Bayesian Hierarchical Models for Detection and Attribution

Matthias Katzfuss, Texas A&M

Dorit Hammerling, NCAR

Richard Smith, Univ. of North Carolina and SAMSI

IDAG, March 14, 2017

Outline

1. Current approaches and limitations
2. A Bayesian hierarchical approach
3. Application to temperature data
4. Conclusions

Aim of Detection and Attribution (D&A)

General definition from the climate literature:

> *The aim of Detection and Attribution is to assess whether observed changes are consistent with internal climate variability only, or with an expected response to a combination of external forcings and internal climate variability.*

Example from IPCC report

All forcings (including anthropogenic)

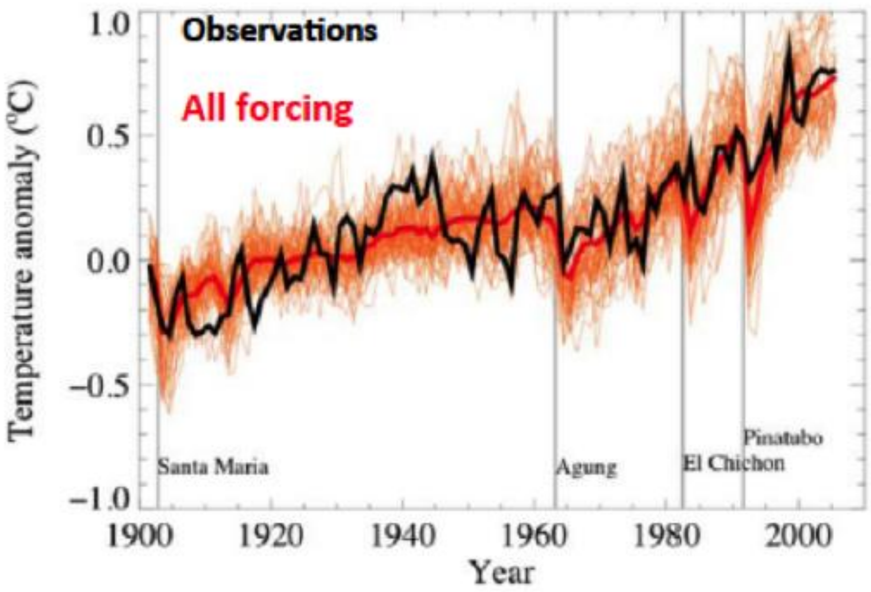

Only natural forcings

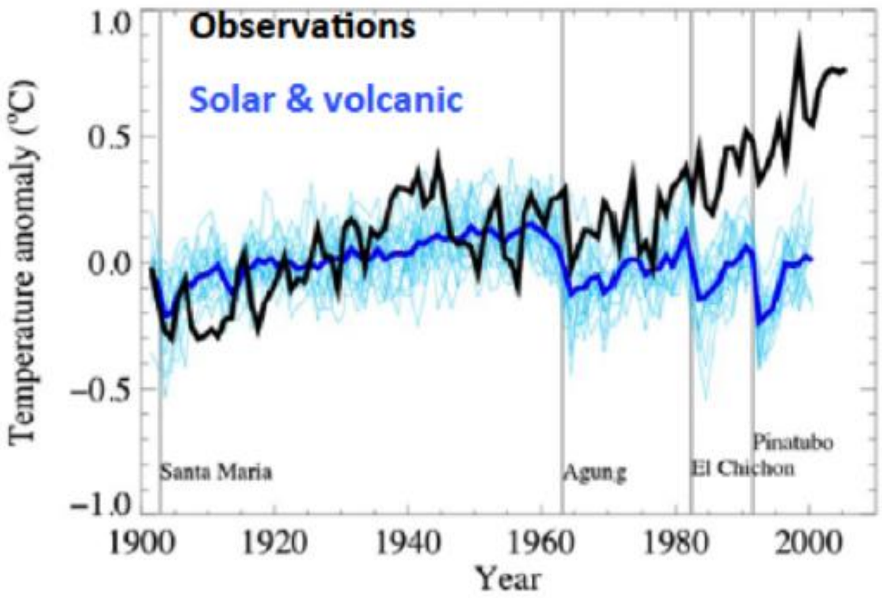

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

IPCC statements (loosely quoted)

Warming of the climate system is unequivocal, as is now evident from observations of increases in average air and ocean temperatures, widespread melting of snow and ice, and rising global average sea level. [..]

Most of the observed increase in global average temperatures since the mid-20th century is **likely** *(TAR 2001) due to the observed increase in anthropogenic greenhouse gas concentrations.*

. . . mid-20th century is **very likely** *(AR 2007) due to . . .*

. . . mid-20th century is **extremely likely** *(AR 2013) due to . . . .*

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

Regression-based D&A

$$\mathbf{y} = \sum_{j=1}^{m} \beta_j \mathbf{x}_j + \mathbf{u} = \mathbf{X}\beta + \mathbf{u}$$

where

- $\mathbf{y}$: vector of observed signals
  Here: observed temperature changes at spatial grid cells
- $\mathbf{x}_1, \ldots, \mathbf{x}_m$: responses to the m different forcings
  Here: temperature changes that would have happened under each forcing scenario ($\mathbf{x}_1$: anthropogenic; $\mathbf{x}_2$: natural)
- $\mathbf{u}$: internal climate variability
  Usual assumption: $\mathbf{u} \sim N_n(\mathbf{0}, \mathbf{C})$

GLS solution:

$$\hat{\beta} = (\mathbf{X}'\mathbf{C}^{-1}\mathbf{X})^{-1}\mathbf{X}'\mathbf{C}^{-1}\mathbf{y}$$

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

Detection and Attribution

If a particular coefficient β_j is significantly greater than 0, we say that the jth forcing factor has been *detected*.

Among those forcing factors that are detected, the corresponding β_js are then interpreted as the *attribution* of the observational signal to the different forcing factors

Challenges

$$\mathbf{y} = \sum_{j=1}^{m} \beta_j \mathbf{x}_j + \mathbf{u}, \quad \mathbf{u} \sim N_n(\mathbf{0}, \mathbf{C})$$

- The $\mathbf{x}_j$s are not actually known (errors-in-variables problem) — climate scientists have addressed this using the *total least squares* algorithm, but there are problems with this approach
- Recently, climate scientists have started to realize that the $\mathbf{y}$s aren't actually known either
  $\Rightarrow$ uncertainty in temperatures are characterized through an ensemble of possible temperatures
- Uncertainty in parameter estimation not taken into account in current approaches

Challenges with the estimation of C

Main challenge: y and $x_1, \ldots, x_m$ are high dimensional (typically thousands) but the number of independent "observations" is comparatively small. This makes estimation of C difficult.

Solutions used by climate scientists:

- Estimate C from *control runs* of the climate model
- Expand in *empirical orthogonal functions* (principal components) and then truncate (in an ad-hoc fashion)

We will address these challenges by formulating and fitting a Bayesian hierarchical model

Bayesian D&A regression model

Bayesian regression model:

$$\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}, \mathbf{C} \sim N_n\Big(\sum_{j=1}^{m} \beta_j \mathbf{x}_j, \mathbf{C} \Big)$$

D&A consists of determining the posterior distribution of the β_j (mainly, $P(\beta_j > 0|\mathbf{y}, \mathbf{X})$).

Challenge: $\mathbf{y}$, $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_m)$, $\boldsymbol{\beta}$, and $\mathbf{C}$ are all unknown

NCAR

NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

Uncertainty in observed temperature changes

True temperature changes in grid cells over the globe are unknown
But: We have an ensemble of N temperature time series, which can be converted to an ensemble of N temperature changes

We assume that

$$\mathbf{y}^{(i)} | \mathbf{y}, \mathbf{W} \overset{iid}{\sim} N_n(\mathbf{y}, \mathbf{W}), \quad i = 1, \ldots, N,$$

where $\mathbf{W}$ is a covariance matrix describing the variability of the ensemble members around the true temperature change

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

Uncertainty in temperature under forcing scenarios

The (true) temperature changes due to forcing are also unknown, but we have an ensemble of GCM outputs for each forcing scenario:

$$\mathbf{x}_j^{(l)} | \mathbf{x}_j, \mathbf{C} \overset{ind}{\sim} N_n(\mathbf{x}_j, \mathbf{C}), \quad l = 1, \ldots, L_j, \; j = 1, \ldots, m,$$

where L_j is the number is the number of GCM runs under the jth forcing scenario, and climate variability is assumed to have covariance matrix $\mathbf{C}$.

Model parameters

- Climate variability: Typically, $\mathbf{C}$ is expanded in *empirical orthogonal functions* and then truncated:
  $\mathbf{C} = \mathbf{BKB}'$, where $\mathbf{B}$ contains the first r principal components estimated from control runs, $\mathbf{K} = diag\{e^{\lambda_1}, \ldots, e^{\lambda_r}\}$, and $r << n$
- Observation uncertainty: Currently, $\mathbf{W} = \sigma^2 \widetilde{\mathbf{W}}$, where $\widetilde{\mathbf{W}}$ is a diagonal matrix containing the empirical variances of $\{\mathbf{y}^{(i)}\}$
- Priors:
  - Noninformative priors for β and σ
  - Vaguely informative priors for $\lambda_1, \ldots, \lambda_r$

Inference

MCMC with adaptive Metropolis-Hastings updates

High-dimensional problem $\rightarrow$ Integrate out $\mathbf{y}$ and $\mathbf{X}$

MCMC computations only rely on low-dimensional quantities and are very fast, even for almost a million data points

Bayesian model averaging

Previous slides assumed r, the number of EOFs, to be fixed.

The number of variables in the model depends on $r \rightarrow$ standard MCMC sampler cannot be used to make inference on θ and r simultaneously.

Instead we perform Bayesian model averaging (BMA) to average the posterior results for each value of r using weights automatically chosen by the data.

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

Bayesian model averaging (cont.)

The posterior of β averaged over the posterior of r (i.e., taking the uncertainty about the value of r into account) is given by

$$[\beta|\mathcal{Y},\mathcal{X}] = \sum_{i=\min}^{\max} [\beta|r_i,\mathcal{Y},\mathcal{X}][r_i|\mathcal{Y},\mathcal{X}],$$

Due to the uniform prior on r, the posterior probability of $r = r_i$ is given by $[r_i|\mathcal{Y},\mathcal{X}] \propto [\mathcal{Y}|r_i,\mathcal{X}]\,[r_i] \propto [\mathcal{Y}|r_i,\mathcal{X}]$.

Fortunately, a good estimate of marginal likelihood $[\mathcal{Y}|r_i,\mathcal{X}]$ can be obtained using the evaluations of the likelihood already performed in the MCMC procedure as

$$[\mathcal{Y}|r_i,\mathcal{X}] = \frac{1}{M}\sum_{j=1}^{M}[\mathcal{Y}|r_i,\boldsymbol{\theta}^{(j)},\mathcal{X}][\boldsymbol{\theta}^{(j)}|r_i]$$

Computational considerations

Bayesian modeling averaging approach is ideally suited for parallelization.

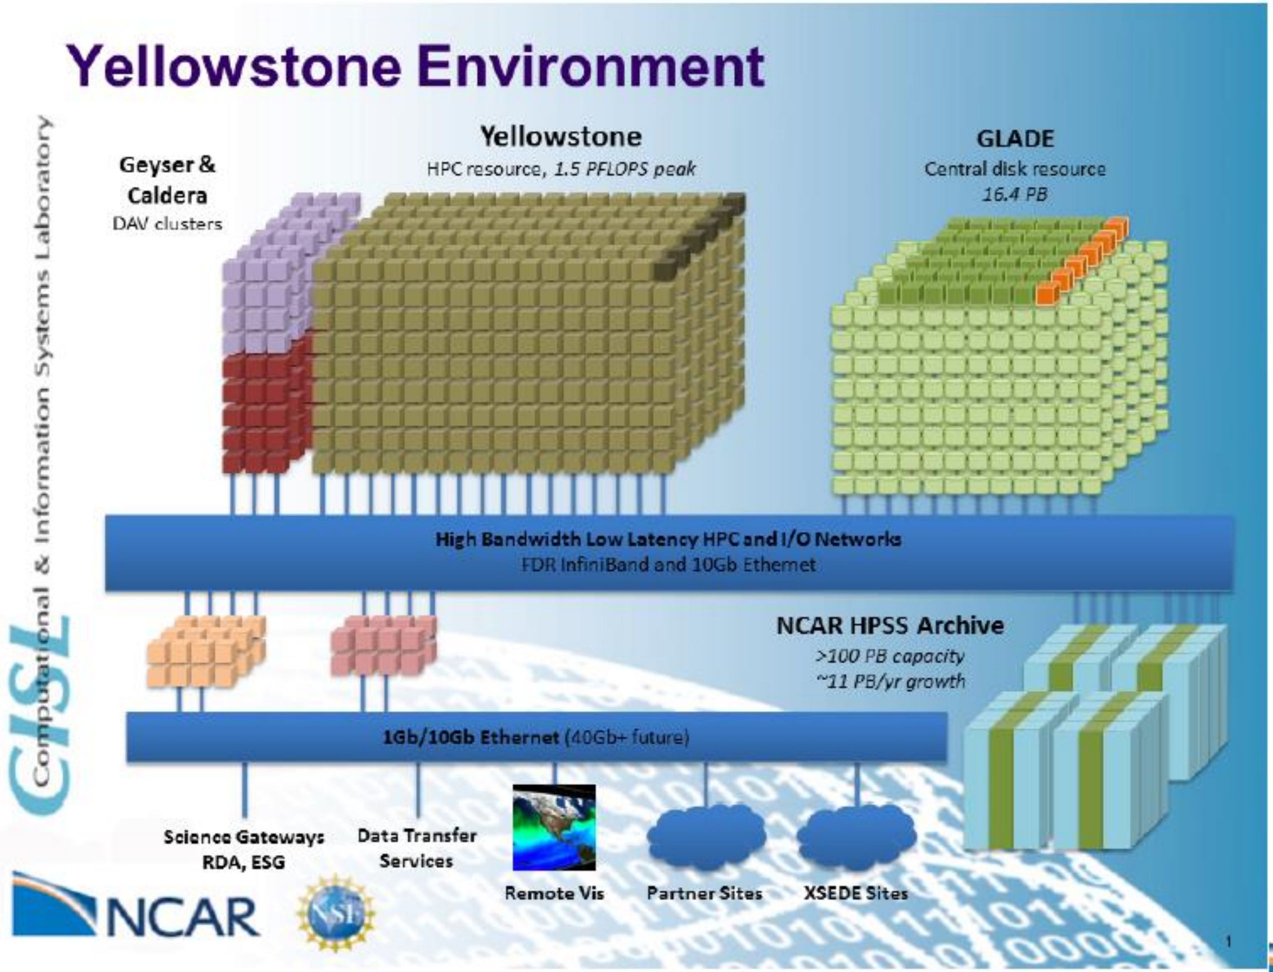

Parallelizing over r (161EOFs): 40 hours on laptop $\rightarrow$ 2 hours on Geyser

The data

- Climate Model Intercomparison project (CMIP5) models: suite of more than 20 models, of which we use as subset (BCC CSM1, CAN ESM2, CSIRO, GISS, IPSL, GFDL)
- Remote Sensing Systems temperature retrievals based on microwave sounding units (MSUs): $N = 396$ realizations

Based on these sources, we consider the linear trends (slopes) of annual lower-tropospheric temperatures between 1979 and 2005 in $n = 2107$ $5^\circ \times 5^\circ$ grid cells on the globe (between -70° and 80° latitude with an altitude lower than 3km)

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

Linear trends 1979–2005: Natural-only forcing

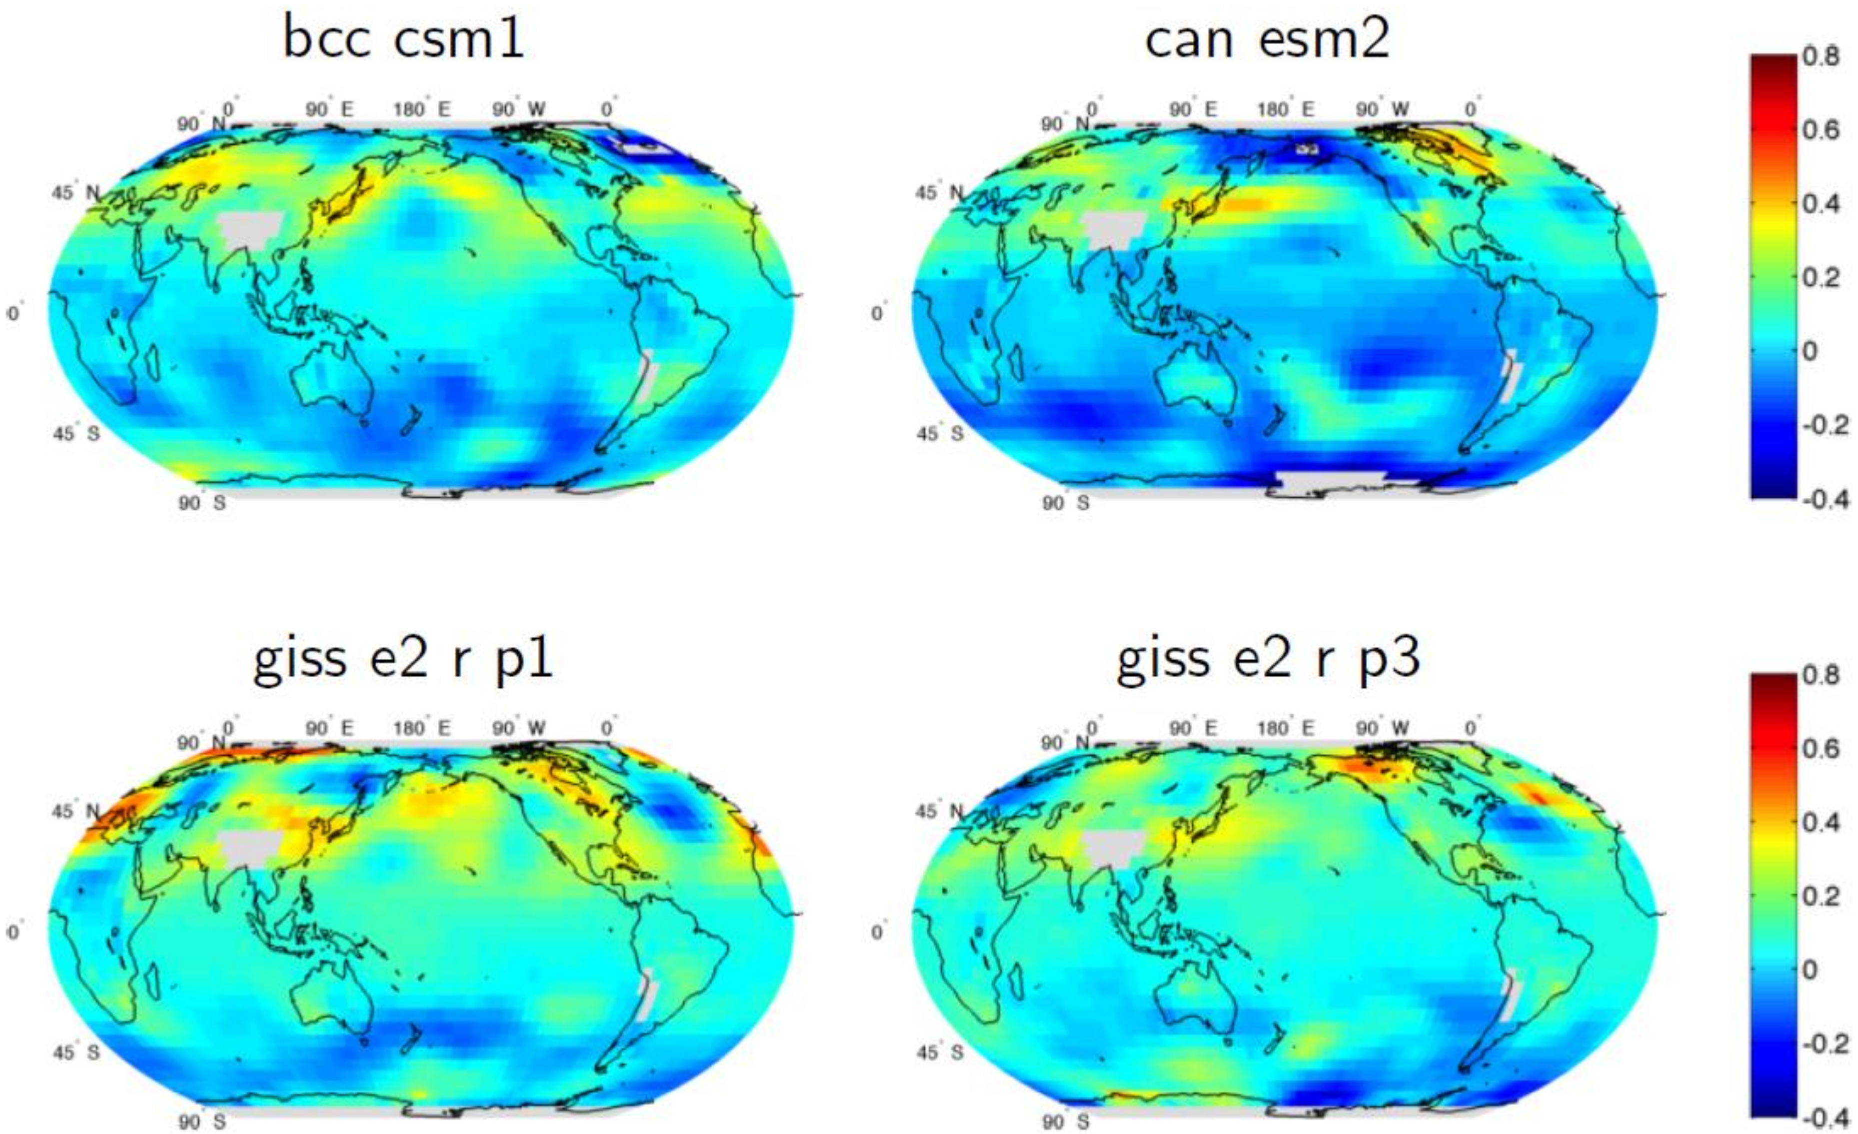

Units are °C per decade

Linear trends 1979–2005: Anthropogenic-only forcing

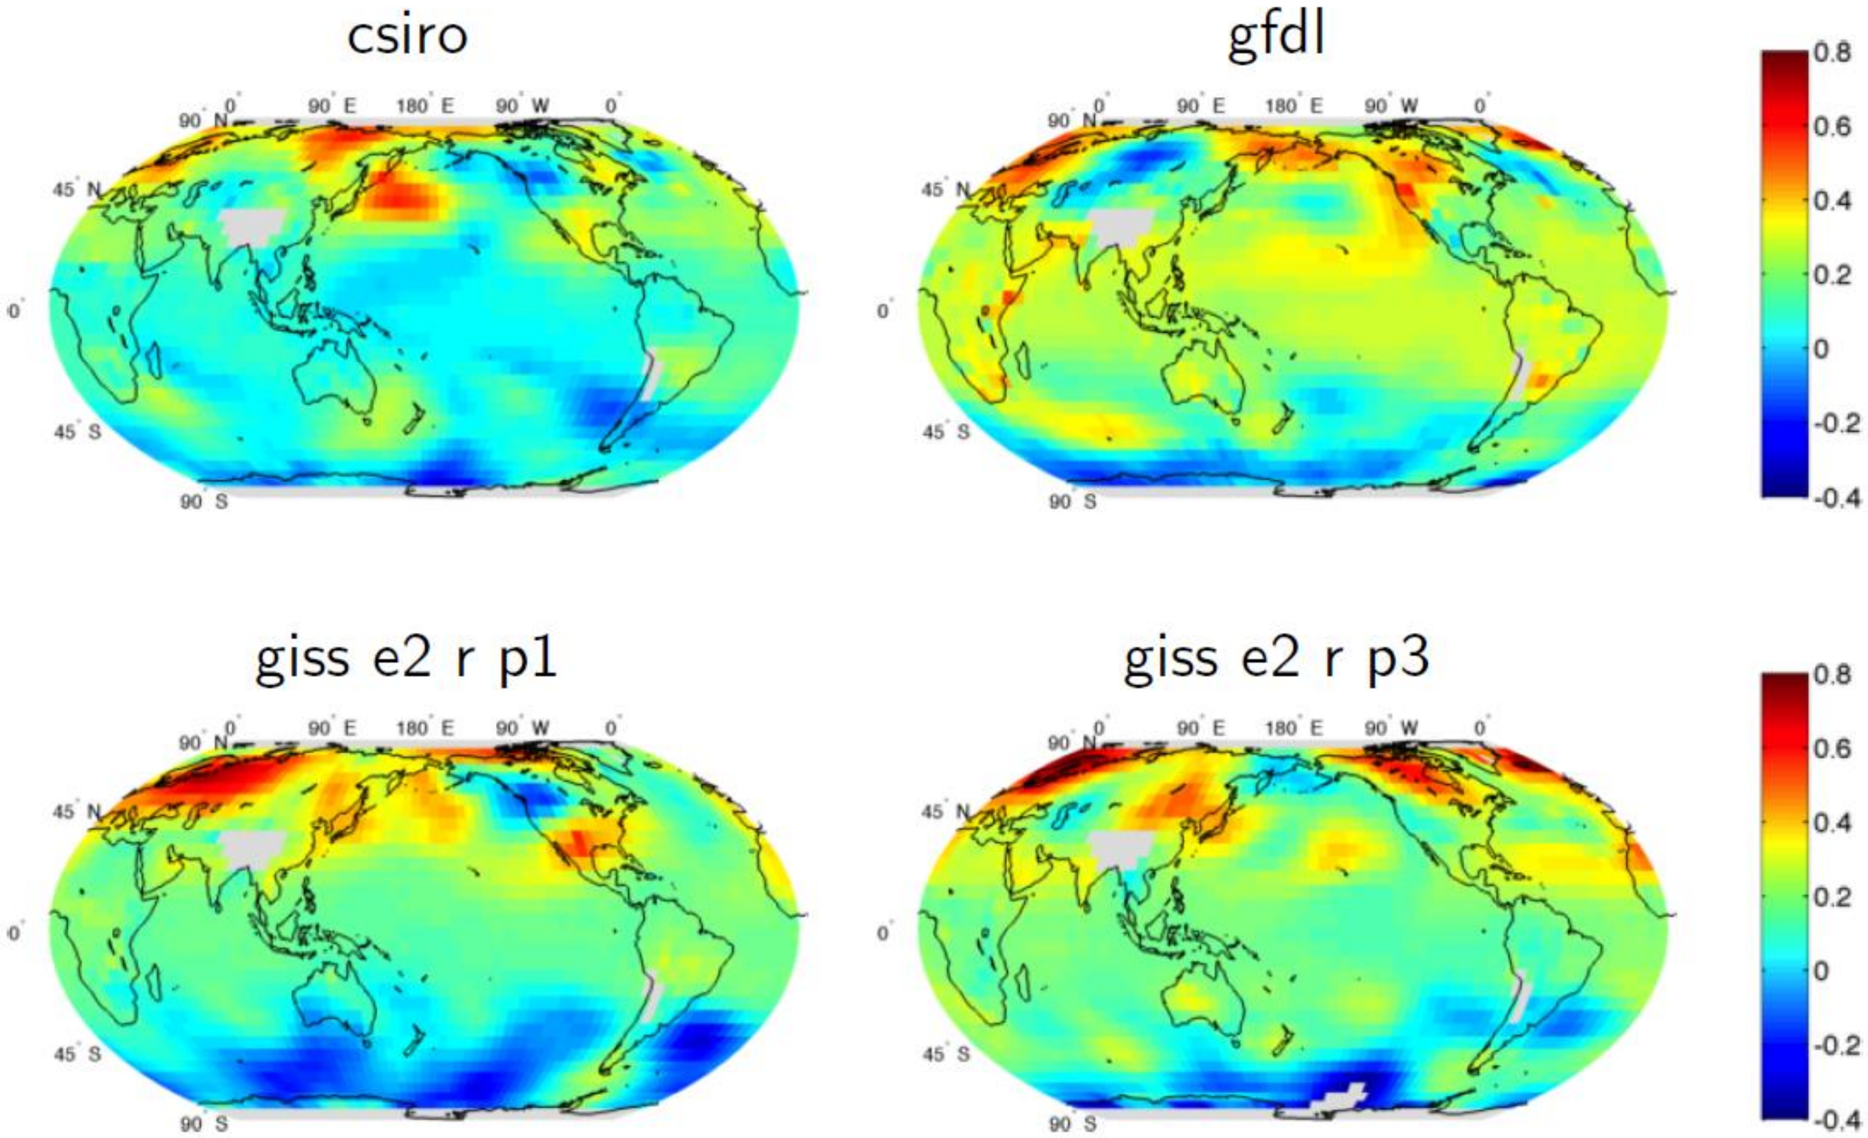

Units are °C per decade

Linear trends 1979–2005: giss models

Natural-only forcing | Anthropogenic-only forcing

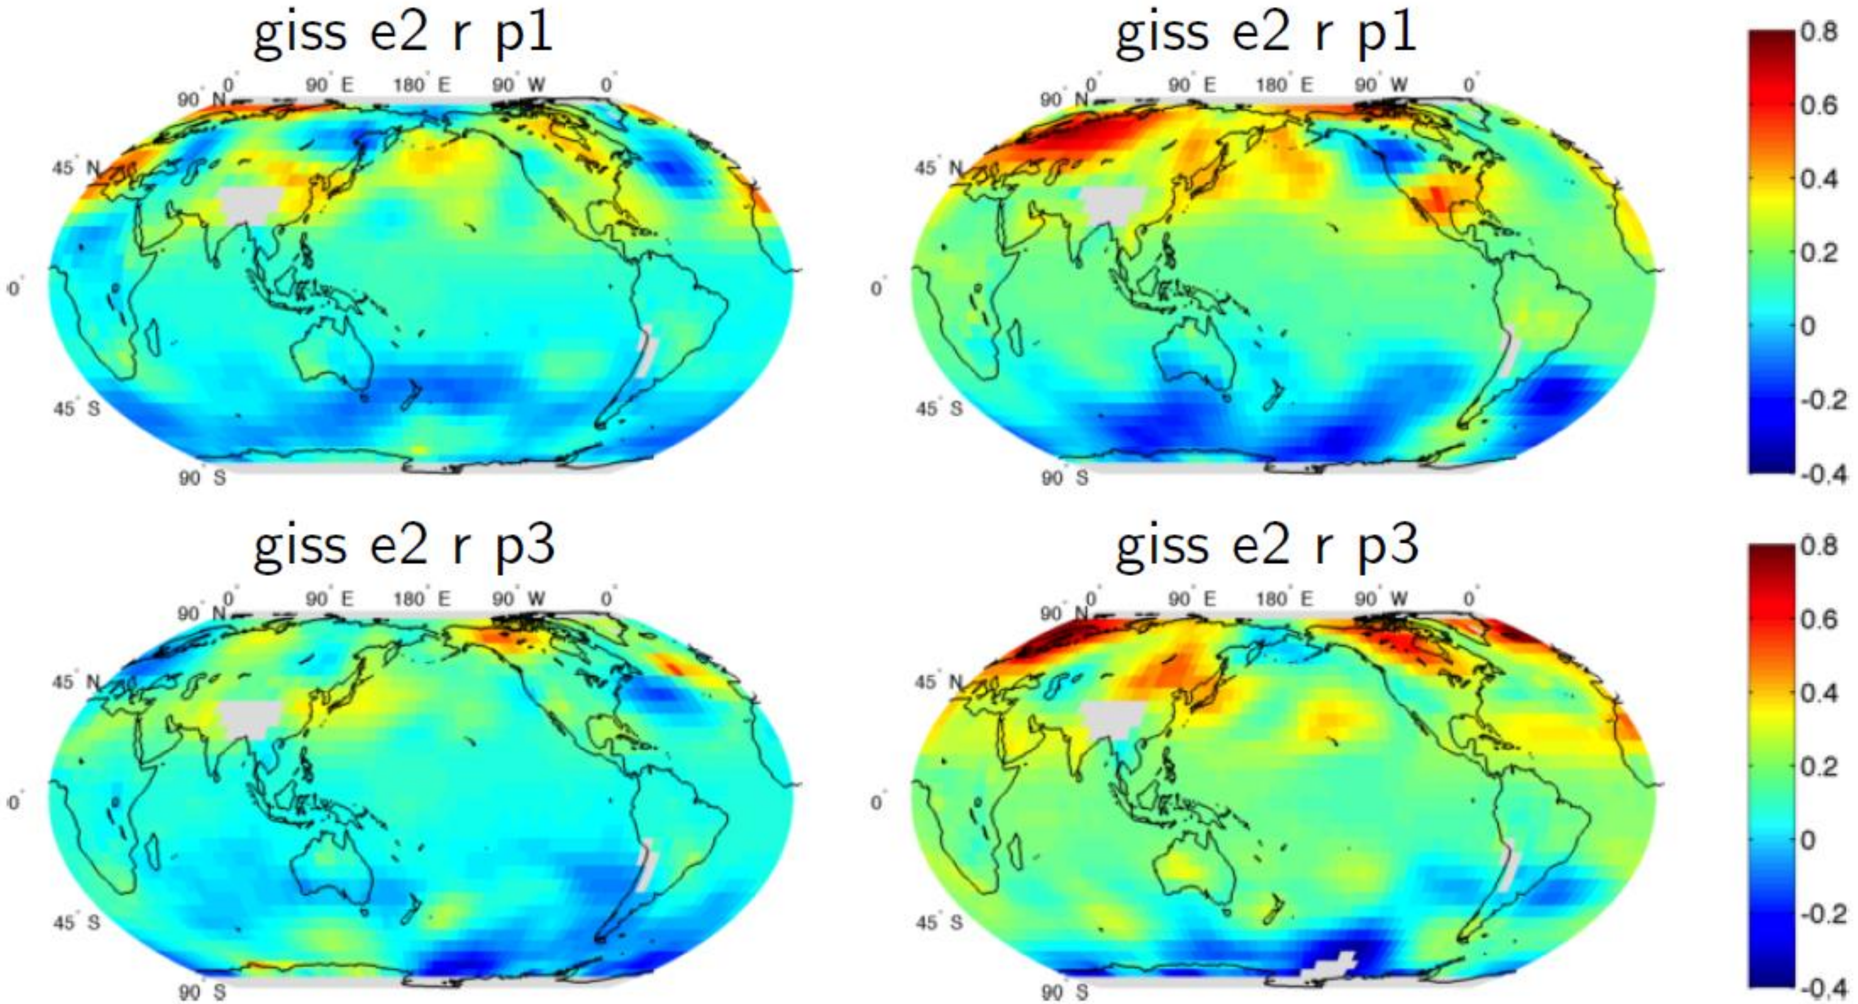

Units are °C per decade

Linear trends in satellite observations 1979–2005

Average of 396 ensemble members

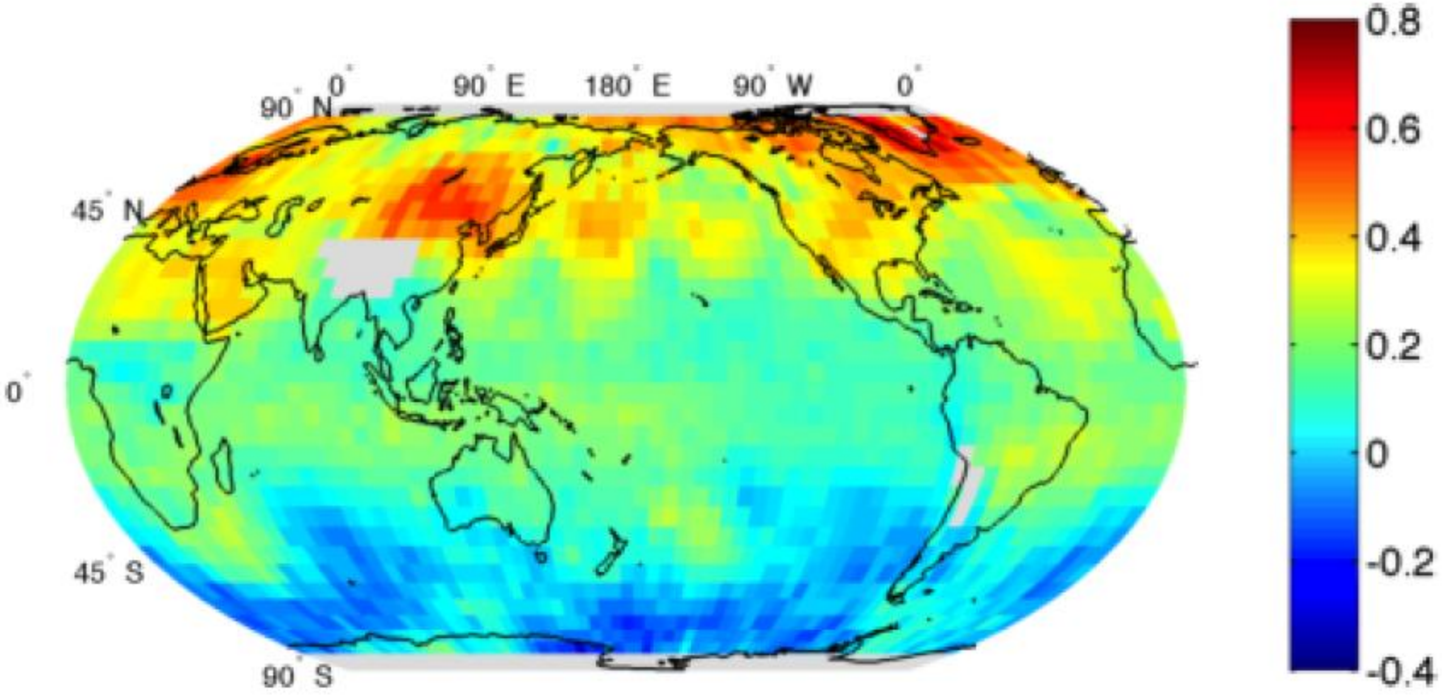

Posterior densities for βs for all values of r

All available GCM models for forced runs, bcc model for control (18 runs)

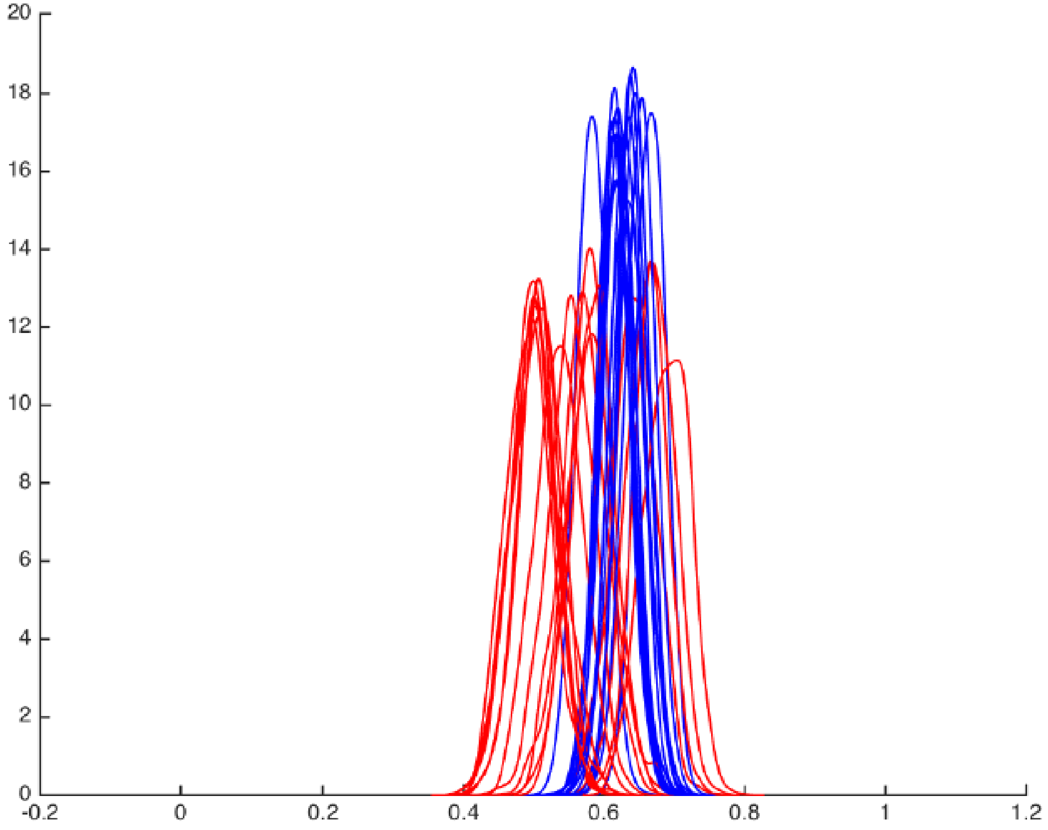

blue(β_1) corresponds to anthropogenic forcings, red(β_2) to natural forcings

Weights for all values of r

All available GCM models for forced runs, bcc model for control (18 runs)

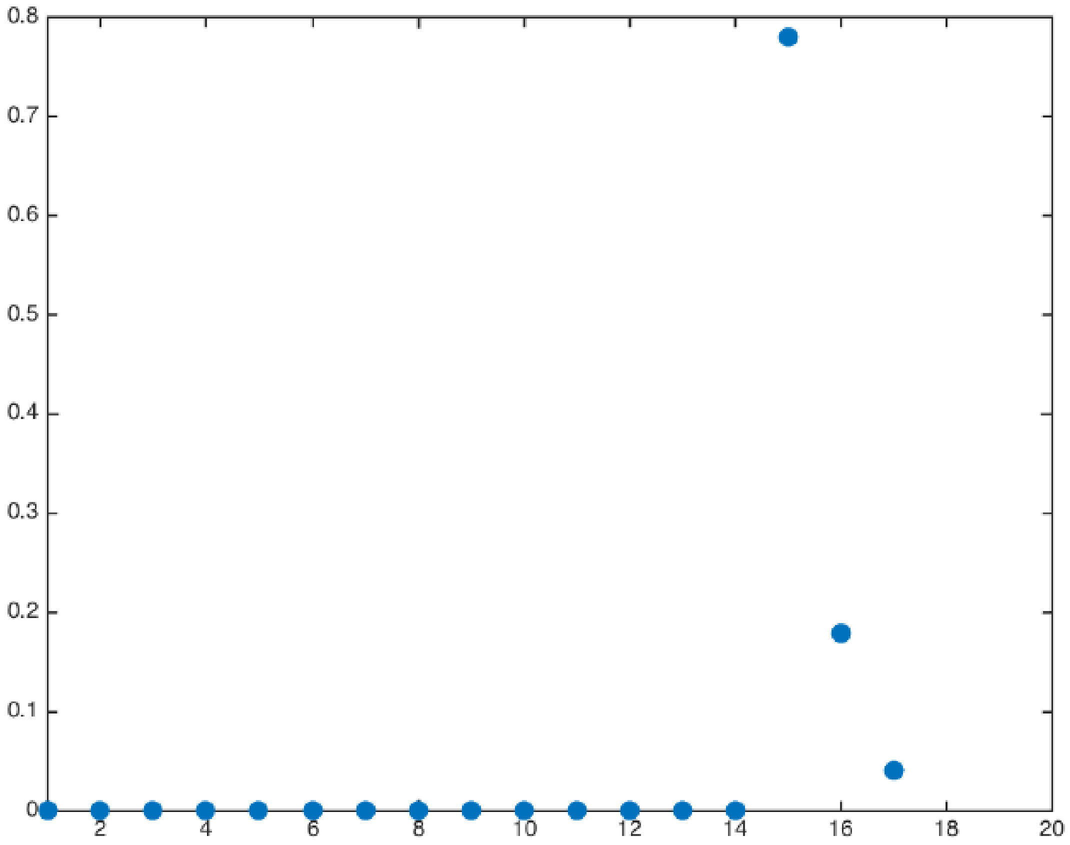

Bayesian model averaged posterior densities for βs

All available GCM models for forced runs, bcc model for control (18 runs)

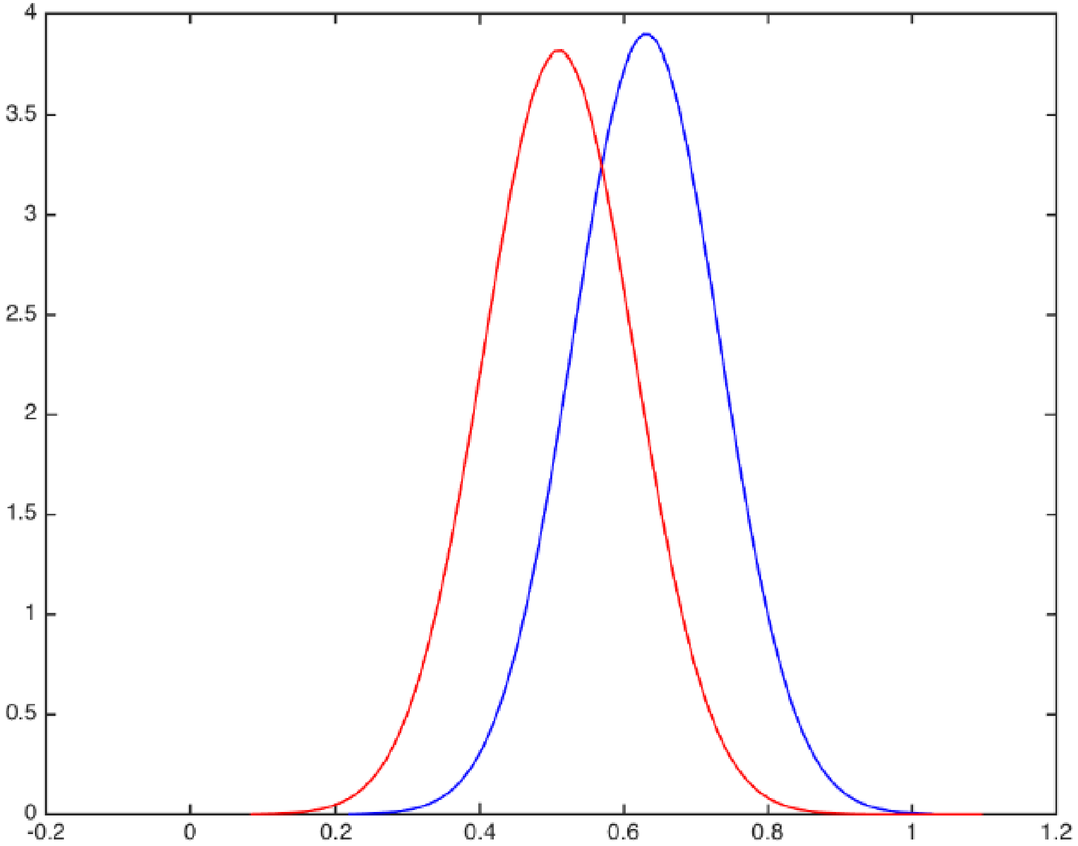

blue(β_1) corresponds to anthropogenic forcings, red(β_2) to natural forcings

Posterior densities using different control runs

Control runs: bcc

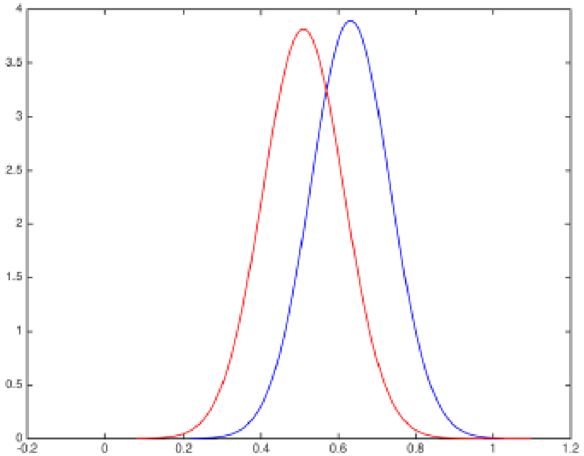

Control runs: gfdl

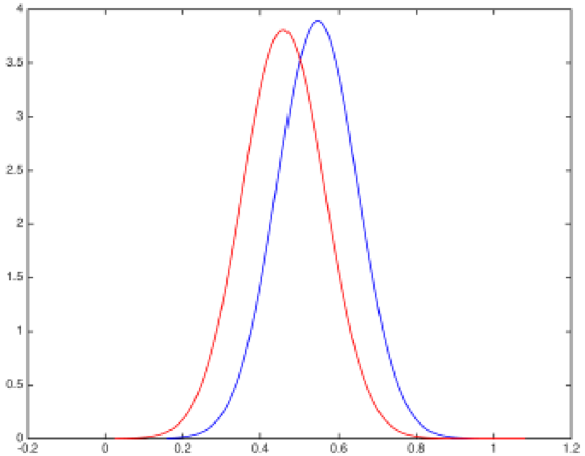

Control runs: csiro + ipsl

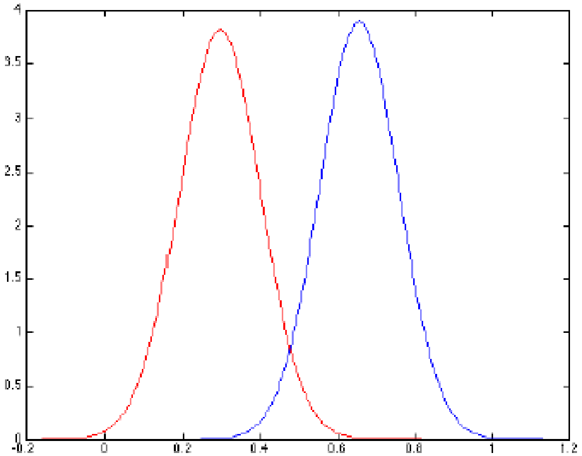

Control runs: from nine models

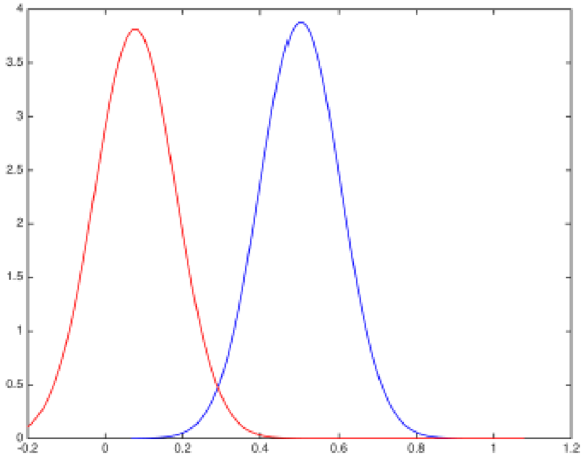

Relation to Method of Hannart (J. Climate 2016)

Hannart presented an "integrated Optimal Fingerprinting" approach that has several overlaps with the current method

- Not expicitly Bayesian but uses several elements derived from Bayesian theory, in particular, *integrated likelihoods*
- Initial $\widehat{C} = S$ where S is sample covariance matrix from control model runs
- Improved estimate $\widehat{C}_\alpha = \alpha\Delta + (1-\alpha)S$ for suitably chosen α, Δ
- Inverse Wishart "prior distribution" on C; integrate out C from Likelihood
- Didn't take account of observational uncertainty
- Open question which method performs better

Summary

- BHM allows natural modeling of uncertainty in all quantities in the D&A regression model
- Posteriors take all (modeled) uncertainties into account
- Results not sensitive to priors
- BUT results are sensitive to choice of control runs

Future work:

- Inference on EOFs themselves
- Or completely different approach to estimating covariance

Acknowledgements

- Originally started as a SAMSI working group
- Data provided by Ben Santer (Lawrence Livermore National Lab) and Carl Mears (Remote Sensing Systems)
- Paper submitted to GRL

SAMSI Program on Mathematical and Statistical Methods for Climate and the Earth System 2017-18

- **Program Leaders**
  - Amy Braverman, Jet Propulsion Laboratory
  - Peter Challenor, Exeter
  - Doug Nychka, National Center for Atmospheric Research
  - Michael Wehner, Lawrence Berkeley National Lab
  - Mary Lou Zeeman, Bowdoin College, Mathematics,
- **Directorate Liaison**
  - Richard Smith, SAMSI
- **Local Scientific Coordinators**
  - Montse Fuentes, North Carolina State University
  - Christopher Jones, University of North Carolina at Chapel Hill
- **National Advisory Committee Liaisons**
  - Habib Najm, Sandia National Laboratory
  - Michael Stein, University of Chicago