Well-estimated global surface warming in climate projections selected for ENSO phase

JAMES S. RISBEY

CSIRO Oceans and Atmosphere
Hobart, Australia *

STEPHAN LEWANDOWSKY

Univ. of Bristol and Univ. of Western Australia

CLOTHILDE LANGLAIS, DIDIER P. MONSELESAN
TERENCE J. O'KANE

CSIRO Oceans and Atmosphere

NAOMI ORESKES

Harvard University

---

*Corresponding author address: James Risbey, CSIRO Oceans and Atmosphere, Box 1538, Hobart, Tas 7001, Australia.
E-mail: james.risbey@csiro.au ⋄ Phone: +61 3 6232-5086

1

The question of how climate model projections have tracked the actual evolution of global mean surface air temperature (GMST) is important in establishing the credibility of their projections. Some studies and the Intergovernmental Panel on Climate Change (IPCC) AR5 suggest that the most recent 15 year period provides evidence that models are overestimating current temperature evolution. Such comparisons are not evidence against model trends because they represent only one realisation where the decadal natural variability component of the model climate is generally not in phase with observations. We present a more appropriate test of models where only those models with natural variability (represented by El Niño-Southern Oscillation [ENSO]) largely in phase with observations are selected from multi-model ensembles for comparison with observations. These tests show that climate models have provided good estimates of 15 year trends, including for recent periods and for Pacific spatial trend patterns.

The differences between climate model *forecasts* and *projections* [1] have come to prominence over interpretation of model simulations of recent temperature trends. A key difference between a climate forecast and a climate projection is that the former attempts to account for the correct phase of natural internal climate variations while the latter does not and can not.

A weather forecast attempts to account for the growth of particular synoptic eddies and is said to have lost skill when model eddies no longer correspond one to one with those in the real world. Similarly, a climate forecast of seasonal or decadal climate attempts to account for the growth of disturbances on the time scale of those forecasts. This means that the model must be initialized to the current state of the coupled ocean-atmosphere system

and the perturbations in the model ensemble must track the growth of ENSO [2, 3] and other subsurface disturbances [4] driving decadal variation. Once the coupled climate model no longer keeps track of the current phase of modes like ENSO, it has lost forecast skill for seasonal to decadal time scales. The model can still simulate the statistical properties of climate features from this point, but that then becomes a projection, not a forecast.

The CMIP5 series of coupled climate models have been run in climate projections mode [5] and, to a limited extent, in decadal forecasting mode [6]. The models run as climate projections apply best estimates of the historical sequence of radiative forcing of climate for the past (until 2005), followed by specified future forcing scenarios. That means that until 2005 the models attempt to stay in sequence only with the year to year and decade to decade fluctuations in climate caused by the historical variation in radiative forcing (and not with internal variations).

Decadal variations in surface climate are due to a range of factors [7, 8]. These factors include radiative forcing variations, but they are not the only, or even the most important, factors. Natural internal variations in the climate system also drive decadal variations [9, 10, 11]. These can occur for example through variations in the rate at which the ocean circulation takes up the additional heat added to the atmosphere from greenhouse forcing [12, 13]. Long term variations in the preference of the coupled system for La Niña and El Niño ENSO states (the Pacific Decadal Oscillation [14]) change the rate of ocean heat uptake and are a key driver of decadal variability [13, 15].

In the CMIP5 models run using historical forcing there is no way to ensure that the model has the same sequence of ENSO events as the real world. This will only occur by chance and only for limited periods, since natural variability in the models is not constrained to occur in the same sequence as the

real world. For any 15 year period the rate of warming in the real world may accelerate or decelerate depending on the phase of ENSO predominant over the period. That means that for a set of model projections well calibrated to the range of natural variability, there will be some 15 year periods where the observed rate of warming is in the low tail of the distribution of model trends for that 15 year period, and some 15 year periods where the observed rate of warming is in the high tail of the model distribution. These cases are illustrated by the 15-year observed trend and CMIP5 model trend distribution for 1998–2012 (Fig. 1a) and 1984–1998 (Fig. 1b). These two periods are no more meaningful in evaluating projections than any other 15-year periods, and we focus on the former here only to evaluate claims made about model projections in the most recent 15-year period.

The precise degree to which the observed trend in Fig. 1a is in the extreme low tail is open to question. When one accounts for a range of observational trend estimates by using series that include a representation of the Arctic region [16, 17], the observed estimates fall further inside the model distribution (blue and red lines in Fig. 1a) than for the HadCRUT4 trend (dashed red).

Regardless of just where in the tail they fall, the fact that the observed trends in Fig. 1a for the recent 15-year period are in the low tail of the model distribution [7, 18] does not constitute evidence against the fidelity of model trend projections because it is only one realisation of 15-year periods. In a properly constructed model forecast ensemble [19], a single observed realisation lying in the tail of the model ensemble provides no more evidence against the model forecast than a single observed realisation lying on the ensemble median. Forecast ensemble members have similar likelihood regardless of where they lie and the real world may follow any of them.

4

In order to see how representative the two 15-year periods in Fig. 1 are of the model's ability to simulate 15 year temperature trends we need to test many more 15-year periods. Using data from CMIP5 models and observations for the period 1880–2012, we have calculated sliding 15-year trends in observations and models over all 15-year periods in this interval (Fig. 2). The 2.5–97.5 percentile envelope of model 15-year trends (grey) envelops within it the observed trends for almost all 15-year periods for each of the observational datasets. There are several periods when the observed 15 year trend is in the warm tail of the model trend envelope ($\sim 1925, 1935, 1955$), and several periods where it is in the cold tail of the model envelope ($\sim 1890, 1905, 1945, 2005$). In other words, the recent 'hiatus' centred about 2005 (1998–2012) is not exceptional in context. One expects the observed trend estimates in Fig. 2 to bounce about within the model trend envelope in response to variations in the phase of processes governing ocean heat uptake rates, as they do.

The variations in the grey envelope of model trends in Fig. 2 represent primarily variations in forcing, since the multi-model ensemble tends to average out the internal variations (which are not phase locked between models or with observations). On the other hand, the observed 15 year trends respond to both variations in forcing and internal variations. Any mismatch in forcing (between the models and real world) or any large natural excursion in the rate of ocean heat uptake moves the observed GMST trend up or down within the model distribution represented by the grey envelope.

The histogram of observed temperature trends for all 15-year periods over 1880–2012 is shown in Fig. 3. Plotted over this histogram is the distribution of 15-year trends from all the CMIP5 models for the same set of 15-year periods over 1880–2012. The observed and model distributions are broadly

similar, consistent with the above results. The observed trend median is slightly higher (warmer) than the model median, and the models have a slightly larger spread.

Taken as a collective of 15-year periods over the period since 1880, the CMIP5 models provide a remarkably good representation of 15-year observed trends. If we wish to evaluate the model trends over just the single 15-year period, 1998–2012, we have only one observed realisation to compare with models that are generally not in phase with the internal variability of the observed system. There are several different approaches that have been used to bring the models in to phase with the real world. We discuss two popular methods and present an alternative third approach that we employ.

1. Climate forecasts. The desired solution to this problem is to do a climate forecast, not a projection. For a climate forecast one would test models that were initialised in 1998 and run with ensemble perturbations designed to track the major decadal instability modes (ENSO, PDO) in the ocean. Decadal forecasting efforts to date show some improvements due to initialisation, but are still in relative infancy [21].

2. Forced/restored projections. Another approach to ensure that the models have the same ENSO phase as the real world is to impose observed sea surface temperatures (sst) [22] or observed winds [23] in the tropical Pacific region of the model. This approach can be applied to ocean-only models. For coupled models this method requires that the model sst or wind stress field be restored rapidly to the observed fields. The forced/restored projection forces the models to lock phase with observed ENSO, though this partly specifies a temperature outcome consistent with observations. When applied in coupled models the results are also unphysical in that part of the model surface field is effectively specified and part is free. Other approaches include post-hoc

removal of a regression estimate of the ENSO phase contribution from model ensemble mean temperature [24].

3. Phase selected projections. Here we develop and apply an alternative simple, but natural (unforced), method to analyse model projections in phase with the real world. The models are not given any information about the observed state in this approach. The method takes advantage of the fact that, by chance, some of the CMIP5 model runs will be at least partially in phase at any given time with internal variability in the real world. One of the major processes driving variability in the rate of ocean heat uptake is ENSO [13]. The long term cycle of ENSO in switching between El Niño preferred periods (slower ocean heat uptake decades) and La Niña preferred periods (faster ocean heat uptake decades) will by chance line up in some of the CMIP5 models. To select this subset of models for any 15-year period, we calculate the 15-year trend in Niño3.4 index [25] in observations and in CMIP5 models and select only those models with a Niño3.4 trend within a tolerance window of $\pm .01 K/year$ of the observed Niño3.4 trend. This approach ensures that we select only models with a phasing of ENSO regime and ocean heat uptake largely in line with observations. In this case we select the subset of models in phase with observations from a reduced set of 18 CMIP5 models where Niño3.4 data were available [26] and for the period since 1950 when Niño3.4 indices are more reliable in observations.

Fig. 4 shows the running 15-year trends for observations in red (GISS top row; Cowtan & Way bottom row) and for the subset of CMIP5 models that fell within the Niño3.4 trend tolerance window in blue in the left column. The right column shows model 15-year trends for the subset of models in grey that were furthest from the observed Niño3.4 slope (least in phase with observations), where the subset of models is constrained so that it contains

7

the same number of models that fell within the best fit tolerance window. The size of plotting symbol for model trends is proportional to the number of models that fell within the Niño3.4 tolerance window. The solid lines are loess smoothed fits to the trend points. The loess smoothing is weighted by the number of models that contributed to each observation. The shaded areas surrounding each loess line represent approximate 95% confidence intervals.

The results in Fig. 4b,d show that selecting models with ENSO well out-of-phase with observations (grey curves) yields 15-year trends that can be very different from observations, and which are mostly not in accord with the observed trends. By contrast, selecting models with ENSO largely in-phase with observed ENSO (Fig. 4a,c; blue curves) yields 15-year trend estimates that are mostly in accord with observed trends. Further, even the 1998–2012 trend in the models lines up well with the observed trend when only models with a similar ENSO regime phase are used to estimate the trend. The selection of models in-phase with ENSO regimes thus removes much of the apparent discrepancy that occurs when model projections are compared with observational realisations irrespective of model natural internal variability phase.

The composite pattern of spatial 15-year trends in the selection of models in/out of phase with ENSO regime is shown for the 1998–2012 period in Fig. 5. The models in phase with ENSO (Fig. 5a) display a PDO-like pattern of cooling in the eastern Pacific, whereas the models least in phase (Fig. 5b) show more uniform El Niño-like warming in the Pacific. The set of models in-phase with ENSO produce a spatial trend pattern broadly consistent with observations (Fig. 5c) over the period. This result is in contrast to the full CMIP5 multi-model ensemble spatial trends, which display broad warming [27] and cannot reveal the PDO-like structure of the in-phase model trends because the internal modes of variability are averaged out.

8

To put the subset of phase-selected model projections in context with the full multi-model projection ensemble we compare them together in Fig. 6. The figure shows how the selection of models in phase with ENSO regimes changes the expected model warming trend relative to the full multi-model warming trend envelope. As noted, the grey all-model envelope moves up and down mostly due to variations in forcing (as the variations due to internal variability tend to average out across many model realisations), whereas in the observations the 15-year trends move up and down due to forcing and the cycles of natural variability. By selecting a subset of models in phase with the observed ENSO regime (blue curve) we introduce some variation due to ENSO in the model 15-year trends. Critically, this variation brings the expected model trends down toward the bottom part of the all-model trend envelope for recent 15-year trends. Models in phase with ENSO warm less rapidly during this period.

This method of phase aligning to select appropriate model trend estimates won't be perfect as the models contain errors in the forcing histories [28] and errors in the simulation of ENSO [26, 29] and other processes. Further, ENSO is not the only process generating natural variability on these time scales and so the method used here can only be approximate. Nonetheless, the phase-selection method provides a fairer and more appropriate basis to compare model projection trends over decadal-scale periods than use of the entire multi-model envelope. When the phase of natural variability is taken into account, the model 15-year warming trends in CMIP5 projections well-estimate the observed trends for all 15-year periods over the past half century.

9

Methods

Observational and model data. Three observational data sets are used to represent annual GMST: GISS [17]; HadCRUT4 [30]; and Cowtan & Way [16]. GMST is also taken from 38 CMIP5 models with multiple runs per model generating a multi-model ensemble with 82 runs. The CMIP5 ensemble runs span 1880–2012 with historical forcing for the period 1880–2005 and RCP8.5 [31] forcing from 2006–2012. Because the latter time period is so short, the choice of RCP scenario is not critical here. Results were repeated using RCP4.5 for the latter period and are very similar. A subset of 18 of the 38 CMIP5 models were available to us with sea surface temperature data to compute Niño3.4 [25] indices. The subset of models are indicated in bold in the set below. The reduced set of 18 models provides a similar model ensemble trend distribution to the larger model set. Niño3.4 values for observations were downloaded from NOAA at `http://www.esrl.noaa.gov/psd/data/climateindices/list/`.

The set of CMIP5 models used are: **ACCESS1-0**, **ACCESS1-3**, bcc-csm1-1, bcc-csm1-1-m, BNU-ESM, **CanESM2**, **CCSM4**, CESM1-BGC, CESM1-CAM5, CMCC-CM, CMCC-CMS, **CNRM-CM5**, **CSIRO-Mk3-6-0**, EC-EARTH, FGOALS-s2, FIO-ESM, **GFDL-CM3**, **GFDL-ESM2G**, GFDL-ESM2M, GISS-E2-H, GISS-E2-H-CC, **GISS-E2-R**, GISS-E2-R-CC, HadGEM2-AO, **HadGEM2-CC**, **HadGEM2-ES**, **INMCM4**, **IPSL-CM5A-LR**, **IPSL-CM5A-MR**, IPSL-CM5B-LR, MIROC-ESM, MIROC-ESM-CHEM, **MIROC5**, **MPI-ESM-LR**, MPI-ESM-MR, **MRI-CGCM3**, **NorESM1-M**, and NorESM1-ME.

Phase selected climate projection. The method of phase selected climate projection is based on selecting a feature of the climate that one wishes to align to the real world phase of the feature. Since we are interested in tracking

decadal fluctuations in GMST, we need to follow natural variations in ocean heat uptake on this time scale. These are driven by a range of processes in tropical and high latitude oceans [11, 32], but ENSO is one of the important ones and is at least broadly simulated by the coupled models.

Climate models cannot simulate every ENSO phase transition, but that is not critical here. To approximate the phase of ocean heat uptake in a 15-year period one may need to capture only the general sense of whether the models are El Niño or La Niña dominated over the period. To this end we use Niño3.4 as an index of ENSO phase and calculate the trend in Niño3.4 over the period to indicate the general tendency of the system toward regimes dominated by either ENSO state (similar to PDO phases). Models with a similar Niño3.4 trend to observations are then selected to represent the subset of models from the projection ensemble that just happen to have a similar response of ENSO over the 15-year period. This approach selects models with a statistical similarity that is related to the desired feature (ocean heat uptake). The approach does not guarantee any kind of dynamical consistency between models for each 15-year period.

The choice to fit 15-year trends to Niño3.4 is not the only way to estimate the phase of ocean heat uptake rates. We also tested a method based on trends of low pass Niño3.4 values to more directly mimic the PDO. Results were repeated where we detrended the Niño3.4 series in observations and models and compared the low pass Niño3.4 slopes to select in-phase models. The results for the low pass method are very similar to those for the direct method shown.

The pdf of observed and modelled 15 year GMST trends in Fig. 3 is based on 82 model runs over 1880–2012, but only 3 observational data sets, and thus contains different sample sizes. To control for sample size, the model pdfs

11

were recalculated with 500 random samples of model 15 year trends with the same sample size as the number of observed trends. This yields very similar model distributions to those in Fig. 3, indicating that any differences between the model and observed pdf are mostly not due to sample size differences.

References

[1] Bray D. and H. von Storch. Prediction or projection: The nomenclature of climate science. *Sci. Commun.*, 30(4):534–543, 2009.

[2] M. Peña and E. Kalnay. Separating fast and slow modes in coupled chaotic systems. *Nonlinear Proc. Geoph.*, 11:319–327, 2004.

[3] S.-C. Yang, C. Keppenne, M. Rienecker, and E. Kalnay. Application of coupled bred vectors to seasonal-to-interannual forecasting and ocean data assimilation. *J. Clim.*, 22(11):2850–2870, 2009.

[4] T. O'Kane, R. Matear, M. Chamberlain, and P. Oke. ENSO regimes and the late 1970s climate shift: The role of synoptic weather and south Pacific ocean spiciness. *J. Comp. Phys.*, 2014. http://dx.doi.org/10.1016/j.jcp.2013.10.058.

[5] K. Taylor, R. Stouffer, and G. Meehl. An overview of CMIP5 and the experiment design. *Bull. Amer. Met. Soc.*, 93(4):485–498, 2012.

[6] G. Meehl and Coauthors. Decadal prediction: can it be skillful? *Bull. Amer. Met. Soc.*, 90(10):1467–1485, 2009.

[7] T. Stocker et al. *Technical summary. 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*, pages 33–115. Cambridge Univ. Press, Cambridge, UK, 2013. 1535pp.

[8] J. Tollefson. The case of the missing heat. *Nature*, 505:276–278, 2014.

[9] E.N. Lorenz. Climatic determinism. *Meteor. Monographs*, 8(30):1–3, 1968.

[10] I. James and P. James. Ultra-low-frequency variability in a simple atmospheric circulation model. *Nature*, 342:53–55, 1989.

[11] T. O'Kane, R. Matear, M. Chamberlain, J. Risbey, B. Sloyan, and I. Horenko. Decadal variability in an OGCM Southern Ocean: Intrinsic modes, forced modes and metastable states. *Ocean Model.*, 69:1–21, 2013.

[12] S. H. Schneider and S. Thompson. Atmospheric co_2 and climate: Importance of the transient response. *J. Geophys. Res.*, 86:3135–3147, 1981.

[13] G. Meehl, J. Arblaster, J. Fasullo, A. Hu, and K. Trenberth. Model-based evidence of deep-ocean heat uptake during surface-temperature hiatus periods. *Nature Climate Change*, 1:360–364, 2011.

[14] N. Mantua and S. Hare. The Pacific Decadal Oscillation. *J. Oceanogr.*, 58(1):35–44, 2002.

[15] G. Foster and S. Rahmstorf. Global temperature evolution 1979–2010. *Environ. Res. Lett.*, 6(4):1–8, 2011.

[16] K. Cowtan and R. Way. Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends. *Q. J. R. Meteor. Soc.*, 2014. DOI: 10.1002/qj.2297.

[17] J. Hansen, R. Ruedy, M. Sato, and K. Lo. Global surface temperature change. *Rev. Geophys.*, 48(4):1–29, 2010. DOI: 10.1029/2010RG000345.

[18] J. Fyfe, N. Gillett, and F. Zwiers. Overestimated global warming over the past 20 years. *Nature Climate Change*, 3:767–769, 2013. doi:10.1038/nclimate1972.

[19] Z. Toth and E. Kalnay. Ensemble forecasting at NCEP and the breeding method. *Mon. Weather Rev.*, 125(12):3297–3319, 1997.

[20] T. Ault, C. Deser, M. Newman, and J. Emile-Geay. Characterizing decadal to centennial variability in the equatorial Pacific during the last millennium. *Geophys. Res. Lett.*, 40(13):3450–3456, 2013.

[21] G. Meehl and Coauthors. Decadal climate prediction: An update from the trenches. *Bull. Amer. Met. Soc.*, 2014. doi: http://dx.doi.org/10.1175/BAMS-D-12-00241.1.

[22] Y. Kosaka and S.-P. Xie. Recent global-warming hiatus tied to equatorial Pacific surface cooling. *Nature*, 501:403–407, 2013.

[23] M. England, S. McGregor, P. Spence, G. Meehl, A. Timmermann, W. Cai, A. Sen Gupta, M. McPhaden, A. Purich, and A. Santoso. Recent intensification of wind-driven circulation in the Pacific and the ongoing warming hiatus. *Nature Climate Change*, 2014. online: doi:10.1038/nclimate2106.

[24] G. Schmidt, D. Shindell, and K. Tsigaridis. Reconciling warming trends. *Nature Geoscience*, 7:158–160, 2014.

[25] K. Trenberth. The definition of El Niño. *Bull. Amer. Met. Soc.*, 78(12):2771–2777, 1997.

[26] J.N. Brown, C. Langlais, and C. Maes. Zonal structure and variability of the Western Pacific dynamic warm pool edge in CMIP5. *Clim. Dynam.*, 2014. doi: 10.1007/s00382-013-1931-5.

[27] J. Fyfe and N. Gillett. Recent observed and simulated warming. *Nature Climate Change*, 4:150–151, 2014.

[28] R. Allen, J. Norris, and M. Wild. Evaluation of multidecadal variability in CMIP5 surface solar radiation and inferred underestimation of aerosol direct effects over Europe, China, Japan, and India. *J. Geophys. Res.*, 118(12):6311–6336, 2013.

[29] E. Guilyardi, A. Wittenberg, A. Fedorov, M. Collins, C. Wang, A. Capotondi, G. van Oldenborgh, and T. Stockdale. Understanding El Niño in ocean-atmosphere general circulation models: progress and challenges. *Bull. Amer. Met. Soc.*, 90(3):325–340, 2009.

[30] C. Morice, J. Kennedy, N. Rayner, and P. Jones. Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The HadCRUT4 data set. *J. Geophys. Res.*, 117:1–22, 2012.

[31] Moss and coauthors. The next generation of scenarios for climate change research and assessment. *Nature*, 463:747–756, 2010.

[32] B. Kirtman et al. *Near-term climate change: projections and predictability. 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*, chapter 11, pages 953–1028. Cambridge Univ. Press, Cambridge, UK, 2013. 1535pp.

Acknowledgements

Author contributions

J.S.R. and S.L. conceived the study and initial experimental design. All authors contributed to experiment design and interpretation. S.L. provided analysis of models and observations. C.L. and D.P.M. analysed Nino3.4 in models. J.S.R. wrote the paper and all authors edited the text.

Additional information

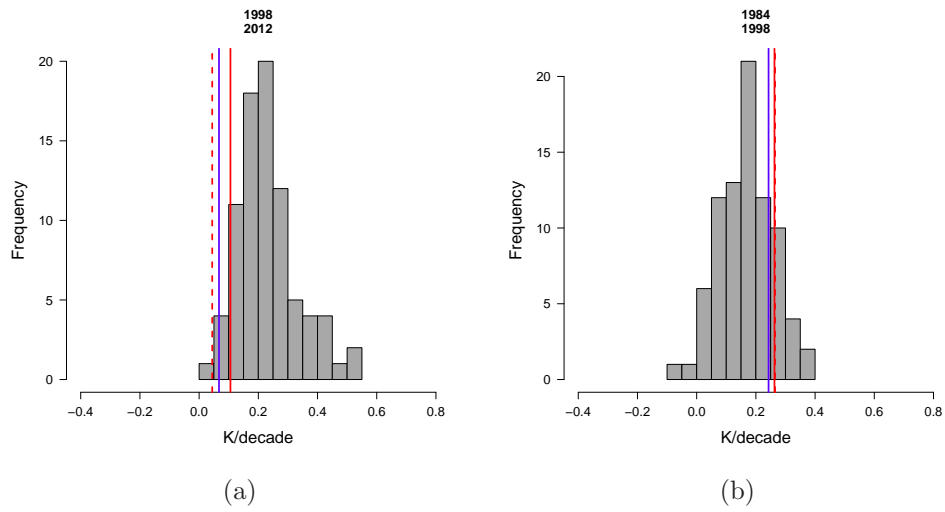The authors declare no competing financial interests.

Figure 1: Ensemble trends and observational trends. Histogram of trends for a) 1998–2012 and b) 1984–1998 for the CMIP5 multi-model ensemble. The thin vertical lines show the observed trend for these periods from HadCRUT4 (dashed red), GISS (blue), and Cowtan & Way (red).
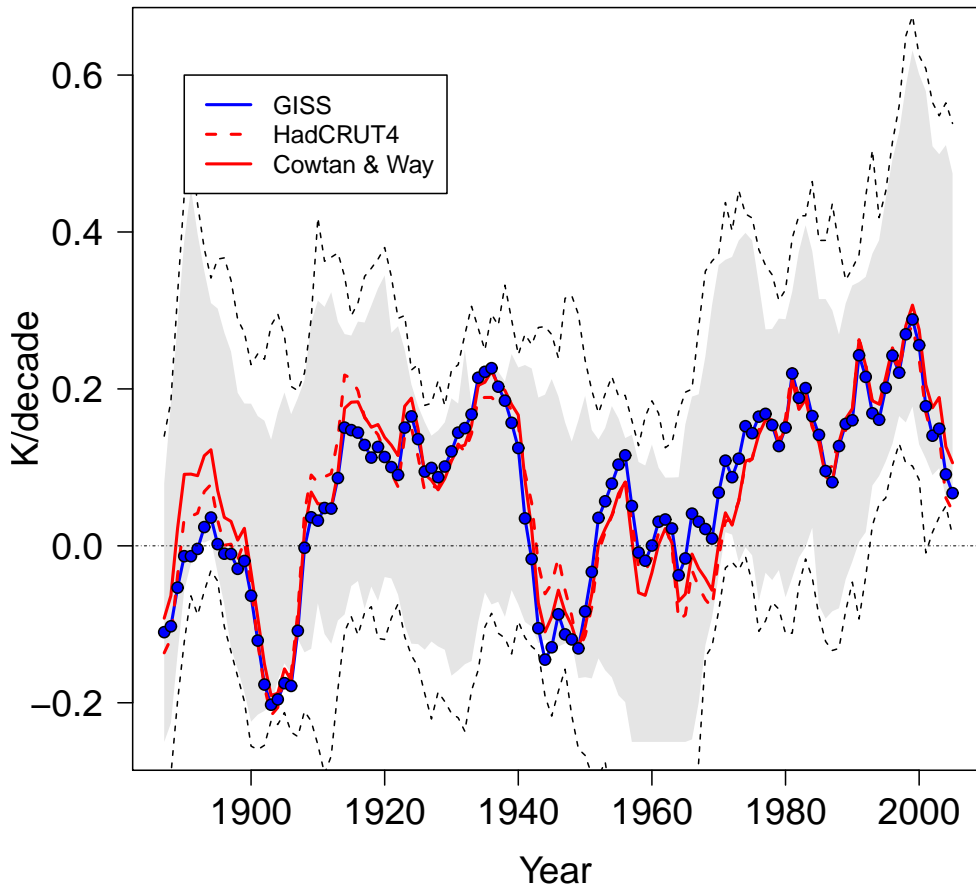
Figure 2: Running 15-year trends. The lines show the observed estimate of 15-year trends plotted at the central year of the 15-year period for GISS (blue), HadCRUT4 (dashed red) and Cowtan and Way (red). The grey envelope spans the 2.5 percentile and 97.5 percentile estimate of the distribution of CMIP5 ensemble trends over the same 15-year periods. The black dashed lines encompass the full CMIP5 model ensemble.
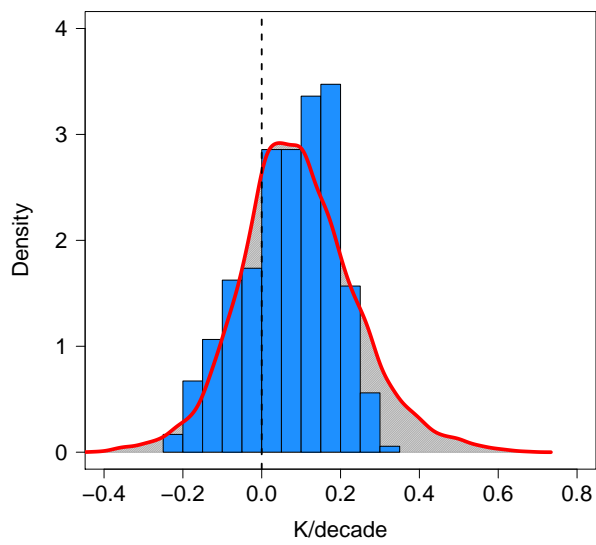
Figure 3: Trend histograms. Histogram of the 15-year trend estimates for all 15-year periods plotted in Fig. 2 for all three observational datasets. The pdf over the histogram is the kernel density estimate of the distribution of all CMIP5 model trends for all 15-year periods.
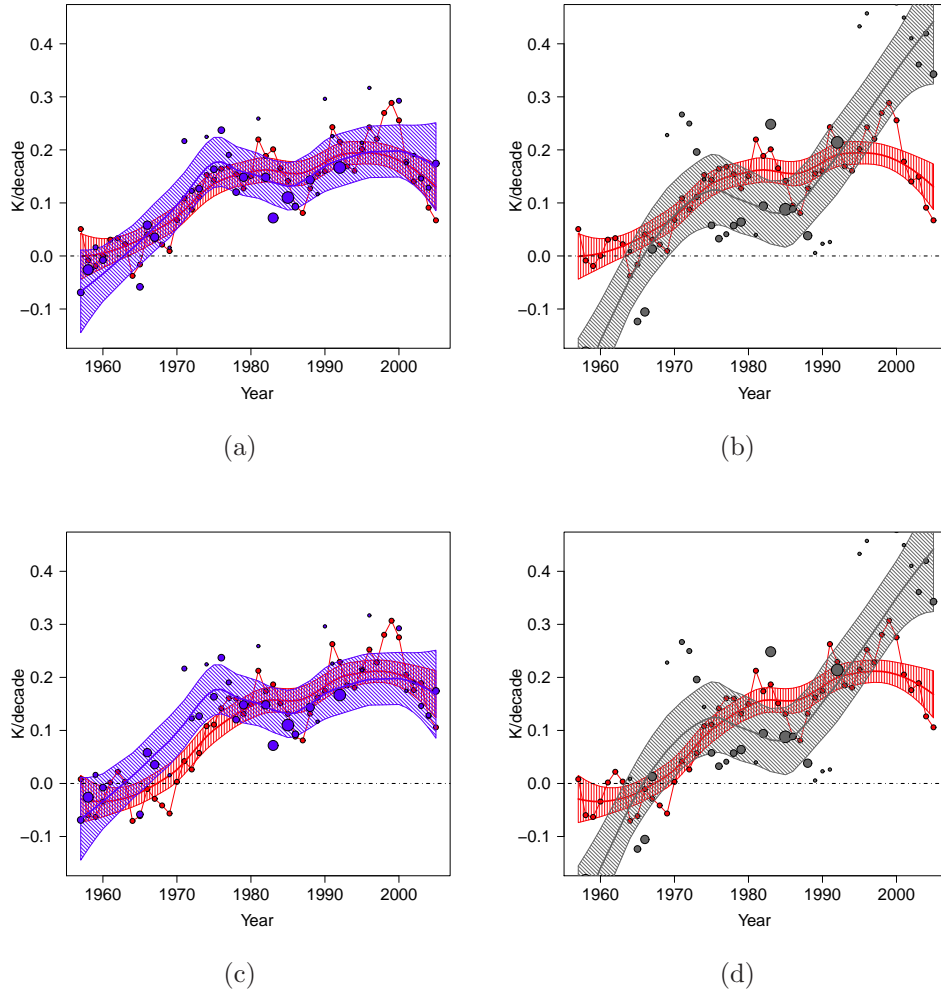
Figure 4: Model phase-selected trends and observed trends. The red dots on the thin red line correspond to the 15-year observed trends for each 15-year period in a), b) GISS and c), d) Cowtan and Way. The red envelope shows a 2.5–97.5 percentile loess smoothed fit to the observed 15-year trends. The blue dots (a, c) show the 15-year average trends from only those CMIP5 runs in each 15-year period where the model Niño3.4 trend is close to the observed Niño3.4 trend. The size of the blue dot is proportional to the number of models selected. If fewer than two models are selected in a period, they are not included in the plot. The blue envelope is a 2.5–97.5 percentile loess smoothed fit to the model 15-year trends weighted by the number of models at each point. Panels b) and d) contain the same observed trends in red for GISS and Cowtan & Way. The grey dots show the average 15-year trends for only the models with the worst correspondence to the observed Niño3.4 trend. Results for HadCRUT4 (not shown) are broadly similar to those of Cowtan & Way.
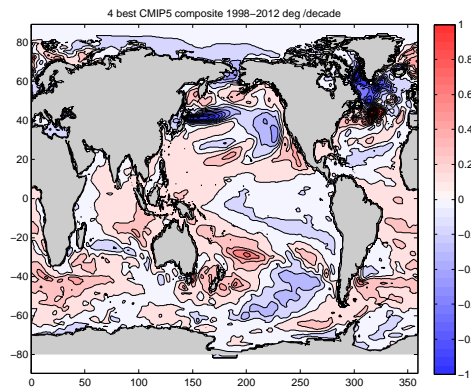
(a)
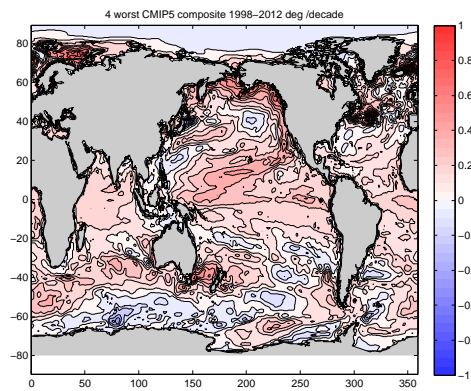
(b)
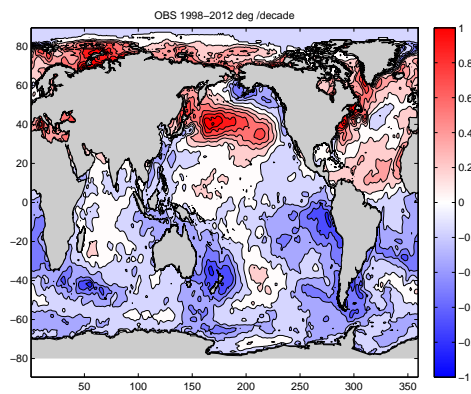
(c)

Figure 5: Composite sst spatial trends. Composite of 15-year sst trends over 1998–2012 for a) CMIP5 models best in phase with Niño3.4 trends, b) CMIP5 models least well in phase with Niño3.4 trends, and c) HadISST observations.
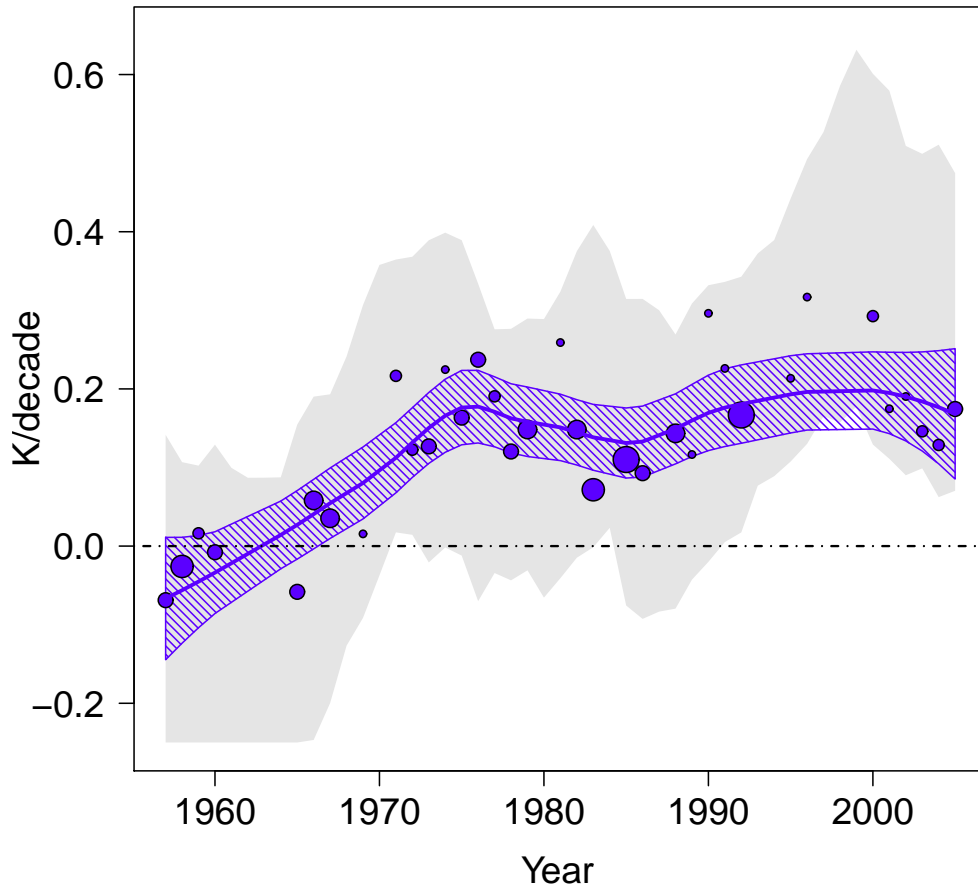
Figure 6: Model ensemble and phase-selected trends. The grey envelope represents the 2.5–97.5 percentile estimates of 15-year GMST trends from all models runs. The blue dots are the 15-year trend estimates from the models best in phase with Niño3.4 trends shown in Fig. 4a,c. The blue envelope is a 2.5–97.5 percentile loess smoothed fit to the model 15-year trends weighted by the number of models at each point.