

**THE PERCEPTION
OF CLIMATE SENSITIVITY:
REVEALING PRIORS FROM POSTERiors**

Masako Ikefuji
Jan R. Magnus

December 2020

The Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

The perception of climate sensitivity: Revealing priors from posteriors

Masako Ikefuji

Faculty of Humanities and Social Sciences, University of Tsukuba
and ISER, Osaka University

Jan R. Magnus

Department of Econometrics and Data Science,
Vrije Universiteit Amsterdam and Tinbergen Institute

December 4, 2020

Abstract: A Bayesian typically uses data and a prior to produce a posterior. In practice, the data and the posterior are often observed but not the prior. We shall follow the opposite route, using data and the posterior information to reveal the prior. We then apply this theory to (equilibrium) climate sensitivity as reported by the Intergovernmental Panel on Climate Change in an attempt to get some insight into their prior beliefs. It appears that the IPCC scientists have agreed *a priori* on a value for the climate equilibrium between 3 and 4 degrees Celsius, while judging the occurrence of a real disaster much more likely than the previous report predicts.

JEL Classification: C11, Q54.

Keywords: Revealed prior; climate sensitivity; IPCC.

Corresponding author: Masako Ikefuji, Faculty of Humanities and Social Sciences, University of Tsukuba, 1-1-1 Tennodai, Tsukuba, Ibaraki 305-8577, Japan.

Email addresses:

ikefuji.masako.gn@u.tsukuba.ac.jp (Ikefuji),
jan@janmagnus.nl (Magnus).

1 Introduction

Everybody has priors, Bayesians and non-Bayesians alike. The priors may be vague and difficult to make explicit, but they are there and they may be important. The purpose of this paper is to show that we can make priors explicit from our knowledge of the data and the posterior, and to apply this theory to the perception of climate sensitivity.

Imagine a group of people (the ‘committee’) with a collective prior, perhaps based on knowledge and experience, perhaps on political beliefs, perhaps on short-term profit. The committee meets privately and we have no information about their discussions. But we do have scientific data (official ‘objective’ statistics and scientific results) and we do have access to their published predictions or policy recommendations, which they present to the public. In other words, we have the data and the posterior, but not the prior which the committee does not reveal and possibly may not even be able to formulate or quantify. Can we recover the prior from the data and the posterior? Yes, this is indeed possible and we shall study the properties of the recovered prior in some detail.

We apply this theory to the estimation of the equilibrium climate sensitivity (ECS), which is an important diagnostic in climate modeling. The data come from various studies as listed and reported by the Intergovernmental Panel on Climate Change (IPCC). The IPCC then formulates its own conclusions (the posterior). What we are interested in is to recover the prior beliefs of the IPCC.

The idea of reversing Bayesian thought and — rather than obtain a posterior from data and prior — recover the prior from data and posterior, does not seem to have received much attention. The current paper attempts to fill this gap. Of course, the list of possible applications is endless. A political party uses scientific data and publishes reports. From these two sources we can recover their priors. Do these conform to the party program? Scientists use data and write papers. The results in these papers may well be influenced by prior beliefs or non-scientific reasons. Can this influence be quantified? Such questions can, in principle, be studied by the theory developed in this

paper.

In Section 2 we introduce the ECS and summarize the findings of the IPCC: the posterior. In Section 3 we analyze how to recover the prior from the data and the posterior within the framework of the normal distribution. In Section 4 we consider the case when there is only one parameter of interest, and extend the framework to monotonic reparametrizations of the normal distribution, like the lognormal distribution. Next, in Section 5, we discuss the data which result from fifteen studies used by the IPCC. Given the posterior and the data, we then recover and discuss the prior mean and standard deviation (Section 6). Section 7 concludes. An Appendix discusses the role of monotonic reparametrizations in general, and the lognormal distribution in particular.

2 Climate sensitivity, the reported posterior

When the radiation balance of the Earth is perturbed, the temperature will change. By how much is measured by the *equilibrium climate sensitivity* (ECS): the long-term temperature rise that is expected to result from a doubling of the atmospheric CO₂ concentration, usually relative to the pre-industrial level (around 1750). It is a prediction of the new global mean near-surface air temperature once the CO₂ concentration has stopped increasing and most of the feedbacks have had time to have their full effect. The ECS is an important diagnostic in climate modelling, but it cannot be measured directly and forms a large source of uncertainty. CO₂ levels rose from 280 parts per million (ppm) in the eighteenth century (IPCC, 2013, p. 100) to about 416 ppm by 2020, an increase of almost 50%. In the same period, the Earth's temperature rose by a little over one degree Celsius. The ECS will be our parameter of interest, and we shall denote it by β .

In estimating β we rely exclusively on the Intergovernmental Panel on Climate Change (IPCC) reports. So far, five so-called Assessment Reports have appeared, the first in 1990, the fifth in 2013. In these reports we find estimates (and precisions) of studies on the ECS (our data) and the IPCC's own estimates (our posterior). From this information we will attempt to

recover the IPCC’s priors.

In the fifth report, more precisely the Working Group I contribution (IPCC, 2013), hereafter IPCC5, the authors state that ‘no best estimate for equilibrium climate sensitivity can now be given because of a lack of agreement on values across assessed lines of evidence and studies’ (IPCC5, p. 16, footnote). But later in the same report they do provide estimates, as follows:

‘...ECS is likely in the range 1.5°C to 4.5°C with high confidence. ECS is positive, extremely unlikely less than 1°C (high confidence), and very unlikely greater than 6°C (medium confidence)’.

(IPCC5, pp. 83–84)

The IPCC also provides a precise interpretation of terms like ‘extremely unlikely’ and ‘medium confidence’ (IPCC5, p. 36), which differs slightly from the interpretation in the previous Assessment Report (IPCC, 2007, p. 22) by explicitly taking into account the degree of ‘agreement’ in the team about the evidence provided by each study. Given this interpretation, IPCC5 concludes that

$$\begin{aligned}\Pr(1.5 < \text{ECS} < 4.5) &= 0.67, \\ \Pr(\text{ECS} < 1.0) &< 0.05, \text{ and} \\ \Pr(\text{ECS} > 6.0) &< 0.10.\end{aligned}$$

In addition (pp. 75 and 817), they summarize information of experiments by the Coupled Model Intercomparison Project Phase 5 (CMIP5) who report a range 2.1–4.7 for the ECS, without however stating the likelihood of this range.

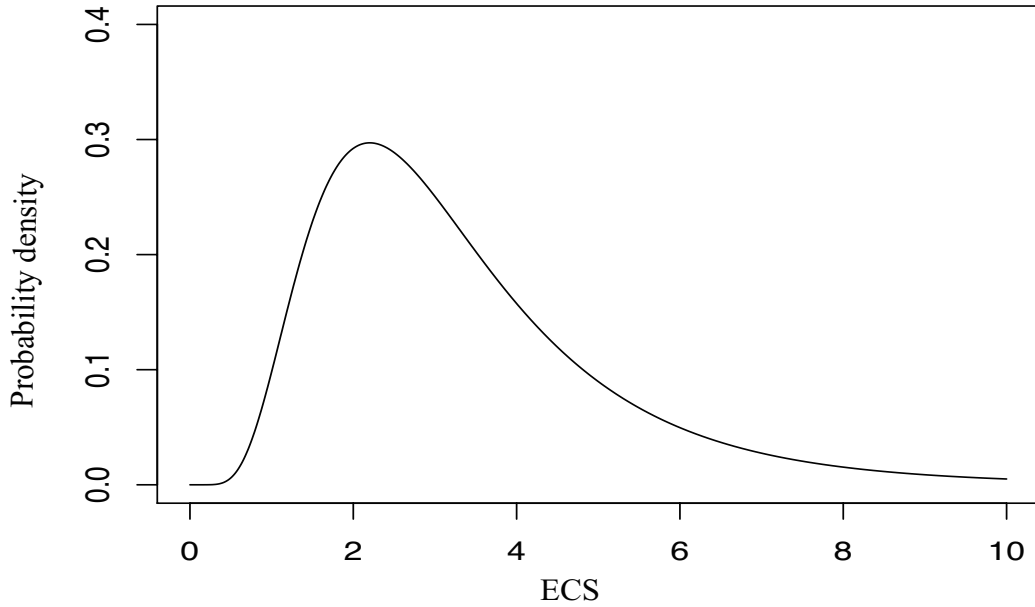
Assuming the ECS β to be lognormally distributed, so that $\log \beta \sim N(b_2, \sigma_2^2)$, we seek combinations (b_2, σ_2) such that the posterior probabilities closely match the probabilities in the IPCC report. This leads to Table 1, where we have calculated the required possibilities for various credible values of b_2 and σ_2 . There is no unique lognormal distribution that fits our data, but $b_2 = 1.07$ and $\sigma_2 = 0.53$ (in bold) seems a reasonable approximation and

Table 1: Posteriors on climate sensitivity, Fifth Assessment Report

Lognormal approx.		Probabilities for ECS (%)			
b_2	σ_2	1.5–4.5	< 1.0	> 6.0	2.1–4.7
0.95	0.57	66.5	4.8	7.0	49.5
1.00	0.55	68.0	3.5	7.5	52.1
1.05	0.53	69.2	2.4	8.1	54.6
1.06	0.53	69.1	2.3	8.4	54.7
1.07	0.53	68.9	2.2	8.7	54.8
1.08	0.52	69.5	1.9	8.6	55.8
1.09	0.52	69.3	1.8	8.9	55.9
1.10	0.51	69.9	1.6	8.7	56.9
1.15	0.49	70.1	0.9	9.5	58.9

is also in line with Hwang, Reynès, and Tol (2013, Figure 4) where $b_2 = 1.071$ and $\sigma_2 = 0.527$.

Figure 1: Posterior lognormal distribution with $b_2 = 1.07$ and $\sigma_2 = 0.53$



The selected posterior distribution, plotted in Figure 1, satisfies

$$\begin{aligned}\Pr(1.5 < \beta < 4.5) &= 68.9\% \text{ (about 67\%),} \\ \Pr(0 < \beta < 1.0) &= 2.2\% \text{ (less than 5\%), and} \\ \Pr(\beta > 6.0) &= 8.7\% \text{ (less than 10\%),}\end{aligned}$$

in accordance with the IPCC report's conclusions. In addition, the interquartile range is $\Pr(2.04 < \beta < 4.17) = 50\%$ and there is a 1% probability of $\beta > 10.0$. The skewness of the distribution is well illustrated by the fact that mode and median of β are quite different: The mode is $e^{b_2 - \sigma_2^2} = 2.20$, while the median is $e^{b_2} = 2.92$.

3 From posterior to prior under normality

In order to highlight the issue under simple conditions, we shall first assume normality, which we shall modify later. We consider a parameter vector of interest β and suppose that data are generated from a normal distribution

$$y|\beta \sim N(X\beta, \Omega), \quad (1)$$

where X is a given $n \times k$ matrix of rank k and Ω is a positive definite $n \times n$ matrix. A non-Bayesian frequentist would estimate β using the generalized least-squares (GLS) estimator (which is also the maximum likelihood estimator)

$$b_0 = (X'\Omega^{-1}X)^{-1} X'\Omega^{-1}y \quad (2)$$

with variance

$$\Sigma_0 = (X'\Omega^{-1}X)^{-1}. \quad (3)$$

A Bayesian, on the other hand, would wish to take prior knowledge about β into account. Suppose this prior information is given by

$$\beta \sim N(b_1, \Sigma_1), \quad (4)$$

where Σ_1 is positive definite. Then the posterior distribution of β is

$$\beta|y \sim N(b_2, \Sigma_2), \quad (5)$$

where

$$b_2 = Wb_1 + (I_k - W)b_0, \quad \Sigma_2 = (\Sigma_1^{-1} + \Sigma_0^{-1})^{-1}, \quad (6)$$

and $W = \Sigma_2 \Sigma_1^{-1}$ is a $k \times k$ weight matrix.

Although W is, in general, not symmetric, its eigenvalues are real and lie between zero and one. In fact, letting $Z = \Sigma_1^{1/2} \Sigma_0^{-1} \Sigma_1^{1/2}$ with eigenvalues $\lambda_i(Z) > 0$ ($i = 1, \dots, k$), we see that

$$\lambda_i(W) = \lambda_i(\Sigma_1^{-1/2} \Sigma_2 \Sigma_1^{-1/2}) = \frac{1}{\lambda_i(\Sigma_1^{1/2} \Sigma_2^{-1} \Sigma_1^{1/2})} = \frac{1}{1 + \lambda_i(Z)}. \quad (7)$$

Note that when the prior becomes uninformative, that is when $\Sigma_1^{-1} \rightarrow 0$, then $b_2 \rightarrow b_0$ and $\Sigma_2 \rightarrow \Sigma_0$.

This is well-established basic Bayesian theory. But now consider the opposite situation where the data and the posterior are available but not the prior. Can we reveal the prior from the data and the posterior? In general we can, and in the special case of normality we obtain the prior moments as

$$b_1 = W^{-1}b_2 + (I_k - W^{-1})b_0, \quad \Sigma_1 = (\Sigma_2^{-1} - \Sigma_0^{-1})^{-1}, \quad (8)$$

with

$$W^{-1} = \Sigma_1 \Sigma_2^{-1} = \Sigma_0(\Sigma_0 - \Sigma_2)^{-1}, \quad (9)$$

which assumes implicitly an upper bound to the posterior variance, namely $\Sigma_2 < \Sigma_0$ in the usual sense that $\Sigma_0 - \Sigma_2$ is positive definite. The prior mean is thus a ‘weighted average’ of b_2 and b_0 , but the eigenvalues of W^{-1} do not lie between zero and one. In fact $\lambda_i(W^{-1}) = 1 + \lambda_i(Z) > 1$ and $\lambda_i(I_k - W^{-1}) = -\lambda_i(Z) < 0$ for all $i = 1, \dots, k$.

The restriction $\Sigma_2 < \Sigma_0$ does not play a role in the usual Bayesian framework where we go from data plus prior to posterior, because the underlying variances Σ_0 and Σ_1 are unrestricted (apart from being positive definite) and

Σ_2 will automatically satisfy the restriction. But it does play a role when we go from data plus posterior to prior, because now the restriction is not automatically satisfied. This has practical consequences as we shall see later.

4 One parameter of interest under lognormality

In the special but important case where we have only one parameter β of interest, we write σ_0^2 , σ_1^2 , and σ_2^2 instead of Σ_0 , Σ_1 , and Σ_2 . From the data (without a prior) we obtain an unbiased estimator of β : $b_0 \sim N(\beta, \sigma_0^2)$. If we add a prior $\beta \sim N(b_1, \sigma_1^2)$, then we obtain the posterior $\beta \sim N(b_2, \sigma_2^2)$, where

$$b_2 = \frac{\sigma_0^2 b_1 + \sigma_1^2 b_0}{\sigma_0^2 + \sigma_1^2}, \quad \sigma_2^2 = \frac{\sigma_0^2 \sigma_1^2}{\sigma_0^2 + \sigma_1^2}. \quad (10)$$

In the reversed case that we are interested in we have an unbiased estimator $b_0 \sim N(\beta, \sigma_0^2)$ from the data and the posterior moments of $\beta \sim N(b_2, \sigma_2^2)$. From these two ingredients we obtain the prior as $\beta \sim N(b_1, \sigma_1^2)$, where

$$b_1 = \frac{\sigma_0^2 b_2 - \sigma_2^2 b_0}{\sigma_0^2 - \sigma_2^2}, \quad \sigma_1^2 = \frac{\sigma_0^2 \sigma_2^2}{\sigma_0^2 - \sigma_2^2}, \quad (11)$$

under the restriction that $\sigma_2^2 < \sigma_0^2$.

So far we have assumed normality because of its transparency in handling Bayesian problems. But it happens frequently that the normality assumption on β is not plausible. Then we need to search for a credible alternative, and the simplest alternative is normality of a monotonic transformation, say $h(\beta)$. Given this transformation the same theory shows that if $b_0 \sim N(h(\beta), \sigma_0^2)$ is an unbiased estimator of $h(\beta)$ from the data and $h(\beta) \sim N(b_2, \sigma_2^2)$ is the posterior, then the prior can be recovered as

$$h(\beta) \sim N(b_1, \sigma_1^2), \quad (12)$$

where b_1 and σ_1^2 are given in (11). From the prior distribution of $h(\beta)$ we

then obtain the prior distribution of β . The most important example is $h(\beta) = \log \beta$ in which case the posterior and the prior are both lognormally distributed; see the Appendix for details. This is the route that we shall follow.

Defining α_m and α_v implicitly by

$$b_2 = \alpha_m b_0, \quad \sigma_2^2 = \alpha_v \sigma_0^2, \quad (13)$$

we can rewrite (11) as

$$b_1 = \kappa_m b_0, \quad \sigma_1^2 = \kappa_v \sigma_0^2, \quad (14)$$

where

$$\kappa_m = \frac{\alpha_m - \alpha_v}{1 - \alpha_v}, \quad \kappa_v = \frac{\alpha_v}{1 - \alpha_v} \quad (15)$$

measure how far the prior is removed from the data and their effect on the prior mean and variance, respectively. Note that α_m is unrestricted but that α_v is restricted by $0 < \alpha_v < 1$.

The two fractions κ_m and κ_v capture the essence of our story. First consider κ_v , which relates to the prior variance. What matters here is whether κ_v is small (strong prior information) or large (weak prior information). This depends only on α_v , not on α_m . When α_v is close to one, then the variance σ_0^2 in the data and the variance σ_2^2 in the posterior are approximately equal, so that the prior has only a small effect. This is represented by a large value of κ_v and hence a large value of the prior variance σ_1^2 . The prior is then uninformative. But when α_v is close to zero, then the data variance and the posterior variance are not close at all, and the prior has a big effect. This is represented by a small value of κ_v and hence a small value of the prior variance σ_1^2 . The prior is then informative.

The situation is quite different with κ_m . What matters here is not whether κ_m is small or large, but rather whether κ_m is close to one or not. This will depend on both α_m and α_v . It is clear that $\kappa_m = 1$ when $\alpha_m = 1$, irrespective

of the value of α_v . Writing

$$1 - \kappa_m = \frac{1 - \alpha_m}{1 - \alpha_v}, \quad (16)$$

we see that the deviation of κ_m from one depends on the deviation of α_m from one *relative* to the deviation of α_v from one. When α_m is close to one but α_v is not, then the mean b_0 in the data and the mean b_2 in the posterior are approximately equal, but the variance σ_0^2 in the data and the variance σ_2^2 in the posterior are not approximately equal. In that case $\kappa_m \approx 1$ and the prior mean agrees with the data and the posterior. But when α_v is close to one but α_m is not, then the variances σ_0^2 and σ_2^2 are approximately equal, but the means b_0 and b_2 are not. In that case κ_m is large (in absolute value). Naturally, for people with a very strong prior ($\sigma_1^2 \approx 0$) we have $\alpha_v \approx 0$, and hence $\kappa_m \approx \alpha_m$ and $b_1 \approx b_2$.

5 The underlying data: Estimating b_0 and σ_0

In addition to the posterior we need information on the data. Our data consist of $n = 15$ studies mentioned in the IPCC5 report. Fourteen of these are contained in Figure 10.20 (p. 925) and Box 12.2 (p. 1110) of the report, and one study (Huber et al., 2011) is included in Figure 2 of Knutti, Rugenstein, and Hegerl (2017) and referred to in various places of IPCC5.

The i th study produces a range (l_i, u_i) (lower and upper bound) with an associated probability p_i (typically 90% or 95%) for the ECS β , our parameter of interest. Given the range (l_i, u_i) and the associated probability $\Pr(l_i < \beta_i < u_i) = p_i$ we can identify the parameters of the associated lognormal distributions, as follows. Since $\log \beta_i \sim N(y_i, \sigma_{0i}^2)$ in the i th study, we have

$$\Pr\left(\frac{\log l_i - y_i}{\sigma_{0i}} < z_i < \frac{\log u_i - y_i}{\sigma_{0i}}\right) = p_i, \quad z_i \sim N(0, 1), \quad (17)$$

and hence

$$y_i = \log(l_i u_i)^{1/2}, \quad \sigma_{0i} = \frac{\log(u_i/l_i)^{1/2}}{q_i}, \quad \Phi(q_i) = \frac{p_i + 1}{2}, \quad (18)$$

where Φ denotes the c.d.f. of the standard-normal distribution. In this way we end up with fifteen observations y_i with associated standard deviations σ_{0i} ; see Table 2. Our task is to find an estimator $b_0 \sim N(\log \beta, \sigma_0^2)$ based on the fifteen data $y_i \sim N(\log \beta, \sigma_{0i}^2)$. This task is not trivial as we shall see.

Table 2: Data on climate sensitivity, Fifth Assessment Report

Study	p_i	Bounds		Lognormal approx.	
		l_i	u_i	y_i	σ_{0i}
1 Lindzen and Choi, 2011	95%	0.5	1.1	−0.30	0.20
2 Schmittner et al., 2011	67%	1.7	2.6		
land and ocean data	90%	1.4	2.8	0.68	0.21
only ocean data	67%	1.5	2.5		
	90%	1.3	2.7		
only land data	67%	2.8	4.1		
	90%	2.2	4.6		
3 Aldrin et al., 2012	90%	1.2	3.5	0.72	0.33
4 Hargreaves et al., 2012	90%	1.0	4.2	0.72	0.44
5 Lewis, 2013	90%	2.0	3.6	0.99	0.18
incl. uncertainties	90%	1.0	3.0		
6 Bender et al., 2010	95%	1.7	4.1	0.97	0.22
7 Otto et al., 2013 (2000–2009)	90%	1.2	3.9		
Otto et al., 2013 (1970–2009)	90%	0.9	5.0	0.75	0.52
8 Schwartz, 2012	90%	1.2	4.9	0.89	0.43
9 Lin et al., 2010	90%	2.8	3.7	1.17	0.08
10 Libardoni and Forest, 2011	90%	1.2	5.3	0.93	0.45
11 Köhler et al., 2010	90%	1.4	5.2	0.99	0.40
12 Olson et al., 2012	95%	1.8	4.9	1.09	0.26
13 Huber et al. 2011	67%	2.9	4.0	1.23	0.17
14 Holden et al., 2010	67%	2.6	4.4		
	90%	2.0	5.0	1.15	0.28
15 Palaeosens, 2012	67%	2.2	4.8		
Palaeosens, 2012	95%	1.1	7.0	1.02	0.47

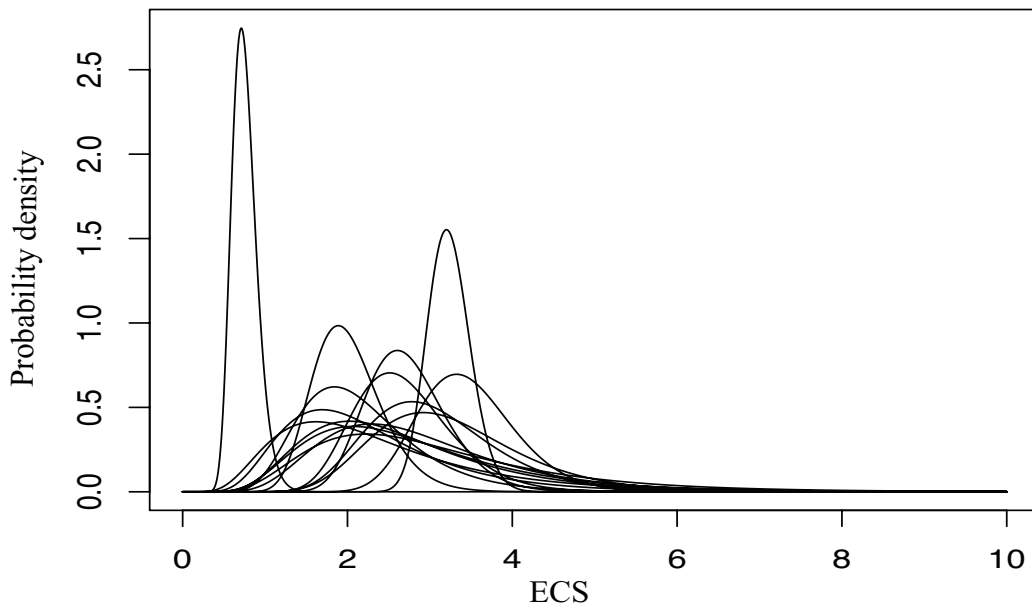
The IPCC5 report distinguishes between an instrumental period and a palaeoclimatic period. The instrumental period is the short period in the Earth’s long history where direct instrumental records on climate are available, while the palaeoclimatic period is the long period preceding such records. Of our fifteen studies some only use data from the instrumental

period (studies 1, 3, 5–9, 12, 13), some only from the palaeoclimatic period (2, 4, 11, 14, 15), and some combine different lines of evidence (3, 10, 12). Studies 3 and 12 appear in both the groups instrumental and combination. This highlights the first of several problems: comparability.

The second problem is that the first study (Lindzen and Choi, 2011) has a big impact and is in essence an outlier. The IPCC raises doubts about the reliability of this study (IPCC5, pp. 923–924), but it has not removed the study from their report. Not only is the revealed value of y_i much lower than in the other studies, but the effect is much strengthened by the fact that the reported precision is high.

Third, we should take into account that the studies are correlated with each other. They all estimate the same parameter, probably using similar highly correlated data sets.

Figure 2: The fifteen studies, lognormal distributions



To gain further insight, let us draw the fifteen lognormal curves; see Figure 2. From the figure several things become clear. First, that the first study (with the highest peak) is an outlier. And second, almost beyond belief, that what the data tell us is ambiguous in the extreme: The modes range from 0.71 (1.89 if we exclude the first study) to 3.31, and the medians

from 0.74 (1.98 if we exclude the first study) to 3.41. To make any sense of these data requires a strong prior.

Not only do we have to find an estimator b_0 but also a standard deviation σ_0 . One of the reasons why this is difficult is that σ_0 should always be *smaller* than each of the individual standard deviations σ_{0i} (more information leads to higher precision), but *larger* than the posterior standard deviation σ_2 . These two facts are not in agreement with each other, since σ_0 must be smaller than 0.08 (the smallest of the underlying standard deviations σ_{0i} in Table 2), while at the same time it must be larger than σ_2 , which we have set at 0.53; see Table 1.

In fact, this is an intriguing problem: we combine information and we thus expect the standard deviation to decrease, but in fact it doesn't. To illustrate, suppose we have one observation $x_1 = 50$ on an unknown parameter θ with standard deviation 4. We use this as our guideline until a new observation appears: $x_2 = 25$ with standard deviation 3. A naive approach would be to conclude that the updated estimate of θ is a weighted average of x_1 and x_2 , in this case 34 with standard deviation 2.4. This would conform statistical theory, at least in the absence of correlation between x_1 and x_2 . But it would not convince, because there appears to be a conflict between the two estimates. The naive approach works if we draw from the same population, but in our case each study has used different (though correlated) data, different methods, a different time period, and so on. Hence, there is not only *intra*-study noise (captured by the reported standard deviations), but also *inter*-study noise, which is not captured and not reported. Each study behaves like the frog in the well-known Japanese proverb 'a frog on the bottom of a well,' which relates the story of a frog living happily at the bottom of a dark well, full of sweet water and buzzing insects, until a turtle tells him about something outside, called the ocean. The problem is closely related to the 'prior-likelihood conflict' in the Bayesian literature and to the ideas of 'scientific ambiguity' (Miller, Dietz, and Heal, 2013; Brock and Xepapadeas, 2019), 'negative learning' (Oppenheimer, O'Neill, and Webster, 2008), and the combination of disparate lines of evidence (Cooke and Wielicki, 2018).

To understand the problem better let us realize that each of the fifteen

estimates comes from a different population, so the reported means and standard deviations are conditional on a ‘state’ S_i . In fact, therefore, the studies report the conditional moments $E(y_i|S_i)$ and $\text{var}(y_i|S_i)$ instead of the unconditional moments

$$E y_i = E E(y_i|S_i) \quad \text{var}(y_i) = E \text{var}(y_i|S_i) + \text{var}(E(y_i|S_i)). \quad (19)$$

To work with $E(y_i|S_i)$ and $\text{var}(y_i|S_i)$ instead of $E E(y_i|S_i)$ and $E \text{var}(y_i|S_i)$ may be reasonable, but to ignore the variation in the conditional mean $\text{var}(E(y_i|S_i))$ is not reasonable. This explains, at least in part, why the reported variance seriously underestimates the true variance.

Since we have no knowledge of the states S_i , we take a different route based on credible scenarios. Since theory prescribes that $\sigma_0 > 0.53$, we consider three values of σ_0 : 0.6, 0.7, and 0.8. In addition, since $\bar{y} = 0.87$ if we include the first study and $\bar{y} = 0.95$ if we exclude it, we consider four values of b_0 : 0.8, 0.9, 1.0, and 1.1.

Table 3: Estimates of b_0 and σ_0 from the data

b_0	σ_0	mode	quantiles		
			5%	50%	95%
0.8	0.6	1.55	0.83	2.23	5.97
	0.7	1.36	0.70	2.23	7.04
	0.8	1.17	0.60	2.23	8.30
0.9	0.6	1.72	0.92	2.46	6.60
	0.7	1.51	0.78	2.46	7.78
	0.8	1.30	0.66	2.46	9.17
1.0	0.6	1.90	1.01	2.72	7.29
	0.7	1.67	0.86	2.72	8.60
	0.8	1.43	0.73	2.72	10.13
1.1	0.6	2.10	1.12	3.00	8.06
	0.7	1.84	0.95	3.00	9.50
	0.8	1.58	0.81	3.00	11.20

Thus we shall consider twelve scenarios, where each scenario obeys the theoretical restrictions and the combination of scenarios covers what we be-

lieve is a credible range of the data parameters. Table 3 presents the mode and three quantiles (5%, 50%, and 95%) for the implied distributions. We see that the median ranges from 2.23 to 3.00, which is a little narrower than in the fourteen studies where the median ranges from 1.98 to 3.41 (if we exclude the first study where the median is 0.74). The scenarios also allow for a wide range of (right) tail behavior. The value β^* of β which defines the right tail (that is, where 5% of the distribution lies to the right of β^*) ranges from 5.97 to 11.20. Given the data (particularly, $\bar{y} = 0.95$) and the restriction $\sigma_0 > \sigma_2 = 0.53$, our preferred estimate is $(b_0, \sigma_0) = (1.0, 0.7)$ (in bold), but we shall consider the whole spectrum of twelve scenarios when we derive and analyze the priors in the next section.

6 Analysis of the prior moments

Given the moments (b_0, σ_0) from the data and (b_2, σ_2) from the posterior, we can now discuss the revealed prior moments. The prior mean is given by (11),

$$b_1 = \frac{b_2 - \alpha_v b_0}{1 - \alpha_v}, \quad \alpha_v = \sigma_2^2 / \sigma_0^2, \quad (20)$$

and thus depends only on the posterior mean b_2 (which we set at 1.07), the data mean b_0 (which ranges from 0.8 to 1.1), and on the ratio of $\sigma_2 = 0.53$ and σ_0 (which ranges from 0.6 to 0.8).

In contrast, the prior standard deviation, also given in (11), is

$$\sigma_1 = \frac{\sigma_0 \sigma_2}{\sqrt{\sigma_0^2 - \sigma_2^2}} = \frac{\sigma_2}{\sqrt{1 - \alpha_v}}, \quad (21)$$

which, given $\sigma_2 = 0.53$, depends only on σ_0 or alternatively on α_v .

Table 4 shows the revealed prior means and standards deviations for each of the twelve scenarios (four values of b_0 , three values of σ_0). Given $(b_2, \sigma_2) = (1.07, 0.53)$ and the twelve selected pairs (b_0, σ_0) , we present the induced values of κ_m , κ_v , b_1 , and σ_1 .

Since b_2 must lie in-between b_0 and b_1 , it follows that $b_1 > 1.07$ if $b_0 < 1.07$ and $b_1 < 1.07$ if $b_0 > 1.07$. The lower is b_0 , the higher must be b_1 . If $b_0 = 0.8$

Table 4: Estimates of prior moments b_1 and σ_1 for given posterior moments $(b_2, \sigma_2) = (1.07, 0.53)$

data		κ_m	κ_v	prior		mode	quantiles		
b_0	σ_0			b_1	σ_1		5%	50%	95%
0.8	0.6	2.54	3.55	2.03	1.13	2.12	1.18	7.61	48.84
	0.7	1.79	1.34	1.43	0.81	2.17	1.10	4.19	15.91
	0.8	1.60	0.78	1.28	0.71	2.18	1.12	3.60	11.53
0.9	0.6	1.86	3.55	1.67	1.13	1.48	0.83	5.33	34.24
	0.7	1.44	1.34	1.30	0.81	1.90	0.96	3.66	13.91
	0.8	1.34	0.78	1.20	0.71	2.02	1.04	3.33	10.66
1.0	0.6	1.32	3.55	1.32	1.13	1.04	0.58	3.74	24.01
	0.7	1.16	1.34	1.16	0.81	1.66	0.84	3.20	12.16
	0.8	1.12	0.78	1.12	0.71	1.87	0.96	3.08	9.86
1.1	0.6	0.88	3.55	0.96	1.13	0.73	0.41	2.62	16.83
	0.7	0.94	1.34	1.03	0.81	1.45	0.74	2.80	10.64
	0.8	0.95	0.78	1.05	0.71	1.73	0.89	2.85	9.12

then we need a high prior mean, leading to a prior median of 3.6°C or more. If the standard deviation in the data is relatively large ($\sigma_0 = 0.8$), then the standard deviation in the prior will be relatively small. But even then the right tail is substantial with a 5% probability that the ECS will exceed 10°C or worse. In our bolded scenario there is a 5% probability that the ECS will exceed 12°C.

As discussed in Section 4, the story can be summarized in terms on the key parameters κ_m and κ_v . What matters regarding κ_m is whether or not it is close to one. When the mean b_0 in the data and the mean b_2 in the posterior are approximately equal, but the variance σ_0^2 in the data and the variance σ_2^2 in the posterior are not approximately equal, then $\kappa_m \approx 1$ and the prior mean agrees with the data and the posterior. But when the variances σ_0^2 and σ_2^2 are approximately equal, but the means b_0 and b_2 are not, then κ_m is large (in absolute value). This is well illustrated in Table 4 where $\kappa_m = 1.12$ is close to one when $b_0 = 1.0$ is close to $b_2 = 1.07$, while $\sigma_0 = 0.8$ is not close to $\sigma_2 = 0.53$. At the other extreme, $\kappa_m = 2.54$ is not close to one, because

$b_0 = 0.8$ is not close to $b_2 = 1.07$, while $\sigma_0 = 0.6$ is close to $\sigma_2 = 0.53$.

The prior standard deviation satisfies $\sigma_1 > \sigma_2 = 0.53$ (as it must), but σ_1 can be larger or smaller than σ_0 . In our set up, σ_1 can only take three values: 0.71 (small), 0.81 (medium), or 1.13 (large), depending on whether σ_0 equals 0.8 (large), 0.7 (medium), or 0.6 (small). Regarding κ_v what matters is whether it is small ($\kappa_v = 0.78$, strong prior information) or large ($\kappa_v = 3.55$, weak prior information). When the standard deviations σ_0 in the data and σ_2 in the posterior are approximately equal, then we see from Table 4 that the prior has only a small effect. This is represented by a large value of κ_v (3.55) and hence a large value of the prior variance σ_1 (1.13). The prior is then uninformative. But when the standard deviation in the data and the posterior are not close, then the prior has a big effect. This is represented by a small value of κ_v (0.78) and hence a small value of the prior standard deviation σ_1 (0.71). The prior is then informative.

So, what do these results tell us about the unknown prior of the IPCC? And in particular, how strong is this prior in view of the fact that a strong prior seems to be required to make any sense of the data.

Let us compare our prior with the conclusions (the posterior) of the previous IPCC report (IPCC, 2007). One would expect that the posterior of the previous study serves as the prior in the next study, at least approximately. The fourth report concludes about the ECS:

‘It is likely to be in the range 2.0°C to 4.5°C with a best estimate of about 3.0°C, and is very unlikely to be less than 1.5°C. Values substantially higher than 4.5°C cannot be excluded, but agreement of models with observations is not as good for those values.’
(IPCC, 2007, p. 12)

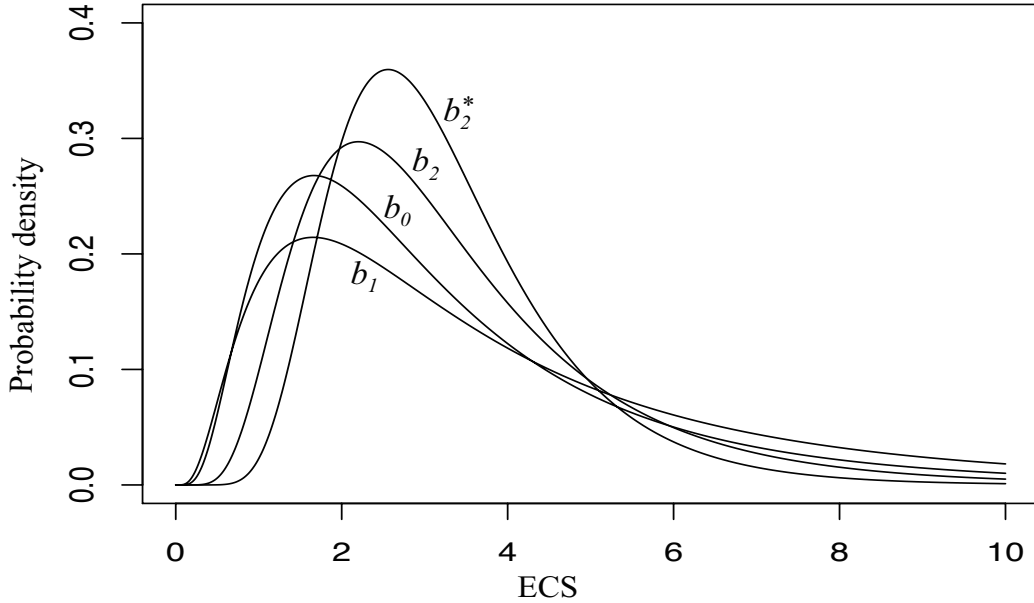
If we fit a lognormal distribution with mean 1.1 and standard deviation 0.4, then these statements translate to a median of 3.0 and

$$\begin{aligned}\Pr(2.0 < \text{ECS} < 4.5) &= 0.69 > 0.67 \text{ (likely),} \\ \Pr(\text{ECS} < 1.5) &= 0.04 < 0.10 \text{ (very unlikely), and} \\ \Pr(\text{ECS} > 4.5) &= 0.16,\end{aligned}$$

which seems about right.

The posterior mean 1.10 from the previous report is a little lower than the prior mean in the current report (say, 1.16), but it is comparable. The same is true for the implied median, which is 3.0 in the previous report, and now a little higher, say 3.2. The main difference is in the standard deviation, which is 0.4 in the posterior distribution of the previous report and about 0.7 in the prior distribution of the current report. It appears that the IPCC scientists have agreed *a priori* on a value for the ECS between 3°C and 4.0°C, while judging the occurrence of a real disaster much more likely than the posterior in the previous report predicts. For example, while $\Pr(\text{ECS} > 6.0)$ equals 4.2% for the posterior in the previous report and 8.7% for the posterior in the current report, the revealed prior gives a probability range between 14.8% ($b_0 = 1.05, \sigma_0 = 0.71$) and 58.3% ($b_0 = 2.03, \sigma_0 = 1.13$). These high numbers reflect the *a priori* view of the IPCC scientists; they are not based on new information becoming available in the current report.

Figure 3: Distributions of β based on data b_0 , prior b_1 , posterior b_2 , and previous report b_2^*



In Figure 3 we plot the four representative lognormal distributions of β for the data ($b_0 = 1.00, \sigma_0 = 0.70$), the prior ($b_1 = 1.16, \sigma_1 = 0.81$), the

posterior ($b_2 = 1.07$, $\sigma_2 = 0.53$), and the posterior from previous report ($b_2^* = 1.10$, $\sigma_2^* = 0.40$). These four graphs summarize our story. The smaller the variance of the lognormal distribution, the higher the top. The posterior b_2^* from the previous reports has the smallest variance, hence the highest top and the smallest probability of a really large temperature. There is a large discrepancy between the posterior b_2^* of the previous report and the prior b_1 of the current report, showing dynamic inconsistency. The main difference between the conclusions in the previous and the current report is that uncertainty has increased. Why has uncertainty increased in the presence of many new scientific studies, often claiming rather precise predictions? This can only be because the IPCC's prior judgement has become more cautious and less optimistic, perhaps rightly so, than the previous posterior justifies.

7 Conclusions and further work

In this paper we have tried to reveal the prior, given information about the data and the posterior. We do so in the context of the normal distribution and an extension (reparametrization) based on a monotonic transformation of the normal distribution, in our case the lognormal distribution. The (log)normality assumption can be relaxed without changing the underlying theory, but the expressions would become less transparent.

We developed the theory and we applied this theory to climate sensitivity, more precisely to the prior views of the IPCC panel towards the value of this parameter. It appears that the IPCC scientists have agreed *a priori* on a value for the ECS between 3°C and 4.0°C, while judging the occurrence of a real disaster much more likely than the posterior in the previous report predicts. Our application assumes that the IPCC behaves in a Bayesian way and thinks in terms of priors and posteriors. This may not be so. But it doesn't matter because we are all Bayesians in the sense that we all have priors. Scientific Bayesians have both explicit and implicit priors; scientific frequentists only have implicit priors. But we all have priors and the purpose of this paper has been to make these priors explicit. Of course, the field of application of the ideas developed in this paper is not confined to climate is-

sues. It can be used more generally to determine people’s biases, for example politicians or scientists.

There are at least two further closely related questions worth investigating. First, suppose that our observations stretch over several, say two, periods. In both periods we have data and posteriors, and we can recover the priors in periods 1 and 2. Dynamic consistency requires that the prior in period 2 is the posterior in period 1. But is it? If it is, then the agent is consistent or rational in this Bayesian updating scheme. But if it isn’t, then the agent is not consistent. One can easily imagine a situation where the agent remains too loyal to their original prior, which one may call ‘prior stubbornness’ or ‘bunching’. This stubbornness may be politically motivated and is related to the theory of learning. It may continue until some bound has been reached (a tipping point), after which the prior is adjusted and moves to a new level. In fact, it should be possible to derive a measure for such stubbornness.

Second, we may consider the situation where there is not one but several, say two, agents over several, say two, periods. The data available to the two agents are the same, but their posteriors are not, and hence their priors are also different. Under what conditions would their priors converge in this setting?

Appendix: The lognormal distribution

Recall the fundamental formula for ‘completing the square’,

$$\frac{(y - \beta)^2}{\sigma_0^2} + \frac{(\beta - a_1)^2}{\sigma_1^2} = \frac{(\beta - a_2)^2}{\sigma_2^2} + \frac{(y - a_1)^2}{\sigma_0^2 + \sigma_1^2}, \quad (22)$$

where

$$a_2 = \frac{\sigma_0^2 a_1 + \sigma_1^2 y}{\sigma_0^2 + \sigma_1^2}, \quad \sigma_2^2 = \frac{\sigma_0^2 \sigma_1^2}{\sigma_0^2 + \sigma_1^2}.$$

Multiplying both sides by $-1/2$ and taking exponentials gives

$$f(y; \beta, \sigma_0^2) f(\beta; a_1, \sigma_1^2) = f(\beta; a_2, \sigma_2^2) f(y; a_1, \sigma_0^2 + \sigma_1^2), \quad (23)$$

where

$$f(x; \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp -\frac{1}{2} \left(\frac{x - \mu}{\sigma} \right)^2$$

denotes the normal density. This, of course, is just Bayes' formula

$$f(y|\beta) f(\beta) = f(\beta|y) f(y)$$

and shows that a normal likelihood plus a normal prior results in a normal posterior.

Now let g denote the lognormal density

$$g(x; \mu, \sigma^2) = \frac{h'(x)}{\sigma\sqrt{2\pi}} \exp -\frac{1}{2} \left(\frac{h(x) - \mu}{\sigma} \right)^2 \quad (x > 0).$$

where $h(x) = \log x$ and consequently $h'(x) = 1/x$. Then it follows immediately from (22) that, for $\beta > 0$,

$$f(y; h(\beta), \sigma_0^2) g(\beta; a_1, \sigma_1^2) = g(\beta; a_2, \sigma_2^2) f(y; a_1, \sigma_0^2 + \sigma_1^2). \quad (24)$$

Hence, a normal likelihood plus a lognormal prior results in a lognormal posterior, so that prior and posterior are conjugate distributions. The function h can be any monotonic transformation of x , not necessarily $\log x$.

Another way of arriving at this result is to realize that the lognormal distribution is not really a new distribution but rather a reparametrization. The parameter of interest remains β , but the analysis is performed on $\log \beta$ (more generally on $h(\beta)$). So we have a normal likelihood with mean $\log \beta$ and a normal prior on $\log \beta$, resulting in a normal posterior on $\log \beta$.

Acknowledgements

We are grateful to Peter Boswijk, Siem Jan Koopman, Gideon Magnus, Mark Steel, and Rick van der Ploeg for helpful discussions. Earlier versions of this paper were presented at the Vrije Universiteit Amsterdam (April 2020), the NESG Conference (May 2020), and at ISER, Osaka University (December

2020). This work was financially supported by the JSPS KAKENHI Grants Number JP20H01477 and the program of the Joint Usage/Research Center for Behavioral Economics at ISER, Osaka University.

References

- Aldrin, M., M. Holden, P. Guttorp, R. B. Skeie, G. Myhre, and T. K. Berntsen (2012). Bayesian estimation of climate sensitivity based on a simple climate model fitted to observations of hemispheric temperatures and global ocean heat content. *Environmetrics*, 23, 253–271.
- Bender, F. A.-M., A. M. L. Ekman, and H. Rodhe (2010). Response to the eruption of Mount Pinatubo in relation to climate sensitivity in the CMIP3 models. *Climate Dynamics*, 35, 875–886.
- Brock, W., and A. Xepapadeas (2019). Regional climate policy under deep uncertainty: Robust control, hot spots and learning. DEOS Working Papers #1903, Athens University of Economics and Business.
- Cooke, R. M. and B. Wielicki (2018). Probabilistic reasoning about measurements of equilibrium climate sensitivity: Combining disparate lines of evidence. *Climatic Change*, 151, 541–554.
- Hargreaves, J. C., J. D. Annan, M. Yoshimori, and A. Abe-Ouchi (2012). Can the Last Glacial Maximum constrain climate sensitivity? *Geophysical Research Letters*, 39, L24702.
- Holden, P. B., N. R. Edwards, K. I. C. Oliver, T. M. Lenton, and R. D. Wilkinson (2010). A probabilistic calibration of climate sensitivity and terrestrial carbon change in GENIE-1. *Climate Dynamics*, 35, 785–806.
- Huber, M., I. Mahlstein, M. Wild, J. Fasullo, and R. Knutti (2011). Constraints on climate sensitivity from radiation patterns in climate models. *Journal of Climate*, 24, 1034–1052.

- Hwang, C., F. Reynès, and R. S. J. Tol (2013). Climate policy under fat-tailed risk: An application of Dice. *Environmental and Resource Economics*, 56, 415–436.
- IPCC (2007). *Climate Change 2007: The Physical Science Basis*. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- IPCC (2013). *Climate Change 2013: The Physical Science Basis*. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Knutti, R., M. A. A. Rugenstein, and G. C. Hegerl (2017). Beyond equilibrium climate sensitivity. *Nature Geoscience*, 10, 727–736.
- Köhler, P., R. Bintanja, H. Fischer, F. Joos, R. Knutti, G. Lohmann, and V. Masson-Delmotte (2010). What caused Earth’s temperature variations during the last 800,000 years? Data-based evidence on radiative forcing and constraints on climate sensitivity. *Quaternary Science Reviews*, 29, 129–145.
- Lewis, N. (2013). An objective Bayesian, improved approach for applying optimal fingerprint techniques to estimate climate sensitivity. *Journal of Climate*, 26, 7414–7429. doi:10.1175/JCLID-12-00473.1.
- Libardoni, A. G., and C. E. Forest (2011). Sensitivity of distributions of climate system properties to the surface temperature dataset. *Geophysical Research Letters*, 38, L22705. Correction in: *Geophysical Research Letters* (2013), 40, 2309–2311. doi:10.1002/grl.50480.
- Lin, B., L. Chambers, P. Stackhouse Jr., B. Wielicki, Y. Hu, P. Minnis, N. Loeb, W. Sun, G. Potter, Q. Min, G. Schuster, and T.-F. Fan (2010). Estimations of climate sensitivity based on top-of-atmosphere radiation imbalance. *Atmospheric Chemistry and Physics*, 10, 1923–1930.

- Lindzen, R. S., and Y. S. Choi (2011). On the observational determination of climate sensitivity and its implications. *Asia-Pacific Journal of Atmospheric Sciences*, 47, 377–390.
- Miller, A., S. Dietz, and G. Heal (2013). Scientific ambiguity and climate policy. *Environmental and Resource Economics*, 55, 21–46.
- Olson, R., R. Sriver, M. Goes, N. M. Urban, H. D. Matthews, M. Haran, and K. Keller (2012). A climate sensitivity estimate using Bayesian fusion of instrumental observations and an Earth system model. *Journal of Geophysical Research: Atmospheres*, 117, D04103.
- Oppenheimer, M., B. C. O’Neill, and M. Webster (2008). Negative learning. *Climatic Change*, 89, 155–172.
- Otto, A., F. E. L. Otto, O. Boucher, J. Church, G. Hegerl, P. M. Forster, N. P. Gillett, J. Gregory, G. C. Johnson, R. Knutti, N. Lewis, U. Lohmann, J. Marotzke, G. Myhre, D. Shindell, B. Stevens, and M. R. Allen (2013). Energy budget constraints on climate response. *Nature Geoscience*, 6, 415–416.
- Palaeosens Project Members (2012). Making sense of palaeoclimate sensitivity. *Nature*, 491, 683–691.
- Schmittner, A., N. M. Urban, J. D. Shakun, N. M. Mahowald, P. U. Clark, P. J. Bartlein, A. C. Mix, and A. Rosell-Melé (2011). Climate sensitivity estimated from temperature reconstructions of the last glacial maximum. *Science*, 334, 1385–1388. Response to comment in: *Science* (2012), 337, 1294.
- Schwartz, S. E. (2012). Determination of Earth’s transient and equilibrium climate sensitivities from observations over the twentieth century: Strong dependence on assumed forcing. *Surveys in Geophysics*, 33, 745–777.