict future climate, as most recently represented by the Coupled Model Intercomparison Project 3 (CMIP3) (Meehl et al. 2007) that collected together model runs used in the Intergovernmental Panel on Climate Change Fourth Assessment Report (IPCC AR4). Although these model runs were not designed to span the full range of uncertainty, are not fully independent, and are not identically forced (see Knutti et al. 2008), they do offer some indication of the range of future climate states. If we are to correctly interpret such an ensemble of opportunity, it is first necessary to determine the interdependence between the models and what range of uncertainty is covered by the ensemble.

An important interdependence was identified between the radiative forcing and climate sensitivity across the CMIP3 models by Schwartz et al. (2007). Schwartz et al. (2007) noted that while 20th century changes in radiative forcing differs by a factor of four (0.6 to 2.4 W/m^2, 5-95% confidence limits) across the models, the resulting temperature spread differs by only a factor of two, suggesting compensation between various model components. Kiehl (2007) then presented evidence that this narrow temperature range results from an anticorrelation between radiative forcing and climate sensitivity, and Knutti (2008) demonstrated that this anticorrelation holds for the CMIP3 models in particular. Differences in radiative forcing arise from how aerosols are treated. Thus, the CMIP3 models approximate the 20th century warming through differing balances between radiative forcing and climate sensitivity.

Inter-model compensation between climate sensitivity and radiative forcing (Schwartz et al. 2007; Kiehl 2007; Knutti 2008) underscores that the models are not based purely on theory but are also conditional upon observations and expectations. That is, the models are tuned. It has been argued that the aerosol tuning manifest in the 20th century simulations does not lead to appreciable bias in the spread of future climate predictions because the radiative forcing from atmospheric CO_2 comes to dominate over aerosols in the emissions scenarios (Kiehl 2007; Knutti 2008). But the question arises whether other features of the models are also tuned and whether these influence the spread in climate predictions.

Webb et al. (2006) observed that the radiative forcing associated with a doubling of CO_2 and climate sensitivity are anticorrelated across those models in the Cloud Feedback Model Intercomparison Project (McAvaney and Le Treut 2003). This suggests that the tuning of radiative forcing extends beyond aerosols. Furthermore, Raper et al. (2002) noted that differences in the efficiency of heat uptake across the models in CMIP2—the second Coupled Model Intercomparison Project—leads to a more similar transient climate sensitivity across models than is expected from purely physical considerations. In light of indications that the radiative forcing and ocean heat uptake of models are adjusted so as to narrow the spread in climate sensitivity, here I analyze the variations in the strength of feedbacks across the CMIP3 models to explore whether these also act to curtail the spread in climate sensitivity.

2. Feedbacks and their covariance

Feedbacks are variously defined in the literature, making it useful to recap the notation used here, which follows the standard electronics literature definition. The relationship between changes in radiative forcing and temperature can be represented as a linear feedback system, $\triangle T = \lambda_o \triangle R + f \triangle T$, where perturbations in radiative forcing ($\triangle R$ in units of W/m^2)

lead to direct changes in temperature ($\triangle T$ in units of °C) according to the basic climate sensitivity (λ_o in °C per W/m^2), as well as through feedback factors (f_x, which are unitless). The feedback factors are linearly additive and those associated with water vapor, the vertical lapse rate, albedo, and clouds are considered: $f_{\text{net}} = f_{wv} + f_{lr} + f_a + f_c$. The mean and variance of f_{net} then depends on the joint probability distribution relating each feedback to one another, a topic returned to later. Solving for $\triangle T$ yields the expression,

$$\triangle T = \frac{\lambda_o \triangle R}{1 - f_{\text{net}}}. \tag{1}$$

This representation is based upon the assumption that the change in Earth's temperature can be modeled as a linear perturbation, but the representation obviously breaks down as f_{net} approaches one.

I rely upon the feedbacks estimated for the CMIP3 models by Soden and Held (2006, also see Figure 1 and Table 1 here), where they considered results using the A1B emission scenario. Note that Soden and Held (2006) did not compute a climate sensitivity for the GISS AOM and GISS EH models because these were only run out to 2100 AD: these models are excluded from the present analysis. Soden and Held (2006) define feedback parameters as $\gamma_x = \triangle R_x / \triangle T_x$ and include the basic model response to changes in radiative forcing as a feedback. To convert to the formulation introduced above, basic climate sensitivity is obtained as $\lambda_o = 1/\gamma_p$, where γ_p is Soden and Held's (2006) Plank feedback. The feedback parameters for each model are then obtained as $f_x = \lambda_o \gamma_x$ (see Figure 1 and Table 1, as well as Bony et al. 2006 and Roe and Baker 2007 for a more detailed discussion).

Anti-correlation between the water vapor feedback and the lapse rate feedback is expected on physical grounds (e.g. Cess 1975; Held and Soden 2000). For example, a decrease in the lapse rate (a negative feedback) implies relatively greater warming aloft and, by the Clausius-Clapeyron relationship, more upper tropospheric water vapor (a positive feedback). Thus, as is common, these two feedbacks are added together to form a single water vapor plus lapse rate feedback, f_{wv+lr}. Note, however, that it can be questioned whether the anticorrelation between these feedbacks is an artifact of the models (Bony et al. 2006), possibly as a result of changes in water vapor being poorly represented by models with coarse vertical resolution (Tompkins and Emanuel 2000). Whether other feedbacks ought to covary in one or another direction is less clear, and will be taken up in greater detail below.

The variance in the net feedbacks across the twelve CMIP3 models, $var(f_{\text{net}})$, is 0.0082, whereas the variance in the individual feedbacks are $\text{var}(f_a) = 0.0004$, $\text{var}(f_{wv+lr}) = 0.0014$, and $\text{var}(f_c) = 0.014$. The variance in cloud feedbacks, f_c is almost double the net variance, f_{net}, indicating that the other feedbacks somehow compensate for variability in f_c. Indeed, the cross-correlation between f_c and f_{wv+lr} is is -0.7 and the correlation between f_c and f_a is -0.4 (see Fig. 1). The anti-covariance between f_c and f_a and between f_c and f_{wv+lr} is actually larger than the variance associated with either f_a or f_{wv+lr} individually (see Table 2). If the covariance between individual feedbacks is suppressed and the individual feedback variances simply added together, the variance of f_{net} becomes 0.016, double the value obtained when covariance is included.

Clouds appear the principal source of uncertainty in the models (e.g. Soden and Held 2006), as follows from the variance in f_c being more than an order of magnitude larger than the variance in f_a or f_{wv+lr}, but variance alone is apparently an insufficient description of the uncertainty. The covariance between clouds and the other feedbacks sums to -0.0082. Thus, the cloud covariance compensates for more than half of the cloud variance and, by coincidence, is very nearly equal in magnitude, albeit opposite in sign, to the net variance, var$f_{\rm net}$=0.0082. Thus, both the variance in f_c and the covariance between f_c and other feedbacks are leading order terms in determining $f_{\rm net}$.

Colman (2003a) also collected estimates of climate feedbacks from various models, but because the feedbacks were estimated using a variety of different approaches, their inter-model variance is more difficult to interpret than the results of Soden and Held (2006). It is nonetheless notable that Colman's (2003) results indicate that the inter-model variance of $f_{\rm net}$ is nearly three times larger when the covariance between f_c, f_a, and f_{wv+lr} is suppressed. Apparently, substantial compensation also occurs between the estimated feedbacks in Colman's (2003) collection of models.

3. Climate sensitivity

The CMIP3 ensemble of models is not designed to capture the full range of uncertainty in climate predictions (e.g. Knutti et al. 2008), but it is still instructive to examine the implications of this ensemble for the distribution of climate sensitivity. Climate sensitivity is defined as $\triangle T/\triangle R_{2\times}$, with $\triangle R_{2\times}$ representing the radiative forcing expected from a doubling of atmospheric CO_2. An indication of the distribution of climate sensitivity can be obtained from the distribution of the $f_{\rm net}$.

For illustrative purposes,$f_{\rm net}$ is assumed to follow a normal distribution characterized by the sample mean and variance obtained from the twelve CMIP3 models. (A Lilliefors test for the normality of the twelve net feedback values yields a p-value of 0.34; thus, normality cannot be rejected, but this is a weak result given the small amount of data.) The assumption of normality is not ideal because it implicitly assumes that an infinite climate sensitivity has non-zero probability. Weitzman (2009b) discusses the implications of very large climate sensitivity under more reasonable assumptions regarding the probability distribution, and Frame et al. (2005) discuss how the choice of priors and distributional forms can influence the resulting estimates of climate sensitivity.

Assuming normality, the distribution of f_{net} can be converted into a distribution for climate sensitivity following Roe and Baker (2007) (see Figure 2). The observed mean and variance of the net feedback (mean(f_{net}) = 0.6, $var(f_{\text{net}}) = 0.008$) gives a distribution of the climate sensitivity with a 95% confidence range between 2.0 to 5.5°C, whereas the net variance obtained without feedback covariance (mean(f_{net}) = 0.6, $var(f_{\text{net}}) = 0.016$) gives a range from 1.9 to 8.0°C. The wider distribtuion of climate sensitivity is more consistent with the climateprediction.net results (Stainforth et al. 2005) and parallels how Roe and Baker (2007) estimated uncertainty across the CMIP3 models. Note that the length and fatness of the tail of the climate sensitivity distribution is particularly sensitive to changes in feedback uncertainty because of how feedback variance asymmetrically maps into climate sensitivity (Roe and Baker 2007), with the upper 95% bound increasing by 2.5°C. Often climate sensitivity is reported with a 90% confidence itnerval, but 95% is also a standard statistical choice and is perhaps of greater societal relevance (Weitzman 2009a).

The two distributions of climate sensitivity considered here are illustrative of the importance of the covariance terms but neither are acceptable estimates. They comes with all the limitations of the CMIP3 ensemble, some of which were noted earlier. Additionally, the CMIP3 models are not independent of one another—both specifically (Tebaldi and Knutti 2007) and generally in that the assumptions, numerical approaches, and training of the modelers widely overlap—biasing the feedback variance low relative to that expected from independent realization. Further, the ensemble spread is curtailed by omission of ice shelf, carbon cycle, and other processes; and, arguably, is widened by ignoring many observational constraints upon climate sensitivity (e.g. Edwards et al. 2007; Knutti and Hegerl 2008). Nonetheless, the enormous attention given to the model indications of climate sensitivity and the spread between these predictions, coupled with a sensitivity to the degree of covariance between feedbacks, suggests inquiring into the origins of feedback covariance is worthwhile.

Below I analyze the covariance between cloud and other feedbacks using some simple statistical tests. A more complete analysis would involve diagnosing the origins of feedback covariance within and across multiple general circulation models and would require serious computational and research commitments.

4. Origins of the covariance

There appear four possible explanations for how the overall negative covariance between feedbacks could arise: by chance, because of how the feedbacks are estimated, model parametrization of the physics inherently resulting in negative covariance, or model tuning. These possibilities are not exclusive of one another.

a. Covariance by chance

What are the odds that the covariance observed between the feedbacks is truly zero and merely arises from chance fluctuations? An analytical approach to assessing these odds would involve modeling the covariance matrix and require assumptions regarding the underlying feedback distributions. Instead, it seems preferable to use a bootstrap method that takes advantage of the sample distribution.

Bootstrapping is performed by shuffling the feedbacks across the different models. For example, the NCAR CCSM3 albedo is randomly reassigning to any one of the albedos in the twelve models, including the NCAR CCSM3 model itself. This shuffling preserves the distribution of the feedbacks across models while destroying the expected covariance between different sets of feedbacks (e.g. Chernick 2007), in accord with a null-hypothesis of zero covariance. The covariance matrix associated with the feedbacks is then recomputed from the shuffled feedback matrix, and summing across the rows and columns gives a realization of the net feedback variance. Note that the diagonal of the covariance matrix is unaffected because only covariance, not variance, depends on the ordering the feedbacks.

Repeating the bootstrap procedure 100,000 times indicates a 0.3% probability for variance to be equal to or lower than the observed value of 0.008 by chance alone. It is thus safe to reject the null-hypothesis and conclude that the small variance between model feedbacks arises from an actual negative covariance between the feedbacks. Now the question becomes why such negative covariance exists.

b. Feedback estimation artifacts

The least interesting explanation of the negative covariance between clouds and the other feedbacks is as an artifact of the manner in which the cloud feedback was estimated. The estimates used here (Soden and Held 2006) were acquired using the partial radiative perturbation approach (Wetherald and Manabe 1988; Held and Soden 2000). For each of twelve models, Soden and Held (2006) computed the change in a climate variable—water vapor, vertically average temperature, lapse rate, or albedo as a function of latitude, longitude, and altitude—relative to the change in mean surface temperature between two decade long control periods. The resulting ratios were then multiplied by the partial derivatives of top of the atmosphere radiation with respect to each climate variable, again as a function of latitude, longitude, and altitude, as estimated within the Geophysical Fluid Dynamics Laboratory (GFDL) model to yield sensitivity fields. These fields of radiative sensitivity to temperature changes were then integrated from the surface to tropopause and averaged globally. The partial radiative perturbation approach is less prone to introducing correlation between clouds and other feedbacks than the other commonly used method—the so-called cloud forcing approach—but is by no means guaranteed to be free of artifacts (Aires and Rossow 2003; Soden et al. 2004; Bony et al. 2006).

One issue is that cloud feedback could not be directly estimated because of changes in their vertical overlap (Soden and Held 2006). Cloud feedbacks were instead found as the residual between the estimated net feedback and the individual feedback estimates, $f_c = f_{\text{net}} - f_a - f_{wv+lr}$. Uncertainties in the estimation of these parameters could, in the limit, lead to f_{net} being unrelated to both f_a and f_{wv+lr}, yielding $f_c = -(f_a + f_{wv+lr}) +$

ϵ, where ϵ is uncorrelated with both f_a and f_{wv+lr}. The expected covariances is then, $cov(f_c, f_a) = -var(f_a) - cov(f_a, f_{wv+lr}) = -0.0004 - 0.0004 = -0.0008$; and $cov(f_c, f_{wv+lr}) = -var(f_{wv+lr}) - cov(f_a, f_{wv+lr}) = -0.0014 - 0.0004 = -0.0018$, where the values for the variance and covariance are taken from the sample values (see Table 2). The case of negative sample covariance imposed by the estimation procedure considered here seems an upper bound, and are higher than the errors described by Soden and Held (2006) imply, yet the resulting covariances are still less negative than the sampled values, $cov(f_c, f_a) = -0.0010$ and $cov(f_c, f_{wv+lr}) = -0.0031$. A scenario in which random draws of feedbacks happen to produce negative covariance, as described in the foregoing section, and the estimation procedure then accentuates this negativity cannot be ruled out, but such a compound explanation seems unsatisfying. Other artifacts could also be present, but whose nature is unclear.

It also notable that the manner in which the cloud feedbacks are calculated absorbs *all* processes which influence each model's sensitivity into this term, excepting those feedbacks which are directly estimated (Soden and Held 2006). It is thus not possible to fully determine which model elements contribute to the variance and covariance associated with f_c . Direct estimation of cloud feedbacks would permit more conclusive results.

c. Inherent covariance between feedbacks in the models

If a climate model's response to increased CO_2 can be linearly separated into distinct feedbacks upon temperature, no covariance will exist between these feedbacks, but is such a degree of separability likely?

Colman et al. (1997) analyzed the feedbacks present in a single model and found evidence for significant nonlinearity in the longwave response of lapse rates, clouds, and water vapor to perturbations in sea surface temperature ranging between $\pm 2°C$. Though interactions between feedbacks were not explicitly diagnosed, these radiative nonlinearities suggest that a linear representation of the feedbacks is incomplete and that covariance should be expected between feedbacks (Colman et al. 1997). A more recent study by Colman (2003b) indicates that the strength of feedbacks also varies over the course of the seasons, further supporting the notion of nonlinearity in model responses. Likewise, Aires and Rossow (2003) highlight nonlinear interactions between feedbacks in the context of a simple model using a neural network approach, and Wohlfahrt et al. (2004) discuss interaction between ocean and vegetation feedbacks.

The nonlinearities inherent to the climate system, and the handful of studies addressing changes in the strengths of feedbacks in models, suggest that it is unlikely for any feedback to be truly independent. Yet the general expectation of interaction is distinct from determination of the magnitude or even the expected sign of the relationship between feedbacks. The more poignant question is whether there is a physical basis by which to expect cloud feedbacks to be anticorrelated with the strength of albedo and water vapor feedbacks?

While the controls upon feedbacks have begun to be parsed(e.g. Bony and Dufresne 2005; Webb et al. 2006), there remains substantial uncertainty (e.g. Bony et al. 2006). Physical attribution of the compensation between feedbacks observed across the CMIP3 models does not appear immediately possible. While a physical basis for the anticorrelation between cloud and other feedbacks cannot be ruled out, the absence of a compelling mechanism for this apparent anticorrelation—coupled with indications that other model components display substantial tuning (Schwartz et al. 2007; Kiehl 2007; Knutti 2008)—leaves open the possibility that tuning may account for the observed compensation between feedbacks.

d. Feedback tuning

Covariance could arise between feedbacks from conditioning the models upon observations or expectations. Some sense of how this would work can be expressed in the context of a game, played initially with two six-sided dice. The dice are fair, and comparison of successive values obtained from each throw would show no correlation. But if a rule is made that the value of the dice must sum to, say, 7, then a perfect anti-correlation exists between the acceptable pairs (i.e., 1, 2, 3... for one die and 6, 5, 4... for the other). Now considering introduction of a twelve-sided die and the three dice to sum to fourteen. An expected cross-correlation of -0.7 will then exist between the values of the twelve-sided die and those of either of the six-sided dice, but the values of the two six-sided dice will have no correlation between them. The summation rule forces the two six-sided dice to compensate for the greater range of the twelve-sided die. This illustrates how placing constraints upon the output of a system can introduce covariance between its components.

An analogous situation may hold for the CMIP3 models, with variations in f_{lr} and f_{wv} compensating for the larger variations in f_c. For example, if $\triangle R_{2x}/\triangle T$ is made to have a specific value or range of values, it follows from Eq. 1 that only certain combinations of feedback values will be acceptable [1], $f_c + f_a + f_{lr+wv} = 1 - \lambda_o \triangle R_{2x}/\triangle T$. For admissible

[1]Of course, the magnitude of $\triangle R_{2x}$ or λ_o could also be adjusted, as also seems to have been the case for the CMIP3 models (Schwartz et al. 2007; Kiehl 2007; Knutti 2008), but only feedbacks are focused on here

models, large values of f_c would tend to be associated with smaller values of f_a and f_{lr+wv}.

The degree of covariance expected to arise from tuning[2] the CMIP3 models can be explored using a more detailed version of the dice game. Consider the case in which feedbacks are drawn from a normal distribution having a mean corresponding to the CMIP3 feedbacks (see Table 1) a standard deviation that is twice the observed value—presumably, untuned model parameters will have a wider spread. Model realization are only accepted if they have a climate sensitivity between the lowest (2.3°C for NCAR CCSM3) and highest CMIP3 values (4.2°C for MPI ECHAMl), where climate sensitivity is calculated according to Eq. 1 with a $\triangle R_{2x}$ of 3.7 W/m^2 and a λ_o of 0.31. Drawing 100,000 realizations according to the above rules resulted in 36,430 models having a climate sensitivity between 2.3 and 4.2°C.

The 36,340 acceptable models have feedback covariances and correlations resembling the CMIP3 values (see Table 3). In particular, anti-correlations exist between f_c and f_a of -0.3 and between f_c and f_{lr+wv} of -0.6, leading to more than a factor of three reduction in the variance of f_{net}. The individual magnitudes of the correlations and covariances in the random realizations are similar to those observed for CMIP3 with the one exception that there is no appreciable correlation between f_a and f_{lr+wv}. The observed correlation between f_a and f_{lr+wv} is 0.6 and may arise from a different form of tuning or one of the other mechanisms described earlier.

The basic result is that specifying climate sensitivity to lie within the observed spread across the CMIP3 models is a sufficient explanation for the origins of the compensation

[2]For the present purposes, no distinction is made between only accepting random model realization that meet a certain criteria and actively adjusting model parameters those criteria are met: both are referred to as tuning.

between f_c and the other feedbacks. Note that for such tuning to occur the actual model feedbacks never need to be explicitly calculated. Indeed, there is little reason to believe that model feedbacks are intentionally adjusted to compensate one another. Rather, it seems more likely that the standards by which models are deemed acceptable leads to an implicit tuning of the parameters. In this sense, the covariance between the CMIP3 model feedbacks may be symptomatic of the uneven treatment of outlying model results.

The tuning in the simple example provided here is more explicit than would be expected in actual model development, but procedures similar in spirit are routine. For example, of the 414 stable model versions Stainforth et al. (2005) analyzed, 6 versions yielded a negative climate sensitivity. Those 6 versions appear to have been subjected to greater scrutiny and were excluded because of non-physical interactions between the model's mixed-layer ocean and tropical clouds. Scrutinizing models that fall outside of an expected range of behavior makes them less likely to be included in an ensemble of results and, therefore, is apt to limit the spread of a model ensemble. Even a subtle limitation in climate sensitivity would be expected to introduce covariance between feedbacks.

It should also be noted that tuning is not restricted to objective or even conscious adjustment of parameters to better fit data, but includes the more nebulous adjustment of parameters to fit conceptions of what is correct. Such tuning is difficult to account for because it depends not only on observations, but also on psychological phenomena such as anchoring[3]. While most climate models must be tuned because some model parameters cannot be directly observed nor their values derived from theoretical considerations alone, such

[3]The 3 ± 1.5°C range of climate sensitivity suggested by the ad hoc study group on carbon dioxide and climate (Charney et al. 1979) may be acting as a heavy anchor.

tuning has potentially important consequences for the resulting spread in climate predictions.

5. Discussion and conclusions

Numerical climate models are indispensable tools for predicting climate. If we are to correctly interpret their results and optimally design future model studies, we must carefully track what assumptions and observations are incorporated into them. Evidence has accumulated that inter-model differences in climate forcing (Webb et al. 2006; Schwartz et al. 2007; Kiehl 2007; Knutti 2008), ocean heat uptake (Raper et al. 2002), and the individual feedbacks that contribute to climate sensitivity (this study) act to reduce the spread in warming realized across models. These compensating model features may have a sound physical basis, but the specter of systematic curtailment of the inter-model spread in climate sensitivity is difficult to dismiss.

Knutti (2008) argued that parameter covariance across models are neither unexpected nor problematic if the models are interpreted as conditional on observations. A problem does arise, however, when model results are used in conjunction with observations to constrain climate sensitivity (see the reviews by Edwards et al. 2007; Knutti and Hegerl 2008), as this runs the risk of doubly calling upon the data. Furthermore, comparison between model results and climate of the 20th century may then be circular (also see Rodhe et al. 2000). Ultimately, we need to know what exactly goes into models, if we are to correctly interpret the output.

While it seems a large undertaking, a more objective tuning approach may be warranted. Standard data sets could be agreed upon for tuning climate models, with other data explicitly withheld for testing. Or perhaps a more readily undertaken course of action is to test model results against less closely monitored aspects of the climate, such as the amplitude (Knutti et al. 2006) or phase (Stine et al. 2009) of the seasonal cycle in surface temperature. The paleoclimate record is also useful in this manner (e.g. Braconnot et al. 2007), in that it can be more safely assumed that models have not been tuned to reproduce these more distant and, during many epochs, dramatically different climates. Convergence between model results, if not truly driven by a decrease in model uncertainty or clearly understood as a result of tuning, could have the unfortunate consequence of lulling us into too great a confidence in model predictions or inferences of too narrow a range of future climates.

As a final note, the CMIP3 archive can be characterized as an ensemble of opportunity, not specifically designed to span the range of uncertainty in future climates. A better indication of the range of possible future climates may be obtained through more exhaustive searches of the behavior of simpler models under perturbation of their parameters (e.g. Stainforth et al. 2005). It may also be sensible to push the most sophisticated models towards generating realizations of future climate which are as *inconsistent* as possible with current predictions, while still being physically sound. Focusing on maximally inconsistent possibilities seems more likely to lead to scientific discoveries and to uncover climatic surprises[4]. A maximally inconsistent ensemble of state-of-the-art model realizations would also have the advantage of suggesting outer bounds upon the range of climate sensitivity and, therefore, be complimentary to existing estimates.

[4]Such promotion of scientific discord may be contrasted with the IPCC process, which tends to prize consensus, albeit for political rather than scientific reasons.

Acknowledgments.

Without implying their agreement, James Hammit, Gerard Roe, Alexander Stine, Martin Weitzman, and Carl Wunsch are thanked for the discussion and insightful suggestions they provided.

REFERENCES

Aires, F. and W. Rossow, 2003: Inferring instantaneous, multivariate and nonlinear sensitivities for the analysis of feedback processes in a dynamical system: Lorenz model case-study. *Quarterly Journal of the Royal Meteorological Society*, **129 (587)**.

Bony, S. and J. Dufresne, 2005: Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models. *Geophys. Res. Lett*, **32**, 20.

Bony, S., et al., 2006: How well do we understand and evaluate climate change feedback processes? *Journal of Climate*, **19 (15)**, 3445–3482.

Braconnot, P., et al., 2007: Results of PMIP2 coupled simulations of the Mid-Holocene and Last Glacial Maximum–Part 1: experiments and large-scale features. *Climate of the Past*, **3 (2)**, 261–277.

Cess, R., 1975: Global climate change- An investigation of atmospheric feedback mechanisms. *Tellus*, **27 (3)**, 193–198.

Charney, J., A. Arakawa, D. Baker, B. Bolin, R. Dickinson, R. Goody, et al., 1979: Carbon dioxide and climate: a scientific assessment. *Washington, DC: National Academy of Sciences.*

Chernick, M., 2007: *Bootstrap methods: a guide for practitioners and researchers.* Wiley-Interscience.

Colman, R., 2003a: A comparison of climate feedbacks in general circulation models. *Climate Dynamics*, **20 (7)**, 865–873.

Colman, R., 2003b: Seasonal contributions to climate feedbacks. *Climate Dynamics*, **20 (7)**, 825–841.

Colman, R., S. Power, and B. McAvaney, 1997: Non-linear climate feedback analysis in an atmospheric general circulation model. *Climate Dynamics*, **13 (10)**, 717–731.

Edwards, T., M. Crucifix, and S. Harrison, 2007: Using the past to constrain the future: how the palaeorecord can improve estimates of global warming. *Progress in Physical Geography*, **31 (5)**, 481.

Frame, D., B. Booth, J. Kettleborough, D. Stainforth, J. Gregory, M. Collins, and M. Allen, 2005: Constraining climate forecasts: The role of prior assumptions. *Geophysical Research Letters*, **32 (9)**, L09 702.

Held, I. and B. Soden, 2000: Water Vapor Feedback and Global Warming. *Annual Review of Energy and the Environment*, **25 (1)**, 441–475.

Kiehl, J., 2007: Twentieth century climate model response and climate sensitivity. *Geophys Res Lett*, **34**, L22 710.

Knutti, R., 2008: Why are climate models reproducing the observed global surface warming so well? *Geophys. Res. Lett*, **35**.

Knutti, R. and G. Hegerl, 2008: The equilibrium sensitivity of the Earth's temperature to radiation changes. *Nature Geoscience*, **1 (11)**, 735–743.

Knutti, R., G. Meehl, M. Allen, and D. Stainforth, 2006: Constraining climate sensitivity from the seasonal cycle in surface temperature. *Journal of Climate*, **19 (17)**, 4224–4233.

Knutti, R., et al., 2008: A review of uncertainties in global temperature projections over the twenty-first century. *Journal of Climate*, **21 (11)**, 2651–2663.

McAvaney, B. and H. Le Treut, 2003: The cloud feedback model intercomparison project: CFMIP. *CLIVAR Exchanges*, **26**, 1–4.

Meehl, G., C. Covey, T. Delworth, M. Latif, B. McAvaney, J. Mitchell, R. Stouffer, and K. Taylor, 2007: The WCRP CMIP3 multimodel dataset. *Bull. Am. Meteorol. Soc*, **88**, 1383–1394.

Raper, S., J. Gregory, and R. Stouffer, 2002: The role of climate sensitivity and ocean heat uptake on AOGCM transient temperature response. *Journal of Climate*, **15 (1)**, 124–130.

Rodhe, H., R. Charlson, and T. Anderson, 2000: Avoiding circular logic in climate modeling. *Climatic Change*, **44 (4)**, 419–422.

Roe, G. and M. Baker, 2007: Why Is Climate Sensitivity So Unpredictable? *Science*, **318 (5850)**, 629.

Schwartz, S., R. Charlson, and H. Rodhe, 2007: Quantifying climate changetoo rosy a picture? *Nature Reports Climate Change*, 23–24.

Soden, B., A. Broccoli, and R. Hemler, 2004: On the use of cloud forcing to estimate cloud feedback. *Journal of Climate*, **17 (19)**, 3661–3665.

Soden, B. and I. Held, 2006: An assessment of climate feedbacks in coupled ocean–atmosphere models. *Journal of Climate*, **19 (14)**, 3354–3360.

Stainforth, D., et al., 2005: Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature*, **433**, 403–406.

Stine, A., P. Huybers, and I. Fung, 2009: Changes in the phase of the annual cycle of surface temperature. *Nature*, **457 (7228)**, 435–440.

Tebaldi, C. and R. Knutti, 2007: The use of the multi-model ensemble in probabilistic climate projections. *Philosophical Transactions A*, **365 (1857)**, 2053.

Tompkins, A. and K. Emanuel, 2000: The vertical resolution sensitivity of simulated equilibrium temperature and water-vapour profiles. *Quarterly Journal of the Royal Meteorological Society*, **126 (565)**.

Webb, M., et al., 2006: On the contribution of local feedback mechanisms to the range of climate sensitivity in two GCM ensembles. *Climate Dynamics*, **27 (1)**, 17–38.

Weitzman, M., 2009a: On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, **91 (1)**, 1–19.

Weitzman, M. L., 2009b: Additive damages, fat-tailed climate dynamics, and uncertain discounting. *Economics: The Open-Access, Open-Assessment E-Journal*, **3 (2009-26)**, URL `http://www.economics-ejournal.org/economics/discussionpapers/2009-26`.

Wetherald, R. and S. Manabe, 1988: Cloud feedback processes in a general circulation model. *Journal of the Atmospheric Sciences*, **45 (8)**, 1397–1416.

Wohlfahrt, J., S. Harrison, and P. Braconnot, 2004: Synergistic feedbacks between ocean and vegetation on mid-and high-latitude climates during the mid-Holocene. *Climate Dynamics*, **22 (2)**, 223–238.

List of Tables

TABLE 1. Columns are the basic response of the system to a change in radiative forcing, λ_o; the albedo, cloud, and combined lapse rate plus water vapor feedback; and the sum of the feedbacks. All values are adapted from Soden and Held (2006, Table 1).

model	λ_o	albedo	clouds	wv+lr	net feedback
NCAR CCSM3	0.31	0.11	0.04	0.33	0.49
GISS ER	0.31	0.05	0.20	0.25	0.49
NCAR PCM1	0.31	0.11	0.06	0.34	0.51
MRI	0.31	0.08	0.07	0.37	0.53
INMCM3	0.31	0.10	0.11	0.33	0.54
GFDL CM2-1	0.31	0.06	0.25	0.26	0.58
GFDL CM2-0	0.31	0.10	0.21	0.32	0.63
CNRM	0.31	0.10	0.25	0.29	0.64
UKMO HADCM3	0.31	0.07	0.34	0.29	0.70
IPSL	0.32	0.07	0.33	0.31	0.70
MIROC MEDRES	0.32	0.10	0.34	0.28	0.72
MPI ECHAM5	0.31	0.09	0.37	0.27	0.73

Table 2. The covariance between feedbacks multiplied by 10,000, the sums of variance (right column and bottom row), and the net variance (bottom right entry). All variances and covariances are multiplied by 10,000. Also shown in parenthesis are the cross-correlations between pairs of feedbacks. Note that albedo and the combined water vapor plus lapse rate feedback each have a covariance with the cloud feedback that exceeds their individual variance: models with higher albedo or combined lapse rate plus water vapor feedbacks thus actually tend to have a lower climate sensitivity.

	albedo	clouds	wv+lr	net
albedo	4 (1)	-10 (-0.4)	4 (0.6)	-1
clouds	-10 (-0.4)	139 (1)	-31 (-0.7)	97
wv+lr	4 (0.6)	-31 (-0.7)	14 (1)	-13
net	-1	97	-13	82

TABLE 3. Similar to Table 2, but for random models that are only accepted if they have a climate sensitivity ranging between 2.3 and 4.2°C.

	albedo	clouds	wv+lr	net
albedo	16 (1)	-13 (-0.3)	-1 (0.0)	-1
clouds	-13 (-0.3)	98 (1)	-44 (-0.6)	40
wv+lr	-1 (0.0)	-44 (-0.6)	50 (1)	4
net	-1	40	4	46

List of Figures

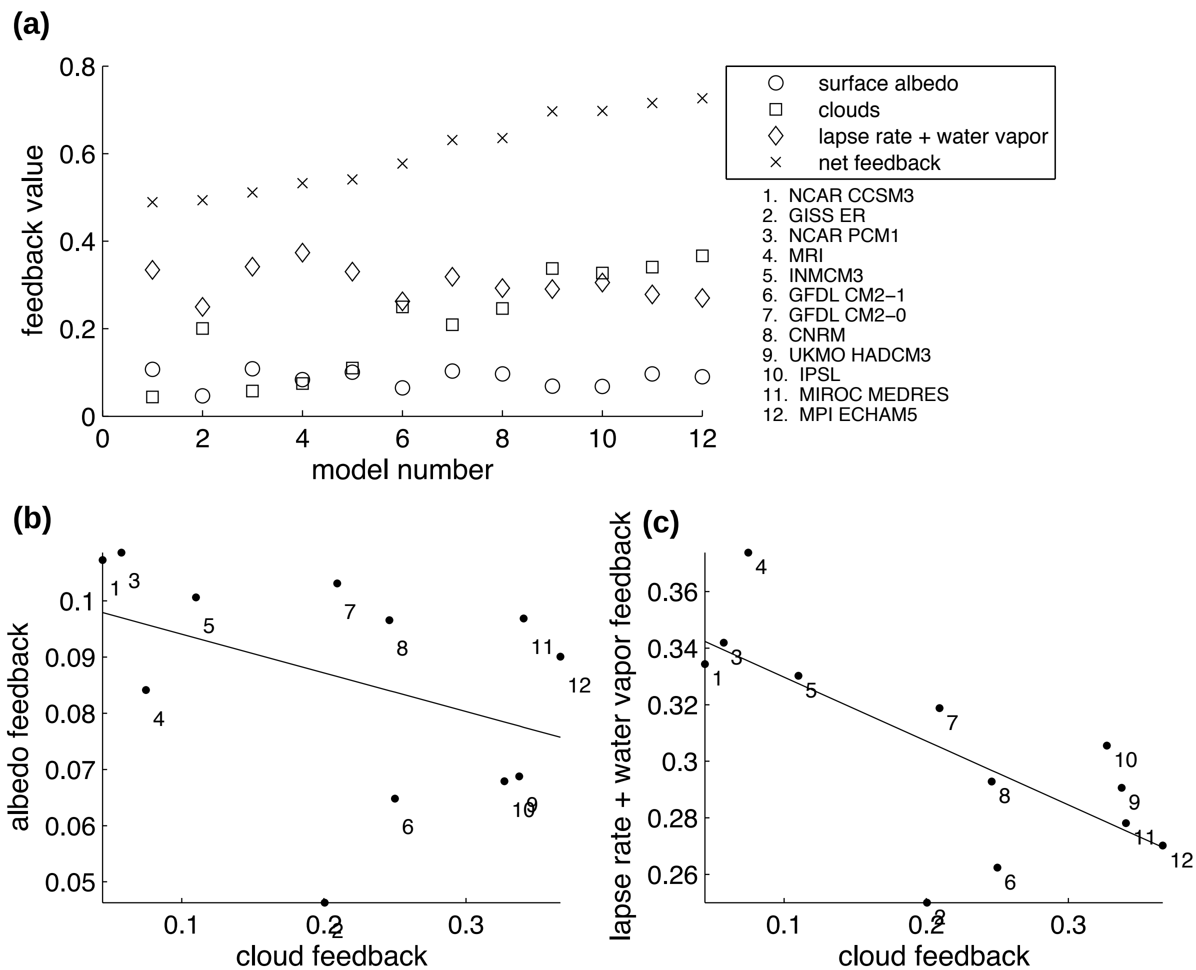

FIG. 1. Feedback values from the CMIP3 collection of models (Soden and Held 2006). **(a)** The individual and net feedback factors for twelve climate models, ordered according to the strength of the net feedback. The cloud feedback plotted against **(b)** the albedo feedback and **(c)** the combined lapse rate and water vapor feedback.

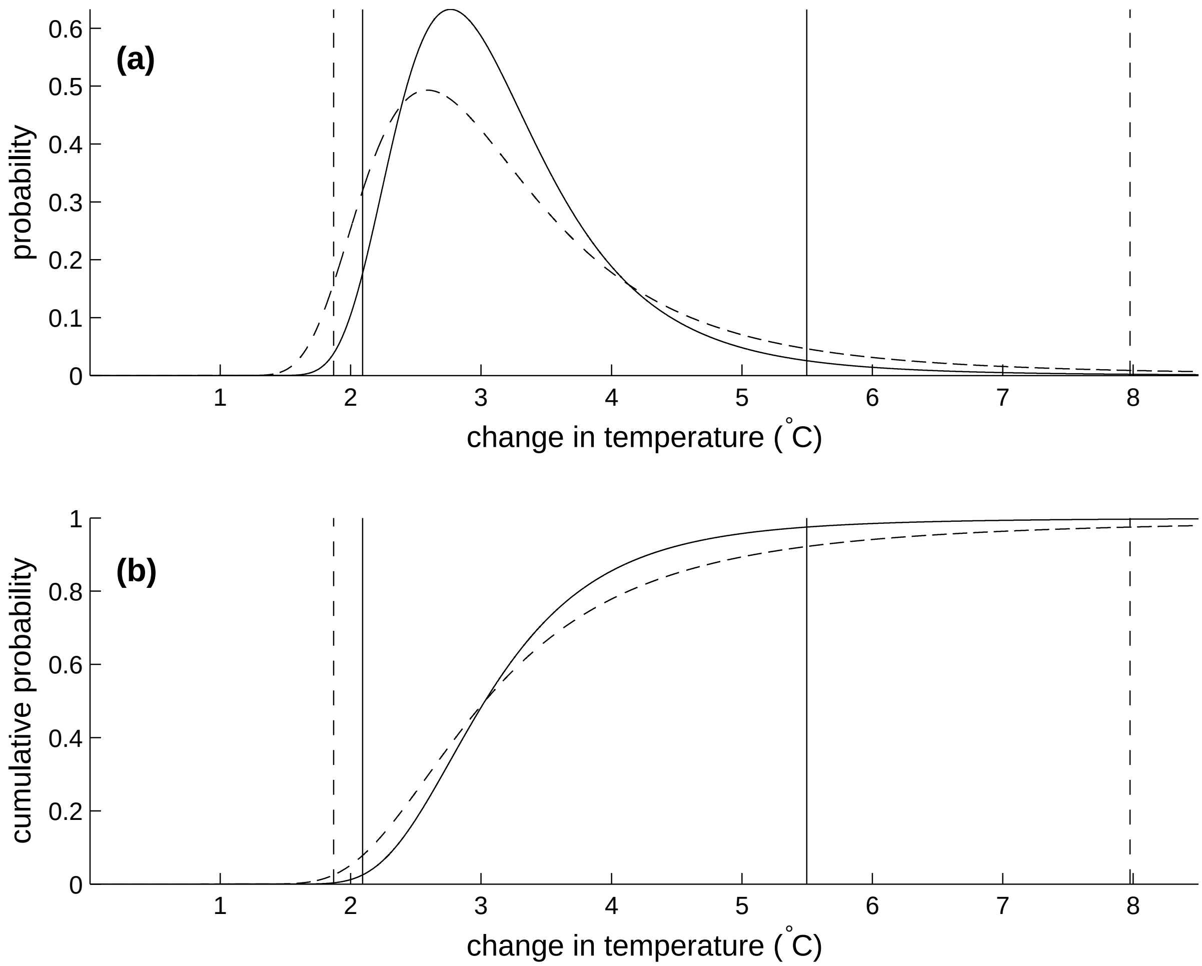

FIG. 2. Climate sensitivity distribution. **(a)** The probability distribution density function for climate sensitivity associated with a mean feedback of 0.6 and a variance of 0.008 (solid lines) or 0.016 (dashed lines). The higher variance results from assuming that the cloud, albedo, and combined water vapor and lapse rate feedbacks are independent. Vertical lines indicate the 95% intervals for each distribution. The positive skew of the probability distribution leads a larger 2.5°C shift in the upper 95% bound. **(b)** Similar to (a) but for the cumulative probability.