

Global Warming Is Likely Affecting Regional Drought across Eurasia

Kate Marvel¹, Benjamin I. Cook¹, and Edward Cook²

¹NASA Goddard Institute for Space Studies

²Tree Ring Laboratory, Lamont-Doherty Earth Observatory of Columbia University

Key Points:

- We present a flexible Bayesian modeling framework for detecting regional hydroclimate responses to rising temperatures.
- We learn the spatiotemporal characteristics of internal variability from tree-ring based paleoclimate records in the pre-industrial era.
- We find that the influence of global warming is likely present in the twenty-first century hydroclimate of many regions.

Corresponding author: Kate Marvel, kate.marvel@nasa.gov

Abstract

While rising global temperatures have altered global drought risk and are projected to continue to change large-scale hydroclimate, it has proved difficult to detect the influence of warming on drought-relevant variables at regional scales. In addition to the inherent difficulty in identifying signals in noisy data, detection and attribution studies generally rely on general circulation models, which may fail to accurately capture the characteristics of naturally forced and internal hydroclimate variability. Here, we use a long tree-ring based paleoclimate record of drought to estimate pre-industrial variability in the Palmer Drought Severity Index (PDSI), a commonly used metric of drought risk. Using a Bayesian framework, we estimate the temporal and spatial characteristics of hydroclimate variability prior to 1850. We assess whether observed twenty-first century PDSI is compatible with this pre-industrial variability or is better explained by a forced response that depends on global mean temperature. Our results suggest that global warming likely contributed to dry PDSI in Eastern Europe, the Mediterranean, and Arctic Russia and to wet PDSI in Northern Europe, East-central Asia, and Tibet.

Plain Language Summary

Are twenty-first century drought conditions due to global warming, or can they just as well be explained by natural climate cycles? Data from tree rings gives us a record of previous drought conditions that stretches all the way back to the year 1000 CE. We use this long paleoclimate record to learn the complex structure of natural climate variability before the Industrial Revolution. Recent conditions in many regions are not well explained by previous patterns of natural variability, but are compatible with an external factor: the influence of rising global temperatures.

1 Introduction

As the planet warms, the risk of drought is expected to change in many regions (Seneviratne et al., 2023; B. I. Cook et al., 2020). Previous studies have identified the influence of rising temperatures on global drought patterns (Marvel et al., 2019; Bonfils et al., 2017) and regional droughts of particular severity, including the 2000-2022 southwest US megadrought (Swain et al., 2014; Williams et al., 2015, 2020). But the identification of novel or unprecedented drought conditions, as well as attribution to specific drivers, usually depends on the use of coupled general circulation models (GCMs). GCMs are used to identify fingerprints of external forcing (e.g. N. Gillett et al. (2002); Hegerl et al. (1996); Allen and Stott (2003); Tett et al. (2002); Stott et al. (2000); Santer, Painter, Bonfils, et al. (2013)) as well as to simulate and quantify pre-industrial climate variability (Santer et al., 2011; Santer, Painter, Mears, et al., 2013). However, the state-of-the-art GCMs participating in the Coupled Model Intercomparison Project, Phase 6 (CMIP6, Eyring et al. (2016)) exhibit many biases in their representation of global (Tokarska et al., 2020; Hausfather et al., 2022) and regional (Richter & Tokinaga, 2020) temperature, precipitation (Yazdandoost et al., 2021), extremes (Kim et al., 2020), and land surface properties that may affect the credibility of their estimates of pre-industrial variability. Moreover, while GCM projections of the future appear coherent over some regions, there is great uncertainty in the magnitude or even sign of future changes in drought risk in some regions (B. I. Cook et al., 2020; Marvel et al., 2021).

In the case of drought risk, we can circumvent many of the challenges associated with the GCMs by drawing upon long reconstructions of last-millennium hydroclimate derived from tree ring measurements. These “drought atlases” provide a record of internal and naturally forced climate variability that stretch back centuries. They allow us to learn about the spatial and temporal properties of this natural variability and provide a GCM-independent means of identifying unusual or unprecedented states or patterns in the present day (e.g. Marvel and Cook (2022)).

Here, we present a flexible, extendable Bayesian method for learning about past and present drought conditions. We use this framework to demonstrate that in many regions, it is likely that rising global temperatures have affected drought conditions. The paper is structured as follows: in section 2, we discuss the data and methods used. We show how the drought atlases may be used to “learn” the parameters of the spatial covariance (i.e., how different regions naturally change in relation to one another) and the temporal autocorrelation (how much drought risk in a particular region depends on previous years). We describe a simple model for recent hydroclimate variability, and show how Bayesian posterior predictive distributions can be used to separate the signal of a forced response to global warming from the noise of pre-industrial variability. In section 3 we present results for the spatiotemporal structure of pre-industrial variability, the fingerprint of regional response to global temperature, and attribution results. In section 4 we discuss the limitations of this method and possible future extensions.

2 Methods

2.1 Drought atlas description

We use the new Great Eurasian Drought Atlas (GEDA, B. Cook et al. (2024)), a tree-ring based reconstruction of past hydroclimate variability that updates existing drought atlases (E. R. Cook et al., 2010, 2015, 2020). The GEDA, which targets summer (JJA) self-calibrating Palmer Drought Severity Index (PDSI, Wells et al. (2004)), spans the 1,021-year period 1000CE–2020CE. Tree-ring based reconstructions are used from 1000–1989 CE and instrumental observations from the University of East Anglia Climate Research Unit (CRU) (van der Schrier et al., 2013) based on the CRU TS gridded dataset (Harris et al., 2020) are used from 1990–2020. Full details on the development and validation of the GEDA can be found in B. Cook et al. (2024).

We average the GEDA spatially over land regions used in the IPCC Sixth Assessment Report (hereafter AR6, Iturbide et al. (2020)). The GEDA provides full coverage over all European and Asian regions with the exception of Southeast Asia (SEA), where coverage extends over only the northern half of the region (Figure 1.) We split the GEDA into “preindustrial” (1000–1849) and recent (1850–2020) components. 1850 is chosen as the dividing line because all Coupled Model Intercomparison Project (CMIP) “historical” simulations begin on this date (Eyring et al., 2016). We standardize PDSI in all regions by subtracting the pre-industrial mean and dividing by the pre-industrial standard deviation.

2.2 Bayesian methods

Bayesian methodology has long been applied to the problem of climate change detection and attribution (e.g. (Annan, 2010; Katzfuss et al., 2017; Berliner et al., 2000)) as well as other problems in climate science (e.g. (Sherwood et al., 2020; Tierney et al., 2020)). In this section, we describe the basics of the Bayesian methodology used in our analysis. Suppose we have data D that we wish to interpret using a model characterized by a set of parameters Θ . If we begin with a set of prior beliefs $P(\Theta)$ about these parameters, we can use Bayes’ Theorem to update these beliefs in light of the evidence D :

$$P(\Theta|D) = \frac{P(D|\Theta)P(\Theta)}{P(D)}. \quad (1)$$

Here, $P(\Theta|D)$ is the posterior distribution, which can be thought of as representing our updated knowledge about the parameters given the evidence. The term $P(D|\Theta)$ is the likelihood of observing the evidence given some value of the parameters. The denominator $P(D)$ is a normalization constant that makes the posterior a true probability distribution.




Figure 1. Tree-ring based reconstructions: spatial extent and regions. Self-calibrating summer (JJA) average Palmer Drought Severity Index for 1500CE from the GEDA, along with selected AR6 regions.

The goal of a Bayesian analysis is to use available evidence to update our priors (Gelman et al., 1995). But what, exactly, are those parameters? The answer depends on the model we use to interpret the evidence. Here, we will use “GCM” to refer to complex general circulation models and reserve the term “model” for this interpretive framework. It is important to clearly specify this model, as we do in the next section.

2.3 Modeling the preindustrial period

In this section, we will show how the Bayesian framework described above can be applied to pre-industrial drought variability as represented by the GEDA. The interpretive model we specify will determine the parameters we fit and the inferences we can make. For example, if we believe pre-industrial variability in PDSI to be pure white noise whose standard deviation is independent of location, then our model would contain a single parameter: the standard deviation σ . Of course, we know that this is not likely to be a very good model for pre-industrial variability: summer soil moisture is known to exhibit strong year-to-year persistence (B. I. Cook et al., 2022). The PDSI in a given year depends on the PDSI in the year before, and perhaps in years prior to that. Moreover, we know that certain modes of internal variability cause PDSI in different regions to co-vary positively or negatively with one another (Baek et al., 2017). This means we should use a more complex model to interpret the pre-industrial period that takes into account the spatio-temporal structure of natural variability.

Here, we assume that pre-industrial regional PDSI in one year depends on regional PDSI in the n_{lag} previous years. We also assume that the spatial relationships between r different AR6 regions are described by a $r \times r$ covariance matrix Σ . We assume the r -dimensional vector of regional PDSI at time t , $\mathbf{D}(t)$, is drawn from a multivariate normal distribution :

$$\mathbf{D}(t) \sim MN(\mu, \Sigma) \quad (2)$$

where the mean depends on the time-varying response to external forcing $F(t)$ and the value of D in the n_{lag} previous years:

$$\mu(t) = \mathbf{F}(t) + \sum_{j=1}^{n_{lag}} \ell_j \mathbf{D}(t-j).$$

That is, the PDSI in any given region depends in some unknown way on what happened in that region in previous years, while the overall PDSI pattern is constrained by (unknown) covariance relationships between different spatial regions. The model allows for n_{lag} lagged correlation coefficients ℓ_j , calculated separately for each lagged region.

We assume the forced response $F(t) = 0$ in the pre-industrial period. This neglects volcanic and solar forcing known to have been present and influencing climate prior to 1850 (e.g. (Schmidt et al., 2011; Schurer et al., 2013; Lücke et al., 2023; Jungclaus et al., 2017)). However, this has the effect of inflating the estimated covariance parameters, and therefore may render our subsequent detection analysis more conservative.

The parameters in this model are $\Theta = (\ell_j, \Sigma)$, where ℓ_j are the $n_{lag} \times r$ lag coefficients and Σ the $r \times r$ covariance matrix. By fitting the Cholesky decomposition of the covariance matrix

$$\Sigma = LL^T, \quad (3)$$

where L is a lower-triangular matrix, we can reduce the number of parameters in the covariance matrix to $r(r-1)/2$. The model (Eq. 2) specifies the likelihood of observing the data $\mathbf{D}(t)$ given values of these parameters:

$$P(\mathbf{D}(t)|\Theta) = (2\pi)^{-r/2} \det(\Sigma)^{-1/2} \exp\left(-\frac{1}{2} [\mathbf{D}(t) - \mu(t)]^T \Sigma^{-1} [\mathbf{D}(t) - \mu(t)]\right) \quad (4)$$

where μ is given by Eq. 2.3.

Now, we must specify prior beliefs $P(\Theta)$ about these parameters. Adopting a lag-2 model ($n_{lag} = 2$), we place Gaussian priors on each lag coefficient:

$$\ell_j \sim N(0, 1).$$

We use the Lewandowski-Kurowicka-Joe (LKJ, (Lewandowski et al., 2009)) prior for the spatial correlation matrix. Combined with priors on the standard deviations (which we set as Exponential(1.0), this yields a prior for the Cholesky matrix L (from which we can recover the full covariance matrix Σ). We can then use Markov Chain Monte Carlo (MCMC) sampling to estimate the posterior distributions for all parameters (Abril-Pla et al., 2023). These are presented in Sections 3.1 and 3.2.

2.4 Modeling recent variability

We consider two different models for recent (post-1850) PDSI variability in the GEDA.

- Model A, in which the recent variability is identical to pre-industrial variability and there is no forced response.
- Model B, in which recent PDSI variability is modeled as pre-industrial variability plus a nonzero, time-dependent forced response $\mathbf{F}(t)$ that differs across each region.

Model A is as described in Section 2.3. In Model B, the data at time t is:

$$\mathbf{D}(t) \sim MN(\mu_F(t), \Sigma) \quad (5)$$

where

$$\mu(t) = \mathbf{F}(t) + \sum_{j=1}^{n_{lag}} \ell_j \mathbf{D}(t-j) \quad (6)$$

and the covariance matrix Σ and the lagged coefficients ℓ_j are as in Eq. 2.3.

We now require a model for the forced response $\mathbf{F}(t)$ in each region over time. Here, we use

$$\mathbf{F}(t) = \beta T(t)$$

where $T(t)$ is the global mean temperature anomaly relative to the 1850-1900 average. β is a vector of scaling constants which are assumed to differ regionally: rising global temperatures may make some regions wetter, some drier, and have no effect on others.

2.5 Hierarchical modeling: incorporating uncertainty in ΔT

The global temperature anomaly $T(t)$ is well-constrained but not precisely known. There is substantial agreement among multiple datasets (e.g. HadCRUT (Morice et al., 2021), Berkeley Earth (Rohde & Hausfather, 2020), and GISTEMP (Lenssen et al., 2019), Figure 2(a)), but they do not match one another exactly. Moreover, the uncertainty in T depends on time: temperatures earlier in the post-industrial period are less well-measured than more recent anomalies. While we expect the uncertainty in T to be a minor component of our analysis, we still would like our results to incorporate the fact that we do not *exactly* know the global mean temperature anomaly.

One of the major advantages to a Bayesian framework is that it is relatively simple to incorporate and propagate uncertainties through a hierarchy of sub-models. Here, we use a random-effects model (see, e.g. (Gronau et al., 2021)) to estimate the “true” global mean temperature anomaly from three observational datasets and their reported uncertainties. We assume the reported temperature anomaly time series from dataset k , denoted \hat{T}_k , differs from the (latent) true temperature anomaly T_k for that dataset, and that all dataset anomalies T_k are drawn from a normal distribution whose mean is the underlying *real-world* temperature anomaly T and whose spread is controlled by an inter-dataset homogeneity parameter τ . In the special case where $\tau = 0$, this reduces to a “fixed effect” model, in which all datasets are assumed to differ only because of sampling error. If τ is allowed to be positive definite, then this becomes a “random effects” model, in which uncertainty due to possible inhomogeneity between datasets is taken into account. Here, we use such a random effects model, which can be written as

$$\begin{aligned}\hat{T}_k &\sim MN(T_k, \Sigma_k) \\ T_k &\sim N(T, \tau) \\ T &\sim g(.) \\ \tau &\sim h(.)\end{aligned}$$

where $g(.)$ and $h(.)$ are priors on the true real-world temperature anomaly T and the inter-dataset spread τ , respectively, which we set to $N(0, 10)$ and $HalfNormal(10)$. The dataset covariance matrices are $\Sigma_k = \text{diag}(\sigma_1^2 \dots \sigma_t^2)$, where σ_t is the reported standard deviation at time t .

Figure 2(b) shows the resulting 95% highest posterior density interval for T . This is the (uncertain) real-world temperature anomaly upon which our assumed forcing βT depends. By incorporating this sub-model within a Bayesian hierarchical structure, we can easily take into account the uncertainty in the global temperature anomaly and propagate this uncertainty through our results. The inter-dataset spread parameter τ is small relative to the rise in global average temperatures (Figure 2 c), reflecting the high degree of agreement between datasets.

2.6 Detecting the influence of global warming

In frequentist detection and attribution, a “fingerprint” (Hegerl et al., 1996) of the expected response to external forcing is generally multiplied by a scaling factor β (e.g.




Figure 2. Estimating the real-world temperature response and its uncertainties.

(a): Global mean temperature relative to 1850-1900 (1880-1900 for GISTEMP) in three observed datasets. Shading represents the reported 95% confidence intervals. (b): Posterior distribution for the “true”, real-world temperature anomaly T . Shown is the 95% highest posterior density interval. (c): Posterior for the inter-dataset spread parameter τ .

(N. P. Gillett et al., 2021)). The goal of the analysis is to calculate the true underlying value of the scaling parameter β and its uncertainty. If β is shown to be incompatible with 0 in a statistical sense, the fingerprint it multiplies is said to have been “detected”. If β is compatible with 1, the observations are said to be attributable to external forcing.

From a Bayesian perspective, there is no such thing as a true value of β . The scaling parameter is just that: a *parameter* in our model about which we hold some prior beliefs based on previous information. Given the evidence, we can update these priors to arrive at a posterior that expresses our confidence in the possible range of β . Hence, we do not base claims of detection or attribution on the value of β .

Moreover, the detection of any external influences is complicated by the temporal structure of pre-industrial variability. In Model B, the scaling parameter multiplies the global mean temperature change, and $\beta T(t)$ is an addition to the expectation value of the PDSI $\mathbf{D}(t)$ at every time step. But if the PDSI in any given year depends on the PDSI in the previous year (or before), then a small wetting or drying arising due to random chance will make the next year more likely to be wet or dry, which will in turn affect the next year, and so on. We must identify the extent to which a persistent trend can be explained by an external driver as opposed to the natural “memory” of the system, as reflected in the temporal autocorrelation.

Instead, we consider two explanatory models for 1850-2020 PDSI variability in the GEDA (Figure 3). In Model A, recent variability is explained by natural variability, as parameterized by $\Theta_A = (\ell_1, \ell_2, \Sigma)$ inferred from the pre-industrial (1000-1849) GEDA.

In Model B, recent variability is explained by this pre-industrial variability plus a forced response that depends on the (uncertain) global mean temperature T , itself estimated from multiple observational datasets with spread τ . Model B therefore has more parameters than Model A: $\Theta_B = (\ell_1, \ell_2, \Sigma, \beta, T, \tau)$.

In statistical modeling, we balance two competing imperatives. On one hand, we want to avoid over-fitting with too many parameters. On the other, we want a model that explains the data well. This means adding parameters to a model is “worth it” only if those parameters have additional explanatory power. In our analysis, detection is a question of model comparison. Does Model B, in which recent variability is explained




Figure 3. Comparing two models of recent PDSI variability. Summary graphs of Model A, in which recent PDSI is assumed to be explained by pre-industrial variability, and Model B, in which it is explained by pre-industrial variability plus a forcing term that depends on the global mean temperature anomaly T . Model A is parameterized by the temporal lag coefficients ℓ and the Cholesky decomposition L of the spatial covariance matrix Σ . Model B is a hierarchical model, in which the global mean temperature T is estimated from three observational datasets with spread τ and the forced response is βT . Variables labeled “Deterministic” are functions of random variables estimated by the models. Shaded ovals are the observed data (GEDA and the global temperature datasets). Because GISTEMP begins in 1880 while HadCRUT and Berkeley Earth begin in 1850, we model 1850–1880 GISTEMP as unobserved values.




Figure 4. Year-to-year persistence in PDSI. (a) Lag-1 coefficients (posterior mean of ℓ_1) for each region. (b): As in (a), but for lag-2 coefficients ℓ_2

by pre-industrial variability plus a temperature-dependent response, fit the data better than Model A, in which it is explained by pre-industrial variability alone? And to what extent?

To answer these questions, we use posterior predictive distributions (PPDs), which allow us to predict out-of-sample data using the posterior distributions for the parameters of each model (Gelman et al., 1995). If $\mathbf{D}(t)$ is the PDSI in the r regions at time t and the PDSI at previous times $\mathbf{D}(t-1), \mathbf{D}(t-2) \dots \mathbf{D}(t=0)$ are known, then

$$P(\mathbf{D}(t)|\mathbf{D}(t-1), \mathbf{D}(t-2) \dots \mathbf{D}(t=0)) = \int P(\Theta|\mathbf{D})P(\Theta)d\Theta. \quad (7)$$

The posterior predictive distribution depends on the parameters Θ , which are set by the model. To compare Model A and Model B, we draw samples from the posteriors for each model $P(\Theta_A|\mathbf{D})$ and $P(\Theta_B|\mathbf{D})$ and use them to “predict” the PDSI in each recent year as if we had never seen it before. Comparing the PPD for the no-forcing model to PPD for the model with a temperature-dependent term allows us to calculate how regional PDSI trends differ, and hence to “attribute” observed trends to natural variability or regional forcing. Essentially, we are asking: is it “worth it”, in terms of predictive power, to include the influence of global warming? Using this framework, we can then quantify the extent to which global mean temperature change influences regional PDSI while taking into account the natural persistence of the system.

3 Results

3.1 Temporal autocorrelation in reconstructed PDSI

Figure 4 shows the posterior mean lag-1 (ℓ_1) and 2 (ℓ_2) coefficients for each region. There is substantial one-year “memory” in each region, with the lag-1 autocorrelation largest in Siberia and smallest in western central Asia. Posteriors for the lag-2 autocorrelation in many regions are not strongly shifted away from zero, indicating weak or no dependence of PDSI on its value two years before. However, in the Arabian Peninsula, West Central Asia, and East Asia, over 98% of the posterior density for ℓ_2 is greater than zero, suggesting that PDSI in these regions is correlated with its value two years before. In western Siberia and south-east Asia, the PDSI in year t appears to be anti-correlated with PDSI two years prior.

3.2 Spatial covariance in reconstructed PDSI

Figure 5 shows the posterior mean of the spatial covariance matrix Σ . For visual clarity, we have excluded terms on the diagonal matrix: that is, we do not show the variance of PDSI in each region. Because the PDSI has been standardized, in the absence




Figure 5. The spatial covariance structure of pre-industrial variability. Posterior mean covariance matrix Σ for the drought atlas data. Redder colors indicate the PDSI in two regions co-varies with one another, while bluer colors indicate the PDSI in two regions is anti-correlated. Also shown are the posterior distributions for the covariance between Northern European PDSI and all other regions.

of temporal autocorrelation these terms would be equal to 1. The larger the autocorrelation, the smaller the diagonal term in the covariance matrix, since more variability is explained by PDSI in prior years. For example, the fact that PDSI in Northern Europe in any given year is positively correlated with PDSI in the year before means that the non-lagged variance is estimated to be less than unity (top left distribution, Figure 5). The posterior for Σ represents the spatial covariance structure between regions. For example, if PDSI in Northern Europe decreases, PDSI in West Central Europe does too, while PDSI in the Mediterranean increases. This reflects the well-understood hydroclimate response to the North Atlantic Oscillation (NAO) (E. R. Cook et al., 2015).

To compare our results with more standard methods of covariance estimation, we calculate the eigenvector of Σ (posterior mean, shown in Figure 6(b)) associated with the largest eigenvalue. We also calculate the leading EOF (EOF1) of the preindustrial GEDA (Figure 6(b)). The eigenvector for the posterior mean Σ resembles EOF1 in many regions: the covariance between European regions is particularly strong in both. Differences in sign or magnitude are likely related to the fact that Σ is estimated from a method that takes temporal covariance into account, whereas EOF1 does not. This is one advantage of our Bayesian approach; other perks include a full estimation of uncertainties in the covariance matrix, as well as avoiding the arbitrary truncation in representing the covariance matrix with a smaller number of EOFs.




Figure 6. Comparing methods of covariance estimation. (a): Leading eigenvector of the posterior mean covariance matrix Σ . (b): EOF1 calculated from 1000-1849 drought atlas data.




Figure 7. The sign of PDSI change with global warming. “Fingerprint” of regional PDSI response to global temperature rise, defined as the posterior mean of the parameter β . When temperature rises, the model predicts

3.3 Fingerprints of temperature increase

The posterior mean for the regional scaling parameters β is shown in Figure 7. Here, β represents the estimated sign and magnitude of any regional PDSI change that scales with global mean temperature, and can be thought of as the calculated “fingerprint” of global warming on regional PDSI. According to this model, northern Europe, northern Europe, Tibet, East Central Asia get wetter as the planet warms; Eastern Europe, Arctic Russia, the Arabian Peninsula and the Mediterranean get drier, and changes are smaller in other regions.

3.4 Comparing with preindustrial drought atlas variability

Temporally autocorrelated and spatially correlated variability is capable of explaining *some* wetting or drying trends. If a region is dry in any given year, it is more likely to be dry the next year, and so on. And long-term wetting or drying trends in some regions are associated with trends in other regions because of teleconnections arising from known modes of variability. Natural variability is not pure white noise, in which long-term trends would be extremely unlikely; we expect to see (and, indeed, we do see, in




Figure 8. What difference does global warming make? This plot shows the mean difference (as a function of time) between the posterior predictive distributions for the Global T model, in which drought responses are assumed to depend on T , and AR2, a model in which they are represented by preindustrial variability alone.




Figure 9. How well do different statistical models explain 21st century PDSI?

Light blue distributions show the posterior predictive distribution for regional 2000-2020 mean PDSI assuming it is explained by natural variability inferred from the 1000-1849 drought atlas. Dark blue distributions show the PPD for regional 2000-2020 mean PDSI assuming it is explained by natural variability plus a global temperature-dependent forced response. Black lines indicate quartiles. Orange dots represent the 2000-2020 mean PDSI in the GEDA.

the preindustrial GEDA) multi-decadal trends in PDSI even in the absence of external forcing. The attribution question is then: to what extent does adding a temperature-dependent forcing to this complex natural variability increase a model's explanatory power?

Figure 8 shows the mean difference between the posterior predictive distribution for Model B (which incorporates a the global temperature response) and the PPD for Model A (in which recent variability is modeled as pre-industrial variability) as a function of time. This represents the mean wetting or drying explained by the inclusion of a T -dependent forced response relative to the wetting or drying that can be explained by natural variability (as inferred from the preindustrial GEDA) alone.

Figure 8 does not incorporate the uncertainty, a crucial step for confident detection or attribution. To illustrate the full posteriors, we compare twenty-first century (2000-2020) mean regional PDSI in both models. The light blue distributions in Figure 9 show the PPD for 21st century PDSI assuming Model A. These reflect the ability of natural variability (as inferred from the preindustrial GEDA) to explain 21st century mean PDSI anomalies. Consider, for example, Eastern Europe (EEU). Pre-industrial variability alone can explain a dry anomaly of a certain magnitude; three-quarters of the PPD mass lies below zero. However, the observed twenty-first century EEU PDSI (orange dot) lies in the tail of the light-blue PPD, indicating that such a large dry anomaly is difficult to explain with natural variability alone. The dark blue distributions in Figure 9 show the PPD for 21st century PDSI assuming Model B. The 21st century EEU anomaly lies near the center of the PPD for Model B, indicating that a temperature-dependent forced response is useful for explaining the observed PDSI.

By contrast, both Model A and Model B appear to be about equally as able to capture the 21st century mean PDSI in East Asia (EAS), indicating that an additional temperature-dependent forced response is not necessarily required to explain the dry PDSI in this region.




Figure 10. Attributing twenty-first century PDSI to global warming. The mean difference between the posterior predictive distribution assuming a temperature-dependent forced response and the PPD assuming natural variability for 2000-2020 mean regional PDSI. Boxes show the quartiles, while whiskers show the “likely” (13-83%) range.

3.5 Where have rising global temperatures likely affected drought?

We summarize the attribution analysis in Figure 10. The box-and-whisker plots show the difference between the Model B PPD and the Model A PPD for 21st century mean regional PDSI. The observed 21st century PDSI for Eastern Europe appears to lie directly at the center of the relevant box, indicating that a temperature-dependent response explains essentially all of the recent drying in this region. The IPCC defines “likely” as within the 66% confidence interval; in our Bayesian framework we will define a “likely” contribution from global warming as one in which the 66% highest-posterior density interval excludes zero. Using this terminology, we assess that global warming likely contributed to dry PDSI in Eastern Europe, the Mediterranean, and Arctic Russia and to wet PDSI in Northern Europe, East-central Asia, and Tibet.

In most regions, the inferred contribution from the temperature-dependent forced response (or at least, the posterior mean) is of the same sign as the observed 21st century mean PDSI. The exceptions are Southeast Asia (SEA), West Central Asia (WCE), and the Russian Far East (RFE), indicating that natural variability is *more* able to explain the observed PDSI than the inferred T -dependent response.

4 Discussion and Conclusions

All detection and attribution studies are model-dependent, and ours is no exception. Although we do not rely on coupled atmosphere-ocean general circulation models, we use simple models to interpret and characterize pre-industrial variability, to estimate the global mean temperature from multiple datasets, and to explain recent PDSI variations. We treat detection and attribution in a unified framework of model *comparison*: which of these models best explains the observed data? Our results suggest that a temperature-dependent forcing term better explains recent variability in many regions than pre-industrial variability, at least as characterized by our spatiotemporal model. Thus, we conclude that global warming is likely making eastern and southern Europe drier, while it is making northern Europe and parts of Asia wetter. This result is contingent on the two models

we compare: it may be that some other model is better able to both characterize pre-industrial variability and explain recent trends. Still, we can be confident in stating that given a choice between pre-industrial variability alone and variability added to the influence of global warming, twenty-first century PDSI in many regions is best explained by the latter.

The flexibility of Bayesian methods opens up the possibilities of many future analyses. The number of sub-models in a Bayesian hierarchy is unlimited, which allows for attribution on multiple levels. For example, one might further model the global mean temperature T as a response to natural and anthropogenic forcing agents, and trace the influence of anthropogenic forcing to regional PDSI via its impact on global mean temperature. Other, more complex models for the PDSI response are also possible: we might go beyond the global mean temperature to consider the effects of, for example, different SST patterns. Finally, the properties of reconstructed pre-industrial hydroclimate variability might be used to evaluate and constrain the output of GCMS, leading to more confident attribution and more coherent projections.

These results reinforce that regional drought risk is, to a certain extent, predictable. The year-to-year persistence in soil moisture is an important source of predictability even in the absence of anthropogenic forcing. We show that, in many regions, another, stronger source of predictability is already emerging: the rising global temperature. In the absence of drastic emission cuts, the planet will continue to warm, and this will become an even more important determinant of drought risk. Our statistical analysis highlights the urgent necessity to understand the underlying physical drivers shaping this relationship, as well as the need for action to adapt to altered drought risk in a warmer world.

5 Open Research

The Great Eurasian Drought Atlas is available at

<https://zenodo.org/records/11059894>.

Global mean temperature datasets and uncertainties may be downloaded at the following links:

- GISTEMP: <https://data.giss.nasa.gov/gistemp/uncertainty/>
- HadCRUT: <https://www.metoffice.gov.uk/hadobs/hadcrut5/data/HadCRUT5.0.2.0/download.html>
- Berkeley Earth: https://berkeley-earth-temperature.s3.us-west-1.amazonaws.com/Global/Land_and_Ocean_summary.txt

Analysis was performed with the PyMC probabilistic programming environment available at <https://www.pymc.io/>. Code to reproduce all figures and analyses is available at https://github.com/netzeroasap/GEDA_BAYES/.

Acknowledgments

K.M and B.I.C. were supported by the NASA Modeling, Analysis and Prediction program.

References

- Abril-Pla, O., Andreani, V., Carroll, C., Dong, L., Fonnesbeck, C. J., Kochurov, M., ... others (2023). Pymc: a modern, and comprehensive probabilistic programming framework in python. *PeerJ Computer Science*, 9, e1516.
- Allen, M., & Stott, P. (2003). Estimating signal amplitudes in optimal fingerprint-

- ing, part i: Theory. *Climate Dynamics*, 21(5-6), 477–491.
- Annan, J. (2010). Bayesian approaches to detection and attribution. *Wiley Interdisciplinary Reviews: Climate Change*, 1(4), 486–489.
- Baek, S. H., Smerdon, J. E., Coats, S., Williams, A. P., Cook, B. I., Cook, E. R., & Seager, R. (2017). Precipitation, temperature, and teleconnection signals across the combined north american, monsoon asia, and old world drought atlases. *Journal of Climate*, 30(18), 7141–7155.
- Berliner, L. M., Levine, R. A., & Shea, D. J. (2000). Bayesian climate change assessment. *Journal of Climate*, 13(21), 3805–3820.
- Bonfils, C. J., Anderson, G., Santer, B. D., Phillips, T. J., Taylor, K. E., Cuntz, M., ... others (2017). Competing influences of anthropogenic warming, enso, and plant physiology on future terrestrial aridity. *Journal of Climate*, 30(17), 6883–6904.
- Cook, B., Cook, E., KJ, A., & Singh, D. (2024). Characterizing the 2010 russian heatwave-pakistan flood concurrent extreme over the last millennium using the great eurasian drought atlas. *Journal of Climate*.
- Cook, B. I., Mankin, J. S., Marvel, K., Williams, A. P., Smerdon, J. E., & Anchukaitis, K. J. (2020). Twenty-first century drought projections in the cmip6 forcing scenarios. *Earth's Future*, 8(6), e2019EF001461.
- Cook, B. I., Smerdon, J. E., Cook, E. R., Williams, A. P., Anchukaitis, K. J., Mankin, J. S., ... others (2022). Megadroughts in the common era and the anthropocene. *Nature Reviews Earth & Environment*, 3(11), 741–757.
- Cook, E. R., Anchukaitis, K. J., Buckley, B. M., D'Arrigo, R. D., Jacoby, G. C., & Wright, W. E. (2010). Asian monsoon failure and megadrought during the last millennium. *Science*, 328(5977), 486–489.
- Cook, E. R., Seager, R., Kushnir, Y., Briffa, K. R., Büntgen, U., Frank, D., ... others (2015). Old world megadroughts and pluvials during the common era. *Science advances*, 1(10), e1500561.
- Cook, E. R., Solomina, O., Matkovsky, V., Cook, B. I., Agafonov, L., Berdnikova, A., ... others (2020). The european russia drought atlas (1400–2016 ce). *Climate Dynamics*, 54(3), 2317–2335.
- Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. (2016). Overview of the coupled model intercomparison project phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development*, 9(5), 1937–1958.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (1995). *Bayesian data analysis*. Chapman and Hall/CRC.
- Gillett, N., Zwiers, F., Weaver, A., Hegerl, G., Allen, M., & Stott, P. (2002). Detecting anthropogenic influence with a multi-model ensemble. *Geophysical Research Letters*, 29(20), 31–41.
- Gillett, N. P., Kirchmeier-Young, M., Ribes, A., Shiogama, H., Hegerl, G. C., Knutti, R., ... others (2021). Constraining human contributions to observed warming since the pre-industrial period. *Nature Climate Change*, 11(3), 207–212.
- Gronau, Q. F., Heck, D. W., Berkhout, S. W., Haaf, J. M., & Wagenmakers, E.-J. (2021). A primer on bayesian model-averaged meta-analysis. *Advances in Methods and Practices in Psychological Science*, 4(3), 25152459211031256.
- Harris, I., Osborn, T. J., Jones, P., & Lister, D. (2020). Version 4 of the cru ts monthly high-resolution gridded multivariate climate dataset. *Scientific data*, 7(1), 1–18.
- Hausfather, Z., Marvel, K., Schmidt, G. A., Nielsen-Gammon, J. W., & Zelinka, M. (2022). Climate simulations: Recognize the ‘hot model’ problem. *Nature*, 605(7908), 26–29.
- Hegerl, G. C., von Storch, H., Hasselmann, K., Santer, B. D., Cubasch, U., & Jones, P. D. (1996). Detecting greenhouse-gas-induced climate change with an

- optimal fingerprint method. *Journal of Climate*, 9(10), 2281–2306.
- Iturbide, M., Gutiérrez, J. M., Alves, L. M., Bedia, J., Cerezo-Mota, R., Gimenez, E., ... Vera, C. S. (2020). An update of ipcc climate reference regions for subcontinental analysis of climate model data: definition and aggregated datasets. *Earth System Science Data*, 12(4), 2959–2970. Retrieved from <https://essd.copernicus.org/articles/12/2959/2020/> doi: 10.5194/essd-12-2959-2020
- Jungclauss, J. H., Bard, E., Baroni, M., Braconnot, P., Cao, J., Chini, L. P., ... others (2017). The pmip4 contribution to cmip6—part 3: The last millennium, scientific objective, and experimental design for the pmip4 past1000 simulations. *Geoscientific Model Development*, 10(11), 4005–4033.
- Katzfuss, M., Hammerling, D., & Smith, R. L. (2017). A bayesian hierarchical model for climate change detection and attribution. *Geophysical Research Letters*, 44(11), 5720–5728.
- Kim, Y.-H., Min, S.-K., Zhang, X., Sillmann, J., & Sandstad, M. (2020). Evaluation of the cmip6 multi-model ensemble for climate extreme indices. *Weather and Climate Extremes*, 29, 100269.
- Lenzen, N. J., Schmidt, G. A., Hansen, J. E., Menne, M. J., Persin, A., Ruedy, R., & Zys, D. (2019). Improvements in the gistemp uncertainty model. *Journal of Geophysical Research: Atmospheres*, 124(12), 6307–6326.
- Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of multivariate analysis*, 100(9), 1989–2001.
- Lücke, L. J., Schurer, A. P., Toohey, M., Marshall, L. R., & Hegerl, G. C. (2023). The effect of uncertainties in natural forcing records on simulated temperature during the last millennium. *Climate of the Past*, 19(5), 959–978.
- Marvel, K., & Cook, B. I. (2022). Using machine learning to identify novel hydroclimate states. *Philosophical Transactions of the Royal Society A*, 380(2238), 20210287.
- Marvel, K., Cook, B. I., Bonfils, C., Smerdon, J. E., Williams, A. P., & Liu, H. (2021, aug). Projected changes to hydroclimate seasonality in the continental united states. *Earth's Future*. Retrieved from <https://doi.org/10.1029/2021ef002019> doi: 10.1029/2021ef002019
- Marvel, K., Cook, B. I., Bonfils, C. J., Durack, P. J., Smerdon, J. E., & Williams, A. P. (2019). Twentieth-century hydroclimate changes consistent with human influence. *Nature*, 569(7754), 59–65.
- Morice, C. P., Kennedy, J. J., Rayner, N. A., Winn, J., Hogan, E., Killick, R., ... Simpson, I. (2021). An updated assessment of near-surface temperature change from 1850: the hadcrut5 data set. *Journal of Geophysical Research: Atmospheres*, 126(3), e2019JD032361.
- Richter, I., & Tokinaga, H. (2020). An overview of the performance of cmip6 models in the tropical atlantic: mean state, variability, and remote impacts. *Climate Dynamics*, 55(9), 2579–2601.
- Rohde, R. A., & Hausfather, Z. (2020). The berkeley earth land/ocean temperature record. *Earth System Science Data*, 12(4), 3469–3479. Retrieved from <https://essd.copernicus.org/articles/12/3469/2020/> doi: 10.5194/essd-12-3469-2020
- Santer, B. D., Mears, C., Doutriaux, C., Caldwell, P., Gleckler, P., Wigley, T., ... others (2011). Separating signal and noise in atmospheric temperature changes: The importance of timescale. *Journal of Geophysical Research: Atmospheres* (1984–2012), 116(D22).
- Santer, B. D., Painter, J. F., Bonfils, C., Mears, C. A., Solomon, S., Wigley, T. M., ... others (2013). Human and natural influences on the changing thermal structure of the atmosphere. *Proceedings of the National Academy of Sciences*, 201305332.

- Santer, B. D., Painter, J. F., Mears, C. A., Doutriaux, C., Caldwell, P., Arblaster, J. M., ... others (2013). Identifying human influences on atmospheric temperature. *Proceedings of the National Academy of Sciences*, 110(1), 26–33.
- Schmidt, G. A., Jungclaus, J. H., Ammann, C., Bard, E., Braconnot, P., Crowley, T., ... others (2011). Climate forcing reconstructions for use in pmip simulations of the last millennium (v1. 0). *Geoscientific Model Development*, 4(1), 33–45.
- Schurer, A. P., Hegerl, G. C., Mann, M. E., Tett, S. F., & Phipps, S. J. (2013). Separating forced from chaotic climate variability over the past millennium. *Journal of Climate*, 26(18), 6954–6973.
- Seneviratne, S. X. Z., Adnan, M., Badi, W., Dereczynski, C., Luca, A. D., Ghosh, S., ... B. Zhou, . (2023, July). Weather and climate extreme events in a changing climate. In *Climate change 2021 – the physical science basis* (pp. 1513–1766). Cambridge University Press. Retrieved from <http://dx.doi.org/10.1017/9781009157896.013> doi: 10.1017/9781009157896.013
- Sherwood, S., Webb, M. J., Annan, J. D., Armour, K., Forster, P. M., Hargreaves, J. C., ... others (2020). An assessment of earth’s climate sensitivity using multiple lines of evidence. *Reviews of Geophysics*, 58(4), e2019RG000678.
- Stott, P. A., Tett, S., Jones, G., Allen, M., Mitchell, J., & Jenkins, G. (2000). External control of 20th century temperature by natural and anthropogenic forcings. *Science*, 290(5499), 2133–2137.
- Swain, D. L., Tsiang, M., Haugen, M., Singh, D., Charland, A., Rajaratnam, B., & Diffenbaugh, N. S. (2014). The extraordinary california drought of 2013/2014: Character, context, and the role of climate change. *Bulletin of the American Meteorological Society*, 95(9), S3.
- Tett, S. F., Jones, G. S., Stott, P. A., Hill, D. C., Mitchell, J. F., Allen, M. R., ... others (2002). Estimation of natural and anthropogenic contributions to twentieth century temperature change. *Journal of Geophysical Research*, 107(D16), 4306.
- Tierney, J. E., Zhu, J., King, J., Malevich, S. B., Hakim, G. J., & Poulsen, C. J. (2020). Glacial cooling and climate sensitivity revisited. *Nature*, 584(7822), 569–573.
- Tokarska, K. B., Stolpe, M. B., Sippel, S., Fischer, E. M., Smith, C. J., Lehner, F., & Knutti, R. (2020). Past warming trend constrains future warming in cmip6 models. *Science Advances*, 6(12), eaaz9549.
- van der Schrier, G., Barichivich, J., Briffa, K., & Jones, P. (2013). A scpsdi-based global data set of dry and wet spells for 1901–2009. *Journal of Geophysical Research: Atmospheres*, 118(10), 4025–4048.
- Wells, N., Goddard, S., & Hayes, M. J. (2004). A self-calibrating palmer drought severity index. *Journal of climate*, 17(12), 2335–2351.
- Williams, A. P., Cook, E. R., Smerdon, J. E., Cook, B. I., Abatzoglou, J. T., Bolles, K., ... Livneh, B. (2020). Large contribution from anthropogenic warming to an emerging north american megadrought. *Science*, 368(6488), 314–318.
- Williams, A. P., Seager, R., Abatzoglou, J. T., Cook, B. I., Smerdon, J. E., & Cook, E. R. (2015). Contribution of anthropogenic warming to california drought during 2012–2014. *Geophysical Research Letters*, 42(16), 6819–6828.
- Yazdandoost, F., Moradian, S., Izadi, A., & Aghakouchak, A. (2021). Evaluation of cmip6 precipitation simulations across different climatic zones: Uncertainty and model intercomparison. *Atmospheric Research*, 250, 105369.