



# Estimating trends in the global mean temperature record

Andrew Poppick<sup>1</sup>, Elisabeth J. Moyer<sup>2</sup>, and Michael L. Stein<sup>3</sup>

<sup>1</sup>Department of Mathematics and Statistics, Carleton College, Northfield, MN, USA

<sup>2</sup>Department of the Geophysical Sciences, University of Chicago, Chicago, IL, USA

<sup>3</sup>Department of Statistics, University of Chicago, Chicago, IL, USA

*Correspondence to:* Andrew Poppick (apoppick@carleton.edu)

Received: 29 December 2016 – Revised: 1 May 2017 – Accepted: 8 May 2017 – Published: 9 June 2017

**Abstract.** Given uncertainties in physical theory and numerical climate simulations, the historical temperature record is often used as a source of empirical information about climate change. Many historical trend analyses appear to de-emphasize physical and statistical assumptions: examples include regression models that treat time rather than radiative forcing as the relevant covariate, and time series methods that account for internal variability in nonparametric rather than parametric ways. However, given a limited data record and the presence of internal variability, estimating radiatively forced temperature trends in the historical record necessarily requires some assumptions. Ostensibly empirical methods can also involve an inherent conflict in assumptions: they require data records that are short enough for naive trend models to be applicable, but long enough for long-timescale internal variability to be accounted for. In the context of global mean temperatures, empirical methods that appear to de-emphasize assumptions can therefore produce misleading inferences, because the trend over the twentieth century is complex and the scale of temporal correlation is long relative to the length of the data record. We illustrate here how a simple but physically motivated trend model can provide better-fitting and more broadly applicable trend estimates and can allow for a wider array of questions to be addressed. In particular, the model allows one to distinguish, within a single statistical framework, between uncertainties in the shorter-term vs. longer-term response to radiative forcing, with implications not only on historical trends but also on uncertainties in future projections. We also investigate the consequence on inferred uncertainties of the choice of a statistical description of internal variability. While nonparametric methods may seem to avoid making explicit assumptions, we demonstrate how even misspecified parametric statistical methods, if attuned to the important characteristics of internal variability, can result in more accurate uncertainty statements about trends.

## 1 Introduction

The physical basis of climate change is understood through a combination of theory, numerical simulations and analyses of historical data. Climate change is driven by radiative forcing, a change in net radiation (downwelling minus upwelling, often specified at the top of atmosphere) resulting from an imposed perturbation of a climate in equilibrium, for example by increasing the atmospheric concentration of a greenhouse gas. The Earth's response to forcing is complex and not fully understood, in part due to physical uncertainties in important feedbacks such as cloud responses (see the assessment reports of the Intergovernmental Panel on Climate Change (IPCC), e.g., IPCC, 2013).

Given the physical uncertainties inherent in all climate simulations, the observed temperature record since the late nineteenth century is often used as a source of empirical information about the Earth's systematic response to forcing. (Figure 1 shows one estimate of annually averaged global mean surface temperatures from the past 136 years, along with estimates of radiative forcings from various constituents during that period, with the data sources described in Sect. 2.) Analysis of the observed temperature record is complicated, however, by the short available record of direct measurements, by uncertainties in the historical radiative forcings themselves, and by the internal temperature variability that exists even in the absence of forcing. Statistical methods are therefore required to quantify the information in the histori-

cal record about the response to forcing: given the data, what do we know about how global temperatures have warmed in response to forcing, how much warming can we expect in plausible future forcing scenarios, and how do we expect uncertainties to change as we continue to observe the Earth's temperatures?

It can be helpful to divide approaches to using the observed temperature record to understand aspects of mean climate change into two categories. One common approach involves assuming a physical model of the system. (Here the term “model” encompasses anything from a simple energy balance model to a very complicated atmosphere–ocean general circulation model (GCM).) Analysis then may involve estimating statistical parameters in order to best fit the observed record. Estimated parameters have statistical uncertainty because of the finite observational record and the internal variability inherent in the climate system, even were the model a perfect representation of reality. Analyses of observed temperatures using simple or moderately complex physical models include Wigley et al. (1997), Gregory et al. (2002), Knutti et al. (2002), Forest et al. (2002), Forest et al. (2006), Gregory and Forster (2008), Padilla et al. (2011), Aldrin et al. (2012), Otto et al. (2013), Rypdal and Rypdal (2014), Masters (2014), Lewis and Curry (2015), Zeng and Geil (2016) and many others. Analyses using GCMs are summarized in IPCC (2013) chap. 10, and references therein.

Another category of approaches involves analyses that are more empirical and appear to de-emphasize assumptions about the underlying physics generating the observed temperatures. Many studies use regression models that treat time rather than radiative forcing as the covariate. This practice is often used, for example, to test for significant warming (e.g., Bloomfield, 1992; Smith and Chen, 1996; Løvstetten and Rypdal, 2016) or for changes in warming trends (e.g., Foster and Rahmstorf, 2011; Cahill et al., 2015; Rajaratnam et al., 2015); in general, regressions in time are widespread in the literature (see also IPCC, 2013, especially chap. 2 Box 2.2, and references therein).

In both categories, analyses require characterizing internal variability for the purpose of quantifying uncertainty, a task that also involves assumptions. A typical approach is to assume a statistical model for the dependence structure of the noise, such as assuming an autoregressive moving average (ARMA) noise model with a small number of parameters, which is fit to the residuals of whatever trend model is being used. Some authors, however, argue for nonparametric (resampling or subsampling) methods for time series rather than parametric approaches (e.g., Gluhovsky, 2011; Rajaratnam et al., 2015)<sup>1</sup>. The argument is again that these approaches

are advantageous because they are ostensibly objective and require fewer assumptions.

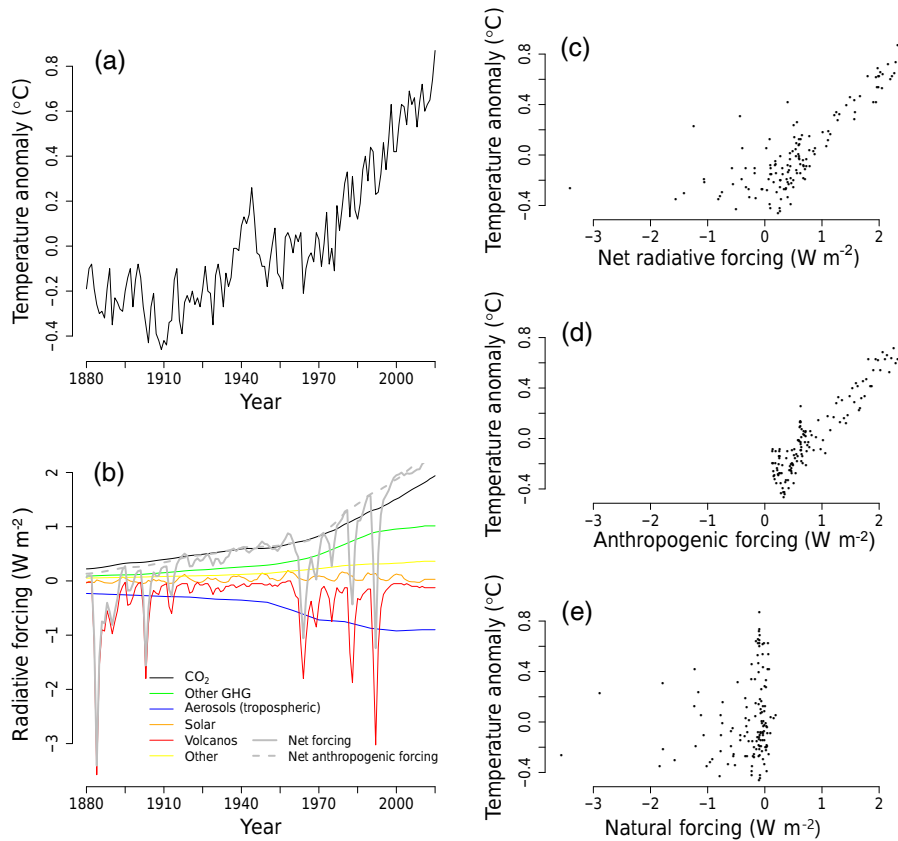
However, methods that de-emphasize assumptions, be they physical or statistical, can be problematic in the climate setting. While regressions in time are simple to apply and do not appear to make explicit assumptions about how temperatures should respond to forcing, these models both limit what can be learned from the data and can result in misleading inferences. Regressions in time are sensitive to arbitrary choices (such as the start and end date of the data analyzed), cannot be expected to apply over even modestly long time frames and cannot in general reliably separate forced trends from internal variability. Furthermore, in accounting for internal variability, nonparametric methods for time series often require long data records to work well, and can be seriously uncalibrated in data-limited settings with strong temporal correlation, such as the setting we are discussing.

In the following, we illustrate two primary points. First, we show that targeted parametric mean models that incorporate even limited physical information can provide better-fitting, more interpretable, and more illuminating descriptions of the systematic response of interest compared to approaches that de-emphasize assumptions. Second, we show that parametric models for residual (i.e., internal) variation can provide for safer and more accurate uncertainty quantifications in this setting than do approaches that de-emphasize assumptions, even if the parametric model is misspecified, as long as the parametric modeling is done with particular attention towards the representation of low-frequency internal variability. We believe that the analysis that we present is informative, even if not maximally so, and we attempt to highlight both complications with our analysis as well as important sources of information about global warming that are ignored in our approach. Parts of our analysis share similarities with others listed above, especially with Rypdal and Rypdal (2014) and Zeng and Geil (2016). In distinction to previous papers, here we are primarily interested in contrasting what can be learned using physically motivated models vs. those that de-emphasize assumptions, as well as in emphasizing the role that accounting for internal variability plays in inferring uncertainties in mean trends. Our goals are to indicate directions in which statisticians can incorporate explicit modeling to positive effect and to highlight what we view are some of the important sources of uncertainty and information in this problem.

This article is organized as follows. In Sect. 2, we introduce the data sources used in our analysis. In Sect. 3, we provide some background on modeling the historical global mean temperature record and contrast a simple, minimally physically informed model with a more empirical approach.

<sup>1</sup>In this work we consider a parametric method to be any that makes explicit assumptions about the functional form of the probability distribution that the data come from, described with a finite number of statistical parameters. In particular, for the purposes of

this work, we consider the class of low-order ARMA models to be a parametric class. We consider a nonparametric method to be one that attempts to make fewer distributional assumptions about the data and does not involve a parametrized statistical model.



**Figure 1.** (a) Estimates of annually averaged global mean surface temperature anomalies (relative to a base period of 1951–1980) from the years 1880 to 2015. (b) Estimates of (top-of-atmosphere) effective radiative forcings from different constituents over this time period (the “other” category includes O<sub>3</sub>, H<sub>2</sub>O, black carbon, contrails, and land use changes). Right, temperature anomalies vs. net radiative forcings (c), vs. anthropogenic forcings (d), and vs. natural forcings (e). Data sources are described in Sect. 2. Despite the correlation in the plots of temperature vs. radiative forcing, temperatures will depend on the full past trajectory of radiative forcings in a potentially complex way, as we discuss in Sect. 3.

In Sect. 4, we highlight insights that can be gained from the more informed approach, with an emphasis on probing different aspects of uncertainty in trends. In Sect. 5, through synthetic simulations, we compare the performance of various parametric and nonparametric methods of uncertainty quantification in the presence of temporal correlation in settings similar to that of the historical temperature record. In Sect. 6, we give some concluding remarks.

## 2 Background and data

This analysis requires estimates of historical global mean temperatures and radiative forcings. To the extent that we are interested in how temperatures may evolve in the future (and how uncertainty in the response to radiative forcing evolves as more data are observed), we also need radiative forcings associated with a plausible future scenario.

### 2.1 Temperature

We use the Land–Ocean Temperature Index from the NASA Goddard Institute for Space Studies (GISS) (Hansen et al., 2010) in our primary analysis. The index combines land and sea surface temperature measurements to estimate annual average global mean surface temperature anomalies (relative to a base period from 1951 to 1980), extending from the year 1880 to the present (comprising 136 years in total). Any dataset of global mean temperature anomalies represents an estimate of that quantity and is subject to some uncertainty. Sources of uncertainty include the spatial coverage of the network of measurements, interpolation schemes used to estimate temperatures at unobserved locations, methods used to incorporate different sources of data (e.g., land- vs. satellite- vs. surface buoy- vs. ship-based data), and instrumental errors. Sources of uncertainty in the GISS dataset are discussed in Hansen et al. (2010). NASA GISS has made some attempt to provide pointwise uncertainty estimates for their data (e.g., Fig. 9a of Hansen et al., 2010), but it is important to realize

that errors will be correlated in time. That said, uncertainties in the global mean temperature record are relatively small compared to the changes in temperatures observed over the 20th century.

To evaluate the effects of uncertainty in the temperature record, we repeat a small portion of our analysis using the HadCRUT4 global annual temperature ensemble (Morice et al., 2012), designed for that purpose (Sect. 4.5). (Some sources of uncertainty are, however, common to both data sources; a brief comparison of these two sources, as well as one from the National Oceanographic and Atmospheric Administration (NOAA), can be found in IPCC, 2013, Sect. 2.4.3.)

## 2.2 Radiative forcing

The primary driver of climate change during the historical period is changing atmospheric CO<sub>2</sub> concentrations; radiative forcing due to these changes scales approximately with the logarithm of the ratio of the CO<sub>2</sub> concentration at any given time to its preindustrial level (e.g., Arrhenius, 1896). However, variations in other agents also have non-negligible forcing effects that must be taken into account to interpret the historical record (aerosols from human sources or volcanoes, other greenhouse gases, etc.). In this work, we aggregate the effects of different forcing agents by using their estimated *effective radiative forcing*, the radiative imbalance after rapid atmospheric adjustments. These adjustments are intended to partially compensate for differing efficacies of forcing agents. In practice, effective forcings are often treated as though they may be combined additively. For effective radiative forcings from 1750 to 2011, we use the estimates in IPCC (2013) Table AII.1.2 (Fig. 1 only shows the forcings after 1880, but the full available record is used in our analysis). From 2011 to 2015, we use the global CO<sub>2</sub> concentrations from NOAA and treat radiative forcing from other sources as constant during this period.

While historical concentrations of CO<sub>2</sub> are relatively well known, since CO<sub>2</sub> is well-mixed and long lived, those of other forcing agents, such as tropospheric aerosols, are more difficult to measure because they are spatially heterogeneous and short lived. Uncertainties in forcings associated with tropospheric aerosols are important because aerosol effects can be negatively confounded with greenhouse gas effects (see Fig. 1, bottom left). Generally, uncertainties vary by constituent, as does the extent to which the estimates are derived from model output vs. observations; see IPCC (2013) chap. 8. The focus of this paper is on the information content of the observed temperature record assuming known forcings, but we make a limited attempt to discuss the effect of uncertainty in the forcings (Sect. 4.5).

For a plausible future radiative forcing scenario, we use the extended Representative Concentration Pathway scenario 8.5 (RCP8.5) (Riahi et al., 2011; Meinshausen et al., 2011), where the change in radiative forcing from the preindustrial

level is 8.5 W m<sup>-2</sup> by the year 2100 and levels off at around 12 W m<sup>-2</sup> in the year 2250. In our simulations, we slightly increase (by about 0.07 W m<sup>-2</sup>) the radiative forcings from the RCP8.5 scenario in the 21st century to match what we take to be the historical value in 2015, and we assume that natural forcings remain constant after 2015.

## 2.3 Ocean heat uptake

The analysis here focuses on information provided solely by the global mean temperature record and assumed known forcings. We do not use additional potential sources of empirical information, including estimates of ocean heat uptake (discussed in, e.g., Forest et al., 2002 and Knutti et al., 2002). Many empirical analyses of the historical record do incorporate information about ocean heat content (e.g., Gregory et al., 2002; Otto et al., 2013; Masters, 2014, and Lewis and Curry, 2015). We compare results with these studies in Sect. A1 in the Appendix.

## 3 Modeling trends in the observed global mean temperature record

Evaluating the systematic response of global mean surface temperatures to forcing is complicated by the long timescales for warming of the Earth system. Because the Earth's climate takes time to equilibrate, the near-term (transient or centennial-scale<sup>2</sup>) climate response will be less than the long-term (equilibrium or millennial-scale) response. The evaluation is also complicated by the fact that historical radiative forcings are not constant but rather evolve in time (e.g., atmospheric CO<sub>2</sub> increases). The physical lags in response imply that the Earth's global mean temperature at any given time depends on the past trajectory of radiative forcings (because the climate does not instantly equilibrate to the present forcing).

A common framework is to decompose observed temperatures into two components: a systematic component changing in response to past forcings and a residual component representing sources of internal variability. That is, for global mean temperatures  $T(t)$  at time  $t$ ,

$$T(t)|\{F(t'); t' \leq t\} = f(F(t'); t' \leq t) + \epsilon(t), \quad (1)$$

where  $f$  is an unknown functional of  $\{F(t'); t' \leq t\}$ , the collection of past radiative forcings associated with each forcing agent, and  $\epsilon(t)$  is a residual process that has mean zero and is correlated in time. Here we emphasize in our notation that the systematic response of interest is the response *conditional on a given forcing trajectory*. The problem, then, is how to

<sup>2</sup>Since the term *transient climate response* has a specific definition in the literature (see the Appendix for a discussion), we use the term *centennial-scale response* to describe the systematic response of temperatures to forcings on the mixing timescale of the mixed layer of the ocean but not of the deep ocean.

estimate the systematic response  $f$  using the historical temperatures and forcings.

### 3.1 Regression models in time

Estimating a model like Eq. (1) is intractable without additional assumptions. As discussed above, one approach is to resort to physical models. But if instead a more empirical analysis of the observed data is desired, it is common practice to consider a surrogate regression on time itself, as stated in Sect. 1. The implicit assumption here is that, when viewed as a function of time, the systematic response to historical radiative forcings is approximately linear in time, at least over the considered time frame:

$$f(F(t'); t' \leq t) \approx \alpha + \beta t, \quad t \in [t_0, t_1], \quad (2)$$

where  $\alpha$  and  $\beta$  are unknown parameters and the trend is considered over the interval  $[t_0, t_1]$ . The linear time trend approach arguably involves assumptions about the forcing history in addition to the systematic response  $f$ : that the forcing itself evolves approximately linearly in time (else the approximation would not be appropriate).

The linear time trend model is widely used, and the general sense is that such a model offers a way of testing for statistically significant changes in mean temperature without having to make physical assumptions and without having to believe that the true forced response is linear in time (e.g., Bloomfield, 1992 and Løvsetten and Rypdal, 2016). The IPCC, accounting for the apparent appeal of the linear time trend model, states that it “is relatively simple, transparent and easily comprehended, and is frequently used in the published research assessed here,” (IPCC, 2013 chap. 2, Box 2.2), but suggests that linearity in time can at best be viewed as an approximation expected to hold over a relatively short period of time. It is well understood that neither the observed temperature record nor the forcing history appear to evolve linearly over the full range of the data record (Fig. 1, left), and most users of the linear time trend approach confine their analysis to only the past few decades.

While the time trend model may be routine to apply, appear objective, and provide a good fit to the data, its use can be precarious. A proper accounting of uncertainty in mean temperature changes relies on distinguishing internal variability from systematic responses. The time trend model is problematic in this respect. If the chosen time interval is short, it can be difficult to distinguish between trends and sources of internal variability that are correlated over longer timescales than the chosen interval (implicitly recognized in, e.g., Easterling and Wehner (2009) and Santer et al. (2011) and explicitly discussed in a broader context in Wunsch, 1999). If the chosen interval is long and the systematic trend is actually nonlinear in time, then assuming a linear model in time will shift part of the systematic response to the residual process and can therefore give the impression of excessive in-

ternal variability over long timescales (and hence excessive uncertainty in trends).

Because the time trend model cannot be applied over long time intervals for arbitrary forcing scenarios, it also does not have a property that may be considered important for making inferences: that we can learn more about the systematic trend of interest by collecting more observations. There will be only a finite amount of information about the systematic response within the interval  $[t_0, t_1]$  (this because sources of internal variability will be positively correlated in time). While this on its own does not invalidate the use of such a model over some narrow time frame, it does mean that what can be learned from the linear time trend model is necessarily limited. More broadly, since the linear time trend model does not map to a physical understanding of the relationship between radiative forcing and global mean temperatures, either during the time interval  $[t_0, t_1]$  or extending beyond it, the questions that can be asked with this model are narrow.

Some argue that many of these problems may be overcome by using a model that is nonlinear in time, such as a spline or other nonparametric regression method. (The IPCC, for example, appears to view nonparametric extensions as more generically appropriate than the linear model.) Nonparametric regressions in time will appear to provide an even better fit to the data than the linear trend model, but many of the above arguments carry over to this setting. Such models have limited interpretational value or ability to capture systematic (non-internal) trends, since they cannot generically be expected to distinguish between the systematic trends of interest and other, internal sources of long-timescale variation in the data. Collectively, these arguments suggest that it is advisable to seek better motivated models if one is interested in understanding the systematic response of global temperatures to forcing.

### 3.2 A simple, physically based model for the centennial-scale response to forcing

A typical approach is to use more complex models, including full GCMs, to explain the systematic response of interest. (Model output is also used in concert with observations in the context of “detection and attribution” studies; see, e.g., chap. 10 of IPCC, 2013.) Some may object to this approach, however, out of a worry that the climate model has already been tuned to match the observed historical temperature trend or is otherwise conditioned on past temperature observations (Knutti, 2008; Huybers, 2010; Mauritsen et al., 2012). There is therefore value in a compromise approach between the linear time trend model and very complex numerical simulations. In this work, we discuss a statistical model that is easy to apply but that encodes some physical intuition for the problem that makes the model interpretable and hopefully applicable over longer time periods. The goal is to show that even simple models incorporating limited physical information can provide more insight about temperature

trends and their uncertainties given the observed data than can regression models in time.

A commonly used, very simplified physical model for the response to an instantaneous change in radiative forcing is that temperatures approach their new equilibrium in exponential decay. That is, writing  $F_{\text{inst}}(t)$  for a step function that changes at time  $t = 0$ ,

$$f(F_{\text{inst}}(t'); t' < t) \approx \mu_0 + \lambda(1 - \rho^{t'})\mathbf{1}\{t \geq 0\}, \quad (3)$$

where  $\lambda$  is the change in equilibrium temperature,  $\mu_0$  is the mean temperature in the baseline state, and  $\rho$  controls the rate at which the changes in temperatures approach  $\lambda$ , taking values between zero (instantaneous response time) and 1 (infinite response time). Equation (3) represents the response function for a linear model of temperature change, so is a natural first approximation for the evolution of temperature in the case of small perturbations from steady state (e.g., MacKay and Ko, 1997). In particular, Eq. (3) can be interpreted as the solution to a simple energy balance model that makes two assumptions: first, that the equilibrium temperature change is linear in the forcing (i.e., that the *linear forcing feedback model* holds) and, second, that the rate of warming is approximately proportional to the heat uptake. The model is overly simplified because the Earth shows responses at multiple timescales (e.g., Held et al., 2010; Olivié et al., 2012; Geoffroy et al., 2013b) but can provide a reasonable approximation for the response on timescales shorter than those associated with full equilibration. In any case, when convolved with a time-varying forcing trajectory, the resulting model for the systematic response is an infinite distributed lag model in the forcing trajectory with weights decaying exponentially (e.g., Caldeira and Myhrvold, 2013; Castruccio et al., 2014). Models based off of Eq. (3) have previously been considered for analyses of the observed global mean temperature record (e.g., Rypdal and Rypdal, 2014; Zeng and Geil, 2016).

We also use a model based off of Eq. (3) for the systematic temperature response in the observed data:

$$f(F(t'); t' \leq t) \approx \mu_0 + \lambda_A h\left(\rho_A, \frac{F_A(t')}{F_{2\times}}; t' \leq t\right) + \lambda_N h\left(\rho_N, \frac{F_N(t')}{F_{2\times}}; t' \leq t\right), \quad (4)$$

where

$$h(\rho, x(t'); t' \leq t) = (1 - \rho) \sum_{k=0}^{\infty} \rho^k x(t - k).$$

In model Eq. (4),  $\lambda_A$  and  $\lambda_N$  represent “sensitivities” to anthropogenic and natural forcings,  $F_A$  and  $F_N$ , respectively, and have units of degrees Celsius temperature change per forcing change  $F_{2\times}$  (the forcing associated with a doubling of atmospheric  $\text{CO}_2$ , approximately  $3.7 \text{ W m}^{-2}$ ). The param-


eter  $\lambda_A$  is similar to the *equilibrium climate sensitivity*<sup>3</sup>, but will be estimated as somewhat lower than that quantity, in part because the proposed model contains only a single timescale of response to anthropogenic forcing (see the Appendix). Caldeira and Myhrvold (2013) and Castruccio et al. (2014) used multiple timescales in modeling longer series from GCM output, but we cannot distinguish these with only 136 years of data and given a smooth past trajectory of anthropogenic forcings. Response timescales to anthropogenic and natural forcings are set by the parameters  $\rho_A$  and  $\rho_N$  (taking values between zero and one). Model (4) should approximate temperature trends reasonably well up to centennial timescales, but not at the millennial timescales at which the deep ocean mixes.

Our approach differs from those in Rypdal and Rypdal (2014) and Zeng and Geil (2016) in some important ways. Zeng and Geil (2016) assumed that net radiative forcing is a constant multiple of  $\text{CO}_2$  forcing, treat the  $\rho$  parameter as known, and do not account for uncertainty due to internal variations that are correlated in time. Rypdal and Rypdal (2014) suggested replacing the standard exponential decay response used in model (4) with a power-law decay response. The basis for this suggestion is an implicit assumption that we do not make: that the same simplified model for the climate’s systematic response to radiative forcing should be used to model residual, internal variability as a response to white noise forcing. More details on our uncertainty quantification for model (4), including our model for internal variability, are given in Sect. 4.

In building model (4), we chose to separate natural and anthropogenic forcings because they seem not strictly comparable (Fig. 1, right). First, there may be some evidence that aerosol forcing from volcanic eruptions is less efficacious than  $\text{CO}_2$  forcing (Sokolov, 2006; Hansen et al., 2005; Marvel et al., 2015). While we use estimates of effective radiative forcing, which should compensate for efficacy, these estimates do not include an adjustment for the volcanic forcing. Moreover, the timescales of response associated with these forcings may also be different, possibly because ocean heat content responds differently to sudden and/or negative changes in forcing (as produced by volcanic eruptions) compared to more gradual and/or positive changes (as in continued anthropogenic emissions of  $\text{CO}_2$ ) (e.g., Gregory and Forster, 2008; Padilla et al., 2011). We combine solar and volcanic forcings out of convenience; the solar forcings do change more rapidly than the anthropogenic forcings, and in any case the magnitude of the changes in solar forcing is small.

To illustrate the use of model (4), we fit it to different segments of the observed global mean surface temperature record and we compare it to the linear fits estimated over the

<sup>3</sup>The change in mean temperature associated with a doubling of  $\text{CO}_2$  concentration, after a sufficiently long time that the climate has reached a new equilibrium.



**Figure 2.** Comparison of the fitted values for model (4) and the linear time trend model fitted to the global mean temperature record over different time frames. (The fitted model Eq. 4 trends are extended back to 1880 regardless of the time frame used to fit the model.) The linear model appears in agreement with model (4) roughly after 1970, but not before. By contrast, model (4) produces fairly stable estimates of the mean response during the 20th century, although we note that the apparent fit to the data may be slightly poorer in the earliest part of the record.

same time frames. Figure 2 shows the resulting fitted models using the data from 1970 to 2015, 1950 to 2015, and 1880 to 2015. The estimated trend from model (4) is relatively insensitive to the time frame used. The estimated linear time trends, on the other hand, differ markedly using different time frames, and agree with model (4) only after around 1970, during the time period over which net radiative forcing was evolving approximately linearly in time (see again Fig. 1). The sensitivity of the inferred linear trend in global mean temperature to the starting date has been previously discussed (e.g., Liebmann et al., 2010). These results suggest both that model (4) does indeed capture important aspects of the underlying physical processes driving temperature trends and that it therefore may be used to answer more interesting questions than can the linear time trend model.

#### 4 Trend and uncertainty: what can we learn from applying our simple model to data?

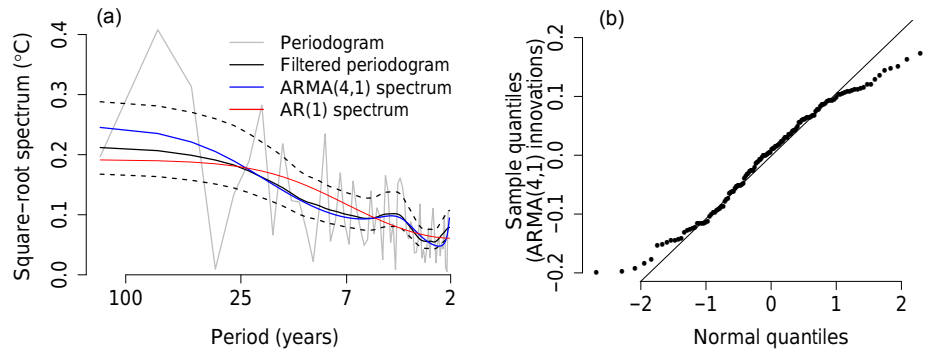
In this section, we illustrate what can be learned by applying the simple model (4) to observed temperatures. To do this, we must introduce an additional model to capture internal variability ( $\epsilon(t)$ ) in the assumed true model Eq. 1). We then use our full model to infer the parameters in model (4), to

evaluate their uncertainties given the data, and to explore the implications for understanding temperature trends.


To diagnose features of internal variability, spectral analysis is an intuitive framework, since the frequency properties of internal variability are tied to uncertainties in trends: uncertainty in smooth trends is more strongly affected by low-frequency than high-frequency internal variability. Figure 3, left, shows the raw periodogram associated with the residuals from model (4) and a smoothed estimate of the spectrum. We also show the spectra associated with fitting the residuals to two models for the internal variability. Both are standard autoregressive-moving average (ARMA) models (Supplement Sect. S3 includes a definition), allowing for dependence between the noise at each time step with the noise for past values. The two models we compare are an ARMA(4,1) and an AR(1) model, which were chosen according to common information criteria used for selecting time series models: respectively, the small sample-corrected Akaike information criterion (AICc) (Hurvich and Tsai, 1989) and the Bayesian information criterion (BIC) (Schwarz, 1978). (See Table S1 in the Supplement for the coefficient estimates, innovation standard deviations, and AICc and BIC associated with these two models.) Both models appear to fit the data reasonably well; the ARMA(4,1) model arguably overfits at the higher frequencies, but the AR(1) model may be underestimating variability at the lowest frequencies. Since low-frequency variability is most important to uncertainties in smooth trends, we adopt the more conservative choice of using the ARMA(4,1) model. Figure 3, right, shows the normal quantile–quantile ( $Q-Q$ ) plot for the sample innovations from this model; there is evidence that the innovations are somewhat more light-tailed than Gaussian, so standard errors based on a Gaussian assumption should not be overoptimistic.

Using the fully parametric model, combining Eq. (4) and a Gaussian ARMA(4,1) noise model, we proceed with uncertainty quantification through a parametric bootstrap<sup>4</sup>. When applied to our model of the historical temperature record, the parametric bootstrap distribution shows, unsurprisingly, that in a relatively short time series and given a smooth past trajectory of forcings, it is difficult to distinguish between a climate with both a high sensitivity (large value of  $\lambda$ ) and slow response (large value of  $\rho$ ) vs. one with a lower sensitivity (smaller value of  $\lambda$ ) but a faster response (smaller value of  $\rho$ ). The estimates  $\hat{\lambda}_A$  and  $\hat{\rho}_A$  are therefore strongly dependent, with  $\hat{\lambda}_A$  increasing explosively as  $\hat{\rho}_A \rightarrow 1$  (Fig. 4 shows their bivariate parametric bootstrap distribution). The strong nonlinear relationship between these two parameters, and the high degree of skewness in the marginal distribution of  $\hat{\lambda}_A$ ,

<sup>4</sup>A parametric bootstrap involves generating repeated, synthetic simulations under the fitted statistical model and then refitting the model to each simulated time series to obtain new estimates of model parameters. The distribution of those estimates then gives a measure of the uncertainty in the original estimates.



**Figure 3.** (a) The raw periodogram of the residuals from model (4), a filtered periodogram (twice applying a moving average of width 5), and the power spectra associated with fitted ARMA(4,1) and AR(1) models. The dashed lines are at  $\pm 2$  standard errors associated with the filtered periodogram. The ARMA(4,1) model appears to more realistically represent low-frequency variation in the residuals, which is crucial for inferences about trends. (b) Normal quantile–quantile plot for the ARMA(4,1) sample innovations. There is some evidence that the innovations are light-tailed compared to the normal distribution.



**Figure 4.** Distribution of the parametric bootstrap estimates of  $\lambda_A$  and  $\rho_A$  from model (4). It is difficult to distinguish between the rate of response and the sensitivity using only global mean temperatures from the recent past.

are the reasons that we rely on bootstrapping for uncertainty quantification, as the typical appeals to asymptotic normality are not viable in this setting.

In the following, we will represent uncertainties using the simple bootstrap percentile method. The percentile method is subject to criticism (e.g., Hall, 1988). We have found, however, that adjusting the percentile method using a nested bootstrap results in narrower confidence limits. For this reason, we believe that the raw percentile-based intervals may be conservative in this setting, so we choose to report the apparently conservative intervals. Our point estimates are based on a two-step procedure wherein model (4) is estimated via least squares and the ARMA(4,1) model is estimated via

maximum likelihood on the residuals. While this procedure may be somewhat suboptimal compared to jointly estimating the mean and covariance structure, the two-step procedure is substantially faster, which is important for carrying out the nested bootstrap.

#### 4.1 Uncertainties in the sensitivity parameters

When using the full 1880–2015 global mean surface temperature record, the point estimate for the centennial-scale sensitivity to anthropogenic forcing is  $\hat{\lambda}_A = 1.8$  °C per doubling, with  $\hat{\rho}_A = 0.80$  (which implies mixing on decadal timescales). The estimated sensitivity to natural forcing is much smaller,  $\hat{\lambda}_N = 0.21$  °C per doubling with  $\hat{\rho}_N = 0.58$ . In this section, we discuss uncertainties in these estimates.

Using our statistical model, the historical data appear to provide a lower bound for  $\lambda_A$  (assuming for now that the forcings are known) but cannot rule out extremely large and implausible values on the order of tens or hundreds of degrees per doubling (the 2.5–97.5th bootstrap percentile interval is 1.5 to 690 °C per doubling). These very large values are not supported by evidence from the paleoclimate record (e.g., IPCC, 2013 chap. 5) and the approximation of the linear forcing feedback model, on which Eq. (4) implicitly relies, breaks down under high sensitivity (Bloch-Johnson et al., 2015). The large upper bound should therefore be interpreted only as a statement about the information in the historical global mean temperature data under the strict assumption that model (4) holds exactly. For the sensitivity to natural forcings, the data cannot rule out  $\lambda_N = 0$  (the 2.5–97.5th bootstrap percentile interval is  $-1.1$  to  $4.1$  °C per doubling). Table 1 gives some intervals at different percentiles for the parameters  $\lambda_A$  and  $\lambda_N$ . (The parameters  $\rho_A$  and  $\rho_N$  are essentially unconstrained by the data.)

The IPCC’s own 66% “likely” interval for equilibrium sensitivity is 1.5 to 4.5 °C per doubling, which subjectively




**Table 1.** Parametric bootstrap percentile intervals for the sensitivities in model (4),  $\lambda_A$  and  $\lambda_N$ , to anthropogenic and natural sources, respectively (in units °C per doubling). The data appear to provide a lower bound for  $\lambda_A$ , but cannot rule out even implausibly large values; the very large values are associated with slow responses (see Fig. 4). The data cannot rule out  $\lambda_N = 0$ , likely because there are few prominent volcanic eruptions in the historical record analyzed and the response to volcanic aerosols may be small.

Percentile interval	$\hat{\lambda}_A = 1.8$		$\hat{\lambda}_N = 0.80$	
	Lower	Upper	Lower	Upper
5–95 (90 %)	1.5	3.0	−0.15	1.5
2.5–97.5 (95 %)	1.5	690	−1.1	4.1
0.5–99.5 (99 %)	1.4	790	−13	490

combines estimates from various sources using multiple lines of evidence, including ensembles of models with different physics, and accounts for other sources of uncertainty that we have so far ignored, such as uncertainty in the forcings themselves (see IPCC, 2013, Sect. 10.8.2 and Box 12.2). The bulk of the distribution of our estimate is somewhat narrower than the IPCC’s, likely because so far we have not accounted for uncertainty in radiative forcings; however, the IPCC rules out the very large and unphysical values in the right tail of our intervals. We here stress again that since we are estimating the centennial-scale response, the estimates of the sensitivity that we provide will tend to be lower than the equilibrium sensitivity estimated in the IPCC’s interval (see the Appendix, which also compares our estimates to other observationally based estimates of a sensitivity parameter). Individual estimates of the equilibrium climate sensitivity and associated uncertainty in the literature are discussed in IPCC (2013) Sect. 10.8.2 (see also again the Appendix).

The main source of uncertainty in the upper bound for  $\lambda_A$  is due to uncertainty in the “equilibration time” of the climate associated with smoothly increasing anthropogenic forcing, controlled by  $\rho_A$ . If we restrict  $\rho_A$  to, say, centennial scales or smaller, then the uncertainty is substantially decreased (see again Fig. 4). One may argue through other lines of evidence and reasoning that extremely large values of  $\rho_A$  and  $\lambda_A$  are implausible, but the statistical model is being used to quantify the information content of the historical temperature record. The fact that additional sources of information are needed to exclude unphysical values suggests that even our minimally informed model may be overly empirical for some purposes. The inability of the data to rule out  $\lambda_N = 0$  is, on the other hand, probably due to the fact that there are few prominent volcanic eruptions in the historical record analyzed and that the response to volcanic aerosols may be small for the reasons discussed above.



**Figure 5.** Projected mean temperature anomalies, and their uncertainties, under the RCP8.5 scenario, based on estimates from model (4) and assuming ARMA(4,1) noise. The black curve shows the observed temperatures. Intervals are pointwise (2.5–97.5)th percentile intervals. Radiative forcing stabilizes in the year 2250, but mean temperatures and especially their uncertainties continue to increase. While uncertainties in the long-term response are quite large, due largely to the inability to rule out implausible values of  $\lambda_A$ , the historical and near-term response is much more certain.

#### 4.2 Uncertainties in near- and long-term trends

The uncertainties in the sensitivity and rate of response parameters imply greater uncertainties in projected longer-term future trends in global mean temperature than in the historical and near-term projected trends. To illustrate this, we examine the implied future trends under the hypothetical (extended) RCP8.5 scenario, in which radiative forcing increases and then stabilizes in the year 2150. We simulate new time series using our estimates of model (4) and the ARMA(4,1) noise model, given radiative forcings from this scenario. The projected trend and associated pointwise uncertainties are shown in Fig. 5.

Projected mean temperatures, and especially their associated uncertainties, continue to increase even after stabilization of forcing. This is a consequence of the joint uncertainty in  $\lambda_A$  and  $\rho_A$ , and in particular of the inability to rule out implausible values of these parameters. If the goal, then, is to provide a long-term projection of mean temperatures given only the historical temperature record, these estimates will unsurprisingly be quite uncertain (even assuming known past and future forcings).

On the other hand, trends in the historical and near-term response are much more certain. The observations strongly suggest that mean temperatures increased in the 20th century; for example, the (2.5–97.5)th percentile interval for the mean

**Table 2.** Evaluation of how uncertainty in the sensitivity parameter,  $\lambda_A$ , decreases as more data are observed, from simulations under the RCP8.5 scenario of future radiative forcing with an assumed value of  $\lambda_A = 1.8^\circ\text{C}$  per doubling. The middle column shows the (2.5–97.5)th percentile intervals for  $\lambda_A$  from simulations under the fitted model (4) and ARMA(4,1) noise. The rightmost column shows the increase in mean temperatures from the preindustrial climate under the fitted model at the year in question. Uncertainties in the upper bound of  $\lambda_A$  decrease relatively slowly as more data are observed.

End year	(2.5–97.5)th percentile ( $^\circ\text{C}$ per doubling)	Change in mean from preindustrial ( $^\circ\text{C}$ )
2025	(1.5, 16)	1.3
2050	(1.6, 11)	2.2
2075	(1.6, 2.8)	3.1
2100	(1.7, 2.0)	3.9

response in the year 2000 (expressed compared to the 1951–1980 average) is well above zero at  $(0.4, 0.6)^\circ\text{C}$ . (The 30 year average of the data around the year 2000 gives an anomaly of  $0.5^\circ\text{C}$ , in line with our estimate using model Eq. 4.) These kinds of distinctions between the uncertainty in the near- and longer-term mean responses are not easily made using a time trend model.

#### 4.3 Decreasing uncertainty in the sensitivity parameter as more data are observed

We have shown that the short historical temperature record alone produces fairly uncertain estimates of the sensitivity parameter,  $\lambda_A$  in model (4) (Fig. 4, Table 1), and therefore longer-term temperature trends (Fig. 5). We now examine how these uncertainties decrease as the temperature record increases (as in, e.g., Kelly and Kolstad, 1999; Ring and Schlesinger, 2012; Padilla et al., 2011; Urban et al., 2014; Myhre et al., 2015). To do this, we artificially extend the temperature record by generating new synthetic time series using the mean and noise models estimated from the historical data and forcings from the same RCP8.5 scenario described above. We then re-estimate model (4) for each synthetic series, using successively longer synthetic datasets. The results suggest that the data will not constrain the upper bound on the sensitivity parameter until another  $\sim 50$  years, by which time (under our estimated model and the RCP8.5 scenario) temperatures will have already risen by about  $3^\circ\text{C}$  from the preindustrial climate. A summary of the evolution of uncertainties is given in Table 2.


These estimates could be more strongly constrained by using additional physical information. As discussed previously, the very high sensitivity estimates in the bootstrap distribution are cases where the estimated response time is unphysically long. Without external information about this timescale, however, long data records are required to rule out the large values of  $\rho_A$  and  $\lambda_A$  that the model entertains.

#### 4.4 Is there evidence of long memory internal variability in global mean temperatures?

One of the complicating factors in estimating trends in climate time series is the question of whether global mean temperatures exhibit *long memory*. Long memory processes have power spectra that behave like  $(2\sin(\omega/2))^{-2d}$  as the frequency  $\omega \rightarrow 0$ , with  $d > 0$  (i.e., infinite power at the frequency zero). When  $d < 1$  the process has finite total variance, as would be expected for a variable like global mean temperature. For more details on long memory processes, see Beran et al. (2013). By contrast, short-memory processes (like the ARMA(4,1) model we assume), have finite power at the origin. Many authors have suggested that internal temperature variability is well-modeled by processes with long memory but finite variance (e.g., Frankignoul and Hasselmann, 1977; Bloomfield, 1992; Smith and Chen, 1996; Gil-Alana, 2005; Lennartz and Bunde, 2009; Fraedrich et al., 2009; Rypdal and Rypdal, 2014; Løvsletten and Rypdal, 2016.). Some authors have moreover claimed that global mean temperatures are well-modeled by a random walk (e.g., Gordon, 1991), as has at least one standard time series textbook (Shumway and Stoffer, 2013), which would imply that global mean temperatures do not have a finite variance over time. In either case, if the Earth's temperatures exhibited long memory, it would be more difficult to estimate trends than in the short memory case, since low-frequency variability would then be more difficult to distinguish from trends.

The evidence for long memory, however, strongly depends on the assumed trend model. Many of the aforementioned authors draw their conclusions by assuming a linear time trend model and applying that model to the temperature record on durations of decades to over a century. (One notable exception is Rypdal and Rypdal, 2014, but they only compare long memory noise models to AR(1) models.) As discussed previously, a linear trend model applied to a time series with a nonlinear trend will imply excessive low-frequency noise. Figure 6 shows the periodograms of the residual global mean temperatures after removing either a linear time trend or a trend of the form of Eq. (4). While the high-frequency behavior of the residuals is not much affected by the choice of trend model, the low-frequency behavior is very much affected. Apparent low-frequency variability is made more severe by assuming that mean temperatures increase linearly in time.

The question of long memory cannot be definitively settled using a dataset of only 136 observations, and other analyses make use of longer climate model runs or the paleoclimate record (e.g., Mann, 2011). Nevertheless, it should be clear that the linear time trend model is especially problematic for this purpose. In general, regression models in time, linear, or otherwise, have a danger either of mistaking systematic trend for apparent low-frequency variability (as just described), or of mistaking low-frequency variability for systematic trend (as would occur, for example, when using a nonparametric



**Figure 6.** Raw periodograms of residuals from models (2) (dashed) and (4) (solid) fit to the full data record. There is substantially more low-frequency variability in the residuals from the linear trend model than in the residuals from model (4). Misspecified mean models will give a misleading impression of low-frequency variability, and therefore misleading uncertainties associated with the mean trend.

regression with too small of a smoothing bandwidth). Either can lead to misstated uncertainties, and therefore can be problematic even if the claim is that the trend model is only being used to test for significant warming and that the model is not believed to be true.

#### 4.5 Implications of uncertain inputs: radiative forcing and temperatures

The analysis thus far has assumed that both radiative forcings and temperatures are known exactly, but uncertainty in the sensitivity and in trends also propagates from uncertainty in these quantities. We therefore discuss at least roughly the potential implications of imperfect knowledge of these inputs.

Of the two factors, uncertainty in radiative forcings, particularly from aerosols, is more consequential, especially for the inferred lower bound of the sensitivity parameter from model (4). The importance of radiative forcing uncertainty for uncertainties in climate sensitivity has been widely noted (e.g., Gregory and Forster, 2008; Padilla et al., 2011; Otto et al., 2013; Masters, 2014; Lewis and Curry, 2015; Myhre et al., 2015). The trajectory of net effective radiative forcing from anthropogenic sources is poorly known; the IPCC states that the difference in net effective radiative forcing from anthropogenic sources between the years 2011 and 1750 was about  $2.29 \text{ W m}^{-2}$  with a 95% confidence interval of 1.13 to  $3.33 \text{ W m}^{-2}$ . We explore the implications of this uncertainty by simply scaling the entire trajectory of net radiative forcings from 1750 to the present such that the uncertainty in 2011 is as stated. Adjusting aerosol forcings to the

high or low ends, respectively, produces sensitivity estimates from Eq. (4) varying by over a factor of 2, from 1.2 to  $3.7 \text{ }^\circ\text{C}$  per doubling (vs. the original estimate  $1.8 \text{ }^\circ\text{C}$ ). The aerosol uncertainty appears important; the value  $1.2 \text{ }^\circ\text{C}$  per doubling is lower than all of the bootstrap values of  $\lambda_A$  generated assuming known forcings (Fig. 4 and Table 1).

Uncertainties in the global mean surface temperature record are comparatively less important. To partially address this issue, we re-estimate model (4) using each of the 100 ensemble members of the HadCRUT4 global mean temperature ensemble. The point estimates of  $\lambda_A$  range from 1.5 to  $2.1 \text{ }^\circ\text{C}$  per doubling, with estimates of  $\rho_A$  ranging from 0.79 to 0.90. The point estimates of  $\lambda_N$  range from 0.71 to  $20 \text{ }^\circ\text{C}$  degrees per doubling, with estimates of  $\rho_N$  ranging from 0.86 to 0.996. The additional uncertainty induced by observational uncertainty in the temperature record is smaller than either the uncertainty induced by internal temperature variability or the uncertainty in radiative forcings.

#### 4.6 Bayesian methods

Some of the uncertainties discussed so far could be addressed in a Bayesian framework (as in, e.g., Forest et al., 2002, 2006; Padilla et al., 2011; Aldrin et al., 2012) although we do not pursue that approach here. For example, the effects of uncertainties in temperatures and radiative forcings could be modeled hierarchically, and outside information could be used to constrain long-timescale responses not informed by the data. We have instead chosen here to focus on the information contained in global mean temperature observations about trends to illustrate some important sources of uncertainty in empirical analyses even assuming known inputs.

### 5 Parametric vs. nonparametric uncertainty quantification

In Sects. 3 and 4, we showed that analyses of the Earth's systematic temperature response are better informed by making physical assumptions than by an ostensibly more empirical approach. In this section, we ask a similar question about characterizing internal variability, in settings similar to that discussed, where data are temporally correlated and limited in length. We used a Gaussian ARMA(4,1) noise model in Sect. 4; while we do not claim a physical motivation for this model, the model does make relatively strong statistical assumptions. Is it more robust to assume a low-order parametric model for the noise, as we have done, or to adopt a nonparametric approach? The answer to this question depends on the length of the data record and the nature of the internal variability.

For the purposes of this illustration, we will use not the actual temperature record but rather some simple synthetic examples. We consider several artificial, trendless time series (the true mean of the process is constant) with temporally correlated noise, and evaluate the results of testing for a lin-

ear time trend (i.e., fitting model Eq. 2 and testing against the null hypothesis that  $\beta = 0$ ) using different parametric (both correctly specified and not) and nonparametric methods for estimating uncertainties. The ordinary least-squares standard errors, which assume uncorrelated noise, tend to be anti-conservative (too small) and time series methods are supposed to ameliorate this overconfidence, but may or may not be successful in this regard depending on the context. The illustrations are simple but our conclusions should be relevant to actual data analysis and to more informed models.

We consider a few parametric approaches common in time series analysis. The typical practice is to assume that the noise follows a low-order model, such as an ARMA model. Ideally, the noise model would be chosen (as in Sect. 4) in consultation with diagnostic plots and an information criterion like the AICc that balances model fit with the number of statistical parameters. However, it is common practice to automatically select a model either solely by minimizing the AICc or another criterion, or just to assume a simple time series model, often the AR(1) model. Uncertainty in the trend parameter(s) can then be estimated in several ways. Usually, this is done assuming asymptotic normality for maximum likelihood estimators, and a test might be carried out using a  $t$  statistic. If asymptotic normality is not viable, as for model (4), an alternative is to resort to using a parametric bootstrap (again as we do in Sect. 4) and to perform the test using bootstrap  $p$  values.

We also evaluate the perhaps most typical nonparametric method for accounting for dependence in a time series, the block bootstrap (Kunsch, 1989). Here, residuals are resampled in blocks to generate bootstrap samples that retain much of the dependence structure in the original data. A popular variant is the circular block bootstrap (Politis and Romano, 1992), in which blocks can be overlapping and blocks starting at the end of the time series wrap back to the beginning. The block bootstrap has been a common method in the climate and atmospheric sciences for at least 2 decades (Wilks, 1997) and has been applied previously to test for time trends (or changes thereof) in the global mean surface temperature record. For example, Rajaratnam et al. (2015) argued that the circular block bootstrap gives better uncertainty estimates than does a parametric analysis in the setting of testing for a trend in a 16-year segment of the global mean temperature record.

While the block bootstrap works very well in some settings, the procedure is not free of assumptions. Like other variants of the nonparametric bootstrap, its justification is based on an asymptotic argument; for the block bootstrap to work well, the size of the block has to be small compared to the overall length of the data but large compared to the scale of the temporal correlation in the data. When the overall data record is short and internal variability is substantially positively correlated in time, as for the historical global mean temperature record, these dual assumptions may not both be

met and we should not expect the block bootstrap to perform well.

In the following, we compare five methods (four parametric and one nonparametric) for generating nominal  $p$  values testing for a significant linear time trend in trendless synthetic data: (a) a  $t$  test assuming independent noise, (b) a  $t$  test assuming Gaussian AR(1) noise, (c) bootstrap  $p$  values from a parametric bootstrap assuming Gaussian AR(1) noise, (d) a  $t$  test assuming Gaussian ARMA( $p, q$ ) noise with the order chosen by minimizing AICc, and (e) bootstrap  $p$  values from a circular block bootstrap. We generate synthetic data from three different time series models, but in cases (b) and (c), the assumed parametric model is always AR(1). (Table S2 summarizes the models from which we simulate, which were chosen in part to share a similar strength of correlation with the observed global mean temperature record.) In cases (a), (b), and (d), the  $t$  test degrees of freedom are approximated by  $n - 2$  minus the number of parameters in the noise model. The parametric models in cases (a) through (d) are all estimated via maximum likelihood. In case (e), we estimate  $p$  values using a few different block lengths and show the results most favorable to the block bootstrap.

After generating nominal  $p$  values, we evaluate the performance of the different methods. Since the null hypothesis is true in this artificial setting (the synthetic series are trendless), a correct  $p$  value would be uniformly distributed. Deviations from the uniform distribution would then be a sign that the procedure generating the nominal  $p$  value is uncalibrated. We therefore use  $Q-Q$  plots to compare the distribution of the simulated nominal  $p$  values with the theoretical uniform distribution. Uncertainties are underestimated when the nominal  $p$  value quantiles are smaller than the theoretical quantiles (i.e., when the  $Q-Q$  plot is below the one-to-one line). In this situation, inferences are anti-conservative and the tests using the selected method will have type I error rates that are larger than the nominal rate.

### 5.1 Parametric vs. nonparametric methods under a correctly specified noise model

We first compare the performance of the five methods for generating nominal  $p$  values in the setting where the assumed AR(1) model is correctly specified (Fig. 7). Unsurprisingly, pre-specified parametric time series methods give reasonably calibrated inferences when the parametric model is correctly specified (Fig. 7, rows 2 and 3), although maximum likelihood gives somewhat anti-conservative estimates of uncertainty with small sample sizes. The anti-conservative bias of the maximum likelihood estimator can be reduced by instead using restricted maximum likelihood (REML) (e.g., McGilchrist, 1989) (see Supplement S4), but we focus here on the performance of maximum likelihood because it is more usually the procedure employed. It is also unsurprising that  $p$  values under assumed independence are quite uncalibrated and anti-conservative in this setting (Fig. 7, top

row) because standard errors are underestimated when positive temporal correlation is ignored.

It may, on the other hand, be surprising that automatically chosen parametric methods and nonparametric methods (blind selection via AICc and the block bootstrap; Fig. 7, rows 4 and 5) can perform *even more poorly than assuming independence* if the sample size is very small. The AICc does improve on the AIC (not shown) by attempting to account for small sample biases, but still performs poorly in very small samples. The (nonparametric) block bootstrap is the worst-performing method at the smallest sample sizes while ostensibly based on the weakest assumptions. For either of these methods to perform comparably to the pre-specified parametric methods, sample sizes must be quite large, indeed in this illustration larger than the available global mean temperature record. This should serve as a warning against using these methods for time series of modest length.

In actual practice, it can be advantageous, as we already discussed, to choose a noise model not automatically but in consultation with diagnostics (such as by comparing theoretical spectral densities with the empirical periodogram). In Sect. 4, we chose a noise model with consideration for the model's representation of low-frequency variability. In that example, the chosen model did minimize the AICc, but gave more conservative inferences than the AR(1) model (which was chosen by BIC). The illustration above shows that this behavior is not generically true and that for small sample sizes it is dangerous to blindly select models by minimizing an information criterion alone.

## 5.2 Parametric vs. nonparametric methods under a misspecified model

The comparisons in the previous section were too favorable to the pre-specified parametric methods because the order of the specified noise model (an AR(1) model) was known to be correct. Now we compare these methods when the assumed noise model is misspecified. The performance of misspecified methods will depend in particular on how the misspecified model represents low- vs. high-frequency variations in the noise process. Models that underestimate low-frequency variability will tend to be anti-conservative for estimating uncertainties in smooth trends, whereas those that overestimate low-frequency variability will tend to be conservative. We therefore repeat the illustrations in the previous section generating the synthetic time series from two different noise models (but still using the pre-specified AR(1) model to generate nominal  $p$  values). The two models are chosen so that their best AR(1) approximations either under- or over-represent low-frequency variability. Figure 8 shows the spectra corresponding to these two noise models and the best AR(1) approximation to each, and Figs. 9 and 10 show  $Q-Q$  plots corresponding to how the various time series methods perform in these two settings. See the supplement Table S2 for the model parameters and Sect. S2 for a comparison un-

der two additional models, including when the true model is the ARMA(4,1) model we assume in our main analysis of the global mean temperature record.

First, we consider an ARMA(1,1) process, whose best AR(1) approximation over-represents low-frequency variability (Fig. 9). The results are similar to those when the model was correctly specified, except that at the largest sample sizes the pre-specified parametric methods slightly over-estimate uncertainties, for the reasons discussed above. As before, both blind selection via AICc and the block bootstrap perform well for large sample sizes but very poorly for the smallest sample sizes.


Second, we consider a fractionally integrated AR(1) process; because this is a long-memory process, the best AR(1) approximation (and indeed any ARMA model) will severely under-represent low-frequency variability (Fig. 10). All of the methods struggle in this setting and produce anti-conservative estimates, but the pre-specified parametric methods still typically perform better than the ostensibly more flexible methods.

These results confirm that approaches to representing noise that appear to weaken assumptions are not guaranteed to outperform even misspecified parametric models. Misspecified parametric models are most dangerous when low-frequency variability is under-represented, but methods like the block bootstrap will also have the most trouble when low-frequency variability is strong because very long blocks will be required to adequately capture the scale of dependence in the data. While it is crucial for the data analyst to scrutinize any assumed parametric model, we believe that in many settings when the time series is not very long relative to the scale of correlation, one will be better served by carefully choosing a low-order parametric model rather than resorting to nonparametric methods.


Note that this illustration uses synthetic simulations that are relatively strongly correlated in time, a feature of the global mean temperature record. Nonparametric methods can work better than illustrated here in settings where correlations are weaker. For example, *local* (rather than global) temperatures, at least over land, tend to be more weakly correlated in time. In general, it is useful to evaluate the performance of statistical methods with simulations that share characteristics with the relevant real data.

## 6 Conclusions

We have sought to show here that targeted parametric modeling of global mean temperature trends and internal variability can provide more informative and accurate analyses of the global mean temperature record than can more empirical methods. Since all analyses involve assumptions, it is important to consider the role that assumptions play in resulting conclusions. In the setting of analyzing historical global mean temperatures, where the data record is relatively short




**Figure 7.** Quantile–quantile plots comparing the distribution of nominal  $p$  values to the theoretical uniform distribution. Simulations are of mean zero, Gaussian AR(1) time series and  $p$  values correspond to a two-sided test for a linear time trend. The length of the time series is given above each column. In the first row, the  $p$  values are from the OLS  $t$  test assuming independent noise; the second row is the  $t$  test assuming AR(1) noise; the third row uses a parametric bootstrap (again assuming AR(1) noise); the fourth row uses a  $t$  test assuming ARMA( $p, q$ ) noise with the order of the model selected by AICc; and the last row uses a circular block bootstrap with block size chosen to be favorable to this method. The  $p$  values calculated assuming the correct parametric model appear approximately correct for modest sample sizes and always outperform both blind selection by AICc and the block bootstrap. The latter two methods can be worse than assuming independence when sample sizes are very small.




**Figure 8.** The power spectra associated with the two models from which we generate synthetic time series in Sect. 5.2, along with spectra of the best AR(1) approximations (in Kullback–Leibler divergence) to both models; see Table S2 for noise model parameters. The corresponding illustrations of the performance of the various time series methods are shown in Figs. 9 and 10. The AR(1) model will tend to overestimate low-frequency variability in the first case and underestimate it in the second.

and temporal correlation is relatively strong, ostensibly more empirical methods can fail to distinguish between systematic trends and internal variability, and can give seriously uncalibrated estimates of uncertainty. While linear-in-time models

can be used for some purposes when applied to moderately narrow time frames (and with careful uncertainty quantification), the demonstrations shown here suggest that they do not have an intrinsic advantage over more targeted analyses. Tar-



**Figure 9.** Same as Fig. 7 but with simulations from an ARMA(1,1) model. The  $p$  values calculated by incorrectly assuming an AR(1) model are increasingly conservative in larger sample sizes in this setting. Both blind selection via AICc and the block bootstrap are anti-conservative for small sample sizes but improve as the sample size increases.



**Figure 10.** Same as Fig. 7 but with simulations from a fractional AR(1) model. While none of the methods perform very well here, the incorrectly specified parametric methods are better, especially in smaller samples.

geted analyses can be used over longer time frames – allowing for better estimates of both trends and noise characteristics – and can address a broader range of questions within a single framework.

The model we use in our analysis provides insights about the information contained in the historical temperature record relevant to both shorter-term and longer-term trend projections. The limited historical record of global mean temperature can provide information about shorter-term trends but unsurprisingly cannot constrain long-term projections very well. The past 136 years of temperatures simply do not alone contain the relevant information about equilibration timescales that would be required to constrain long-term projections, especially when aggregated to a single global value. (Use of spatially disaggregated data may provide additional information.) The distinction between uncertainties in shorter-term and longer-term projections, itself not easily made using a time trend model, serves to further illustrate that while the historical data record is an important source of information, it alone cannot be expected to answer the most important questions about climate change without bringing more scientific information to bear on the problem.

We believe that our discussion is illustrative of broader issues that arise in applied statistical practice, and will have particular relevance to problems involving trend estimation in the presence of temporally correlated data and in relatively data-limited settings, common in climate applications. We suspect that many applied statisticians have personally felt the tension between targeted modeling on the one hand and more empirical analyses on the other. One lucid discussion of the broader issues surrounding this tension can be found in the discussion of model formulation in Cox and Donnelly (2011). Empirical approaches have very successfully generated new insights and predictions in important areas, but there is a wide range of scientific problems where these approaches do not perform well and where targeted, domain-specific modeling is required. Statisticians can bring important insights to scientific problems; a crucial role for the statistician is to consider the modeling choices that will be both the most illuminating and the most reliable given the scientific questions and data at hand.

**Data availability.** The NASA GISS Land–Ocean Temperature index is updated periodically; the data we analyze were accessed on the date 2016-02-03. The current version is available at <http://data.giss.nasa.gov/gistemp/>. The HADCRUT4 data, used in Sect. 4.5, is available at <http://www.metoffice.gov.uk/hadobs/hadcrut4/data/current/download.html>.

Historical radiative forcings until 2011 are available in IPCC (2013) Table AII.1.2. Forcings corresponding to the RCP 8.5 scenario can be found at <http://tntcat.iiasa.ac.at/RcpDb>. NOAA CO<sub>2</sub> concentrations are available at [ftp://aftp.cmdl.noaa.gov/products/trends/co2/co2\\_annmean\\_gl.txt](ftp://aftp.cmdl.noaa.gov/products/trends/co2/co2_annmean_gl.txt).



## Appendix A: Implications for the transient climate response and equilibrium sensitivity

In this paper, we use the historical temperature record to estimate the trend model (4) that treats both the sensitivity parameter,  $\lambda_A$ , and the timescale of response parameter,  $\rho_A$ , as unknown. In this section, we compare our results to those from more commonly used methods of inferring the climate response from projections in climate models and of using the historical temperature record to estimate a climate sensitivity parameter.

Model (4) is justifiable over the relatively short (centennial) timescales over which the *linear forcing–feedback model* is applicable. The linear forcing–feedback model says that the mean temperature response to an instantaneous forcing  $F$  behaves like  $L\Delta T(t) \approx F - Q(t)$ , where  $\Delta T(t)$  is the change in mean temperature at time  $t$  from the baseline state,  $Q(t)$  is the heat uptake (in  $\text{W m}^{-2}$ ) at time  $t$ , and  $L$  is a climate response parameter (with units  $\text{W m}^{-2} \text{K}^{-1}$ ). (In the context of this model, our sensitivity parameter is then  $\lambda_A = F_{2\times}/L$ .) Our method essentially infers  $Q(t)$  by assuming that it decays exponentially until reaching zero again at equilibrium. The linear forcing–feedback model does not hold for all time, however, in part because both feedbacks and rates of warming vary spatially. In models, this typically results in additional warming over the millennial timescales on which the deep ocean mixes (e.g., Winton et al., 2010; Armour et al., 2013; Rose et al., 2014 and others). Additionally, climate feedbacks – and therefore climate sensitivity – may be state dependent; this effect also typically amplifies global mean temperature rise (Bloch-Johnson et al., 2015). Collectively, this suggests that our sensitivity parameter  $\lambda_A$  will be smaller than the equilibrium climate sensitivity that measures the final, equilibrium temperature response.

We can compare the results of our model to reported results from GCMs by estimating the *transient climate response* (TCR), a popular metric of the short-term temperature response to forcing in climate models. The TCR is defined as the change in mean temperature after 70 years of a  $\text{CO}_2$  concentration scenario that increases by 1% per year (so doubles after 70 years). In a multi-model comparison of these centennial-scale projections, IPCC (2013) reports a 66% “likely” interval for the TCR of 1.0 to 2.5 °C. Our results from fitting model (4) to the historical data are fairly consistent with this; our estimate of the TCR is 1.7 °C, with the 2.5–97.5th bootstrap percentile interval of 1.2 to 1.9 °C. These intervals are not strictly comparable because the IPCC’s subjectively combines information from multiple lines of evidence. That said, our interval is shorter in part because it does not account for uncertainties in historical radiative forcings from anthropogenic aerosols. If we repeat the exercise of Sect. 4.5 (scaling the past radiative forcing trajectory to approximate the upper and lower bounds for forcing accounting for aerosol uncertainties), our central estimate

of the TCR would be about 1.2 °C if anthropogenic aerosols were on the high end and about 3.4 °C if on the low end.

We can also compare our estimate of  $\lambda_A$  to prior estimates of a sensitivity parameter that also use the historical temperature record. The most typical approach in the literature shares some commonalities with our method, beginning with the same linear forcing–feedback model that implicitly underlies our analysis, but estimating a sensitivity parameter by using an additional observational estimate of global heat uptake. That is, studies use estimates of changes in forcing,  $F'$ , heat uptake  $Q'$ , and historical temperature change,  $\Delta T'$ , between a base and a final period to compute  $\hat{L} = (F' - Q')/\Delta T'$  and therefore  $\hat{\lambda} = F_{2\times}\Delta T'/(F' - Q')$ . Studies using this method include Gregory et al. (2002), Otto et al. (2013), Masters (2014), and Lewis and Curry (2015). The resulting sensitivity parameter estimate should be similar to our  $\lambda_A$ . Table A1 shows the results from these analyses; the central estimates are indeed similar to our estimate of  $\lambda_A$ . The uncertainty ranges given in these studies, however, also tend to be slightly larger than the intervals we give for  $\lambda_A$ , again because these authors attempt to account for radiative forcing uncertainty in their analysis.

A distinguishing feature of this other, common approach is that it includes data about historical heat uptake, in addition to temperature and radiative forcing data. Ocean heat content is an additional, albeit uncertain, source of information that may improve estimates. On the other hand, since these methods do not involve an explicit trend model and require averaging the inputs over decadal or longer time spans, they cannot use the historical temperature record to estimate internal temperature variability. Most studies therefore estimate internal variability using climate model output, but climate models do not perfectly realistically represent even present-day variability in global annual mean temperature.

By contrast, an advantage of our approach is that it allows one to use the historical data to understand internal variability. Additionally, our approach allows one to answer questions about both historical trends and longer-term projections in the framework of one statistical model, whereas the approaches discussed above do not allow one to infer trends in increasing-in-time forcing scenarios. A disadvantage of our approach is that, as discussed above, we rely on the historical global mean temperature record to estimate the “equilibration” timescales ( $\rho_A$  and  $\rho_N$ ), but the data contain little information about these quantities.

Regardless of the different advantages and disadvantages just discussed, both approaches to using the historical temperature record give similar results concerning the sensitivity parameter, and uncertainties in this parameter remain high. This demonstrates the limitations of the information content of the historical global mean temperature record alone for estimating longer-term projections of mean temperature changes. As noted in Sect. 6, spatially disaggregated data may contain more information.

**Table A1.** Comparison of estimates of a sensitivity parameter from studies that use observational data and a simple energy balance approach. The large best (median) estimate from Gregory et al. (2002) is due to a very fat right tail in their analysis; the mode of their distribution is 2.1 °C per doubling. The estimates given in these studies are similar to our estimates of  $\lambda_A$  (which should be smaller than the equilibrium sensitivity).

	Study	Best estimate (°C per doubling)	90 % Interval (°C per doubling)
Energy balance model using temperatures, forcing and heat uptake	Gregory et al. (2002)	6.1	> 1.6
	Otto et al. (2013)	2.0	1.2–3.9
	Masters (2014)	2.0	1.2–5.2
	Lewis and Curry (2015)	1.6	1.1–4.1
Trend model (4)	this work	1.8	1.5–3.0

**The Supplement related to this article is available online at <https://doi.org/10.5194/ascmo-3-33-2017-supplement>.**

**Competing interests.** The authors declare that they have no conflict of interest.

**Acknowledgements.** The authors thank Jonah Bloch-Johnson, Malte Jansen, Cristian Proistosescu, and Kate Marvel for helpful conversations and comments related to parts of this work. We additionally thank the reviewers of this paper, whose suggestions led to a number of improvements. This work was supported in part by STATMOS, the Research Network for Statistical Methods for Atmospheric and Oceanic Sciences (NSF-DMS awards 1106862, 1106974 and 1107046), and RDCEP, the University of Chicago Center for Robust Decision-making in Climate and Energy Policy (NSF grant SES-0951576). We thank NASA GISS, NOAA, the Hadley Centre, the IPCC, and IIASA for the use of their publicly available data. We acknowledge the University of Chicago Research Computing Center, whose resources were used in the completion of this work.

Edited by: C. Forest

Reviewed by: three anonymous referees

## References

- Aldrin, M., Holden, M., Guttorp, P., Skeie, R. B., Myhre, G., and Berntsen, T. K.: Bayesian estimation of climate sensitivity based on a simple climate model fitted to observations of hemispheric temperatures and global ocean heat content, *Environmetrics*, 23, 253–271, 2012.
- Armour, K. C., Bitz, C. M., and Roe, G. H.: Time-varying climate sensitivity from regional feedbacks, *J. Climate*, 26, 4518–4534, 2013.
- Arrhenius, S.: On the influence of carbonic acid in the air upon the temperature of the ground, *Philos. Mag.* 5, 41, 237–276, 1896.
- Beran, J., Feng, Y., Ghosh, S., and Kulik, R.: *Long-memory processes*, Springer, New York, 2013.
- Bloch-Johnson, J., Pierrehumbert, R. T., and Abbot, D. S.: Feedback temperature dependence determines the risk of high warming, *Geophys. Res. Lett.*, 42, 4973–4980, 2015.
- Bloomfield, P.: Trends in global temperature, *Climatic Change*, 21, 1–16, 1992.
- Cahill, N., Rahmstorf, S., and Parnell, A. C.: Change points of global temperature, *Environ. Res. Lett.*, 10, 084002, <https://doi.org/10.1088/1748-9326/10/8/084002>, 2015.
- Caldeira, K. and Myhrvold, N. P.: Projections of the pace of warming following an abrupt increase in atmospheric carbon dioxide concentration, *Environ. Res. Lett.*, 8, 034039, <https://doi.org/10.1088/1748-9326/8/3/0340>, 2013.
- Castruccio, S., McInerney, D. J., Stein, M. L., Liu Crouch, F., Jacob, R. L., and Moyer, E. J.: Statistical emulation of climate model projections based on precomputed GCM runs, *J. Climate*, 27, 1829–1844, 2014.
- Cox, D. R. and Donnelly, C. A.: *Principles of applied statistics*, Cambridge University Press, New York, 2011.
- Easterling, D. R. and Wehner, M. F.: Is the climate warming or cooling?, *Geophys. Res. Lett.*, 36, L08706, <https://doi.org/10.1029/2009GL037810>, 2009.
- Forest, C. E., Stone, P. H., Sokolov, A. P., Allen, M. R., and Webster, M. D.: Quantifying uncertainties in climate system properties with the use of recent climate observations, *Science*, 295, 113–117, 2002.
- Forest, C. E., Stone, P. H., and Sokolov, A. P.: Estimated PDFs of climate system properties including natural and anthropogenic forcings, *Geophys. Res. Lett.*, 33, L01705, <https://doi.org/10.1029/2005GL023977>, 2006.
- Foster, G. and Rahmstorf, S.: Global temperature evolution 1979–2010, *Environ. Res. Lett.*, 6, 044022, <https://doi.org/10.1088/1748-9326/6/4/0440>, 2011.
- Fraedrich, K., Blender, R., and Zhu, X.: Continuum climate variability: Long-term memory, scaling, and 1/f-noise, *Int. J. Mod. Phys. B*, 23, 5403–5416, 2009.
- Frankignoul, C. and Hasselmann, K.: Stochastic climate models, Part II Application to sea-surface temperature anomalies and thermocline variability, *Tellus*, 29, 289–305, 1977.
- Geoffroy, O., Saint-Martin, D., Bellon, G., Voldoire, A., Olivié, D., and Tytéca, S.: Transient climate response in a two-layer energy-balance model. Part II: Representation of the efficacy of deep-ocean heat uptake and validation for CMIP5 AOGCMs, *J. Climate*, 26, 1859–1876, 2013a.
- Geoffroy, O., Saint-Martin, D., Olivié, D. J. L., Voldoire, A., Bellon, G., and Tytéca, S.: Transient climate response in a two-layer energy-balance model. Part I: Analytical solution and parameter calibration using CMIP5 AOGCM experiments, *J. Climate*, 26, 1841–1857, 2013b.
- Gil-Alana, L. A.: Statistical modeling of the temperatures in the Northern Hemisphere using fractional integration techniques, *J. Climate*, 18, 5357–5369, 2005.
- Gluhovsky, A.: Statistical inference from atmospheric time series: detecting trends and coherent structures, *Nonlin. Processes Geophys.*, 18, 537–544, <https://doi.org/10.5194/npg-18-537-2011>, 2011.
- Gordon, A. H.: Global warming as a manifestation of a random walk, *J. Climate*, 4, 589–597, 1991.
- Gregory, J. M. and Forster, P. M.: Transient climate response estimated from radiative forcing and observed temperature change, *J. Geophys. Res.-Atmos.*, 113, D23105, <https://doi.org/10.1029/2008JD010405>, 2008.
- Gregory, J. M., Stouffer, R. J., Raper, S. C. B., Stott, P. A., and Rayner, N. A.: An observationally based estimate of the climate sensitivity, *J. Climate*, 15, 3117–3121, 2002.
- Hall, P.: Theoretical comparison of bootstrap confidence intervals, *Ann. Stat.*, 16, 927–953, 1988.
- Hansen, J., Sato, M., Ruedy, R., Nazarenko, L., Lacis, A., Schmidt, G., Russell, G., Aleinov, I., Bauer, M., Bauer, S., et al.: Efficacy of climate forcings, *J. Geophys. Res.-Atmos.*, 110, D18104, <https://doi.org/10.1029/2005JD005776>, 2005.
- Hansen, J., Ruedy, R., Sato, M., and Lo, K.: Global surface temperature change, *Rev. Geophys.*, 48, RG4004, <https://doi.org/10.1029/2010RG000345>, 2010.
- Held, I. M., Winton, M., Takahashi, K., Delworth, T., Zeng, F., and Vallis, G. K.: Probing the fast and slow components of global warming by returning abruptly to preindustrial forcing, *J. Climate*, 23, 2418–2427, 2010.

- Hurvich, C. M. and Tsai, C.-L.: Regression and time series model selection in small samples, *Biometrika*, 76, 297–307, 1989.
- Huybers, P.: Compensation between model feedbacks and curtailment of climate sensitivity, *J. Climate*, 23, 3009–3018, 2010.
- IPCC: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, New York, NY, edited by: Stocker, T., Qin, D., Plattner, G.-K., Tignor, M., Allen, S. K., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P., 2013.
- Kelly, D. L. and Kolstad, C. D.: Bayesian learning, growth, and pollution, *J. Econ. Dyn. Control*, 23, 491–518, 1999.
- Knutti, R.: Why are climate models reproducing the observed global surface warming so well?, *Geophys. Res. Lett.*, 35, L18704, <https://doi.org/10.1029/2008GL034932>, 2008.
- Knutti, R., Stocker, T. F., Joos, F., and Plattner, G.-K.: Constraints on radiative forcing and future climate change from observations and climate model ensembles, *Nature*, 416, 719–723, 2002.
- Kunsch, H. R.: The jackknife and the bootstrap for general stationary observations, *Ann. Stat.*, 17, 1217–1241, 1989.
- Lennartz, S. and Bunde, A.: Trend evaluation in records with long-term memory: Application to global warming, *Geophys. Res. Lett.*, 36, L16706, <https://doi.org/10.1029/2009GL039516>, 2009.
- Lewis, N. and Curry, J. A.: The implications for climate sensitivity of AR5 forcing and heat uptake estimates, *Clim. Dynam.*, 45, 1009–1023, 2015.
- Liebmann, B., Dole, R. M., Jones, C., Bladé, I., and Allured, D.: Influence of choice of time period on global surface temperature trend estimates, *B. Am. Meteorol. Soc.*, 91, 1485–1491, 2010.
- Løvstetten, O. and Rypdal, M.: Statistics of regional surface temperatures post year 1900. Long-range versus short-range dependence, and significance of warming trends, *J. Climate*, 29, 4057–4068, 2016.
- MacKay, R. M. and Ko, M. K.: Normal modes and the transient response of the climate system, *Geophys. Res. Lett.*, 24, 559–562, 1997.
- Mann, M. E.: On long range dependence in global surface temperature series, *Climatic Change*, 107, 267–276, 2011.
- Marvel, K., Schmidt, G. A., Miller, R. L., and Nazarenko, L. S.: Implications for climate sensitivity from the response to individual forcings, *Nature Climate Change*, 6, 386–389, 2015.
- Masters, T.: Observational estimate of climate sensitivity from changes in the rate of ocean heat uptake and comparison to CMIP5 models, *Clim. Dynam.*, 42, 2173–2181, 2014.
- Mauritsen, T., Stevens, B., Roeckner, E., Crueger, T., Esch, M., Giorgetta, M., Haak, H., Jungclaus, J., Klocke, D., Matei, D., Mikolajewicz, U., Notz, D., Pincus, R., Schmidt, H., and Tomassini, L.: Tuning the climate of a global model, *J. Adv. Model. Earth Syst.*, 4, M00A01, <https://doi.org/10.1029/2012MS000154>, 2012.
- McGilchrist, C.: Bias of ML and REML estimators in regression models with ARMA errors, *J. Stat. Comput. Sim.*, 32, 127–136, 1989.
- Meinshausen, M., Smith, S. J., Calvin, K., Daniel, J. S., Kainuma, M., Lamarque, J., Matsumoto, K., Montzka, S. A., Raper, S. C. B., Riahi, K., Thomson, A., Velders, G. J. M., and van Vuuren, D. P.: The RCP greenhouse gas concentrations and their extensions from 1765 to 2300, *Climatic Change*, 109, 213–241, 2011.
- Morice, C. P., Kennedy, J. J., Rayner, N. A., and Jones, P. D.: Quantifying uncertainties in global and regional temperature change using an ensemble of observational estimates: The HadCRUT4 data set, *J. Geophys. Res.-Atmos.*, 117, D08101, <https://doi.org/10.1029/2011JD017187>, 2012.
- Myhre, G., Boucher, O., Bréon, F.-M., Forster, P., and Shindell, D.: Declining uncertainty in transient climate response as CO<sub>2</sub> forcing dominates future climate change, *Nat. Geosci.*, 8, 181–185, 2015.
- Olivié, D. J. L., Peters, G. P., and Saint-Martin, D.: Atmosphere response time scales estimated from AOGCM experiments, *J. Climate*, 25, 7956–7972, 2012.
- Otto, A., Otto, F. E. L., Boucher, O., Church, J., Hegerl, G., Forster, P. M., Gillett, N. P., Gregory, J., Johnson, G. C., Knutti, R., Lewis, N., Lohmann, U., Marotzke, J., Myhre, G., Shindell, D., Stevens, B., and Allen, M. R.: Energy budget constraints on climate response, *Nat. Geosci.*, 6, 415–416, 2013.
- Padilla, L. E., Vallis, G. K., and Rowley, C. W.: Probabilistic estimates of transient climate sensitivity subject to uncertainty in forcing and natural variability, *J. Climate*, 24, 5521–5537, 2011.
- Politis, D. N. and Romano, J. P.: A circular block-resampling procedure for stationary data, in: *Exploring the limits of bootstrap*, edited by: LePage, R. and Billard, L., Wiley, New York, 263–270, 1992.
- Rajaratnam, B., Romano, J., Tsiang, M., and Diffenbaugh, N. S.: Debunking the climate hiatus, *Climatic Change*, 133, 129–140, 2015.
- Riahi, K., Rao, S., Krey, V., Cho, C., Chirkov, V., Fischer, G., Kindermann, G., Nakicenovic, N., and Rafaj, P.: RCP 8.5—A scenario of comparatively high greenhouse gas emissions, *Climatic Change*, 109, 33–57, 2011.
- Ring, M. J. and Schlesinger, M. E.: Bayesian learning of climate sensitivity I: Synthetic observations, *Atmos. Clim. Sci.*, 2, 464–473, 2012.
- Rose, B. E. J., Armour, K. C., Battisti, D. S., Feldl, N., and Koll, D. D. B.: The dependence of transient climate sensitivity and radiative feedbacks on the spatial pattern of ocean heat uptake, *Geophys. Res. Lett.*, 41, 1071–1078, 2014.
- Rypdal, M. and Rypdal, K.: Long-Memory Effects in Linear Response Models of Earth's Temperature and Implications for Future Global Warming, *J. Climate*, 27, 5240–5258, 2014.
- Santer, B. D., Mears, C., Doutriaux, C., Caldwell, P., Gleckler, P. J., Wigley, T. M. L., Solomon, S., Gillett, N. P., Ivanova, D., Karl, T. R., et al.: Separating signal and noise in atmospheric temperature changes: The importance of timescale, *J. Geophys. Res.-Atmos.*, 116, D22105, <https://doi.org/10.1038/ngeo1836>, 2011.
- Schwarz, G.: Estimating the dimension of a model, *Ann. Stat.*, 6, 461–464, 1978.
- Shumway, R. H. and Stoffer, D. S.: *Time series analysis and its applications*, Springer, New York, 2013.
- Smith, R. L. and Chen, F.-L.: Regression in long-memory time series, in: *Athens Conference on Applied Probability and Time Series Analysis*, Springer, New York, 378–391, 1996.
- Sokolov, A. P.: Does model sensitivity to changes in CO<sub>2</sub> provide a measure of sensitivity to other forcings?, *J. Climate*, 19, 3294–3306, 2006.
- Urban, N. M., Holden, P. B., Edwards, N. R., Sriver, R. L., and Keller, K.: Historical and future learning about climate sensitivity, *Geophys. Res. Lett.*, 41, 2543–2552, 2014.

- Wigley, T. M. L., Jones, P. D., and Raper, S. C. B.: The observed global warming record: What does it tell us?, *P. Natl. Acad. Sci. USA*, 94, 8314–8320, 1997.
- Wilks, D. S.: Resampling hypothesis tests for autocorrelated fields, *J. Climate*, 10, 65–82, 1997.
- Winton, M., Takahashi, K., and Held, I. M.: Importance of ocean heat uptake efficacy to transient climate change, *J. Climate*, 23, 2333–2344, 2010.
- Wunsch, C.: The interpretation of short climate records, with comments on the North Atlantic and Southern Oscillations, *B. Am. Meteorol. Soc.*, 80, 245–255, 1999.
- Zeng, X. and Geil, K.: Global warming projection in the 21st century based on an observational data-driven model, *Geophys. Res. Lett.*, 43, 947–954, 2016.