An Objective Bayesian Improved Approach for Applying Optimal Fingerprint Techniques to Estimate Climate Sensitivity

DOI: 10.1175/JCLI-D-12-00473.1

Supporting Information

Nicholas Lewis[1]

Bath, UK

S1. Bayesian updating with differing nonlinear data–parameter relationships

It is not widely appreciated that use of a noninformative prior when generating a Bayesian posterior PDF for a parameter will only result in objective inference upon using that posterior as the prior in a second application of Bayes' theorem (Bayesian updating) if the data from the second 'experiment' has the same form of relation to the parameter as that in the first 'experiment'. That is, either the data are linearly related to the parameter in both cases, or they have the same form of nonlinear relationship. Otherwise the noninformative prior for Bayesian inference from each experiment separately will differ, and that for inference from the combined data from the two experiments will be different from either of those separate priors. This is most easily shown by carrying out a large number of trials in a simple case, using noninformative Jeffreys' priors computed from Fisher's information matrix.

[1] *Corresponding author address*: Nicholas Lewis, Walden, Widcombe Hill, Bath, United Kingdom.
E-mail: nhlewis@btinternet.com

Take a univariate parameter and two experiments, A and B. In experiment A, the datum represents the value, measured with error, of the parameter. In experiment B, the datum represents the value, measured with error, of the reciprocal of the parameter. The true parameter value is θ=1, and the errors have independent $N(0, 0.25^2)$ distributions (standard deviation = 0.25). It is easily proven that the noninformative prior for experiment A is uniform, that for experiment B is proportional to θ^{-2}, while that for the combined experiments is proportional to $(1+\theta^{-4})^{1/2}$. 30,000 random draws of the measured datum were made for each experiment and a posterior PDF for the parameter computed on various combinations of prior and experimental data. The percentage point of each CDF (the integral of the posterior PDF) at the true parameter value was then recorded, and the cumulative proportion of instances was plotted against those percentage points. Perfect frequentist probability coverage is represented by, at each CDF percentage point *x*, the parameter value corresponding to that percentage point of the posterior CDF exceeding the true parameter value in *x*% of cases.

The left hand panel of Figure S1 shows that for experiments A and B on their own, using their own noninformative priors (respectively uniform and proportional to θ^{-2}), frequentist coverage probabilities are accurately matched, but that using the each other's noninformative priors that is far from the case.

The right hand panel of Figure S1 shows that for experiments A and B combined, coverage probability matching is inaccurate – in different directions – using either the uniform or θ^{-2} priors. The two cases correspond to analyzing data from one experiment first, using a prior that is noninformative for that experiment on its own, and then incorporating the results of the other experiment using Bayesian updating. However,

coverage probability matching is accurate when the calculated noninformative prior (being $(1+\theta^{-4})^{1/2}$) for the combined experiments is used.

S2. Deriving the $mF_{m,\nu}$ distribution used in the Forest studies for a flat parameter surface

The $mF_{m,\nu}$ distribution that F01 states Δr^2 has follows, when the parameter surface is flat and so everywhere coincident with the tangent hyperplane, from the geometry involved. By definition, the $\kappa F_{\kappa,\nu}$ distribution in (1) is ν times the ratio of a X^2_κ distribution to an independent X^2_υ distribution. The numerator X^2_κ distribution is that of the sum of squares of κ independent random variables each having a $N(0,1)$ distribution. Without variance uncertainty, the whitened differences, whose sum of squares is $\tilde{\boldsymbol{s}}^2$, would fit this description were they zero-mean, as they are at the true parameter settings. The variance uncertainty, which we set aside temporarily, is allowed for by the X^2_υ denominator distribution. Since at the true parameter settings $\tilde{\boldsymbol{s}}$ has a $N(0, \mathbf{I}_\kappa)$ normal distribution, any m-vector $\boldsymbol{y} = \mathbf{E}\tilde{\boldsymbol{s}}$, where $\mathbf{E}$ is any $(m \times \kappa)$ matrix satisfying $\mathbf{E}\mathbf{E}^{\mathrm{T}} = \mathbf{I}_m$, will have a $N(0, \mathbf{I}_m)$ distribution (Mardia et al., 1979, Theorem 3.2.1). The m equations defining a set of orthogonal unit length basis vectors for the tangent hyperplane to the parameter surface at the best-fit point, $\hat{\boldsymbol{s}}$, with origin thereat, form such a matrix $\mathbf{E}$, with $\boldsymbol{y}$ representing the projection of $\tilde{\boldsymbol{s}}$ onto that hyperplane. Define the 'delta error' $\mathbf{s}_\Delta = \mathbf{s} - \hat{\mathbf{s}}$ as the vector in whitened difference space from the best-fit differences to those at the true parameter values. By hypothesis, $\mathbf{s}_\Delta$ lies in the tangent hyperplane and so is orthogonal to $\hat{\boldsymbol{s}}$, which meets that hyperplane at its origin.

Therefore, $\mathbf{s}_\Delta$ represents the projection of **s** onto that hyperplane, and thus has a $N(0, \mathbf{I}_m)$ distribution. Further, since the tangent hyperplane is orthogonal to the best-fit whitened differences vector, at the true parameter values point, where $\tilde{\mathbf{s}} = \mathbf{s}$:

$$\|\tilde{\mathbf{s}}\|^2 = \|\hat{\mathbf{s}}\|^2 + \|\mathbf{s}_\Delta\|^2 \qquad \text{(S1)}$$

Combining (S1) with (2) and (5), at the true parameter settings:

$$\Delta r^2 = \|\tilde{\mathbf{s}}\|^2 - \|\hat{\mathbf{s}}\|^2 = \|\mathbf{s}_\Delta\|^2 = \mathbf{s}_\Delta^{\mathrm{T}}\mathbf{s}_\Delta \qquad \text{(S2)}$$

Since $\mathbf{s}_\Delta \sim N(0, \mathbf{I}_m)$, $\|\mathbf{s}_\Delta\|^2 \sim X_m^2$, ignoring the variance uncertainty. Reintroducing the uncertainty in variance by dividing by X_ν^2/ν, we turn the X_m^2 distribution into a $mF_{m,\nu}$ distribution and obtain (4), confirming the F01 Eq.7 result.

However, the parameter surface will only be flat if the relationships between the model-predictions and some 3D function of the parameters are linear (as equations defining a 3D hyperplane, being a linear subspace of the embedding vector space, are). One would not expect that to be the case here. Therefore, for any parameter combination, including the true combination with whitened differences ***s*** of length $r^2 = \mathbf{s}^{\mathrm{T}}\mathbf{s}$, the delta-error vector $\mathbf{s}_\Delta = \mathbf{s} - \hat{\mathbf{s}}$, with squared length $r_\Delta^2 = \mathbf{s}_\Delta^{\mathrm{T}}\mathbf{s}_\Delta$, that joins the whitened difference vectors for the true and for the best-fit parameter combinations, will not be exactly orthogonal to $\hat{\mathbf{s}}$. Hence $\Delta r^2 = r^2 - r_{\min}^2 \neq r_\Delta^2$. Near the best-fit point, the parameter surface can be expected to be convex viewed from the whitened difference space origin: concavity would be inconsistent with that point being the best-fit. Therefore, at least locally, $\Delta r^2 > r_\Delta^2$. With $\Delta r^2 > r_\Delta^2$, the $\Delta r^2 \sim mF_{m,\nu}$ treatment will undersize parameter value confidence regions. The $mF_{m,\nu}$ distribution in (4) is based on

varying parameter values causing the end-point of $\tilde{\boldsymbol{s}}$ just to move in the tangent plane, where only $m=3$ d.f. are available. As the parameter surface curves away from the tangent hyperplane, the component of $\mathbf{s}_\Delta$ that is orthogonal thereto can access an additional $(\kappa - m)$ d.f., which can increase Δr^2 substantially.

For the univariate *do* diagnostic, by necessity $m = 1$ and $\Delta r^2 = r^2$, since a perfect fit to the single observation can in principle be obtained by varying any climate parameter. The resulting $F_{1,v}$ distribution for r^2 equates to a t_v distribution for r.

S3. Covariance estimation and variance uncertainty

AT99 state that the X^2_m and $mF_{m,v}$ distributions can be regarded as equivalent if $v > 100$, and also that a ~ 50% increase in d.f. is gained where overlapped vectors are used, compared with the number of vectors that have no overlap. F06 used (C. Forest, pers. comm.) a *sfc* diagnostic control matrix derived from the near 1700 year HadCM2 extended control-run (Johns et al., 1997), which provides over 160 overlapping data vectors each comprising four latitudinal band mean temperatures for five non-overlapping decades, as anomalies from the mean for nine decades ending with the last of those five decades. Allowing for the first four decades' data only being usable to anomalise the first vector, 33 non-overlapping five-decade sets of control data exist. A 50% increase implies 49 d.f. For the revised six decade *sfc* diagnostic, the corresponding calculation implies 40 d.f. These two d.f. are used, as appropriate, in all our results. However, the correlations between data from different decades are low. Therefore overlapping should not greatly reduce the effective d.f. in the control data below 160. Forming a 360 column control data matrix from full 72 latitude band resolution HadCM2 data supports this contention. Its log-eigenvalue spectrum declines

only linearly until beyond the 150th EOF. Therefore, there seems no real need to distinguish between F and X^2 distributions in respect of the *sfc* diagnostic. The differences between simulated natural variability in the AOGCM control-run and in 20th century actual climate, insofar as represented in the retained EOFs, are likely to be of more significance.

While the *ua* diagnostic HadCM2 control data, of non-overlapping vectors, has (after centering) only 39 d.f., used in all our results, the *ua* diagnostic constrains probability distribution much less than the *sfc* diagnostic, with its whitened differences generally being substantially smaller that the *sfc* diagnostic ones. Therefore, variance uncertainty in the *ua* diagnostic whitened differences, other than when computing likelihoods, is sufficiently unimportant to be ignorable.

Although the *do* diagnostic gives only one whitened difference, it is important since it discriminates strongly against S_{eq} and K_v both being high. The *do* temperature trend uncertainty estimate is large and has relatively few d.f., implying the trend estimate has a wide-tailed *t*-distribution. A *t*-distribution is equivalent to a weighted mixture of differing variance normal distributions, good emulation being possible with four normals. That computational approach has been adopted, enabling the mathematical analysis to assume a normal distribution; which of the four emulation components is involved only affects the size of the *do* whitened difference. Four separate joint combined diagnostic likelihoods and noninformative priors are computed, and a weighted average of their products taken prior to the final computation of the parameter PDFs.

S4. Derivations of the geometrical error space volume adjustment and the X^2_m distribution

When $m > 2$, the $mF_{m,v}$ density goes to zero as the best-fit point – where the probability of the true parameter values is highest – is approached. This occurs because the *F*-PDF gives the probability density in terms of increments in the normalized squared error sum involved, $\Delta r^2/m$, not of volume elements in error variable space. As the best-fit point is approached and $\Delta r^2 \to 0$, the volume in error space in a shell containing points for which the error sum-of-squares is between Δr^2 and $\Delta r^2 + \delta\Delta r^2$, the fixed increment $\delta\Delta r^2$ being minute, becomes vanishingly small. For the Forest studies, where $m = 3$, 3D Euclidean geometry shows that the volume of an infinitesimally thin shell scales with $\sqrt{\Delta r^2}$. The *F*-distribution density function should therefore be divided by this factor. Without a volume adjustment, the joint probability density for the delta-errors, and the parameters, is underestimated around the best-fit point.

Now consider the general dimension case. We have a *p*-dimensional Euclidean whitened difference space of independent real random variables, (x_1, x_2, …, x_p) or **x** in vector form. By presumption, at the true parameter settings the variables have independent $N(0,1)$ probability distributions. Their joint probability density at any point in the space is therefore

$$p(\mathbf{x}) = \prod_{i=1:p} (2\pi)^{-0.5} \, \mathrm{e}^{x_i^2/2} = (2\pi)^{-p/2} \, \mathrm{e}^{\mathbf{x}^{\mathrm{T}}\mathbf{x}/2} = (2\pi)^{-p/2} \, \mathrm{e}^{\|\mathbf{x}\|^2/2} \qquad \text{(S3)}$$

One can express this probability density as a function of the distance *r* from the origin in the *p*-dimensional space, or of r^2. The volume in that space occupied by a spherical shell of thickness *dr*, centered on the origin, is *dr* multiplied by the surface area of a *p*-dimensional hypersphere of radius *r*, being:

$$2\pi^{p/2} r^{p-1} / \Gamma(\tfrac{p}{2})\, dr \qquad \text{(S4)}$$

As probability density in the whitened difference space is symmetrical around the origin, it is the same throughout the shell. The total probability for the shell is therefore its volume multiplied by $p(\mathbf{x})$. Conversion of probability densities on a change of variables, reflecting relative infinitesimal volumes, then gives the probability density in p-space as a function of r as:

$$(2\pi)^{-p/2}\, \mathrm{e}^{-r^2/2}\;\; 2\pi^{p/2} r^{p-1} / \Gamma(\tfrac{p}{2}) \qquad \text{(S5)}$$

Expressed as a density in terms of r^2, this becomes (since $\frac{\partial r}{\partial r^2} = \frac{1}{2r}$):

$$p_{r^2}(r^2) = \frac{1}{2(r^2)^{1/2}} \mathrm{e}^{-r^2/2}\;\; (r^2)^{p/2-1/2}\;\; \frac{2}{2^{p/2}\Gamma(\frac{p}{2})} = \frac{(r^2)^{p/2-1}\, \mathrm{e}^{-r^2/2}}{2^{p/2}\Gamma(\frac{p}{2})} \qquad \text{(S6)}$$

which is, as it should be, the density function of the Chi-squared distribution with p degrees of freedom. Therefore, comparing (S6) with (S3), in order to convert the density of a Chi-squared distribution with p degrees of freedom into a density for the p variables whose sum-of-squares is the argument of that Chi-squared distribution, one must divide the Chi-squared density by $\pi^{p/2}(r^2)^{p/2-1} / \Gamma(\frac{p}{2})$. For the F06 Δr^2 method, $p = 3,\ \Gamma\left(\frac{3}{2}\right) = \sqrt{\pi}/2$, so the divisor is $2\pi\sqrt{\Delta r^2}$ (it being assumed that $\Delta r^2 = r_\Delta^2$), confirming the result arrived at by considering 3D Euclidian geometry.

S5. Differential geometry approach to deriving a joint PDF for the parameters

The result in (1) and (B16) can be reached directly using differential geometry methods to obtain a conditional probability density in model (parameter) space where a higher dimensional data space, representing data obeying the model subject to independent

errors, is involved. In essence, the PDF conversion factor/noninformative prior equates to the induced metric on the parameter surface as a submanifold of the combined model-parameter, data space. For further details of this approach to probabilistic inference, see Mosegaard and Tarantola (2002). Equation (46) thereof reduces to (B16) in the case here: there is no prior information as to parameter values, so $\rho_{\mathrm{m}}(\mathbf{m}) \propto \sqrt{\det \mathbf{g}_{\mathrm{m}}}$ per their Rule 2, with the model space metric $\mathbf{g}_{\mathrm{m}} \sim 0$ since parameter space in principle extends infinitely; $\rho_d(\boldsymbol{d})$ (the divisor of which is misprinted and should read $\sqrt{\det \mathbf{g}_{\mathrm{d}}}$ not $\sqrt{\det \mathbf{g}_{\mathrm{m}}}$) equates to $p(\tilde{\mathbf{w}}_{\mathrm{c}} \,|\, \boldsymbol{\theta}_{\mathrm{t}})$; and the data space metric $\mathbf{g}_{\mathrm{d}} \propto \mathbf{I}_{\kappa_c}$ since the whitened observations are independent unit variance normally distributed variates. (This metric implies that the data space is flat, with the same probability density everywhere in the absence of any information about even the dependence of the data on parameter values.) The Mosegaard and Tarantola formulation shows that the noninformative prior represents a mapping of probability that adjusts for relative volumes in data space (where the likelihood exists) and parameter space, reflecting their metrics. It conveys no beliefs about model parameter values.

S6. Probability density in 3D delta-error space

Consider the case for the *sfc* or *ua* diagnostic on its own: a different analysis applies for the univariate *do* diagnostic. Suppose a flat parameter surface, so that the $F_{3,\nu}$ PDF used in F06 is valid. Then what that provides, after dividing by first a factor of 3, to convert from an *F*-distribution density to a Chi-squared distribution density, and then by the geometric volume adjustment factor of $2\pi\sqrt{\Delta r^2}$, is the probability density for three orthogonal estimated $N(0,1)$ distributed delta-error differences, taken as lying in the 3D

tangent hyperplane to the parameter surface at the best-fit point. After division by the volume adjustment factor, the PDF for the X_3^2 distribution constituting the numerator of the $F_{3,\nu}$ distribution becomes:

$$p_z(X_3^2(\Delta r^2)) / (2\pi\sqrt{\Delta r^2}) \propto \mathrm{e}^{-\Delta r^2/2} \qquad \text{(S7)}$$

The $F_{3,\nu}$ distribution, after the volume correction, thus gives a probability density in the 3D delta-error space, in successive spherical shells surrounding the best-fit point, which density must from symmetry considerations be uniform within each shell. Each point in any such shell corresponds to a point in parameter space, but the volumes in parameter space corresponding to equal volumes in 3D delta-error (sub)space will vary according to the point in parameter space involved, since the rate at which model estimates vary with each of the parameters depends on the parameter combination involved.

Expressing Δr^2, the squared length of a vector in the 3D subspace, as $\Delta r^2 = a_1^2 + a_2^2 + a_3^2$ in terms of the lengths a_1, a_2 and a_3 of vectors in that subspace orientated in the three directions of any set of orthonormal basis vectors (say, $\mathbf{e}_1$, $\mathbf{e}_2$ and $\mathbf{e}_3$), we have from equation (S7):

$$p_{\Delta r^2}(X_3^2(\Delta r^2)) / \sqrt{\Delta r^2} \;\propto\; \mathrm{e}^{-a_1^2/2}\;\; \mathrm{e}^{-a_2^2/2}\;\; \mathrm{e}^{-a_3^2/2} \qquad \text{(S8)}$$

This is proportional to the joint probability density of three independent $N(0,1)$ delta-error variates defined by the 3D delta-error vector. From the definition of a X_3^2 distribution, with the best-fit point as the origin, Δr^2 is the sum of the squares of three independent $N(0,1)$ random variables. Since variables representing mixtures of such normal variates scaled to unit variance also have a $N(0,1)$ distribution and, if

orthogonal, are independent, a_1, a_2 and a_3 will be independent $N(0,1)$ random variables, whatever set of basis vectors is chosen.

A joint posterior PDF for the parameters on the assumption that the parameter surface is flat, resulting in the 3D delta-error space comprising the tangent hyperplane to that hypersurface at the best-fit point being coincident with it everywhere, can therefore be derived as follows. Compute Δr^2 for the combined whitened errors from all diagnostics and thus a joint likelihood for the 3D delta-errors using a $3F_{3,\nu}$ distribution and equation (S8). The Bayesian step in (B5) then relates to 3D delta-errors, a uniform prior again being noninformative since independent location variables are involved. Here there is a one-to-one relationship between equi-dimensional delta-errors and parameters, and we obtain the joint posterior for the parameters simply by performing a standard change of variables thereto. The conversion factor in the usual formula for converting PDFs upon a change of variables with unchanged dimensionality equals the absolute Jacobian determinant for that transformation, (Mardia et al., 1979, Eq. 2.4.1). Here, that determinant equates to a noninformative prior for the parameters. Its form is as expected: upon a one-to-one reparameterisation, reference (minimum information) priors transform in accordance with the usual conversion of PDFs formula (Bernado and Smith, 1994, Prop. 5.22). Both the best-fit point and the 3D delta-errors must be determined using the combined errors from all three diagnostics to retain the validity of the 3D delta error analysis and of the computation of the Jacobian determinant.

Using the Δr^2 method, and thereby making the parameter surface flatness assumption, but otherwise using our method, marginal posterior PDFs for the parameters are

narrower than when not making that assumption. The effect is larger when the parameters are less well constrained by the data.

S7. Incorporating uncertainties in non-aerosol forcings and in surface temperatures

F06 did not incorporate allowance for uncertainties in forcings other than aerosols (including, implicitly, any forcings having a similar spatio-temporal profile to aerosols). Nor did F06 incorporate allowance for uncertainties in surface temperature observations, on the basis that they were small relative to natural variability. However, neither of these types of uncertainty is negligible. Although it is not possible to make accurate allowance for them, we attempt to do so approximately, as follows.

First, we estimate uncertainty in surface temperature observations. Total measurement error, sampling and bias uncertainty in HadCRUT4 global mean decadally smoothed data averages 0.022 K, as a standard deviation, over 1902–2001, and 0.018 K over 1942–2001. Coverage error is not included here since a gridded masking process is used in deriving the surface diagnostic. Assuming inter-decadal independence of errors, the standard deviation for the 9 decade average climatology data is 0.007 K, which when added in quadrature to 0.018 K for the diagnostic data gives 0.02 K, rounding up. Since four-zone data was used, and the observational error with coverage error excluded is not strongly latitude dependent, that figure is simply doubled to give an approximation to zonal observational error. That brings the total observational uncertainty standard deviation for each diagnostic decade and zone to 0.04 K.

Next, we estimate uncertainty in non-aerosol forcings. We make the following estimates of 5–95% uncertainties for forcings: greenhouse gases ±20%; tropospheric ozone ±0.20 Wm^{-2}, stratospheric ozone ±0.15 Wm^{-2}, land-use ±0.10 Wm^{-2}, and in relation to forcings not explicitly modeled, stratospheric water vapor ±0.05 Wm^{-2} and combined contrails and contrail induced cirrus ±0.07 Wm^{-2}. All these uncertainties are assumed to be independent and normally distributed. Greenhouse gas forcing uncertainty is specified in fractional terms since it relates almost entirely to the radiative effect of accurately known atmospheric concentrations. Applying a ±20% uncertainty range to the greenhouse gas forcing increase of 2.0 Wm^{-2} during 1902–2001 produces a standard deviation of 0.24 Wm^{-2}. The remaining uncertainties are assumed to apply to changes in forcings. Adding them in quadrature and likewise converting to a standard deviation gives a non-greenhouse gas forcing uncertainty standard deviation of 0.17 Wm^{-2}. In principle only part of the greenhouse gas forcing uncertainty need be taken into account when deriving an error range for S_{eq}, since that uncertainty will oppositely affect the calibration of the model run S_{eq} values. However, no reduction is made on that account, since there are undoubtedly other uncertainties, such as regarding solar and volcanic forcings and model error, that have been omitted. Accordingly, adding in quadrature 0.17 Wm^{-2} to 0.24 Wm^{-2} gives a total forcing uncertainty of 0.30 Wm^{-2}.

The approach used to incorporate the observational surface temperature and the forcing uncertainty estimates is to treat the originally estimated posterior PDF for S_{eq} as arising from an energy balance estimation method with normally distributed uncertainties as to changes both in global surface temperature and in forcing (net of radiative imbalance). On that basis, the PDF for S_{eq} is then fitted to a PDF arising as the ratio of two normal distributions, by selecting the means and standard deviations of

those distributions. It is necessary to fix the change in global surface temperature assumed, which forms the mean of the numerator normal distribution. Since the non-GHG forcing uncertainty estimates relate to the change since preindustrial times, the rise in mean global surface temperature from 1851–1860, the first decade in the HadCRUT4 record, to 1992–2001, the final surface diagnostic decade, is used. That rise is fairly similar to the rise between the part of the climatological period before the diagnostic decades and 1992–2001. Since only the standard deviation of the denominator distribution affects the skewness of the PDF, the mean of the denominator normal distribution and the standard deviations of both normal distributions can then be fitted without difficulty to produce a PDF matching the posterior PDF for S_{eq} fairly closely.

Two variants of this method were used to derive an adjusted 5–95% range for S_{eq}. In the first, a ratio of normal distributions was fitted directly to the main results PDF for S_{eq}. The observational surface temperature uncertainty and the aggregate forcing uncertainty were then added in quadrature to the standard deviations of respectively the numerator and denominator normal distributions, the PDF arising from their ratio recomputed and its 5–95% range determined. The estimated zonal observational surface temperature uncertainty was not reduced on account of it being used in a global energy balance situation, as in practice it is the uncertainties in zonal temperatures that affect the actual estimation of S_{eq}. In the second variant, the main results PDF for S_{eq} was recomputed with the zonal observational surface temperature uncertainties added when deriving the whitened surface diagnostic variables and their error sum-of-squares, giving a PDF that allowed for observational surface temperature uncertainty. A ratio of normal distributions was then fitted as before. The aggregate forcing uncertainty was

then added to the standard deviation of the denominator normal distribution, the PDF arising from the ratio of the two normal distributions recomputed and its 5–95% range determined. The second variant of the uncertainty estimation method is theoretically more appropriate, but could turn out to be less robust in practice. Both variants produce fairly similar results, with the PDF for S_{eq} becoming substantially more skewed – slightly more so when using the second variant – and wider. Taking the average change for the two variants of the method, the median estimate for S_{eq} is virtually unchanged as compared with the main results, whilst the mode decreases marginally, reflecting the increased skewness of the PDF. To be conservative, the outer of each pair of 5% and 95% points from the two variants is used, giving an adjusted 5–95% uncertainty range for S_{eq} of 1.0 to 3.0 K. Figure S3 shows the PDF for S_{eq} estimated as the average of the PDFs from the two variants of the uncertainty estimation method.

Data and computer code relating to the paper are available at http://niclewis.wordpress.com/objective-bayesian-sensitivity/ or from the author upon request

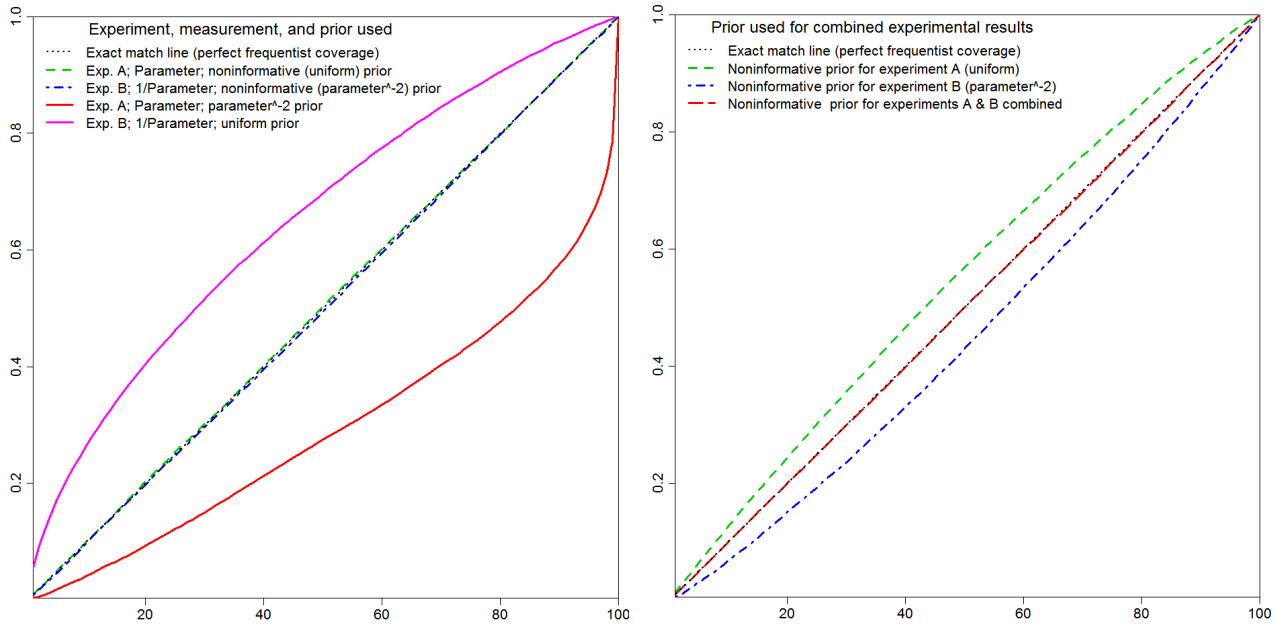

FIG. S1 Coverage probabilities. Measurements, subject to random error, are of variables corresponding to : Experiment A – parameter itself; Experiment B – reciprocal of parameter.

x-axis: 0–100 percentage points of the parameter posterior CDF from 30,000 randomly drawn experimental measurements

y-axis: proportion of random draws for which, at each x-axis CDF percentage point, the parameter value for that percentage point of their posteriors exceeds the true parameter value

Left hand panel: inference based on data from experiments A and B separately, using the correct noninformative prior for that experiment (green and blue lines respectively) or the noninformative prior for the other experiment (red and pink lines respectively).

Right hand panel: inference based on combined experiments, using the noninformative prior for experiment A (green line), that for experiment B (blue line) or using the correct noninformative prior for the combined experiments (red line).

Black dotted lines show perfect frequentist coverage probabilities (both panels).

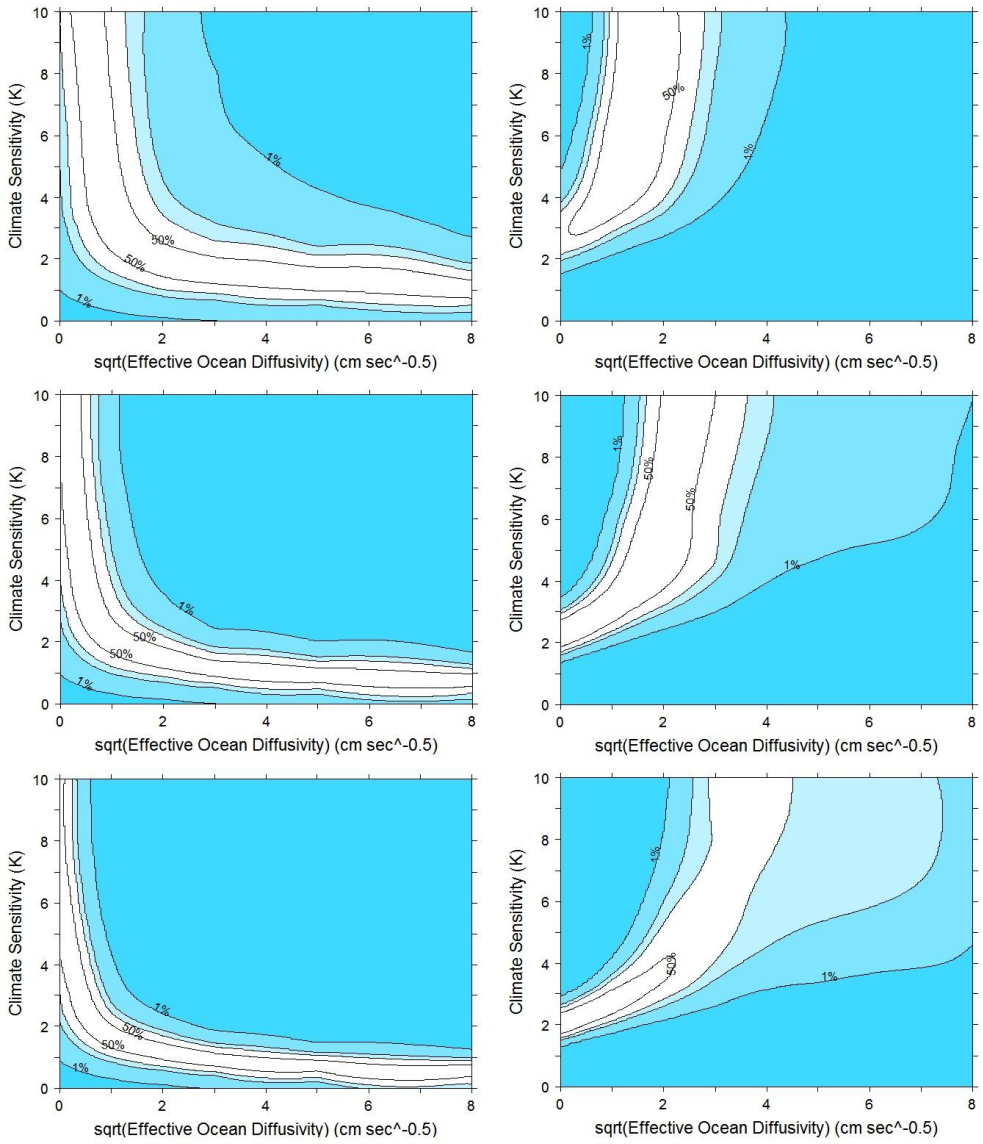

FIG. S2 Joint rejection regions in $S_{eq} - \sqrt{K_v}$ space, using the original F06 diagnostics, based on the likelihood from the deep-ocean diagnostic only (left hand panels) and the surface diagnostic (at $\kappa_{sfc} = 16$) only (right hand panels). The top, middle and bottom panels show CDF regions when probabilities are conditioned on respectively F_{aer} = -0.75, -0.5 and -0.25 Wm^{-2}, using the best-fit $S_{eq} - \sqrt{K_v}$ combination at the relevant F_{aer} level. For the surface diagnostic, where $r^2_{\min}$ varies with F_{aer}, probability regions are therefore not comparable between the upper, middle and lower panels. Shading shows rejection regions at significance levels of 20% (lightest), 10% and 1% (darkest). A labeled contour line shows the rejection boundary at a 50% significance level.

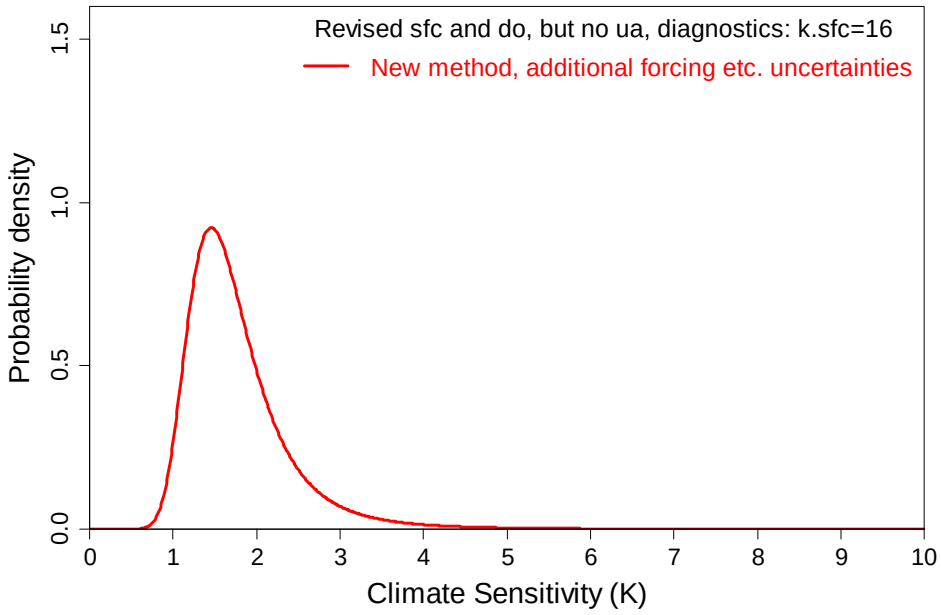

FIG. S3 Marginal posterior PDF for climate sensitivity estimated using the new objective Bayesian method and revised diagnostics, with allowance made for additional uncertainties in forcings and surface temperature observations. Estimated as the unweighted mean of the PDFs arising from the two variants of the uncertainty estimation method used.