On the generation and interpretation of probabilistic estimates of climate sensitivity

J.D. Annan[1] and J.C. Hargreaves
Frontier Research Center for Global Change
Japan Agency for Marine-Earth Science and Technology
3173-25 Showa-machi Kanazawa-ku
Yokohama, Kanagawa 236-0001, Japan

28 May 2009

[1]email: jdannan@jamstec.go.jp

Abstract

The equilibrium climate response to anthropogenic forcing has long been one of the dominant, and therefore most intensively studied uncertainties, in predicting future climate change. As a result, many probabilistic estimates of the climate sensitivity (S) have been presented. In recent years, most of them have assigned significant probability to extremely high sensitivity, such as $P(S > 6C) > 5\%$.

In this paper, we investigate some of the assumptions underlying these estimates. We show that the popular choice of a uniform prior has unacceptable properties and cannot be reasonably considered to generate meaningful and usable results. When instead reasonable assumptions are made, much greater confidence in a moderate value for S is easily justified, with an upper 95% probability limit for S easily shown to lie close to 4°C, and certainly well below 6°C. These results also impact strongly on projected economic losses due to climate change.

1 Introduction

The long-term response of the climate system to anthropogenic forcing, traditionally expressed as the equilibrium sensitivity (S) of the globally-averaged temperature to a doubling of the atmospheric concentration of CO_2, has long been considered as having great significance in terms of our understanding of the climate system. A number of estimates have been presented over recent decades, perhaps the most famous being the assessment of the NRC (1979) that S was believed to lie in the range of 1.5–4.5°C, with that rather vague statement later formally presented as representing a probability somewhere in the range of 66–90% (Houghton et al., 2001). More recently, a proliferation of probabilistic estimates explicitly based on calculations using observational data have also been presented (eg Andronova and Schlesinger, 2001; Gregory et al., 2002; Forest et al., 2002; Hegerl et al., 2006). Many of these results suggest a worryingly high probability of high sensitivity, such as $P(S > 6^{\circ}\mathrm{C}) > 5\%$ (Solomon et al., 2007, Box 10.2). The focus of this paper is to discuss some of the assumptions underlying these estimates, and implications for users.

To avoid possible misunderstandings, we establish at the outset that the notion of probability discussed here is the standard Bayesian paradigm of probability as the subjective degree of belief of the researcher in a proposition (Bernardo and Smith, 1994). In order to calculate a posterior probability distribution function (pdf) for S in the light of observational evidence O, $f(S|O)$, two inputs are required: our prior belief about $f(S)$, and the likelihood $f(O|S)$ which depends on our analysis of the observations which are used. These inputs are combined through Bayes' Theorem:

$$f(S|O) = f(O|S)f(S)/f(O).$$

While the Bayesian approach is not the only possible paradigm for the treatment of epistemic uncertainty in climate science (eg Kriegler, 2005), it appears to be the dominant one in the literature. We do not wish to revisit the wider debate concerning the presentation of uncertainty in climate science (eg Moss and Schneider, 2000; Betz, 2007; Risbey, 2007; Risbey and Kandlikar, 2007) but merely note that despite this debate, numerous authors have in fact presented precise pdfs for climate sensitivity, and furthermore their results are frequently used as inputs for further economic and policy analyses (eg Yohe et al., 2004; Meinshausen, 2006; Stern, 2007; Harvey, 2007). The IPCC 4th Assessment report (Solomon et al., 2007) remains somewhat vague about climate sensitivity, only making the comment that S is likely ($> 66\%$) to lie in the range 2–4.5°C and very unlikely ($< 10\%$) to lie below 1.5°C (a marginal change from previous assessments) but with no further detailed quantitative assessment. In the absence of an authoritative and widely-accepted pdf for S, researchers have generally either used some parameterised distribution fitted to the IPCC's statements, or else sourced estimates directly from the relevant literature.

Given the subjective nature of the assumptions underlying all of these analyses, it is important to properly understand the relationship of the results to the assumptions that underly them. The goal of this paper is to explore one aspect in particular which has received rather limited (and in our opinion rather confusing) treatment in the literature — that is, the choice of prior $f(S)$. In the following section, we outline the methods used in this paper. Then we investigate the prior assumptions which underly most of the results in the literature. These results illustrate serious shortcomings in the approach that has been most widely used. We next present some alternatives that we consider to be more reasonable and practically usable, and then investigate the robustness of our results. Finally, we conclude with some suggestions that researchers may wish to consider in their future work in this area.

2 Methods

We present results both in terms of probabilistic predictions of climate change under a given emissions scenario, and also in terms of the economic implications of this climate change. Since the only climate parameter we are considering in this paper is the equilibrium climate sensitivity, we focus our attention to a long-term stabilisation scenario in which the atmospheric CO_2 concentration is stabilised at 550ppm (double the pre-industrial level, corresponding to the IPCC S550 scenario (Houghton et al., 1996)) and the climate system comes into equilibrium with this forcing. Thus, the global mean temperature change in the long term is precisely the value of the climate sensitivity. To examine the economic consequences of climate change, we use the damage function of the DICE model (Nordhaus, 2008) to provide an estimate of the consequential economic loss due to climate change on the global scale. According to this model, climate change is estimated to cause a loss of the form $C(T) = 0.284T^2$ where $C(T)$ gives the loss in percentage of global GDP as a function of the global mean temperature change T. For this function, a warming of 4°C causes a loss of about 4.5%, which is near the top end of the IPCC range of 1-5% loss for this amount of warming (Parry et al., 2007, p17). At 6°C warming, the cost according to DICE exceeds 10% of global GDP, which is also at the higher end of the range of likely damages estimated by the Stern Review (Stern, 2007, Fig 6.6). As the cited reviews indicate, there is substantial uncertainty among alternative damage functions, although they also suggest that our choice is towards the pessimistic end of the range supported by the literature and could easily be overstated by a factor of 2 or more. Thus, we do not claim here to provide a comprehensive probabilistic assessment of the economic harm associated with climate change. Rather, the economic assessment is intended to demonstrate how the details of the choice of prior may have a substantial downstream impact on users of climate science information. Our results are qualitatively insensitive to the particular choice of the DICE model as the basis for the economic cost. Alternative economic damage functions which also show substantial relative rises in harm across moderate to high temperature rises (say 3–10°C) would support a qualitatively similar analysis and conclusions.

The quadratic curve of the DICE model must be truncated at $T \simeq 19$°C where a complete destruction of the world economy is implied. Extrapolation of this curve to such high temperatures is of course of dubious validity, but our results are qualitatively insensitive to the precise details.

We ignore complications such as the much-debated discount rate, and also risk-aversion (as could be expressed by performing the analysis in terms of a nonlinear utility function of global GDP), as outside the scope of this paper. Instead we directly present the results of our economic analyses as an expected loss of global GDP in percentage terms. For our $2\times CO_2$ scenario, the expected cost of climate change (relative to a pre-industrial baseline), given a probabilistic estimate of the climate sensitivity $f(S)$, is therefore given by

$$\int f(S)C(S)\,dS.$$

We note that, for a more pessimistic stabilisation at 1000ppm CO_2 corresponding to the IPCC S1000 scenario, the temperature rise would be virtually double the value of the climate sensitivity, and the corresponding economic damage would be nearly four times greater than the estimates we provide here.

3 Prior beliefs

According to Bayes Theorem as written in the Introduction, a posterior distribution for S in the light of observational evidence O, which we write as $f(S|O)$, logically requires a prior distribution $f(S)$ which represents the researcher's beliefs concerning S in the absence of these observations. Note that this does not actually require a chronological relationship between the formation of the prior, observation of the data, and calculation of the posterior, although such a relationship may exist. Determining a suitable prior is potentially challenging given that in many or even most cases, the researcher is already aware of the broad implications of the data before the detailed quantitative analysis is undertaken. It is important to be aware of the risk of double-counting the data by accounting for it both in the prior and again through the likelihood, as committing this error would result in over-confident estimates.

3.1 Ignorant priors

In an attempt to avoid the risk of double-counting evidence, researchers have often chosen to use a prior which is uniform in S, which has been described as encapsulating no knowledge about S (eg Lee et al., 2005; Frame et al., 2005). However, it must be recognised that in fact there can be no prior that genuinely represents a state of complete ignorance (Bernardo and Smith, 1994, Section 5.4), and indeed the impossibility of representing true ignorance within the Bayesian paradigm is perhaps one of the most severe criticisms that is commonly levelled at it (eg Walley, 1991, p234). Any proper prior for S must assign a specific level of belief to the proposition that $S > 6°\text{C}$, and it also cannot help but imply a specific prior expectation of loss under the climate change scenario presented in Section 2. Furthermore, the uniform priors which have been widely used represent beliefs that in our opinion are extreme and difficult to justify. For example, the uniform prior U[0C,20C] of Frame et al. (2005) actually represents a prior belief that S is "likely" (70% probability) greater than 6°C, with a mean value for S of 10°C and a 50% probability of exceeding this figure. Even when truncated to U[0C,10°C] as in Hegerl et al. (2006) and Solomon et al. (2007, Figure 9.20), this uniform prior still represents the belief that $P(S > 6°\text{C}) = 40\%$, and also that S is more than twice as likely to lie outside the conventional 1.5–4.5°C "likely" range, as inside it. Perhaps more importantly, the probability assigned to high values of S has a dominant effect on the expected cost of climate change, due to the strong rise in economic losses that accompany increases in temperature. Using the stabilisation scenario and economic analysis described in Section 2, the U[0,10°C] prior implies an expected loss of 9.5% of GDP, with the 40% probability assigned to the interval 6°C < S < 10°C accounting for 87% of that total loss. The broader prior U[0,20°C] implies a massive expected loss of 37% of GDP. Again, the high end of the sensitivity range dominates the loss calculation. It is clear from examining those economic calculations that even low probabilities of high values for for S can have a very strong influence on economic decision making. In fact, changing a deterministic estimate for S from 2°C to 3°C (which covers the range of values where most analyses agree on a highest likelihood for S) has a smaller effect on the total expected loss, than changing our probability of 6°C < S < 20°C by a mere 3%, distributed uniformly across that range. Thus, the notion that such priors (or indeed any prior) can encapsulate the notion of "ignorance" may be superficially attractive but is surely not defensible in detail.

For illustration of how these prior beliefs feed through into posterior estimates of climate sensitivity, we consider the ERBE data which were recently analysed by Forster and Gregory (2006) (henceforth FG). Their analysis is based on fairly recent observational data, and does not rely on the analysis of climate model output, which is valuable in ensuring that we are not double counting data in the analysis presented later. First, we note that their analysis can be interpreted (and was presented by the authors) as generating a likelihood function which is Gaussian in radiative feedback $L = 1/S$. While the authors did not actually present a fully Bayesian analysis of their data, the use of regression-based estimates as likelihood functions is extremely common in climate science and elsewhere, and has some theoretical justification (eg Leroy, 1998). Forster and Gregory's Gaussian regression-based estimate for $L = 2.3 \pm 1.4$ $\mathrm{Wm}^{-2}\mathrm{K}^{-1}$ (at 2 standard deviations) therefore naturally translates into a likelihood function $f(O|L) \propto e^{\frac{-(2.3-L)^2}{2\times 0.7^2}}$. Although the authors acknowledge that their estimate, based on interannual variability, may exclude some feedbacks that are relevant to longer term global warming, they present analysis and arguments that their estimate is accurate to within its stated uncertainties. To further demonstrate the robustness of the results presented here we report on some sensitivity analyses on this likelihood later in the paper.

Figure 1 shows results obtained when different uniform priors are updated with this likelihood function. We have selected two priors that have been used in the recent literature; U[0,10°C] (Hegerl et al., 2006; Solomon et al., 2007), U[0,20°C] (Frame et al., 2005). In contrast to the upper bound, the lower bound of 0°C on the prior is uncontentious: a negative value of sensitivity implies an unstable climate system, and so long as the prior's lower bound is not set so high as to actually exclude reasonable values of S, all results obtained are insensitive to the details of that choice.

In both of these cases, the likelihood we use is exactly the same function presented by FG. That is, the difference in results here is entirely due to the choice of upper bound on the prior, rather than anything relating to the observations or their interpretation. These results, along with others discussed in this paper, are also summarised in Table 1.

It would, we argue, be difficult to claim on theoretical or scientific grounds that either of these prior choices was more objective or defensible than the other one, or any other uniform prior with a different upper bound. Indeed we are unaware of any such arguments in the literature. However, the posterior pdf, and in particular the upper 95% probability threshold, differs greatly between these two results. We can easily explain this by examining the shape of the likelihood function in more detail. A Gaussian likelihood in feedback space has the inconvenient property that $f(O|L=0)$ is strictly greater than zero, and so for all large S, $f(O|S)$ is bounded below by a constant. Therefore, the integral of this function (with respect to S) is unbounded. This means that if the improper semi-unbounded uniform prior $U[0,\infty]$ is used, no proper posterior pdf results. Thus it is necessary to impose bounds on the uniform prior — it would be more appropriate to describe it as *a* uniform prior rather than *the* uniform prior — and the results are strongly dependent on where this upper bound is placed. The inverse Gaussian shape of FG's likelihood function is broadly similar in shape to the marginal likelihood functions for sensitivity that have been obtained from a wide variety of investigations including those cited previously. There are fundamental physical reasons for the skewed shape which have been well understood for many years (Hansen et al., 1985). Therefore, this strong sensitivity to the upper bound of the prior is a generic feature which applies much more widely than just to this specific case.

The expected cost of our stabilisation scenario is also strongly affected by the choice of prior, changing from 3% of global GDP for the U[0,10] prior to 7% for U[0,20]. Increasing the upper bound on the prior still further would generate an even larger posterior expected loss, which would actually tend

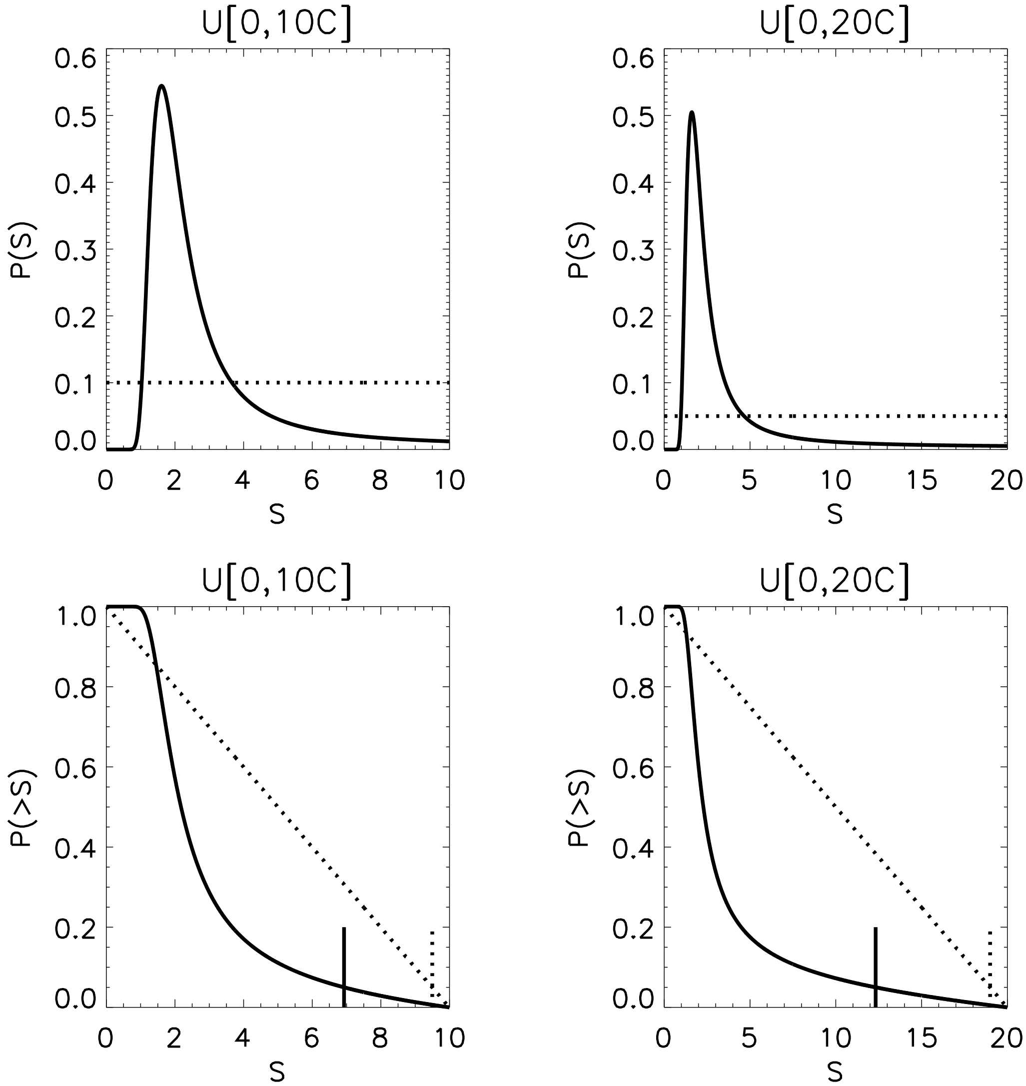

Figure 1: Effect of using different bounds on a uniform prior. Upper row shows marginal pdfs, lower row shows cumulative pds. Dotted lines indicate the priors, solid lines indicates the posterior pdf after updating with the likelihood function of Forster and Gregory (2006). Vertical lines in the lower row indicate the upper 95% bound in each case.

Inputs		Upper 95% limit ($^{\circ}C$)		Cost of $2\times CO_2$ (% GDP)	
Prior	Likelihood	Prior	Posterior	Prior	Posterior
U[0,10]	FG	9.5	6.9	9.5	3.1
U[0,20]	FG	19	12.3	37.4	7.1
Webster	FG	5	3.6	2.3	1.5
Cauchy (2.5,3)	FG	14.5	4.2	9.7	2
Cauchy (3,3)	FG	14.6	4.6	10.4	2.4
Cauchy (3,6)	FG	19.6	5.1	13.9	2.6
Cauchy (3,12)	FG	26.2	5.7	18.8	3.1
Webster	FG+50%	5	4.4	2.3	1.9
Cauchy (2.5,3)	FG+50%	14.5	6.8	9.7	4.4
Cauchy (3,6)	FG+50%	19.6	9.1	13.9	6.4
Cauchy (2.5,3)	Hegerl+	14.5	4.5	9.7	2.4
Cauchy (2.5,3)	AH+	14.5	4.2	9.7	2.8
Cauchy (2.5,3)	Hegerl+, FG	14.5	3.5	9.7	1.6
Cauchy (2.5,3)	AH+, FG	14.5	3.7	9.7	2.0

Table 1: Selected data from prior and posterior analyses discussed in the text. Upper 95% probability limit for S, and expected cost of $2\times CO_2$ scenario are shown. Location and scale parameters of Cauchy distribution are given in parentheses. FG indicates the likelihood from Forster and Gregory (2006), FG+50% indicates 50% increase in the width of their likelihood function. Hegerl+, AH+ are likelihoods based on Hegerl et al. (2006) and Annan and Hargreaves (2006) respectively, with minor modifications as discussed in the text.

to 100% as the bound increases. As a result, choosing between these uniform priors (or any other) could be expected to have strong implications for policy decisions. For example, if stabilising CO_2 at 275ppm compared to 550ppm was estimated to cost an additional 5% of GDP, then switching the prior between U[0,10°C] and U[0,20°C] would reverse the economically optimal decision. Again we emphasise that this difference in outcomes is entirely due to an arbitrary decision to truncate the prior at different points. Moreover, the qualitative behaviour of the specific results presented here will apply to *any* analysis based on a likelihood function that is approximately Gaussian in feedback (or indeed, more generally, bounded away from zero at $L = 0$). Thus we believe the vast bulk of the results in the published literature are sensitive in this way to the arbitrarily-selected upper bound on the uniform prior. Even though we can expect estimates to converge on the true value of the climate sensitivity as more data are obtained and analysed, at any given time our likelihood functions will still be bounded above zero at $L = 0$. That is, if we start with a sufficiently wide uniform prior, analyses of this nature will always present us with an expectation of disaster. In the next section, we discuss alternatives to uniform priors.

3.2 Expert priors

Having demonstrated how the widely-used approach of a uniform prior fails to adequately represent "ignorance" and generates rather pathological results which depend strongly on the selected upper bound, we now consider how to represent reasonable opinion through an expert prior. As we have already mentioned (and proponents of uniform priors are quick to argue), it would be hard today to find a "cloistered expert" who could give an authoritative prior but who was not already aware of the modern data with which we aim to update their prior. Moreover, it is well established that experts often tend to claim more certainty than is appropriate (Morgan and Henrion, 1990). Nevertheless, it certainly seems worthwhile (given the lack of a viable alternative) to attempt this quest, while remaining aware of the potential pitfalls. We minimise the risk of overconfidence by performing some sensitivity analyses to investigate how critically our results depend on the choices made.

In the absence of a cloistered expert, one reasonable approach would be to look back through the literature to see what climate scientists actually wrote prior to the observation and analysis of modern data sets. After what was perhaps the earliest early estimate for S of around 5°C (Arrhenius, 1896), all subsequent model-based estimates have been clearly lower (Manabe and Wetherald, 1967; Hansen et al., 1983), culminating in the "likely" range of 1.5–4.5°C (NRC, 1979). This estimate was produced well in advance of any modern probabilistic analysis of the warming trend and much other observational data, and could barely have been affected by the strong multidecadal trend in global temperature that has emerged since around 1975. Therefore, it could be considered a sensible basis for a credible prior to be updated by recent data.

Simple physically-based arguments also point towards this range as at least having higher probability than much higher or lower values. That is, the radiative forcing effect of a doubling of CO_2 alone is estimated to be roughly 1°C for a suitable grey body via the Stefan-Boltzman law, with the expected water vapour feedback roughly doubling this to around 2°C (Houghton et al., 2001). Cloud feedback is widely acknowledged to be highly uncertain even in sign, and therefore the nonlinear relationship between feedbacks and sensitivity implies that any prior for S should have broad support that extends to high values. However, a prior for S of U[0,20°C] requires the belief that not only is it "very likely" (90%) that the cloud feedback is positive, but furthermore "likely" that it is almost as large as the powerful water vapour effect. Since the sign of this feedback is considered uncertain even now (Solomon et al., 2007, Section 8.6.3.2), it does not seem reasonable to start from such confidently biased beliefs in a prior. We emphasise that we do not propose that the simple theoretical and model-based arguments presented above can provide a precise estimate for S, or even that they justify the prohibition of high values *a priori*. Rather, we merely use them to support our contention that the uniform priors which have been widely used represent an extreme viewpoint which cannot readily be reconciled with any credible scientific opinion, past or present.

A composite expert prior has previously been presented by Webster and Sokolov (2000), which is also broadly consistent with the long-held viewpoint that S is likely to be moderate. This prior (which is shown in Figure 2) is a Beta function with the parameters chosen to approximately fit the range of results found in a survey of experts (Morgan and Keith, 1995), and is also compatible with the NRC report in assigning a probability of 67% to the range 1.5–4.5°C. It produces an expected cost of 2.3% of GDP under our $2\times CO_2$ scenario. It has already been shown by Forest et al. (2002) that updating this expert prior with global temperature data from the 20th century results in greatly increased confidence in a moderate value for S. It is, however, hard to shake off the accusation that the experts who were surveyed by Morgan and Keith were aware of the recent warming rate, and had therefore already accounted for that data in their estimates (albeit informally, as there were no published analyses at this time). Perhaps for that reason, this estimate of Forest et al. appears to have been largely ignored. However, such a criticism of double-counting is no longer tenable if instead of using a likelihood function based on historical temperature data, we consider FG's recent analysis of the ERBE data, which was published more than 10 years after Morgan and Keith performed their expert survey. We can also note further that the raw observational data upon which the FG analysis was based entirely post-dates the NRC (1979) report so cannot possibly have influenced that assessment. As FG also mention, their analysis does not depend on calculations of climate models, and therefore we believe we are fully justified in treating it as independent of the NRC estimate.

We should therefore update the expert prior with the likelihood function arising from FG's analysis of the ERBE data, and present the results in Figure 2. The resulting 5–95% posterior probability interval is 1.2–3.6°C, and the expected cost of $2\times CO_2$ drops from 2.3% to 1.5% of GDP.

4 Sensitivity analysis

In order to investigate the robustness of this result, we next examine its sensitivity to changes in both the prior and the likelihood function.

4.1 Sensitivity to the prior

The Beta function of Webster and Sokolov may be considered rather optimistic in the very low probability that it assigns to high values of S. Also, by being strictly bounded to $S < 15$°C, it prohibits any more extreme values of S irrespective of the data. Therefore, we consider the use of an alternative, more pessimistic prior with greatly exaggerated tails, also illustrated in Figure 2 (right hand plots). This is a Cauchy distribution with location parameter 2.5 and scale 3, giving it the functional shape $f(S) \propto 1/((S-2.5)^2+3)$. We truncate this prior at 0°C and 100°C for numerical convenience but, in contrast to the uniform priors previously discussed, the influence of the upper bound on the results presented here is negligible. The parameters of

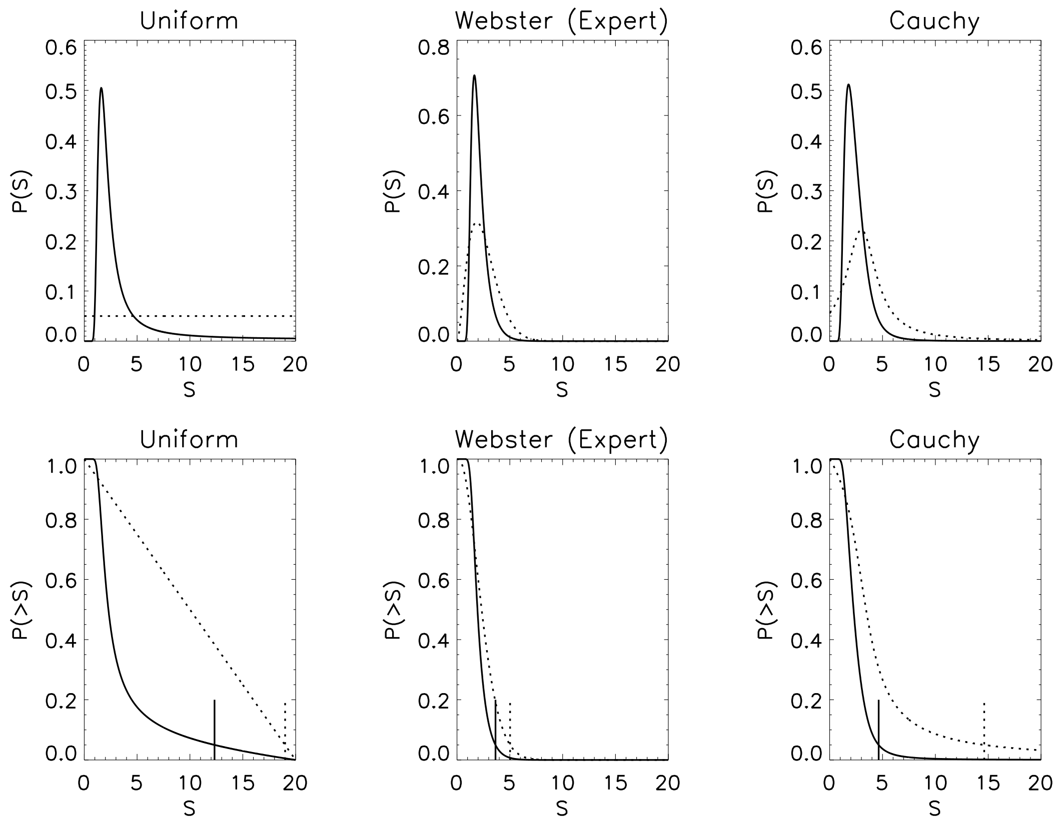

Figure 2: Pdfs arising from different prior distributions. Dotted lines indicate the priors, solid lines indicates the posterior pdf after updating with the likelihood function of Forster and Gregory (2006). Vertical lines indicate the upper 95% bound in each case.

this distribution were chosen to be roughly compatible with both the NRC report and the simple physical arguments presented above, but we have tried to err on the side of pessimism. Thus, this prior only assigns 55% probability to S lying in the traditional range of 1.5–4.5°C, insufficient to satisfy the IPCC's interpretation of "likely". It also implies $P(S > 6°C) = 18\%$ and $P(S > 15°C) = 5\%$. We suspect that if such an estimate had been presented in the NRC report, it would have (rightly) been met with widespread concern from those who accepted it, but also would have provoked hostility and scepticism from others who would have argued that it was rather too pessimistic.

The Cauchy distribution has extremely long and slowly-declining tails, such that it has neither a mean nor a variance, although it does have a well-defined median and finite probability intervals for any probability less than 1. Thus it represents a much more pessimistic outlook than the strictly bounded Beta distribution chosen by Webster and Sokolov (2000) and used by Forest et al. (2002). Indeed, its extreme tail assigns non-zero probability to substantially higher values than any of the uniform priors that appear in the literature. Even in this rather extreme case, however, the posterior 5–95% probability range after updating with FG's results only covers the interval 1.3–4.2°C. Such a result represents a substantial decrease in uncertainty compared to all recent published estimates. Furthermore, it cannot be argued that this has been achieved by using a prior which rules out high S *ab initio*. On the contrary, it assigns a substantial level of prior belief to such a hypothesis, including 8% probability to $S > 10°C$, an event which is deemed impossible by the currently favoured U[0,10°C] prior. If the data had actually indicated a strong likelihood for high sensitivity, say via a hypothetical likelihood function for radiative feedback given by $1/S = L = N(0.4, 0.1)$, then the posterior distribution for S would have a 5–95% probability range of 6.4–14.8°C, comfortably bracketing the maximum likelihood value for S of $3.7/0.4 = 9.2°C$. Therefore, it is clear that the choice of such a prior in no way prevents the posterior from indicating a high probability of high sensitivity, if the data were to actually support this.

The implications for economic analysis are also encouraging. Whereas the Webster and Sokolov prior alone indicates a loss of only 2.3% of GDP for stabilisation at $2\times CO_2$, which is reduced in the posterior to 1.5%, the long-tailed Cauchy prior presented here has a much greater expected loss of 9.7% but this is still reduced to only 2% by the Bayesian updating with the FG likelihood.

We now investigate how much more pessimistic the prior would have to be in order to strongly affect the posterior. We start by shifting the location parameter of the Cauchy prior up to 3 (remembering that the median of the prior will always be higher than this, due to the asymmetric truncation at 0°C and 100°C). This change alone only increases the posterior cost marginally to 2.4%. An additional doubling of the prior's scale parameter, to 6, increases the 95%th percentile of the prior to almost 20°C, and its median to 3.8°C. In this case, the posterior cost of climate change creeps up a little more to 2.6%. Doubling the scale factor again, to 12, leads to a posterior cost of 3.1% of GDP, still fractionally lower than that generated by the U[0,10] prior and less than half of that arising from U[0,20]. In this case, the prior has a median of 4.5°C, which is double that of the Webster and Sokolov prior, and it also assigns 19% probability to $S > 10$°C. We consider that this prior is a very long way removed from any credible expression of prior expert opinion.

4.2 Sensitivity to the likelihood

It is also plausible that FG's interpretation of their data is too optimistic, so we also test the sensitivity of our results to both increasing the uncertainty on their result (combining this modified likelihood function with different priors for S), and also consider replacing their analysis with other analyses of different observational data sets. As before, the results are detailed in the table, and we now describe the basis of these calculations. First we change the uncertainty of FG's likelihood function. A 50% increase on its width in feedback space (to 1.05 at one standard deviation) results in it providing only a rather weak update to the Webster and Sokolov prior, reducing the upper 95% probability limit for S and the cost of a doubling of CO_2 by only a few tenths of a degree and a few tenths of a percentage point respectively. When that broader likelihood is combined with our original Cauchy prior, the posterior cost of a doubling of CO_2 is as high as 4.3% of GDP, and this number rises substantially as the Cauchy prior is broadened and shifted to higher values. Therefore, it seems that the interpretation of the data can be a relatively important factor (at least, if one considers an increase in uncertainty as extreme as 50% to be plausible).

We must, however, also recognise that there are many other data that can also help to inform us on the climate sensitivity. Even considering the possibility that our prior already accounts for that part of the historical warming trend observed prior to NRC (1979) or Morgan and Keith (1995), there has been additional gradual ongoing warming since then, and also distinct analyses such as explorations of paleoclimatic data (Annan et al., 2005; Schneider von Deimling et al., 2006), and simulation of the short-term response to volcanic perturbations (Yokohata et al., 2005; Wigley et al., 2005). If we are not implicitly considering this evidence as part of our prior knowledge about S (and we have deliberately chosen to base the prior on very little evidence), it should rightfully contribute to the likelihood function. The common procedure of using each analysis separately, combined with a "ignorant" prior, cannot by construction generate a credible posterior belief concerning the climate sensitivity, since such an approach makes no attempt to present an integrated analysis of our information. It is well-known that additional data are expected to decrease uncertainty. To be precise, we can only say that additional data cannot be expected to increase uncertainty, but limiting cases where they have no effect whatsoever tend to be somewhat pathological (eg Wynn, 2008). Some analyses of combined data sets have already illustrated the increase in confidence that can arise from such joint analysis (Annan and Hargreaves, 2006; Hegerl et al., 2006). We note that both of these analyses were actually based on an underlying uniform prior, which suggests that a less naive choice would have generated rather sharper results. Since the likelihoods used in these papers in no way depended on the choice of prior, we can illustrate this point by simply post-multiplying their results by a different prior.

The Hegerl et al. analysis was abruptly truncated at 10°C, so we must first extend their likelihood function to higher values. We do this by fitting an inverse Gaussian to the tail of high values, this shape being (as noted above) supported by various theoretical arguments and empirical evidence. With this interpretation of their result, simply extending their uniform prior to cover the range 0–20°C would result in a 95% probability limit of 6.8°C. When we combine their likelihood function with the Cauchy prior, however, the 95% probability limit is much lower at 4.5°C, and the expected cost is 2.4%. These values represent a significant improvement on their published result where an upper 95% probability limit of 6.2°C was reported.

We can also treat the results of Annan and Hargreaves in a similar way. In that work, some strongly truncated likelihood functions (Gaussian and Beta) were chosen for convenience, which may be considered too optimistic in the light of our earlier discussion about extreme tails. Therefore, we impose a lower bound on the likelihood function arising from this work, never allowing it to drop below a value of 1% of its peak. This value was chosen to approximately match the likelihood ratio for the high tail of Hegerl et al.. The resulting posterior pdf, if a U[0,20°C] prior is used, would have a cost of 5.9% of GDP under 2×CO_2 and a 95% probability limit of 9°C. Using instead the Cauchy prior gives a cost of 2.8% and a 95% limit of 4.2°C. Again, the upper limit for S is lower than the originally published result, despite the increased likelihood function for extreme values.

Both of these analyses exclude the FG work, which was based on independent observations. Therefore, we think it would be hard to argue that all of these results are substantially overconfident. If we were to further update each of these analyses with the FG likelihood, the posterior cost would reduce still further to 1.6% of GDP with a 95% limit for S of 3.5°C based on the Hegerl et al. analysis, and 2% of GDP with a 95% limit for S of 3.7°C for Annan and Hargreaves. At such a level of precision, it would probably be worth re-examining the accuracy of assumptions in some detail, such as those regarding the linearity of the climatic response to forcing, and the independence of the analyses of the distinct data sets. Nevertheless, such results may be interpreted as hinting at an achievable upper bound to the precision with which we can reasonably claim to know S, given our current scientific knowledge.

5 Conclusions

We have investigated the assumptions underlying many recent probabilistic analyses of climate sensitivity, and shown how these may influence policy through a simple economic analysis of a climate change scenario. We have two main results. First, we have shown that results based on a uniform prior are sensitive to the selection of the upper bound, and it is surprising that this point has not received more discussion in the literature. We also observe that the uniform priors that have been used represent extremely pessimistic beliefs about climate sensitivity that cannot truly be considered to represent either "ignorance" or plausible prior beliefs of reasonable scientists. For these reasons, we consider uniform priors to be a poor choice for this analysis. Second, we have shown that when an analysis is performed based on a reasonable expert prior, the results are robust with respect to changes in either the details of the choice of prior, or the likelihood function. Even using a Cauchy prior (based on a rather pessimistic interpretation of the NRC report), which has such extremely long and fat tails that it has no mean or variance, gives quite reasonable results. While we are reluctant to openly endorse such a pathological choice of prior, we think it should at least be immune from any criticism that it is overly optimistic. When this prior is updated with the analysis of Forster and Gregory (2006), the long fat tail that is characteristic of all recent estimates of climate sensitivity simply disappears, with an upper 95% probability limit for S easily shown to lie close to 4°C, and certainly well below 6°C. Alternative likelihoods based on Hegerl et al. (2006) and Annan and Hargreaves (2006) generate similar results. Thus it might be reasonable for the IPCC to upgrade their confidence in S lying below 4.5°C to the "extremely likely" level, indicating 95% probability of a lower value. Expected economic losses are also strongly impacted by this reduction in the uncertainty of climate sensitivity. While the economic analysis presented here ignores a substantial additional source of uncertainty by selecting one specific damage function, the use of reasonable expert priors results in a substantial lowering of the expected loss to around 2% of global GDP for the simple $2\times CO_2$ scenario used.

Acknowledgments

This work was partly supported by the Global Environment Research Fund (S-5) of the Ministry of the Environment, and partly supported by the Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan. We are grateful to James Risbey and two anonymous reviewers for their helpful comments.

References

N. G. Andronova and M. E. Schlesinger. Objective estimation of the probability density function for climate sensitivity. *Journal of Geophysical Research*, 108(D8):22,605–22,611, 2001.

J. D. Annan and J. C. Hargreaves. Using multiple observationally-based constraints to estimate climate sensitivity. *Geophysical Research Letters*, 33(L06704), 2006.

J. D. Annan, J. C. Hargreaves, R. Ohgaito, A. Abe-Ouchi, and S. Emori.

Efficiently constraining climate sensitivity with paleoclimate simulations. *SOLA*, 1:181–184, 2005.

S. Arrhenius. On the influence of carbonic acid in the air upon the temperature of the ground. *Philosophical Magazine*, 41:237–276, 1896.

J.M. Bernardo and A.F.M. Smith. *Bayesian Theory*. Wiley, Chichester, UK, 1994.

G. Betz. Probabilities in climate policy advice: a critical comment. *Climatic Change*, 85(1):1–9, 2007.

C. E. Forest, P. H. Stone, A. P. Sokolov, M. R. Allen, and M. D. Webster. Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science*, 295(5552):113–117, 2002.

P. M. Forster and J. M. Gregory. The climate sensitivity and its components diagnosed from earth radiation budget data. *Journal of Climate*, 19(1): 39–52, 2006. doi: doi: 10.1175/JCLI3611.1.

D. J. Frame, B. B. B. Booth, J. A. Kettleborough, D. A. Stainforth, J. M. Gregory, M. Collins, and M. R. Allen. Constraining climate forecasts: The role of prior assumptions. *Geophysical Research Letters*, 32(L09702), 2005.

J. M. Gregory, R. J. Stouffer, S. C. B. Raper, P. A. Stott, and N. A. Rayner. An observationally based estimate of the climate sensitivity. *Journal of Climate*, 15(22):3117–3121, 2002.

J. Hansen, G. Russel, A. Lacis, I. Fung, and D. Rind. Climate response times: Dependence on climate sensitivity and ocean mixing. *Science*, 229: 857–859, 1985.

J. E. Hansen, G. Russell, D. Rind, P. Stone, A. Lacis, S. Lebedeff, R. Ruedy, and L. Travis. Efficient three dimensional global models for climate studies; models I and II. *Monthly Weather Review*, 3:609–662, 1983.

L.D.D. Harvey. Dangerous anthropogenic interference, dangerous climatic change, and harmful climatic change: non-trivial distinctions with significant policy implications. *Climatic Change*, 82(1):1–25, 2007.

G. C. Hegerl, T. J. Crowley, W. T. Hyde, and D. J. Frame. Climate sensitivity constrained by temperature reconstructions over the past seven centuries. *Nature*, 440:1029–1032, 2006.

J. T. Houghton, L. G. Meiro Filho, B. A. Callander, N. Harris, A. Kattenberg, and K.Maskell. *Climate change 1995: the science of climate change. Contribution of Working Group I to the Second Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge University Press, 1996.

J. T. Houghton, Y. Ding, D. J. Griggs, M. Noguer, P. J. Van Der Linden, X. Dai, K. Maskell, and C. A. Johnson. *Climate Change 2001: Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge University Press, 2001.

E. Kriegler. *Imprecise probability analysis for integrated assessment of climate change.* Ph.D. Thesis, Potsdam University, 2005.

T.C.K. Lee, F.W. Zwiers, G.C. Hegerl, X. Zhang, and M. Tsao. A Bayesian Climate Change Detection and Attribution Assessment. *Journal of Climate*, 18(13):2429–2440, 2005.

S.S. Leroy. Detecting Climate Signals: Some Bayesian Aspects. *Journal of Climate*, 11(4):640–651, 1998.

S. Manabe and R. T. Wetherald. Thermal equilibrium of the atmosphere with a given distribution of relative humidity. *Journal of the Atmospheric Sciences*, 24(3):241–259, 1967.

M. Meinshausen. What does a 2C target mean for greenhouse gas concentrations? a brief analysis based on multi-gas emission pathways and several climate sensitivity uncertainty estimates. In H. J. Schellnhuber, editor, *Avoiding dangerous climate change*, chapter 28. Cambridge University Press, 2006.

M. G. Morgan and D. Keith. Subjective judgments by climate experts. *Env. Sci. Tech.*, 29:468–476, 1995.

M.G. Morgan and M. Henrion. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge University Press, 1990.

R.H. Moss and S.H. Schneider. Uncertainties in the IPCC TAR: Recommendations to lead authors for more consistent assessment and reporting. *Guidance Papers on the Cross Cutting Issues of the Third Assessment Report of the IPCC*, pages 33–51, 2000.

W.D. Nordhaus. *A question of balance: weighing the options on global warming policies*. Yale Univ Pr, 2008.

National Research Council NRC. *Carbon Dioxide and Climate: A Scientific Assessment*. National Academy Press, Washington, DC, 1979.

M. Parry, O. Canzaiani, J. Palutikof, P. Van der Linden, and C. Hanson. *Climate change 2007: impacts, adaptation and vulnerability; Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge Univ. Press, 2007.

J.S. Risbey. Subjective elements in climate policy advice. *Climatic Change*, 85(1):11–17, 2007.

J.S. Risbey and M. Kandlikar. Expressions of likelihood and confidence in the IPCC uncertainty assessment process. *Climatic Change*, 85(1):19–31, 2007.

T. Schneider von Deimling, H. Held, A. Ganopolski, and S. Rahmstorf. Climate sensitivity estimated from ensemble simulations of glacial climate. *Climate Dynamics*, 27(2–3):149–163, 2006.

S. Solomon, D. Qin, M. Manning, Z. Chen, et al. Climate change 2007: The physical science basis. Contribution of the Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, 2007.

N. Stern. *The Economics of Climate Change: The Stern Review*. Cambridge University Press, 2007.

P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, 1991.

M. D. Webster and A. P. Sokolov. A methodology for quantifying uncertainty in climate projections. *Climatic Change*, 46(4):417–446, 2000.

T. M. L. Wigley, C. M. Amman, B. D. Santer, and S.C. B. Raper. Effect of climate sensitivity on the response to volcanic forcing. *Journal of Geophysical Research*, 110(D09107), 2005.

H. P. Wynn. Bayesian information-based learning and majorization. *Biometrika*, 2008. (Submitted).

G. Yohe, N. Andronova, and M. Schlesinger. To hedge or not against an uncertain climate future? *Science*, 306:416–417, 2004.

T. Yokohata, S. Emori, T. Nozawa, Y. Tsushima, T. Ogura, and M. Kimoto. Climate response to volcanic forcing: Validation of climate sensitivity of a coupled atmosphere-ocean general circulation model. *Geophysical Research Letters*, 32(L21710), 2005.