Combining Emergent Constraints for Climate Sensitivity

Christopher Bretherton[1] and Peter Caldwell[2]

[1]University of Washington and Vulcan Inc.
[2]Lawrence Livermore National Laboratories

November 22, 2022

Abstract

A method is proposed for combining information from several emergent constraints into a probabilistic estimate for a climate sensitivity proxy \$Y\$ such as equilibrium climate sensitivity (ECS) or the climate feedback parameter \$\lambda\$. The method is based on fitting a multivariate Gaussian PDF for \$Y\$ and the emergent constraints using an ensemble of global climate models (GCMs). For a single perfectly-observed constraint \$X\$, it reduces to a linear regression-based estimate of \$Y\$. The method accounts for uncertainties in sampling this multidimensional PDF with a small number of models, for observational uncertainties in the constraints, and for overconfidence about the correlation of the constraints with the climate sensitivity. Two methods are presented. Method C accounts for correlations between emergent constraints but can fail if some constraints are too strongly related. Method U assumes constraints are uncorrelated except through their mutual relationship to the climate proxy; it is robust to small GCM sample size and is appealingly interpretable. These methods are applied to ECS and \$\lambda\$ using a previously-published set of 11 possible emergent constraints derived from climate models in the Coupled Model Intercomparison Project (CMIP). This study corroborates and quantifies past findings that most constraints predict higher climate sensitivity than the CMIP mean. The \$\pm2\sigma\$ posterior range of ECS for Method C with no overconfidence adjustment is \$4.1 \pm 0.8\$ K. For Method U with a large overconfidence adjustment, it is \$4.0 \pm 1.3\$ K.

Combining Emergent Constraints for Climate Sensitivity

Christopher S. Bretherton*

University of Washington, Seattle, WA

Peter M. Caldwell

Lawrence Livermore National Laboratory, Livermore CA

Submitted to *J. Climate*, Dec. 6, 2019

**Corresponding author address:* Christopher S. Bretherton, Department of Atmospheric Sciences, University of Washington, Seattle, WA 98195-1640

E-mail: breth@uw.edu

ABSTRACT

A method is proposed for combining information from several emergent constraints into a probabilistic estimate for a climate sensitivity proxy Y such as equilibrium climate sensitivity (ECS) or the climate feedback parameter λ. The method is based on fitting a multivariate Gaussian PDF for Y and the emergent constraints using an ensemble of global climate models (GCMs). For a single perfectly-observed constraint X, it reduces to a linear regression-based estimate of Y. The method accounts for uncertainties in sampling this multidimensional PDF with a small number of models, for observational uncertainties in the constraints, and for overconfidence about the correlation of the constraints with the climate sensitivity. Two methods are presented. Method C accounts for correlations between emergent constraints but can fail if some constraints are too strongly related. Method U assumes constraints are uncorrelated except through their mutual relationship to the climate proxy; it is robust to small GCM sample size and is appealingly interpretable. These methods are applied to ECS and λ using a previously-published set of 11 possible emergent constraints derived from climate models in the Coupled Model Intercomparison Project (CMIP). This study corroborates and quantifies past findings that most constraints predict higher climate sensitivity than the CMIP mean. The $\pm 2\sigma$ posterior range of ECS for Method C with no overconfidence adjustment is 4.1 ± 0.8 K. For Method U with a large overconfidence adjustment, it is 4.0 ± 1.3 K.

1. Introduction

Climate change is a defining problem of our time. It is hard to plan for future warming without knowing its magnitude, but our $\pm 1\sigma$ 'likely' confidence range for equilibrium climate sensitivity (ECS, the global-average surface warming due to doubling CO_2 and letting the climate re-equilibrate) is currently 1.5-4.5 K (Stocker et al. 2013) - which is disturbingly large. This uncertainty has persisted for decades despite large advances in our understanding of the climate system (Knutti et al. 2017).

Emergent constraints offer a possible path to narrowing this spread. An emergent constraint is a current-climate quantity which has skill at predicting future changes in climate. Such predictors may be valuable shortcuts to the complex and uncertain process of directly simulating climate change in a general circulation model (GCM) or inferring it from imperfect observational records. Because the physical processes governing climate change are generally the same ones that control present-day seasonal, weather-scale, and diurnal variations, it is likely that real emergent constraints exist. Hall and Qu (2006) was one of the first papers to identify such a constraint. They found that the seasonal cycle of snow albedo over northern-hemisphere land is tightly correlated with snow albedo feedback over this region in 17 model simulations from the 3rd phase of the Coupled Model Intercomparison Project (CMIP). This emergent constraint has an obvious motivation: surface warming reduces snow cover irrespective of whether that warming is due to seasonal changes in insolation or CO_2-induced climate change. Nearly 40 other emergent constraints have been proposed since 2006 (Hall et al. 2019), though few have had such a satisfying physical explanation.

Several limitations and assumptions apply to the use of emergent constraints to predict ECS. First, the ECS simulated by a climate model is generally estimated from an integration of finite length (customarily 150 years in CMIP5) that is not fully equilibrated and keeps certain physics fixed (e. g. vegetation type and land ice). The response of that model (or the real climate system) to a time-varying radiative forcing is not determined purely by the ECS, but also depends upon natural and forced changes in the pattern of surface warming, e.g. Armour et al. (2013); Gregory and Andrews (2016). ECS is also a problematic target for emergent constraints because it arises from interaction between many processes. As a result, it is questionable whether any single current-climate variable would explain a large fraction of ECS variability. This is why Klein and Hall (2015) suggest that emergent constraints should be targeted towards a single climate feedback mechanism (e.g. snow cover) whenever possible. Nevertheless, many studies (including this one) focus on emergent constraints for global climate sensitivity proxies such as ECS or the climate feedback parameter λ (Cess et al. 1989) because of their importance. Lastly, emergent constraints in general derive from a blend of scientific reasoning and a posteriori optimization to maximize their correlation with ECS or λ over a modest set of GCMs, and only the most promising constraints are likely to be published. This suggests a risk of constraints being 'overconfident', i. e. better correlated with ECS or λ over the GCMs on which they were first tested and optimized than in another independent set of GCMs.

Emergent constraints have already been noted to predict larger climate sensitivity than expected from other lines of evidence (Tian 2015; Klein and Hall 2015). If agreement between constraints gives us confidence in their predictions, this is an alarming finding. A goal of this paper is to develop an approach for combining emergent constraints to provide a confidence range for ECS or any other climate sensitivity proxy, while accounting for the issues just raised.

The emergent constraints used and the relevant data are described in Sect. 2, and terminology is described in Sect. 3. In Sect. 4, we derive and apply our method to individual emergent constraints with observational uncertainty. Sect. 5 discusses correlations between the constraints. Sect. 6

presents and applies Method C to derive a PDF of the climate proxy given multiple correlated constraints, including the need to 'prune' constraints when those correlations are too strong. Sect. 7 presents and applies Method U, which neglects any correlations between constraints outside of their mutual correlation with the climate proxy. Method U can be fully analyzed, easily interpreted, and does not require constraint pruning. Sect. 8 presents and applies an overconfidence adjustment to accounting for overfitting. Sect. 9 presents conclusions.

2. Data

a. Choice of emergent constraints

For this study, we rely on 11 emergent constraints evaluated in Caldwell et al. (2018) (hereafter CZK18). These include the four constraints CKZ18 judged to be 'credible' (significantly correlated with ECS and supported by a physical mechanism which correctly identifies dominant physical processes and geographical regions which create this correlation), and seven constraints they judged to be 'uncertain' or 'unclear' (significantly correlated with ECS but not amenable to the above assessment of credibility). We will call these constraints 'possible'. Two other 'unclear' Klein constraints from CKZ18 had to be excluded from our analysis for technical reasons described in the following subsection. We also excluded six constraints assessed not to be credible in Table 4 of CKZ18. Short explanations of each constraint that we used along with original citations and evaluations from CZK18 are provided in Table 1.

b. Observational estimates of constraints

CZK18 focused on the evaluation of emergent constraints using model data, while this study aims to use the observed values of those constraints to make climate sensitivity predictions. This requires observational estimates (including uncertainty) for the constraints. It would be ideal to obtain these estimates directly from the original data sources, but this is impractical given the number and diversity of constraints we use. Thus we rely almost exclusively on values communicated by the papers originally proposing each constraint. These studies employed a variety of approaches and levels of detail in describing observational uncertainty. As a result, we are forced to make approximations to achieve uniformity of observed uncertainty estimates across constraints.

For simplicity, our analysis assumes observational uncertainty is normally distributed. The PDF of observed values for constraint i is specified by its mean $\mu_{i,o}$ and standard deviation $\sigma_{i,o}$. While convenient, this assumption is not appropriate for two 'Klein' constraints discussed by CZK18, which are based on positive semi-definite measures of model skill. Hence these two constraints were excluded from our analysis.

Our observed values and the information used to construct them are summarized in Table 2. Studies which provide mean and some multiple of the standard deviation were trivial to process. For studies which provide bounds for a given confidence level, we compute the number of standard deviations for that confidence level for a normal distribution, and we rescale the quoted range to estimate $\sigma_{i,o}$. Where several estimates of $\mu_{i,o}$ and $\sigma_{i,o}$ were provided, we average the estimated means, and we increase $\sigma_{i,o}$ such that $\mu_{i,o} \pm 1\sigma_{i,o}$ just encompasses all of the individual estimated $\pm 1\sigma$ ranges. For constraints which provide only minimum and maximum credible values (often taken from a pair of observations), we take the average of these values as the mean and 1/2 the distance between these values as the standard deviation. Because two samples provide a very poor sense of uncertainty, we occasionally use extra information from papers and/or expert judgement to modify these values, as noted in Table 2.

3. Terminology and covariance estimation

Our mathematical nomenclature is as follows. Capitalized Latin letters denote random variables, and lowercase versions of the same letter indicate particular values of these variables. Vectors are boldfaced. We define $\tilde{Y}$ to be a climate sensitivity proxy such as the equilibrium climate sensitivity or a climate feedback strength, for which the constraints are derived. A single emergent constraint variable is denoted $\tilde{X}$. A collection of n emergent constraints will be labeled $\tilde{X}_i, i = 1, ..., n$. Versions of these random variables which have been normalized to have zero mean and variance of 1 are similarly denoted, but without the tilde. The PDF of any random variable U is $p(u)$, and similarly for multivariate distributions.

The main mathematical formula that we use is the joint PDF of the components of a column vector $\mathbf{U}$ of m zero-mean Gaussian random variables which are known to have an $m \times m$ covariance matrix C with determinant $|C|$:

$$p(\mathbf{u}) = (2\pi)^{-m/2}|C|^{-1/2}\exp\left(-\frac{1}{2}\mathbf{u}^T C^{-1}\mathbf{u}\right). \tag{1}$$

a. Calculating correlation and covariance from GCM samples

A key input to our analysis is the $(n+1) \times (n+1)$ covariance matrix $\tilde{C}_{GCM}$ between the climate proxy $\tilde{Y}$ and the n constraints $\tilde{X}_i$, derived from the available sample of GCMs. An important complication is that not all GCMs provide the data needed to compute all constraints. Using all available GCMs for calculating each needed covariance, rather than just the 8 models which supplied data for all 11 constraints, is essential to obtaining an adequate sample size. The method used to do this must preserve the positive definiteness of the covariance matrix for the multivariate Gaussian method to provide stable results.

The approach that we settled on is to build a GCM covariance matrix based on the best possible estimates of the correlation coefficients. We compute the correlation coefficient $\tilde{r}_{ij}$ between each pair (i,j) of constraints using all GCMs for which both constraints are available. We use a similar approach for GCM-based correlation coefficients $\tilde{r}_{0j}$ between the climate proxy and the j'th constraint, as well as for calculating the standard deviation of each constraint $\tilde{\sigma}_j$ across the GCM sample. The standard deviation $\tilde{\sigma}_0$ of the climate proxy is computed across all GCMs. The elements of the covariance matrix are computed as:

$$\tilde{C}_{GCM,ij} = \tilde{r}_{ij}\tilde{\sigma}_i\tilde{\sigma}_j, \qquad i,j = 0,...,n \tag{2}$$

Note that in general, a different set of GCMs is used for computing each of the three terms on the right-hand side.

Here and in the rest of the paper, rows and columns of the covariance matrix are indexed starting at 0, index 0 corresponds to the climate proxy, and indices 1-n correspond to the n constraints.

4. Climate Sensitivity PDF from a Single Constraint

In this section we describe our approach for computing a PDF of the climate sensitivity proxy $\tilde{Y}$ from a single constraint $\tilde{X}$. This is a useful first step towards treating multiple constraints. Like Bowman et al. (2018) and others, we first estimate a joint PDF of $\tilde{Y}$ and $\tilde{X}$ from the GCMs, then we apply our observational knowledge about $\tilde{X}$ to derive a constrained pdf of $\tilde{Y}$. A key assumption, questioned by Williamson and Sansom (2019), is that the GCM-derived joint pdf is applicable to the real climate, i. e. is a suitable prior for interpreting an observation of the constraint. Also like Bowman et al. (2018), we make the important simplifying assumption that the multivariate pdfs that we estimate are Gaussian, allowing them to be described in terms of a vector of means and a covariance matrix. See Cox et al. (2018), Brient and Schneider (2016), Bowman et al. (2018), and Williamson and Sansom (2019) for other approaches to deriving a PDF of a climate sensitivity proxy from a single emergent constraint, and for further discussions about issues with applying emergent constraints. Our approach captures all sources of uncertainty without sacrificing simplicity, and it is easily extensible to multiple constraints.

An emergent constraint is based on a GCM-based relationship between $\tilde{X}$ and $\tilde{Y}$. Such a relationship should not be trusted well outside the range of GCM values. Indeed, if the observed value of the constraint $\tilde{X}$ lay well outside the expected GCM range, we might interpret this as a physical shortcoming of the GCMs that requires further attention, rather than a solid basis for inferring that the climate sensitivity proxy $\tilde{Y}$ lies outside its GCM range.

Philosophically, this frames our mathematical representation of emergent constraints. Unlike prior studies, we do not start by performing a GCM-based linear regression to determine $\tilde{Y}$ from an observationally-constrained $\tilde{X}$. Instead we estimate a joint Gaussian pdf between $\tilde{X}$ and $\tilde{Y}$ by substituting their 2×2 sample covariance matrix $\tilde{C}_{GCM}$ into (1). In contrast to linear regression, this retains information about the GCM-preferred range of $\tilde{Y}$. It tacitly assume that relations between the climate proxy and the constraints are nearly linear. It also assumes that all GCMs have equal value in estimating how the emergent constraint is related to the climate proxy, whether or not they predict realistic values of the proxy. That assumption has been reasonably criticized (e.g. Brient 2019) but it is a fundamental premise of emergent constraints that the underlying relationship with the climate proxy should rely on a mechanism sufficiently robust as to be insensitive to details of the GCM physical formulation, even though those details are important to actually obtaining an observationally consistent value of the constraint.

If we could exactly observe that $\tilde{X} = x$, we could substitute into the joint PDF to obtain the conditional pdf of $\tilde{Y}$. In the language of Bayesian analysis, this is the posterior probability for $\tilde{Y}$

based on the GCM-only Gaussian prior and the observation of the constraint. However, there are two further practical complications to consider.

First, we cannot exactly observe the true value of $\tilde{X}$. To handle observational uncertainty, we define a random variable $\hat{X}$, the estimated constraint, which is the sum of $\tilde{X}$ plus a normally-distributed observational error with zero mean and the observed standard deviation σ_o. This observational error is assumed to be independent of $\tilde{Y}$ and the physically-determined (true) value of $\tilde{X}$. Thus $\hat{X}$ is a Gaussian random variable whose variance is the sum of its variance across GCMs and its observational variance. $\tilde{Y}$ and $\hat{X}$ have a Gaussian joint PDF determined by their covariance matrix,

$$\hat{C} = \begin{bmatrix} \mathrm{var}(\tilde{Y}) & \mathrm{cov}(\tilde{Y}, \hat{X}) \\ \mathrm{cov}(\tilde{Y}, \hat{X}) & \mathrm{var}(\hat{X}) \end{bmatrix} = \tilde{C}_{GCM} + \begin{bmatrix} 0 & 0 \\ 0 & \sigma_o^2 \end{bmatrix}, \tag{3}$$

where $\tilde{C}_{GCM}$ is computed using Eq. (2).

Second, the process of formulation and selection of emergent constraints may result in overconfidence, i. e. constraints that are more highly correlated with the climate sensitivity proxy than would be obtained from a different independent random sample of GCMs, if such existed. We counteract overconfidence by artificially reducing the covariance between X and $\tilde{Y}$ without changing the other elements of $\hat{C}$. We have not explicitly included this adjustment in our single-constraint analysis because a goal of that analysis is to evaluate overlap between PDFs for each constraint, and ignoring overconfidence provides a lower bound for that overlap. We consider the sensitivity of a multiple-constraint analysis to an overconfidence correction in Sect. 8.

It is convenient to work with standardized variables with a mean of zero and a standard deviation of 1:

$$
\begin{aligned}
Y &= (\tilde{Y} - \overline{y})/\tilde{\sigma}_Y, &(4)\\
X &= (\hat{X} - \overline{x})/\hat{\sigma}_X, &(5)\\
\hat{\sigma}_X^2 &= \tilde{\sigma}_X^2 + \sigma_o^2. &(6)
\end{aligned}
$$

The normalization of the standardized X accounts for both its variance across GCMs and its observational uncertainty. Overlines indicate averages over the GCM ensemble.

The covariance matrix of Y and X, derived from the sample of GCMs and adjusted for observational constraint uncertainty, is

$$
C = \begin{bmatrix} 1 & r \\ r & 1 \end{bmatrix}, \tag{7}
$$

where $r = \hat{C}_{01}/(\tilde{\sigma}_Y \hat{\sigma}_X)$ is the correlation coefficient between Y and X, or equivalently between $\tilde{Y}$ and $\hat{X}$. Because of the observational uncertainty, r is smaller in magnitude than the GCM-estimated correlation coefficient $\tilde{r}_{01}$ between the climate proxy $\tilde{Y}$ and the constraint $\hat{X}$

We translate the observational estimate of the constraint, $\hat{X} = \hat{x}$, into the standardized form

$$
X = x' = (\hat{x} - \overline{x})/\hat{\sigma}_X. \tag{8}
$$

We condition the joint pdf of Y and X on this known value of X to obtain a Gaussian posterior for Y:

$$
\begin{aligned}
p(y|x') = p(y,x')/p(x') &= (2\pi|C|)^{-1/2} \exp\left(-\frac{1}{2}\begin{bmatrix} y & x' \end{bmatrix} C^{-1} \begin{bmatrix} y \\ x' \end{bmatrix} + \frac{x'^2}{2}\right)\\
&\propto \exp \ell(y), &(9)
\end{aligned}
$$

where the log-probability is

$$
\ell(y) = -\frac{1}{2}\frac{y^2 - 2rx'y + x'^2}{1 - r^2} + \frac{x'^2}{2} = -\frac{1}{2}\frac{(y - rx')^2}{1 - r^2}.
$$

From this formula, we can read off the mean $y^{(1)}$ and standard deviation $\sigma^{(1)}$ of the posterior distribution of Y:

$$
\begin{aligned}
y^{(1)} &= rx', &(10)\\
\sigma^{(1)} &= (1-r^2)^{1/2}. &(11)
\end{aligned}
$$

The superscript (1) denotes that this is a one-constraint estimate. We will define the estimated posterior 'range' of Y as lying within 2 standard deviations of the mean, i. e. $y^{(1)} \pm 2\sigma^{(1)}$. We use a subscript i to denote an estimate based on constraint i.

Past studies of emergent constraints have typically used linear regression to quantify the relationship between the constraint and the climate sensitivity proxy. If we ignore observational uncertainty we could obtain the above result by regressing the climate sensitivity proxy Y on the constraint X. In this case, r would be the GCM-based correlation coefficient between $\tilde{X}$ and $\tilde{Y}$ with no adjustment for the observational uncertainty. The best fit regression line $y = rx$ matches the posterior mean $y^{(1)}$ when evaluated at x'. The residual in Y around that fit has a standard error $(1-r^2)^{1/2}$ that matches $\sigma^{(1)}$.

Unlike this regression approach, our approach naturally incorporates observational uncertainty in the constraint. For a single constraint, this was also done by Bowman et al. (2018) under similar assumptions, but using a different mathematical approach. Reassuringly, after after accounting for our different notation and normalization, our formulas (10) and (11) are isomorphic to eqns. (18) and (23) of Bowman et al. (2018). However, unlike earlier work our approach extends naturally to many constraints.

a. Single-constraint results for ECS

Table 3 gives the correlation coefficients r_i between each constraint i and two choices of climate proxy Y (ECS and climate feedback parameter λ). To simplify the ensuing discussion, we henceforth 'sign-correct' all constraints so that $r_i > 0$, by flipping the sign of those constraints that are negatively correlated with Y (indicated by a '-' in the 'Sign' column). This table also gives normalized constraint values x_i' (in units of standard deviation) for the sign-corrected constraints, such that $x_i' > 0$ favors $y > 0$ (ECS larger than the GCM mean). Eight of the 11 constraints have positive x_i', with values up to 2.4 for Constraint 1 (Sherwood D). Constraints 2 (Brient Shal), 9 (Lipat) and 11 (Cox) have modestly negative x_i' in the range -0.4 to -0.8.

Lastly, Table 3 gives $\sigma_o/\tilde{\sigma}_X$, the ratio of the observational uncertainty to the GCM-based standard deviation for each constraint. For the credible constraints 1-4 and constraint 10 (Siler), this ratio is less than 0.5 and observational uncertainty is relatively unimportant to the posterior range of Y. For the remaining 6 constraints, the ratio is larger than 0.5. For these constraints, observational uncertainty substantially reduces r_i, broadens the posterior range, and moves $y_i^{(1)}$ toward zero, i. e. it moves the posterior mean Y toward the GCM mean. This is most pronounced for Constraint 6 (Qu), with a ratio of 1.7; the ratio lies between 0.64 and 1.01 for the remaining five constraints. In general, we conclude that it is important to account for observational uncertainty in the constraints.

Our approach is illustrated in Fig. 1, using the Zhai constraint as an example. For clarity, the figure is presented in terms of dimensional rather than standardized variables. The cyan ellipse shows a contour of joint probability density between $\tilde{Y}$ and $\tilde{X}$ derived from our collection of CMIP models (dots). The maximum joint density is at the center of this ellipse, which is the centroid of the CMIP data points. The green ellipse corrects this PDF for observational uncertainty in the constraint, the PDF of which is shown along the y-axis in blue. Note that for this constraint, observational uncertainty has a small effect, as evident by the similarity between green and cyan ellipses. This situation is found for other constraints as well. Were we to adjust for overconfidence by using a joint PDF with reduced correlation r between $\tilde{Y}$ and $\hat{X}$, this ellipse would be vertically broadened. The posterior for ECS (red) is the cross section of the bivariate PDF at the best-guess observed constraint value $\hat{x}$ (horizontal blue dashed line).

The posterior PDFs of ECS given each constraint separately are shown in Fig. 2, along with the PDF of ECS from CMIP3+CMIP5 models, shown both as a histogram and a Gaussian fit. Credible constraints are shown in panel (a) and possible constraints are shown in panel (b). Their means vary from 3.14 K (constraint 9 = Lipat) to 3.60 K (constraint 1 = Sherwood), with a standard deviation of 0.27 K (constraint 3 = Zhai) to 0.37 K (constraint 6 = Qu). Eight of the 11 constraints have PDFs peaked at an ECS greater than the GCM mean. Constraint credibility has no systematic effect on peak probabilities or distribution widths.

These Gaussian PDFs are derived from the formulas for nondimensional posterior mean (10) and standard deviation (11) using the information in Table 3. These are redimensionalized by scaling with the GCM standard deviation $\tilde{\sigma}_Y = 0.38$ K and adding the GCM mean $\bar{y} = 3.24$ K, which can be read off the 'GCM Ensemble' line of Table 4.

These PDFs do not account for possible overconfidence, which would further broaden their width. Even so, all of the constraints have substantial overlap with each other. This is reassuring because if two constraints had disjoint PDFs we would be forced to conclude that at least one of them must be wrong, and multi-constraint analysis would be pointless. Results for the climate sensitivity parameter λ (not shown) are similar.

5. Dependence between Constraints

Combining constraints is only useful if those constraints provide independent information. Because agreement between samples is typically used as a proxy for uncertainty, failure to account for redundant samples is dangerous.

Interdependence of constraints was investigated in CZK18 by computing correlations between each pair of constraints. CZK18 found that constraints were more correlated than expected by chance, but identifying pairs of significantly correlated constraints based on related physical explanations was generally unsuccessful. They noted that some pairwise correlation is expected because all constraints are by construction correlated with ECS.

For our 11 constraints, there are 55 pairs of constraints. If we flip the sign of constraints which are negatively correlated with ECS, then 52 of the 55 constraint pairs are positively correlated across the GCMs, which is strong evidence for such mutual dependence; 14 of these have a positive correlation coefficient exceeding 0.4 and none have a correlation coefficient more negative than -0.25.

To correct for the mutual dependence between constraints that is due to ECS, we compute the partial correlation coefficient between constraints X_i and X_j given ECS Y (denoted hereafter by an subscript 0):

$$r_{ij\cdot 0} = \frac{r_{ij} - r_i r_j}{(1 - r_i^2)^{1/2}(1 - r_j^2)^{1/2}}. \tag{12}$$

Here r_i is the sample correlation coefficient between X_i and Y, and similarly for r_j, while r_{ij} is the sample correlation coefficient between X_i and X_j. We call two constraints with a partial correlation of zero ‘conditionally uncorrelated’; for the multivariate Gaussian distributions assumed in this paper, this is equivalent to conditional independence. Positive sign-corrected partial correlation indicates constraints that covary in the same sense as we would expect based on their correlation with Y, but more strongly.

Correcting for the mutual dependence with ECS removes some but not all of the covariation between our 11 constraints. Fig. 3 shows the sign-corrected pairwise partial correlations. Well over half (38 of 55) are positive. This is suggestive, but are these partial correlations statistically significant? As in CZK18, choosing an appropriate number of degrees of freedom (DOF) is difficult because there are complicated structural dependences between models (Masson and Knutti 2011; Knutti et al. 2013; Sanderson et al. 2015a,b). Following CZK18, we handle this issue by using a fairly lax 90% two-sided test and by assuming each GCM that goes into the calculation of a particular correlation coefficient r_{ij} is independent. Both assumptions favor false positives, but partial correlations deemed not to be significant would almost certainly also be deemed insignificant with other reasonable assumptions. For a typical number of contributing GCMs (20), a partial correlation of magnitude 0.35 or larger is significant by this standard.

We found only 6 (1) of 55 constraint pairs have a significantly positive (negative) partial correlation; these are shown by orange (purple) shading on Fig. 3. Some expected relationships between constraints are borne out, like tight correlation between Siler and Volodin. Other constraints are significantly correlated even though their explanations seem to be unrelated, like Brient Shal and Brient Alb. This motivates our Method C (for 'correlated'), which accounting for partial correlation between constraints. However, other constraint pairs, like Zhai and Brient Alb, are weakly correlated even though they share a physical explanation, and only 11% of the constraint pairs have a partial correlation above the 90% significance threshold. Thus it is a reasonable overall assumption to neglect partial correlations among our set of constraints, motivating the simpler Method U (for 'uncorrelated') discussed in Sec. 7.

6. Method C: Climate Sensitivity PDF from Multiple Correlated Constraints

a. Theory

The methodology introduced in Sec. 4 generalizes transparently to the case of multiple constraints $\tilde{X}_1, ..., \tilde{X}_n$ with observational estimates $\hat{x}_1, ..., \hat{x}_n$ having uncertainties $\sigma_{o,1}, ..., \sigma_{o,n}$. We form a column vector $\mathbf{U}$ of length $n+1$ whose components are the climate sensitivity proxy $\tilde{Y}$ and the $\tilde{X}_i' s$. The components of its $(n+1)\times(n+1)$ covariance matrix modified for observational uncertainty are

$$\hat{C}_{ij} = \begin{cases} \mathrm{var}(\tilde{Y}) & i = j = 0 \\ \mathrm{cov}(\tilde{Y}, \tilde{X}_j) & i = 0, j = 1, ..., n \\ \mathrm{cov}(\tilde{X}_i, \tilde{Y}) & j = 0, i = 1, ..., n \\ \mathrm{cov}(\tilde{X}_i, \tilde{X}_j) + \sigma_{o,i}^2 \delta_{ij} & i, j = 1, ..., n \end{cases} \tag{13}$$

Here δ_{ij} is 1 if $i = j$ and 0 otherwise. The matrix $\hat{C}$ is non-negative definite if all covariances are computed with the same set of GCMs, but this is not guaranteed if different elements of $\hat{C}$ are computed with different subsets of GCMs, as we are forced to do in this study.

We standardize $\tilde{Y}$ and the constraints as before, producing a new covariance matrix C between the standardized ECS Y and the standardized constraints X_i, and we derive standardized values x_i' for each observational estimate $\hat{x}_i$. To simplify notation, define

$$A = C^{-1} \tag{14}$$

to be the inverse covariance matrix, which is symmetric and positive definite if $\hat{C}$ is non-singular. To obtain the resulting posterior PDF of Y conditioned on the observational estimates x_i', we define a column vector

$$\mathbf{u} = [y, x_1', ..., x_n']^T \tag{15}$$

and substitute this, together with (13) into the joint PDF (1) :

$$
\begin{aligned}
p(y|x'_1,\ldots,x'_n) &= p(y,x'_1,\ldots,x'_n)/p(x'_1,\ldots,x'_n) \\
&\propto \exp \ell(y), \\
\ell(y) &= -\frac{1}{2}\mathbf{u}^T C^{-1}\mathbf{u} + \text{terms not involving } y \\
&= -\frac{1}{2}\left[A_{00}y^2 + 2\sum_{i=1}^{n} A_{0i}yx'_i\right] + \text{terms not involving } y \\
&= -\frac{(y-y_C^{(n)})^2}{2\sigma_C^{(n)2}} + \text{terms not involving } y, \qquad (16)
\end{aligned}
$$

which describes a Gaussian distribution with:

$$
\begin{aligned}
\text{Method C posterior mean of } Y: &\quad y_C^{(n)} = \sum_{i=1}^{n} a_i x'_i, &(17)\\
\text{Constraint weights:} &\quad a_i = -A_{0i}/A_{00}, &(18)\\
\text{Posterior std. dev. of } Y: &\quad \sigma_C^{(n)} = A_{00}^{-1/2}. &(19)
\end{aligned}
$$

In theory, Method C solves our problem for multiple correlated constraints.

b. Sampling uncertainty in the covariance matrix

In practice, if the covariance matrix is derived from a finite sample of GCMs, it is sensitive to sampling uncertainty, especially if the number of constraints is comparable to the number of GCMs. For instance, if we have 10 constraints and 20 GCMs, we are using 20 samples to estimate an 11-dimensional correlation matrix with $11 \cdot 12/2 = 66$ independent entries, which is a highly underconstrained problem. Thus we anticipate that Method C may fail or give spurious results, especially if partial correlations between the constraints are important. In the next subsection, we apply Method C to ECS estimation and develop a way to test and (if necessary) improve its robustness.

c. *Applying Method C to ECS and the Need to Prune Constraints*

1) 4 CREDIBLE CONSTRAINTS

The red curve in Fig. 4a shows the ECS posterior PDF estimated using Method C based on the 4 sign-corrected credible constraints. It is compared to other PDFs, including the Gaussian CMIP prior (black dash) and the average of the single-constraint PDFs (solid black). Methods U and U3 will be discussed in the following section.

Method C gives a strikingly higher and narrower PDF than the average of the single-constraint PDFs, which in turn is shifted slightly higher than the CMIP mean PDF. This may seem surprising, but can be viewed as the result of strong evidence (3 of 4 constraints) agreeing that ECS is likely more positive than the GCM mean. Mathematically stated, of the four observed x'_i, only $x'_2 < 0$.

Using (17)-(19), Method C predicts that the nondimensional most likely ECS is

$$y_C^{(4)} = 0.22x'_1 + 0.25x'_2 + 0.48x'_3 + 0.28x'_4 = 1.19. \tag{20}$$

The nondimensional ECS standard deviation is $\sigma_C^{(4)} = 0.54$ and does not depend on the constraint values. Redimensionalizing using the mean (3.24 K) and standard deviation (0.75 K) of the GCM ensemble yields a 4-constraint $\pm 2\sigma$ ECS range of 4.14 ± 0.82 K.

Constraint 3 (Zhai) has the strongest weight $a_3 = 0.48$ on the Method C ECS mean. Constraint 4 (Brient Alb) has only 60% of the weight of constraint 3 despite having a similar correlation $r \approx 0.7$ with ECS; this is due to its covariances with the other constraints. All four constraints have positive weights $a_i > 0$. This positivity property is physically reasonable and desirable. For instance, if constraint i has observed value $x'_i > 0$, its correlation with ECS is $r_i > 0$, and all other constraints have observed values of zero, then we expect $y > 0$; this is only possible if $a_i > 0$. However, weight positivity is not guaranteed by Method C.

Fig. 4b assesses the sensitivity to the GCM sampling by redoing the 4-constraint analysis with each of the GCMs omitted in turn when computing the needed correlation matrix. The posterior ECS PDF is quite robust to this test. Removing any two GCMs provides very similar results, with the most likely value of ECS always being 3.9-4.2 K (not shown).

2) ALL CONSTRAINTS

Method C can also be applied to all 11 constraints, since the 12×12 estimated correlation matrix has all positive eigenvalues. It gives a $\pm 2\sigma$ ECS range 4.28 ± 0.69 K, shown as the dashed red curve in Fig. 4c. This PDF has a slightly higher mean and a narrower distribution of ECS than the 4-constraint estimate. Eight of the 11 weights a_i are positive. Two are marginally negative, and one, $a_{10} = -0.47$, is strongly negative. This arises because constraint 10 (Siler) has strong partial correlations with some of the other constraints, exceeding 0.5 with constraints 2 (Brient Shal) and 5 (Volodin). Because the observed value $x'_{10} = 0.35$ is small, the large negative weight doesn't have a strong direct impact on the most probable ECS. Nevertheless this suggests the Method C ECS PDF may be less robust to sampling uncertainty in the correlation matrix with 11 constraints than with 4 constraints. This defeats the point of using more constraints.

Fig. 4d assesses this by removing one GCM from the sample, as in Fig. 4b. The posterior ECS PDFs are mostly robust to this test again, but less so than with the 4 constraints, since there are now three clear outliers among the 11 ranges. The mean ECS ranges from 3.7-4.6 K across these cases, and 3.6-5 K when 2 GCMs are removed.

3) PRUNING CONSTRAINTS WITH LARGE NEGATIVE WEIGHTS

When Method C gives negative sign-corrected a_i's, we do not recommend its use unless there are so many independent GCMs that the correlation matrix is accurately known. Instead, we suggest removing ('pruning') the constraint which has the most negative weight, repeating the analysis, and testing if there are still negative sign-corrected weights. If so, continue the procedure until all weights lie in the desired range. A weaker threshold for pruning a constraint such as $a_i < -0.1$ may still provide satisfactory results, since constraints with weights close to zero will have little impact on the estimate of the most likely ECS. With our set of constraints, certain combinations of as few as three constraints (e. g. 2, 4, 10) lead to negative weights with $a_i < -0.1$, although there are combinations of as many as 10 constraints for which all weights exceed -0.1, and there are combinations of 8 constraints for which all weights are positive.

For the 11-constraint case, pruning just constraint 10 (Siler) brings the most negative sign-corrected weight down from -0.47 to -0.10 and changes the Method C $\pm 2\sigma$ posterior range of ECS to 4.12 ± 0.77 K, very similar to the 4-constraint value. We regard this 10-constraint PDF as more plausible than the 11-constraint PDF, so it is shown as a solid (rather than dashed) red curve in Fig. 4c.

7. Method U: Estimation of Y From Conditionally Uncorrelated Constraints

Overall, Fig. 3 suggests that most of the partial correlations between the constraint pairs are not that large and are statistically insignificant. Thus it is reasonable to consider the special case that all the constraints are mutually uncorrelated when conditioned on a given value of the climate sensitivity proxy Y. Since the constraints have Gaussian PDFs, they are therefore independent when conditioned on Y. This assumption allows a closed-form posterior estimate of the PDF of Y that we call Method U. Method U provides valuable insights into Method C, and our experience is that it also gives results close to Method C when the latter yields all positive sign-corrected weights. Method U requires only estimates of the correlation coefficients r_i between Y and the individual constraints. These can be computed fairly reliably with 10 or more GCMs; for our sample of GCMs, they are given in Table 3. For only one constraint, or when the constraints are conditionally uncorrelated, Method U is identical to Method C, although derived differently.

We work in terms of the standardized variables. With the assumption of conditional independence, the joint PDF of the constraints conditioned on Y can be written as a product:

$$p(x_1,...,x_n|y) = p(x_1|y)...p(x_n|y), \tag{21}$$

where the conditional PDF for constraint X_i is given by (10-11) with x and y swapped:

$$p(x_i|y) \propto \exp\left\{-(x_i - r_i y)^2/2(1-r_i^2)\right\}. \tag{22}$$

Thus the joint multivariate Gaussian PDF simplifies to

$$\begin{aligned} p(y,x_1,...,x_n) &= p(x_1,...,x_n|y)p(y) \\ &\propto \exp\left\{-\frac{1}{2}[y^2 + \sum_{i=1}^{n}(x_i - r_i y)^2/(1-r_i^2)]\right\}. \end{aligned} \tag{23}$$

Setting the constraints equal to their standardized observed values, denoted again by primes, we obtain a posterior PDF

$$\begin{aligned} p(y|x'_1,...,x'_n) &= p(y,x'_1,...,x'_n)/p(x'_1,...,x'_n) \\ &\propto \exp \ell(y) \\ \ell(y) &= -\frac{1}{2}\left[y^2 + \sum_{i=1}^{n}(x'_i - r_i y)^2/(1-r_i^2)\right] \\ &= -\frac{1}{2}\sum_{i=0}^{n} w_i(y'_i - y)^2 \\ &= -\frac{(y-y_U^{(n)})^2}{2\sigma_U^{(n)2}}. \end{aligned} \tag{24}$$

The subscript U refers to the assumption of conditional uncorrelation. Here,

$$\text{Method U posterior mean of } Y: \qquad y_U^{(n)} = \sum_{i=0}^{n} w_i y_i' \tag{25}$$

$$= \sum_{i=1}^{n} a_i x_i', \tag{26}$$

$$\text{Posterior std. dev. of } Y: \qquad \sigma_U^{(n)} = s^{-1/2}, \tag{27}$$

where

$$\text{Proxy estimates} \qquad y_i' = \begin{cases} 0, & i = 0 \\ x_i'/r_i, & 1 \le i \le n \end{cases}, \tag{28}$$

$$\text{Proxy weights} \qquad w_i = v_i/s, \tag{29}$$

$$v_i = \begin{cases} 1, & i = 0 \\ r_i^2/(1-r_i^2), & 1 \le i \le n \end{cases}, \tag{30}$$

$$s = \sum_{i=0}^{n} v_i, \tag{31}$$

$$\text{Constraint weights} \qquad a_i = w_i/r_i. \tag{32}$$

Method U expresses the mean (26) of the posterior PDF as a linear combination of the observed constraint values x_i'. This is the expected form for a special case of Method C, but has three major interpretational advantages. First, it gives explicit formulas for the constraint weights a_i that depend only on the correlation coefficients r_i. The relative weights for constraints with correlation coefficients of 0.3, 0.6 and 0.9 with the climate proxy are 0.3, 0.9, and 4.7; constraints with large correlations have much more impact on the posterior PDF. These explicit formulas are less sensitive to small-sample errors than entries in the inverse covariance matrix. Second, these weights automatically satisfy the desirable positivity condition that $a_i > 0$ for $r_i > 0$.

Third, there is an appealingly interpretable alternate form (25) for $y_U^{(n)}$ in terms of proxy weights $w_i = a_i r_i$ multiplying climate proxy estimates $y_i' = x_i'/r_i$. For each constraint, y_i' is the predicted value of the climate proxy we would obtain by linear regression applied to the single-constraint problem with the dependent and independent variables swapped, i. e. regressing X_i on Y and evaluating at the observed value x'_i.

The proxy weights sum to one if we add a weight w_0 for 'constraint 0', the GCM prior. Thus $y_U^{(n)}$ is a weighted *average* of the y'_i. Each proxy weight is a strongly increasing function of its correlation coefficient with the climate sensitivity proxy Y. For constraints that are poorly correlated with Y, y'_i can be quite large but is multiplied by a very small proxy weight. Nevertheless, if a lot of constraints with modest correlations with Y all suggest large positive anomalies y'_i, they combine to create a large $y_U^{(n)}$, suggestive of a Y much larger than the GCM mean.

The posterior standard deviation $\sigma_U^{(n)}$ decreases as the inverse square root of the sum s of positive contributions v_i from each constraint. More conditionally uncorrelated constraints always narrow the posterior range of Y, especially those having a high correlation with Y.

a. Consistency with single-constraint posterior PDF

When there is only one constraint, Methods C and U are identical, and they must both reduce to the single-constraint posterior PDF. Using the Method U formulas (after dropping the index 1 to match the single-constraint notation):

$$s = 1 + \frac{r^2}{1-r^2} = \frac{1}{1-r^2}, \tag{33}$$

$$y_U^{(1)} = [1 \cdot 0 + \frac{r^2}{1-r^2} y']/s = r^2 y' = r x', \tag{34}$$

$$\sigma_U^{(1)} = s^{-1/2} = (1-r^2)^{1/2}, \tag{35}$$

which indeed match our single-constraint formulas. The interpretation from Method U is that the posterior mean of Y is a weighted average of its GCM prior (zero), and the $y' = x'/r$ predicted by a regression of X on Y, multiplied by a proxy weight $r^2 < 1$. In the single-constraint case, we noted this posterior PDF can alternatively be derived from a regression of Y on X, but that perspective doesn't generalize to multiple constraints.

b. Method U results for ECS

The red curve in Fig. 4a shows the Method U estimate of the ECS posterior given the 4 credible constraints. This Gaussian PDF has a $\pm 2\sigma$ range of 3.3-4.9 K, similar to the 4-constraint estimate from Method C. That is, partial correlation between the 4 credible constraints has almost no impact on the posterior PDF. Indeed, according to Method U, the most likely ECS is

$$y_U^{(4)} = 0.15x_1' + 0.14x_2' + 0.42x_3' + 0.41x_4' = 1.10, \tag{36}$$

The coefficients are quite similar to those in the analogous formula (20) derived from Method C.

The red curve in Fig. 4c shows the Method U ECS posterior for all 11 constraints, which has a $\pm 2\sigma$ range of 3.4-4.6 K. This is similar but slightly narrower than the 4-constraint Method U PDF and the PDF from Method C pruned to 10 constraints per the earlier discussion. That is to be expected, as any conditional correlation between the constraints reduces the independent information provided by adding new constraints, Unlike for Method C, pruning constraint 10 has a negligible impact on the posterior range predicted by Method U (not shown).

Fig. 5a shows the 11 proxy weights w_i, calculated using Eq. (30). Each constraint is represented by a colored vertical line with height proportional to the proxy weight, located at a Y value of y_i'. Since these weights sum to one, the most likely ECS (black circle) is at the horizontal center of gravity (weighted average) of these bars. Most of the constraints have correlation coefficients with ECS of between 0.4 and 0.6 and modest proxy weights of 0.02-0.1. Constraints 3 (Zhai) and 4 (Brient Alb) have larger weights near 0.2 because of their stronger correlation with ECS, so they are particularly important in setting the ECS posterior range. 'Constraint 0', the GCM prior, has a weight of 0.21. Thus with 11 constraints, the GCM range of ECS plays a modest role in shaping the posterior function. Constraint 1 (Sherwood D) has an anomalously large y_1' because its observational estimate is $x_1' = 2.4$ standard deviations from the mean, and because it has a relatively small correlation $r_1 \approx 0.4$ with ECS. However, that small r_1 causes it to have a small proxy weight.

Fig. 5b shows the contribution of individual constraints to changing the posterior mean of ECS away from the GCM mean of 3.2 K ($a_i x_i'$ before redimensionalization). Constraint 3 (Zhai) is most important, because its observational estimate x_3' is over one standard deviation from the GCM mean and it also has the highest correlation with ECS. Three other constraints (1, 4, 8) give substantial positive contributions exceeding 0.1 K, four give modest positive contributions of less than 0.08 K, and three give modest negative contributions with magnitude less than 0.05 K.

8. Adjusting for overconfidence

We present a method to adjust Methods C and U for overconfidence by artificially reducing the correlation coefficients r_i between the constraints and Y, while leaving the partial correlation coefficients between the constraints unchanged.

An obvious way to do this would be to scale all the correlation coefficients r_i by the same factor. We use a related but slightly different approach. A constraint that is nearly perfectly correlated with ECS over 20+ GCMs should also be extremely highly correlated with ECS over a different set of GCMs. Thus we treat overconfidence as leading to a systematic underestimate by a user-specified factor $\alpha^2 \leq 1$ of the variance in each of the constraints that is unexplained by regression onto Y. This amounts to scaling the ratio $v_i = r_i^2/(1-r_i^2)$ of the explained to unexplained variance by the specified factor α^2:

$$v_i^* = \alpha^2 v_i, i = 1, ..., n,$$

for which the correspondingly reduced correlation coefficients are

$$r_i^* = \frac{\alpha r_i}{[1-(1-\alpha^2)r_i^2]^{1/2}}.$$

For small values of r_i, the reduced correlation coefficient is a factor of α as large as the original coefficient; for large r_i, the fractional reduction is less.

To use Method C, we could adjust the correlation matrix to preserve the partial correlations $r_{ij\cdot 0}$ between constraints while reducing the r_i's. However, if we are highly uncertain about r_i, then our empirical estimates of the partial correlations are at least as uncertain and potentially meaningless.

Some overconfidence correction ($\alpha < 1$) seems merited given our earlier arguments about *a priori* selection and optimization of constraints. One compelling basis for such a judgement is testing based on an independent set of GCMs. Some previous examples include some of the discredited constraints in CZK18, and the minimal correlation of the 'credible' Sherwood D constraint with ECS across a perturbed parameter ensemble. Wagman and Jackson (2018) recommend a large uncertainty enhancement for constraint overconfidence even for physically appealing emergent constraints. Further research on this issue is needed.

As a sensitivity test, we apply an extreme overconfidence adjustment $\alpha = 1/3$ (a nine-fold inflation of unexplained variance) to Method U. The magenta curves U3 in Figs. 4a and c show the resulting ECS PDFs for the 4-constraint and the 11-constraint cases. In both cases, the overconfidence correction widens the posterior range (Table 4) and moves the posterior mean ECS somewhat toward the GCM mean. The wider range is expected, because constraints weakly correlated with ECS are less informative.

9. Conclusions

We derive a new 'Method C' of combining correlated emergent constraints to estimate a Gaussian PDF for ECS or any other global climate sensitivity proxy from a finite sample of GCMs. It accounts for observational uncertainty of the constraints and (optionally) for over-confidence in the estimated correlation between the constraints and the proxy. The method is based on approximating the joint PDF of the proxy and the constraints as multivariate Gaussian. With our limited sample of 40 CMIP GCMs, most of which provided inadequate outputs to compute some constraints, this PDF is inadequately sampled and the method may not give robust results as the number of constraints becomes comparable to the number of GCMs. We develop a systematic constraint-pruning method to improve the robustness of Method C. We also present 'Method U', which neglects any partial correlations between emergent constraints that are not related to their joint correlation with the climate proxy; this is more robust when applied to a small sample of GCMs, does not require constraint pruning, and is more interpretable.

We apply these methods to a set of four credible and seven other 'possible' constraints from CZK18 and compare them with PDFs derived from single constraints, for which our method is essentially equivalent to regression of the climate sensitivity proxy on the constraint. Taken singly, the 11 constraints imply ECS PDFs with $\pm 2\sigma$ ranges having means between 3-4 K and $\pm 2\sigma$ widths of $\pm$0.8-1.2 K. Reassuringly, all of these constraints give overlapping PDFs for ECS, with eight of the 11 favoring ECS higher than the GCM mean.

The 4-constraint $\pm 2\sigma$ ECS range estimated by both Methods C and U is close to 4 ± 0.8 K, which lies exclusively in the upper half of the GCM range. To apply method C with all 11 constraints, we had to prune one constraint, after which the ECS range was again similar and higher than the PDFs estimated from most of the constraints taken one at a time. We obtain the same ECS range using Method U with 11 constraints. The interpretation from Method U is that since a large majority of the constraints individually favor ECS values above the GCM mean, combining them more strongly favors that result. This is an important and interesting result of this analysis.

We propose a user-chosen adjustment factor α to account for constraint overconfidence. This factor reduces the correlation coefficients of all the constraints with the climate sensitivity proxy while leaving their conditional correlations with each other unaltered. With a strong overconfidence adjustment $\alpha = 1/3$, the $\pm 2\sigma$ 11-constraint estimated ECS range of Method U doubles in width to 2.7-5.3 K, but its mean remains nearly 4 K.

The same hierarchy of approaches was applied to the climate sensitivity parameter λ, with similar results shown in Tables 3 and 4.

We conclude that climate sensitivity estimated from combining the most reasonable current emergent constraints is very likely above the CMIP3/5 GCM mean of 3.2 K and has roughly even odds of exceeding 4 K. To better interpret and bolster this surprisingly strong result, we should continue to search for more physically-motivated emergent constraints aimed at regime-specific cloud feedbacks, e. g. Qu et al. (2013).

Acknowledgments. We would like to acknowledge the modeling groups, the Program for Climate Model Diagnosis and Intercomparison (PCMDI) and the World Climate Research Program's Working Group on Coupled Modeling for making available the CMIP3 and CMIP5 multi-model datasets, supported by the U.S. Department of Energy (DOE) Office of Science. This work was supported by the Office of Science (BER) at Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and BER's Regional and Global Climate Modeling (RGCM) Program. Steve Klein instigated this study and heavily influenced early drafts. Yuying Zhang provided estimates of the Volodin constraint from multiple satellites.

References

Armour, K. C., C. M. Bitz, and G. H. Roe, 2013: Time-varying climate sensitivity from regional feedbacks. *J. Climate*, **26**, 45184534, doi:10.1175/JCLI-D-12-00544.1.

Bowman, K. W., N. Cressie, X. Qu, and A. Hall, 2018: A hierarchical statistical framework for emergent constraints: Application to snow-albedo feedback. *Geophysical Research Letters*, **45 (23)**, 13,050–13,059, doi:10.1029/2018GL080082.

Brient, F., 2019: Reducing uncertainties in climate projections with emergent constraints: Concepts, examples and prospects, doi:10.1007/s00376-019-9140-8, accepted.

Brient, F., and T. Schneider, 2016: Constraints on climate sensitivity from space-based measurements of low-cloud reflection. *Journal of Climate*, **29 (16)**, 5821–5835, doi:10.1175/JCLI-D-15-0897.1, URL https://doi.org/10.1175/JCLI-D-15-0897.1.

Brient, F., T. Schneider, Z. Tan, S. Bony, X. Qu, and A. Hall, 2015: Shallowness of tropical low clouds as a predictor of climate models' response to warming. *Climate Dynamics*, 1–17, doi:10.1007/s00382-015-2846-0, URL http://dx.doi.org/10.1007/s00382-015-2846-0.

Caldwell, P. M., M. D. Zelinka, and S. A. Klein, 2018: Evaluating emergent constraints on equilibrium climate sensitivity. *Journal of Climate*, **31 (10)**, 3921–3942, doi:10.1175/JCLI-D-17-0631.1, URL https://doi.org/10.1175/JCLI-D-17-0631.1.

Cess, R. D., and Coauthors, 1989: Interpretation of cloud-climate feedback as produced by 14 atmospheric general circulation models. *Science*, **245**, 513–516.

Cox, P. M., C. Huntingford, and M. S. Williamson, 2018: Emergent constraint on equilibrium climate sensitivity from global temperature variability. *Nature*, **553**, 319 EP –, URL http://dx.doi.org/10.1038/nature25450.

Gregory, J. M., and T. Andrews, 2016: Variation in climate sensitivity and feedback parameters during the historical period. *Geophysical Research Letters*, **43 (8)**, 3911–3920, doi:10.1002/2016GL068406.

Hall, A., P. Cox, C. Huntingford, and S. Klein, 2019: Progressing emergent constraints on future climate change, submitted.

Hall, A., and X. Qu, 2006: Using the current seasonal cycle to constrain snow albedo feedback in future climate change. *Geophy. Res. Let.*, **33**, L03 502, DOI: 10.1029/2005GL025 127.

Klein, S. A., and A. Hall, 2015: Emergent constraints for cloud feedbacks. *Current Climate Change Reports*, **1 (4)**, 276–287, doi:10.1007/s40641-015-0027-1, URL http://dx.doi.org/10.1007/s40641-015-0027-1.

Knutti, R., D. Masson, and A. Gettelman, 2013: Climate model genealogy: Generation CMIP5 and how we got there. *Geophy. Res. Let.*, **40**, 1194–1199.

Knutti, R., M. A. A. Rugenstein, and G. C. Hegerl, 2017: Beyond equilibrium climate sensitivity. *Nature Geoscience*, **10**, 727736, doi:10.1038/ngeo3017.

Lipat, B. R., G. Tselioudis, K. M. Grise, and L. M. Polvani, 2017: Cmip5 models' shortwave cloud radiative response and climate sensitivity linked to the climatological hadley cell extent. *Geophysical Research Letters*, **44 (11)**, 5739–5748, doi:10.1002/2017GL073151, URL http://dx.doi.org/10.1002/2017GL073151, 2017GL073151.

Masson, D., and R. Knutti, 2011: Climate model genealogy. *Geophysical Research Letters*, **38 (8)**, doi:10.1029/2011GL046864.

Qu, X., A. Hall, S. A. Klein, and P. M. Caldwell, 2013: On the spread of changes in marine low cloud cover in climate model simulations of the 21st century. *Climate Dynamics*, **42 (9)**, 2603–2626, doi:10.1007/s00382-013-1945-z, URL http://dx.doi.org/10.1007/s00382-013-1945-z.

Read, W. G., and coauthors, 2007: Aura Microwave Limb Sounder upper tropospheric and lower stratospheric H2O and relative humidity with respect to ice validation. *J. Geophys. Res.*, **112**, doi:10.1029/2007JD008 752.

Reynolds, R. W., N. A. Rayner, T. Smith, D. Stokes, and W. W., 2002: An improved in situ and satellite SST analysis for climate. *J. Clim.*, **15**, 16091625.

Rossow, W., and R. A. Schiffer, 1991: Isccp cloud data products. *Bull. Amer. Meteor. Soc.*, **72**, 2–20.

Sanderson, B. M., R. Knutti, and P. Caldwell, 2015a: A representative democracy to reduce interdependency in a multimodel ensemble. *Journal of Climate*, **28 (13)**, 5171–5194, doi: 10.1175/JCLI-D-14-00362.1.

Sanderson, B. M., R. Knutti, and P. Caldwell, 2015b: A representative democracy to reduce interdependency in a multimodel ensemble. *Journal of Climate*, **28 (13)**, 5171–5194.

Sherwood, S. C., S. Bony, and J.-L. Dufresne, 2014: Spread in model climate sensitivity traced to atmospheric convective mixing. *Nature*, **505 (7481)**, 37–42, doi:10.1038/nature12829, URL http://dx.doi.org/10.1038/nature12829.

Siler, N., S. Po-Chedley, and C. S. Bretherton, 2017: Variability in modeled cloud feedback tied to differences in the climatological spatial pattern of clouds. *Climate Dynamics*, **48**, 1–12, doi: 10.1007/s00382-017-3673-2.

Stocker, T., and Coauthors, 2013: Summary for policymakers. in: Climate change 2013: The physical science basis. contribution of working group i to the fifth assessment report of the intergovernmental panel on climate change. Tech. rep., Intergovernmental Panel on Climate Change. 33pp, http://www.climatechange2013.org/images/uploads/WGI_AR5_SPM_brochure.pdf.

Su, H., J. H. Jiang, C. Zhai, T. J. Shen, J. D. Neelin, G. L. Stephens, and Y. L. Yung, 2014: Weakening and strengthening structures in the hadley circulation change under global warming and implications for cloud response and climate sensitivity. *Journal of Geophysical Research: Atmospheres*, **119 (10)**, 5787–5805, doi:10.1002/2014JD021642, URL http://dx.doi.org/10.1002/2014JD021642.

Tian, B., 2015: Spread of model climate sensitivity linked to double-intertropical convergence zone bias. *Geophysical Research Letters*, **42 (10)**, 4133–4141, doi:10.1002/2015GL064119, URL http://dx.doi.org/10.1002/2015GL064119, 2015GL064119.

Volodin, E. M., 2008: Relation between temperature sensitivity to doubled carbon dioxide and the distribution of clouds in current climate models. *Izvestiya, Atmospheric and Oceanic Physics*, **44**, 288–299.

Wagman, B. M., and C. S. Jackson, 2018: A test of emergent constraints on cloud feedback and climate sensitivity using a calibrated single-model ensemble. *Journal of Climate*, **31 (18)**, 7515–7532, doi:10.1175/JCLI-D-17-0682.1.

Williamson, D. B., and P. G. Sansom, 2019: How are emergent constraints quantifying uncertainty and what do they leave behind? *Bulletin of the American Meteorological Society*, **0 (0)**, null, doi:10.1175/BAMS-D-19-0131.1, URL https://doi.org/10.1175/BAMS-D-19-0131.1.

Zhai, C., J. H. Jiang, and H. Su, 2015: Long-term cloud change imprinted in seasonal cloud variation: More evidence of high climate sensitivity. *Geophysical Research Letters*, **42 (20)**, 8729–8737, doi:10.1002/2015GL065911, URL http://dx.doi.org/10.1002/2015GL065911, 2015GL065911.

LIST OF TABLES

TABLE 1. Short description of each emergent constraint tested in this paper along with original citation and evaluation from CZK18. Possible constraints are classified as untestable (lack a physical explanation or don't have an explanation which can be decomposed into feedback and forcing terms) or unclear (ambiguous when evaluated using the CZK18 criteria).

Name	Citation	Credible?	Description
Sherwood D	Sherwood et al. (2014)	Yes	Strength of resolved-scale mixing between BL and lower troposphere in tropical E Pacific and Atlantic
Zhai	Zhai et al. (2015)	Yes	Seasonal response of BL cloud amount to SST variations in oceanic subsidence regions between 20-40°latitude
Brient Shal	Brient et al. (2015)	Yes	Fraction of tropical clouds with tops below 850 mb whose tops are also below 950 mb
Brient Alb	Brient and Schneider (2016)	Yes	Sensitivity of cloud albedo in tropical oceanic low-cloud regions to present-day SST variations
Volodin	Volodin (2008)	Untestable	Difference between tropical and southern-hemisphere midlatitude total cloud fraction
Qu	Qu et al. (2013)	Unclear	BL cloud amount response to SST variations in subtropical stratocumulus regions (after removing EIS contribution)
Su	Su et al. (2014)	Untestable	Error in vertically-resolved tropospheric zonal-average RH between 40°N and 45°S
Tian	Tian (2015)	Untestable	Strength of double-ITCZ bias
Lipat	Lipat et al. (2017)	Unclear	Latitude of the southern edge of the Hadley cell in austral summer
Siler	Siler et al. (2017)	Untestable	Extent to which cloud albedo is small in warm-SST regions and large in cold-SST regions
Cox	Cox et al. (2018)	Untestable	Strength of global-average surface temperature variations and temporal autocorrelation

TABLE 2. Observed values for all constraints used in this study and explanation of how they were obtained. See text for details.

	Mean	σ_o	Explanation
Sherwood D	0.413	0.031	Provides MERRA and ERA Interim values (0.382 and 0.444, respectively)
Brient Shal	44.5%	3.5%	Mean and σ values were provided for ERA Interim over 1979-2012 (45±3%) and Calipso/GOCCP over 2006-20012 (45±3%). Flaws in both estimates were noted.
Zhai	-1.28% K^{-1}	0.187 % K^{-1}	Provides mean and 3σ values computed using CloudSat/CALIPSO cloud fraction and AMSRE SST. Uncertainty was taken as the larger of the northern and southern hemisphere interannual standard deviations for 2006-2010.
Brient Alb	-0.96 % K^{-1}	0.13 % K^{-1}	Computes mean and bootstrapped 90% bounds of 0.96 ± 0.22 % K^{-1} from CERES-EBAF Ed2.8 shortwave fluxes and monthly ERSST sea surface temperature for March 2000 through May 2015. See Brient and Schneider (2016) for details.
Volodin	-25%	5.5%	Suggests mean of -25% with uncertainty <5% based on ISCCP D-2 data. We have computed our own Volodin constraint values of -19%, -30%, and -25% using Calipso, MODIS, and ISCCP data and report the mean and standard deviation over them.
Qu	-2.48% K^{-1}	1.6	Their Table 6 says dLCC/dSST averaged over regions of interest has observed value -2.48% K^{-1} with 95% confidence bounds of ±3.13% K^{-1}. Observed slopes combine ISCCP cloud data (Rossow and Schiffer 1991) and NOAA optimum interpolation monthly SST version 2 (Reynolds et al. 2002).
Su	1	0.25	Our Su metric is the regression slope of model versus observed RH profiles, so matching observations means a slope of 1. RH measurements are taken from August 2004 to December 2012 using AIRS for pressures >300 mb and MLS for lower pressures. Uncertainty of 25% was used based on MLS measurement uncertainty from Read and coauthors (2007).
Tian	0	0.5	Tian measures annual-mean precipitation bias over the Southeast Pacific using monthly-mean data from GPCP between January 1986 to December 2005. Consistency of correlation using TRMM and RH measures is used, but no direct estimate of observational uncertainty is provided. Thus our uncertainty estimate comes completely from expert judgement.
Lipat	35.83	1.75	Table 1 in Lipat et al. (2017) provides the observed mean latitude from ERA Interim for DJF of years 1984-2008 and Figs 3-4 include 95% confidence intervals of 34 to 37.5°.
Siler	0.36	0.05	Their Fig. 9b gives the CERES-EBAF Ed2.8 2003-2008 observed value of 0.36 with 90% bounds of ±0.08 computed assuming the observations follow a normal distribution.
Cox	0.14	0.05	Their extended data table 2 provides mean and σ values from HadCRUT (0.13±0.016), NOAA (0.16±0.034), Berkeley Earth (0.13±0.021), and GISSTEMP (0.12±0.025).

TABLE 3. For each constraint i: Sign of correlation with ECS, standardized best-guess value x_i', ratio of observational uncertainty σ_o to GCM-based standard deviation $\tilde{\sigma}_X$, and sign-corrected GCM-derived correlations r_i with the two climate proxies, after adjusting for observational uncertainty as described in the text.

i	Constraint	Sign	x_i'	$\sigma_o/\tilde{\sigma}_X$	r_i	
	Credible				ECS	λ
1	Sherwood D	+	2.40	0.45	0.40	0.42
2	Brient Shal	+	-0.44	0.18	0.37	0.46
3	Zhai	-	1.23	0.28	0.70	0.76
4	Brient Alb	-	0.68	0.28	0.69	0.80
	Possible					
5	Volodin	-	0.80	0.64	0.45	-0.48
6	Qu	-	0.70	1.70	0.20	-0.14
7	Su	+	0.70	0.88	0.44	0.52
8	Tian	-	1.46	0.81	0.44	-0.38
9	Lipat	-	-0.80	1.01	0.33	-0.22
10	Siler	+	0.35	0.05	0.54	0.64
11	Cox	+	-0.46	0.76	0.50	0.46

TABLE 4. Estimated $\pm 2\sigma$ climate proxy ranges

Method	ECS (K)			λ (K W^{-1} m^2)		
GCM Ensemble	3.24 ± 1.51			-1.14 ± 0.56		
Multiple Emergent Constraints	Credible 4	All 11	No #10	Credible 4	All 11	No #10
Method C	4.14 ± 0.82	4.28 ± 0.69	4.12 ± 0.77	-0.78 ± 0.25	-0.74 ± 0.22	-0.78 ± 0.24
Method U	4.07 ± 0.79	4.01 ± 0.69	4.04 ± 0.72	-0.85 ± 0.29	-0.85 ± 0.25	-0.84 ± 0.26
Method U3	3.84 ± 1.32	3.96 ± 1.30	3.95 ± 1.30	-0.91 ± 0.52	-0.85 ± 0.49	-0.86 ± 0.50

LIST OF FIGURES

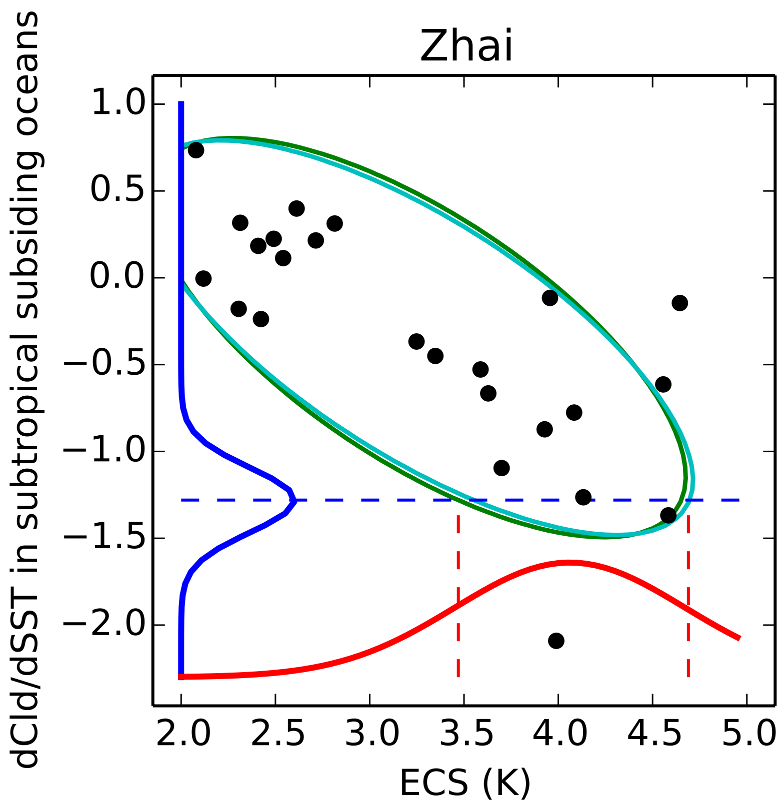

FIG. 1. Example of single-constraint methodology applied to Zhai constraint. Black dots are values for individual GCMs. Ellipses are the 0.94 contour of the bivariate Gaussian between the ECS and the constraint excluding (cyan) and including (green) observational uncertainty. The PDF of the observed constraint value is in blue along the y axis. The red curve is the posterior PDF of ECS given the observed best-guess constraint value (blue dash) and the green bivariate PDF.

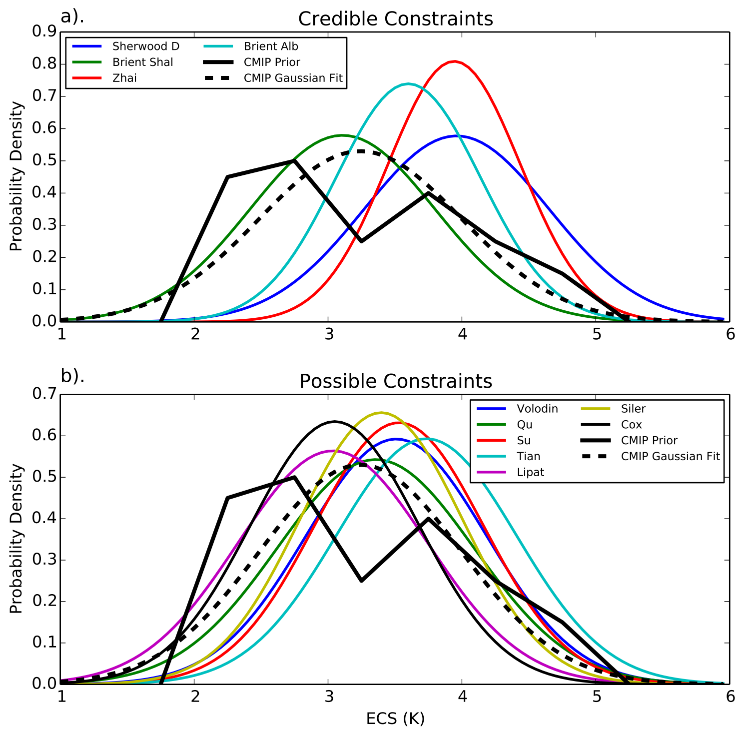

705 FIG. 2. Posterior PDFs for ECS based on individual constraints using the method of Sect. 3.4. Panels (a) and
706 (b) are for credible and possible constraints, respectively.

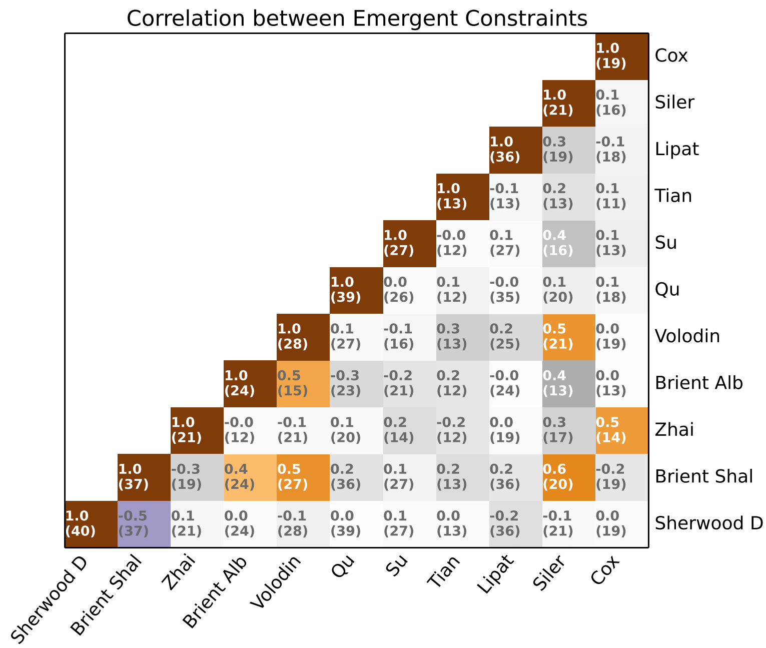

FIG. 3. Partial correlations between pairs of emergent constraints over all available models, removing their joint correlation with ECS. The number of models used for each constraint pair is listed in parentheses in the bottom of each box. Boxes with orange or purple hue are statistically significant at 90% using a 2-tailed T-test. Grayish boxes are not statistically significant.

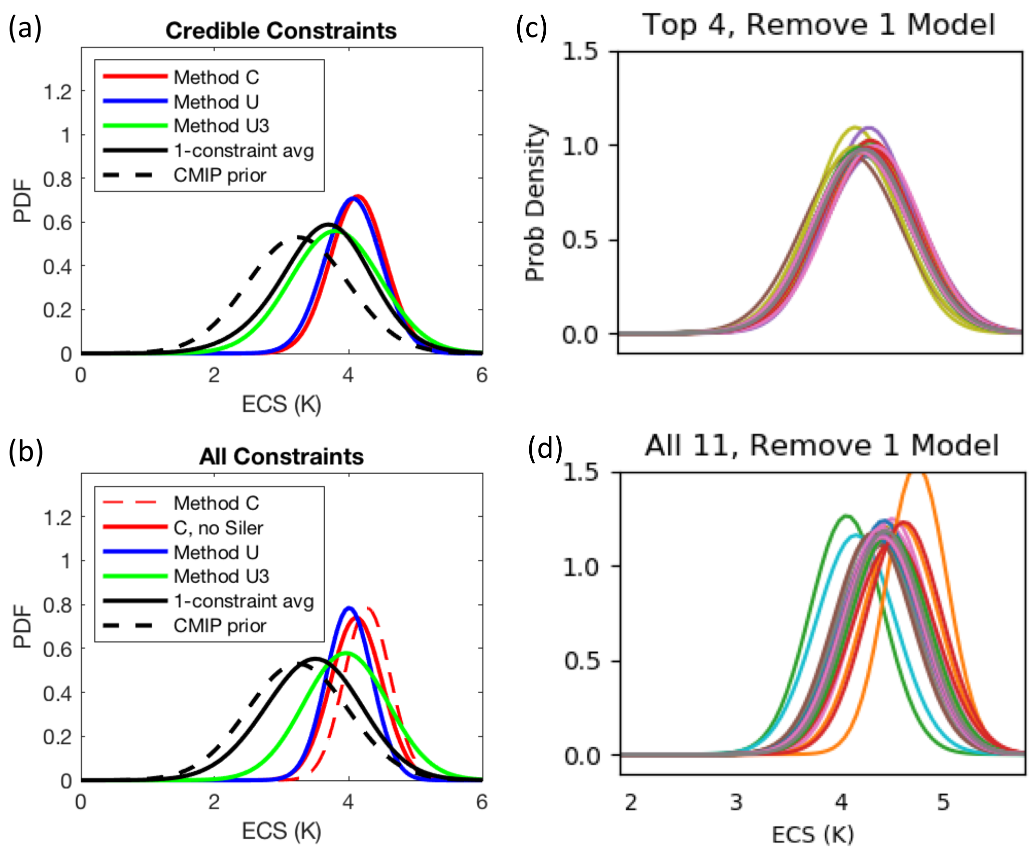

FIG. 4. Posterior PDFs for ECS based on (a) 4 credible constraints and (b) all 11 constraints, using Methods C, U and U3, and (Method C only) the 10 constraints excluding Siler. Black dashed lines in both panels show the CMIP Gaussian prior for comparison. (c) and (d) show Method C results, but each PDF is calculated with one GCM removed to test for robustness.

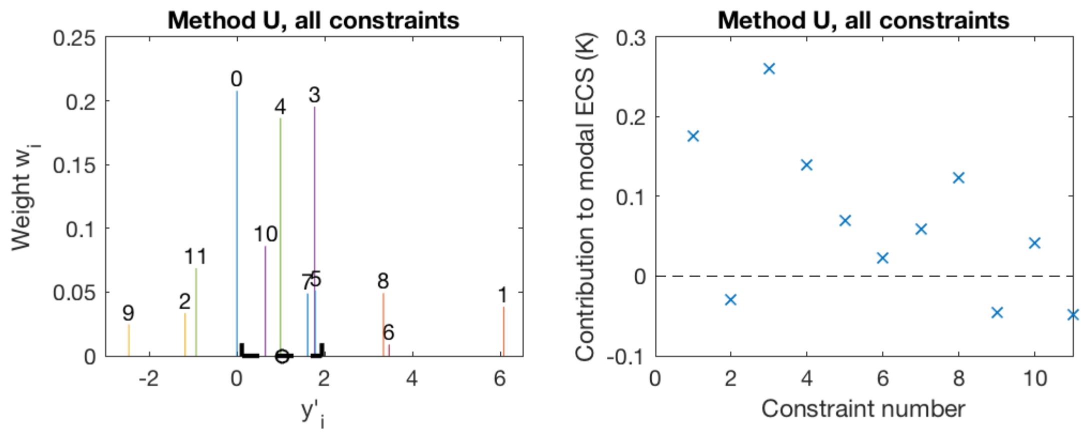

FIG. 5. Method U all-constraint interpretive results: (a) Constraint contributions to posterior mean of ECS, shown as vertical line segments with heights equal to the constraint weights w_i at horizontal positions equal to the nondimensional ECS y_i' that is implied by 'swapped' regression of the constraint on the ECS together with the observed nondimensional constraint value. Each line segment is labelled by its constraint number i; $i = 0$ refers to the GCM prior. The circle and the horizontal dashed black line between the two bars denotes the resulting $\pm 2\sigma$ posterior range of ECS. (b) Contributions of individual constraints to changing the posterior mean of ECS away from the GCM mean of 3.2 K.