

On the Attribution of Weather Events to Climate Change Using Empirically Fit Extreme Value Distributions

PETER SHERMAN¹, PETER HUYBERS^{1,a} AND ELI TZIPERMAN^{1,a,b}

^a *School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts*

^b *Department of Earth and Planetary Sciences, Harvard University, Cambridge, Massachusetts*

(Manuscript received 7 September 2023, in final form 14 February 2025, accepted 4 March 2025)

ABSTRACT: Changes in extreme weather events are a potentially important aspect of anthropogenic climate change (ACC), yet they are difficult to attribute to ACC because the record length is often similar to, or shorter than, extreme-event return periods. This study is motivated by the “World Weather Attribution” (WWA) initiative and, specifically, their approach of fitting extreme value distribution functions to local observations. They calculate the dependence of distribution parameters on global mean surface temperature (GMST) and use this dependence to attribute extreme events to ACC. Applying this method to preindustrial climate simulations with no time-varying greenhouse gas forcing, we still find a strong dependence of distribution parameters on GMST. This dependence results from internal climate variability (e.g., ENSO) affecting both extreme events and GMST. Therefore, dependence on GMST does not necessarily imply an effect of ACC on extremes. We further consider whether an extreme value, normal, or lognormal distribution better represents the data, if a GMST dependence of distribution parameters is justified using a likelihood ratio test, and if a meaningful attribution is possible given uncertainties in GMST dependence. We find, for example, that an attribution of Australia’s 2020–21 bushfires to ACC is difficult due to the effects of internal variability. For the 2019–21 drought in Madagascar, we find that the small number of available data points precludes a meaningful attribution analysis. Overall, we find that the effects of internal climate variability on GMST and the uncertain relationship between GMST and regional extremes may lead to inaccurate attribution conclusions using the part of the WWA approach examined here.

KEYWORDS: Extreme events; Precipitation; Climate change; Temperature; Statistics

1. Introduction

The Intergovernmental Panel on Climate Change (Eyring et al. 2021; IPCC 2023, 2021) found that global mean surface temperature (GMST) between 2011 and 2020 was 1.09°C (0.95–1.20°C) above preindustrial (1850–1900) levels. Although climate change is clearly manifested in GMST warming, and there are reasons to expect extreme events to change with mean climate, it is a complex task to attribute individual extreme weather events to anthropogenic climate change (ACC). Extreme events—i.e., heat waves, droughts, and heavy precipitation—are rare by definition, so the observational record may be too short to detect anthropogenic signals. This issue is especially pervasive in lower-income countries where data are sparse, and the historical record is short (Otto et al. 2020). The situation is further problematic for extreme precipitation event attribution, which is highly sensitive to the specific time scales and percentiles used to define extreme precipitation events (Pendergrass 2018).

The field of attribution of extreme events has progressed significantly over the past decade and encompasses several approaches (see example reviews and cited references within,

Stott et al. 2010; Hegerl and Zwiers 2011; National Academies of Sciences 2016; Clarke et al. 2022; Philip et al. 2020). The approaches that have been applied include the fraction of attributable risk (Allen 2003; Stott et al. 2004), using simulations from climate models, using constraints derived from optimal fingerprinting techniques (Christidis et al. 2015), using climate models to validate observed probability density functions (PDFs; Perkins et al. 2007; Lobeto et al. 2021), and the “analog-based approach” (Lorenz 1969; Yiou et al. 2007; Vautard and Yiou 2009; Faranda et al. 2022). Another methodology fits statistical distributions to the historical record (again, either observed or from climate models) in what Stott et al. (2016) refer to as “empirical-based approaches.” This approach assesses how the return period of particular events has changed over the historical record, with the purpose of detecting potential trends in the frequency of extreme events. These statistical approaches have been used to study trends in extreme events such as flooding in Thailand in 2011 (van Oldenborgh et al. 2012). Approaches using climate models allow for a better estimate of the magnitude and direction of human-induced contributions to an extreme event than is possible using observations alone.

The climate model approaches detailed in the above paragraph can also be performed with single-model initial-condition large ensembles (SMILES), composed of a climate model run several times with slightly different initial conditions. Averaging over multiple climate model ensemble members can reduce the signal from modes of internal variability, which are typically uncorrelated between ensemble members, facilitating the identification of forced climate change signals. This averaging generally improves the detection of trends in the magnitude or frequency

Supplemental information related to this paper is available at the Journals Online website: <https://doi.org/10.1175/JCLI-D-23-0542.s1>.

Corresponding author: Peter Sherman, petersherman@g.harvard.edu

of an extreme event, which are difficult to quantify in a single simulation or in the observational record where the signal from internal variability may be dominant (Yoshida et al. 2017; Yamaguchi et al. 2020). These simulations have been used in the context of tropical cyclone slowdown (Yamaguchi et al. 2020), as well as drought events in South and East Africa (Lott et al. 2013; Pascale et al. 2020). The attribution literature is extensive, and for further detailed discussion, we refer the reader to the above reviews and to recent studies (Stott et al. 2016; Knutson et al. 2017; Seneviratne et al. 2021).

One particular approach that is the focus of this paper and that has been used in recent years by World Weather Attribution (WWA) as a rapid response to extreme events attempts to fit the observed data to an extreme value PDF (The World Weather Attribution Project 2024). The fit allows for a dependence of the distribution parameters on explanatory variables, typically the GMST (e.g., Ciavarella et al. 2021; van Oldenborgh et al. 2021; Harrington et al. 2022; Philip et al. 2020). A GMST dependence is then interpreted as a response of extreme events to ACC. This dependence is then used to attribute extreme events by quantifying the effects of anthropogenic forcing on the probability of their occurrence. Technically, this is achieved by writing the location parameter of the extreme event distribution μ as a function of $\bar{T}$, the (low-pass filtered, via a 4-yr running average) GMST “as a measure of anthropogenic climate change” (Harrington et al. 2022), such that $\mu = \mu_0 + \alpha\bar{T}$. The GMST dependence α is estimated via a fit to observations. The fitted PDF is then used to evaluate the probability of extreme events at 1900 $P[\mu_0 + \alpha\bar{T}(1900)]$ relative to those at 2020 $P[\mu_0 + \alpha\bar{T}(2020)]$. This allows for estimating how the probability of an extreme event changed between these times. A change in the probability of an extreme event is inferred by the WWA approach if α is nonzero, based on the change in GMST.

Although WWA does not carry out formal hypothesis testing, we interpret that they adopt a null hypothesis of α equal to zero, implying no ACC influence on extremes, and an alternative hypothesis that, if α is not equal to zero, ACC does influence extremes. This issue is discussed in detail in section 2f and further addressed throughout the text.

Our focus is on the empirically based approaches used by WWA, but it should be noted that WWA also requires evidence from climate model simulations before making any attribution statements. WWA often uses multiple climate models and is careful to exclude models based on their performance and return periods relative to observations (e.g., Pinto et al. 2023). In this study, we use a single model (CESM) to demonstrate some of our points but focus on the assessment of observational methodologies associated with fitting of extreme value distributions to observations. Our goal is limited to evaluating the observational approach described by WWA and suggesting additional tests to examine its robustness, as opposed to developing alternative attribution approaches.

Motivated by these pioneering efforts based on the fit to extreme event PDFs, we reexamine three test cases analyzed by WWA: the Siberian heat wave analysis of Ciavarella et al. (2021, hereafter CC2021), the analysis of high temperatures associated with the Australian bushfire by van Oldenborgh

et al. (2021, hereafter OK2021), and the Madagascar drought analysis of Harrington et al. (2022, hereafter HW2022). These cases were chosen as they represent a range of applications, including the attribution of heat waves at high and low latitudes (Siberia vs Australia) and of an extreme drought. Each of these studies leverages either a generalized extreme value (GEV) or generalized Pareto distribution (GPD) fit to assess potential contributions to extreme events from ACC. WWA studies use both climate model simulations and observational analysis, and we focus on the observational component. In each of the three cases, we use the CESM Large Ensemble (Kay et al. 2015) to put the observed instrumental record into perspective, and we explore potential caveats concerning the methods examined and suggest how these limitations might be addressed.

The issue highlighted by our analysis is simple: internal variability modes such as ENSO affect both local temperature extremes and GMST. In some cases, the effect on local extremes (say the effects of ENSO on heat waves in Australia) may be large, but the effect on GMST is small. This might lead to a false conclusion that a small GMST signal leads to a comparatively large effect on local extremes. To demonstrate this point, we analyze a preindustrial climate model run with no anthropogenic forcing. We find, for example, that the fit of Australian heat extremes leads to values of α that are substantially greater than zero in the absence of ACC because of the effects of internal climate variability modes such as ENSO on both GMST and on local extremes. Our findings show that α can be strongly affected by internal variability, biasing the attribution results and complicating the WWA approach.

Our study addresses four questions: 1) Is a GMST dependence of the PDF parameters necessarily an indication of the signal of ACC? 2) Is a GEV or GPD necessary to fit the available observations, or are more standard normal or lognormal distributions sufficient? 3) Is the addition of GMST-dependent distribution parameters, which form the basis for the attribution analysis, justified statistically, and is the uncertainty range of these parameters, specifically α , sufficiently small to allow for attribution? 4) What do SMILEs tell us about the amount of data needed to determine the GMST dependence with confidence? WWA uses multiple attribution tools, including analysis of simulations from multiple models and observations (Philip et al. 2020). We focus only on the latter, specifically, on the fitting of extreme value distributions to observations. Our goals are limited to evaluating this observational approach and suggesting additional tests to examine its robustness and do not include developing alternative attribution approaches.

The rest of the paper is structured with section 2 describing the statistical models, observation datasets, and methodologies. Section 3a shows that a GMST dependence of distribution parameters does not necessarily represent the effects of ACC using an analysis of a CESM preindustrial simulation (i.e., without time-varying anthropogenic and natural forcings). We then further examine the three attribution cases in the remainder of section 3. Section 4 provides a summary and discussion of the results.

2. Methods and data

In this section, we briefly describe the extreme value distributions used in the following sections 2a and 2b and discuss uncertainty estimates and empirical cumulative distribution functions (CDFs) that we will use later (section 2c). We describe the deviance statistic test which we use to examine the justification for adding GMST-dependent parameters (section 2d) and the mean residual life plot used to test the justification for using a GPD and its threshold value (section 2e). The data used are described in section 2g.

a. The GEV distribution

Extreme value distributions are a family of PDFs that can be used to represent the statistical behavior of block maxima or minima of a record. An example is an annual time series of the maximum daily temperature at each year, the “block” being 1 year in this case. It can be shown that the distribution of such maxima follows one of three classes of PDFs: Gumbel, Fréchet, or Weibull (Coles et al. 2001), depending on the shape of the tail of the distribution of the events for which block maxima are calculated. Rather than explicitly specifying one of these three classes in fitting data for extreme events, one can combine these distributions into one functional form, represented by the GEV distribution. Large quantities of data are typically needed to accurately fit a GEV and calculate its parameters (e.g., Philip et al. 2020; Trevino et al. 2020).

The PDF of a variable x whose statistics are governed by the GEV is given by

$$\text{GEV}(x, \mu, \sigma, \xi) = \frac{1}{\sigma} t(x, \mu, \sigma, \xi)^{\xi+1} e^{-t(x, \mu, \sigma, \xi)}, \quad (1)$$

where

$$t(x, \mu, \sigma, \xi) = \left(1 + \xi \frac{x - \mu}{\sigma}\right)^{-1/\xi}, \quad (2)$$

under the assumption that $\xi \neq 0$. If $\xi = 0$, $t(x, \mu, \sigma, \xi)$ is defined as

$$t(x, \mu, \sigma, \xi) = e^{-(x - \mu/\sigma)}. \quad (3)$$

In some of the attribution studies we follow for this paper (specifically, CC2021; OK2021), the location parameter μ is assumed to vary linearly with the 4-yr smoothed GMST $\bar{T}$ as

$$\mu = \mu_0 + \alpha \bar{T}. \quad (4)$$

Estimation of the parameter α is the key to the attribution of extreme events. However, this determination may be limited due to the short observational record (Zwiers et al. 2011; De Paola et al. 2018). To estimate the mean μ_0 , scale parameter σ , the GMST dependence α , and the shape parameter ξ , we follow Coles et al. (2001) and WWA (Philip et al. 2020) and maximize the log likelihood of the data points x_i using a GEV distribution as follows:

$$l(\mu_0, \sigma, \xi, \alpha) = \sum_{i=1}^m \log[\text{GEV}(x_i, \mu_0, \sigma, \xi, \alpha)]. \quad (5)$$

We apply this approach, following CC2021 and OK2021, to observations and to data from the CESM Large Ensemble (Kay et al. 2015).

b. GPD

In some applications, one is interested only in points in a time series that exceed a specified threshold, whose statistics are represented by the generalized Pareto distribution (Coles et al. 2001). To study the extreme drought events in Madagascar using a record of precipitation, HW2022 fit a GPD to the bottom 20th percentile of precipitation, i.e., years associated with drought. This low tail was converted to a high tail by multiplying the precipitation data by -1 . Although this converts the lowest precipitation values to maxima, the transformed data have a hard maximum at zero. We note that the validity of using a GPD, in this case, is questionable because no such hard maximum is allowed by—or represented in the functional form of—the GPD.

The PDF of the GPD is given by

$$\text{GPD}(x, u, \sigma, \xi) = \frac{1}{\sigma} \left(1 + \xi \frac{x - u}{\sigma}\right)^{-(1/\xi)+1}. \quad (6)$$

HW2022 assumed that the threshold parameter, which they denote u (denoted μ by Philip et al. 2020) and scale parameter σ both depend exponentially on the 4-yr running-averaged GMST $\bar{T}$ such that

$$\begin{aligned} u &= u_0 \exp(\alpha \bar{T}/u_0), \\ \sigma &= \sigma_0 \exp(\alpha \bar{T}/u_0). \end{aligned} \quad (7)$$

The exponential fit was assumed by HW2022 to ensure that the scaling factor $\exp(\alpha \bar{T}/u_0)$ applied to both u_0 and σ_0 is strictly positive and maintains a fixed ratio between the mean and variance. We obtain the optimized GPD parameters (u_0, σ, ξ, α) by again maximizing the log likelihood [Eq. (4.10), Coles et al. 2001] as follows:

$$l(u, \sigma, \xi) = \sum_{i=1}^m \log[\text{GPD}(x_i, u, \sigma, \xi)], \quad (8)$$

using the trust-region constrained algorithm in the SciPy Python package (“trust-constr” method).

The GPD requires the bound $\sigma > 0$ and consistency relations of the form $x \geq u$ (for $\xi \geq 0$) and $u \geq x$ and $x \leq u - \sigma/\xi$ (for $\xi < 0$). As the optimization iteratively searches for the optimal parameter values, these constraints are occasionally violated, leading to complex numbers or not a number (NaN) in the likelihood. The GPD optimization is notably more prone to such failures and sensitive to the initial guesses of α, u_0, ξ , and σ than the GEV optimization. To avoid such search failures, we replace complex and NaN terms in the log-likelihood sum with large penalty terms whenever they occur due to constraints being violated. In addition, if the optimization fails, we rerun the optimization again with new randomly selected initial guesses until convergence to a solution that maximizes the log likelihood and satisfies the constraints is found. We note that related solutions

have been used to deal with GEV/GPD convergence issues by [Robin and Ribes \(2020\)](#).

c. Uncertainty estimates and empirical cumulative distribution functions

We follow WWA studies and use nonparametric bootstrapping to estimate uncertainty ranges for the optimized parameters. In particular, we use the uncertainty in α as one way to evaluate the validity of the attribution results. We also compare the appropriately fitted distributions (GEV and normal distributions for extreme temperatures or GPD and lognormal for droughts) to the empirical CDFs of the observations to assess the quality of the fit. The empirical CDF is calculated by sorting extreme values (either temperature or precipitation) from smallest to largest, calculating the cumulative sum, and normalizing to a maximum CDF value of one. We calculate error bars in the empirical CDFs based on bootstrap resampling with replacement of the observations. The empirical CDF is calculated for each sampling, and we then calculate the 5th–95th and 2.5th–97.5th percentile ranges from the resulting distribution at each value of the variable whose CDF is estimated. To estimate uncertainty in the parameters estimated from our GEV or GPD maximum likelihood fit, we similarly use bootstrap resampling. We sample with replacement the GMST and corresponding extreme temperature record 5000 times and optimize the parameters for each resampling. We then estimate the 90th and 95th percentiles from 5000 bootstrap samples. For the Madagascar precipitation case, we perform 5000 bootstrap resamples when analyzing the observations but 1000 bootstrap resamples when analyzing the CESM Large Ensemble due to the slow convergence and expensive computations using the GPD log-likelihood maximization.

d. Deviance statistics (likelihood ratio test)

In general, when fitting a model to data, if the number of fit parameters is increased, one expects a better fit because of the greater flexibility of the fitted model. Specifically, in the case considered here, adding a parameter α to include the GMST dependence of the location/threshold parameter will increase the maximum log likelihood of the data. The salient question is whether the improvement in the fit exceeds that expected simply on account of the larger number of model parameters. To evaluate the improvement of the fit, [Coles et al. \(2001, theorem 2.7\)](#) recommend the use of a deviance statistic also known as a likelihood ratio test. This test uses a measure D equal to twice the difference of the sum of log likelihoods of the model with k additional parameters minus that of the simpler model. Because more parameters lead to a better fit and thus higher log likelihood, D is nonnegative. D is then compared to χ_k^2 distribution to assess if k additional parameters significantly improve the likelihood. For example, with $k = 1$, the 95th percentile of χ_1^2 is equal to 3.8. If D is larger than this threshold, we conclude that adding an additional parameter significantly improves the fit. It can then be inferred that the added parameter—in our specific case, the GMST dependence α —meaningfully explains features of the data. On

the other hand, if the deviance statistic is insignificant, one concludes that the addition of another parameter to the statistical fit is not justified by the data and, therefore, that this parameter cannot be used to draw conclusions about the data. Specifically, if one finds that the addition of α is not justified, the implication is that the estimated value should not be used to calculate the effects of climate change on the extreme events under examination. We also estimate the distribution and uncertainty of D through bootstrap resampling of the data 5000 times.

e. Mean residual life plot

In the process of fitting a GPD, a minimum threshold is chosen to select the data that represent extreme values that should be fit. The GPD fit can be sensitive to the selection of this threshold, where too low of a threshold could violate the asymptotic nature of the GPD, and too high of a threshold would provide only a few relevant data points and lead to an unstable fit. To determine an appropriate threshold, one leverages what is known as the mean residual life (MRL) plot. The MRL uses the GPD-based expectation value of the data [Eq. (4.8), [Coles et al. 2001](#)]. For a given threshold u_0 , the expectation value (mean) of the data X that exceed the threshold can be written for $\xi < 1$ as

$$E(X - u_0 | X > u_0) = \frac{\sigma_{u_0}}{1 - \xi}.$$

It can be shown that for all thresholds $u > u_0$, this expectation value is [Eq. (4.9), [Coles et al. 2001](#)]

$$E(X - u | X > u) = \frac{\sigma_{u_0} + u\xi}{1 - \xi}. \quad (9)$$

That is, for $u > u_0$, $E(X - u | X > u)$ is a linear function of u . This expected linearity of Eq. (9) with u —for values of u for which the GPD fit is appropriate—can be used to test the validity of the chosen threshold used in the GPD analysis via a plot which we use in our results section known as the MRL plot.

f. Hypothesis testing

Although WWA does not carry out formal hypothesis testing, we interpret that they adopt a null hypothesis of α equal to zero, implying no ACC influence on extremes, and an alternative hypothesis that, if α is not equal to zero, ACC does influence extremes. This inference is based on the quote, “If the one- or two-sided 95% confidence interval (for the change in probability, which has the same effect as α) excludes zero, the change is statistically significant” (see section 4.3.1 of [Philip et al. 2020](#)). We assess whether rejection of a null hypothesis of α equal to zero permits accepting an alternate of ACC influencing extreme events. In other words, we assess if a non-zero α necessarily means a role for ACC. We note that another null hypothesis for purposes of heat could be that the location of the PDF of heat extremes is shifted by an amount equaling the GMST warming, as would be represented by

$\alpha = 1$. For purposes of consistency with the null implied by WWA, however, we only consider a null of α equal to zero.

WWA uses two variants of their analysis approach: one in which the extreme event in question is incorporated into the estimate of the PDFs, which we follow here, and one in which it is excluded (Philip et al. 2020). The act of choosing to study a class of events immediately following the occurrence of a particularly extreme example is known to induce a bias related to whether it is included or excluded (Miralles and Davison 2023). We choose to include the extreme value in question, which tends to make the parametric dependence on GMST larger and tends to make it easier to reject the null of no change in favor of the alternate hypothesis. Within the context of the WWA hypothesis test, this choice to include the extreme in question makes it more likely to conclude that ACC significantly influences extreme events.

g. Data

We briefly review the observations and our processing used in each of the three case studies reconsidered here. In addition, we describe the model Large Ensemble and preindustrial simulations that we analyze.

1) SIBERIAN HEAT WAVE

Following CC2021, we rely on two sets of observations: 1) daily maximum temperature extremes for each year from station data in Verkhoyansk as discussed in CC2021 and 2) 4-yr smoothed GMST anomalies from GISTEMP (Lenssen et al. 2019). CC2021 also performed a GEV analysis on temperature anomalies over a region in Siberia, which we do not focus on here to study specifically the extreme values detected at the Verkhoyansk station. At the time of the CC2021 study, the daily maximum temperature extreme in 2020 at the Verkhoyansk station was recorded as 38°C. This temperature extreme currently appears as a missing data point on the National Climatic Data Center website (<https://www.ncei.noaa.gov/cdo-web/search>, station RSM00024266, VERHOJANSK, data downloaded August 2023), but we still use this unconfirmed value in our analysis.

We also study GMST anomalies and extreme temperatures over the Siberian region defined in CC2021 within the CESM Large Ensemble (Kay et al. 2015) over the period 1926–2019. The CESM Large Ensemble contains 40 simulations of the climate over the period 1920–2100 under historical and RCP8.5 external forcing, of which we use the 35 that were run on the NCAR supercomputer. These simulations differ only in small perturbations to the initial conditions, which lead to different temporal sequences of modes of internal variability. We use model output only up to the year 2020, and, as a result, there is little sensitivity to the emissions scenario used [Detection and Attribution Model Intercomparison Project (DAMIP); Gillett et al. 2016, uses RCP4.5] because emissions scenarios only significantly diverge later. We do not use some additional runs done on the University of Toronto supercomputer that contain a global mean temperature bias (National Center for Atmospheric Research 2024).

In addition, we study GMST anomalies and extreme temperatures within a 1800-yr CESM preindustrial simulation, yielding 18 segments of 94 years, chosen for consistency with the length of the observed record. This preindustrial simulation uses fixed greenhouse gas concentrations, volcanic activity, anthropogenic aerosols, land use, and solar forcing, all set to their respective levels in the year 1850. Most importantly in the context of this study is the lack of time-varying anthropogenic forcing in this simulation.

2) AUSTRALIAN BUSHFIRE HEAT WAVE

Following OK2021, we again use two sets of observations for the years 1920–2019: 1) annual (July–June) maxima of 7-day moving average daily maximum surface temperature data from the Australian Water Availability Project (AWAP) over the southeastern region in Australia defined in OK2021 and 2) 4-yr smoothed GMST anomalies from GISTEMP. We apply a similar analysis to the CESM Large Ensemble and to the long CESM preindustrial run as for the Siberian heat wave case except, in this case, using 18 segments of 98 years for consistency with the observations.

3) MADAGASCAR DROUGHT

Following HW2022, we use two datasets. The first data are 2-yr means of precipitation from ERA5 (Hersbach et al. 2020) over the region of Madagascar that was defined in HW2022. HW2022 explained that they used a “24-month running mean rainfall data from July to June.” Our approach is to calculate total annual precipitation and then calculate a running 2-yr average of this annual data resulting in a new smoothed annual time series. For 1951–2020, the bottom 20% precipitation 2-yr periods are selected for the analysis, corresponding to the following fourteen 2-yr periods: 1956–57, 1957–58, 1958–59, 1959–60, 1962–63, 1990–91, 1991–92, 1992–93, 2008–09, 2009–10, 2015–16, 2016–17, 2017–18, and 2019–20. The second data are 4-yr smoothed GMST anomalies from GISTEMP. We note that the analysis following HW2022 relies on a fit to only 14 data points. We discuss the issues that are involved in fitting extreme distribution functions to such a small number of data points below. We also study the equivalent variables within the CESM Large Ensemble over the period 1950–2019, consistent with the period of observations. We again apply a similar analysis to the CESM Large Ensemble and to the long CESM preindustrial run, in this case using 18 segments of 70 years.

We note that in each of these three case studies, WWA leverages many datasets, including ERA5 reanalysis, station observations, and the multimodel CMIP5 ensemble. We use a subset of the data that WWA study but also the CESM Large Ensemble data in a consistent manner across the three example cases that we evaluate.

3. Results

We first describe our analyses of how extreme events depend on GMST in the context of the long CESM preindustrial simulation in section 3a. The simulation contains no time-varying

anthropogenic forcing, making it useful for evaluating a null hypotheses associated with no anthropogenic influence. We then consider each of the three case studies in sections 3b–3d. For each case, we assess how GEV or GPD distributions fit the observations relative to normal or lognormal distributions. We then evaluate the justification for the inclusion of a GMST dependence of the distribution parameters using the deviance statistic and estimate the uncertainty level of the parameter α used for the attribution. Finally, we examine the amount of data needed to adequately constrain GMST dependence in a GEV or GPD fit.

a. Extremes analysis of a preindustrial simulation

WWA examines the dependence of the distribution parameters of extreme events on the 4-yr running mean GMST. Commonly, a linear dependence of a location parameter on GMST, as represented by the parameter α , is used to estimate the PDF of extreme events in, for example, 1900 versus 2020. Changes in the PDF between these epochs are interpreted as owing to the effects of anthropogenic climate change, implying that α indicates the extent to which anthropogenic climate change affects extremes. We show here that it is possible, however, for internal climate variability, such as ENSO, to produce covariance between GMST and extremes that is expressed as a nonzero α even without anthropogenic influence. This covariance is demonstrated most clearly in this section by using a preindustrial climate model run, but we also elaborate upon this basic result using the CESM Large Ensemble. For attribution purposes, this result suggests that the value of α calculated in the presence of a greenhouse gas (GHG) increase can be affected by internal variability, biasing the attribution results. It should be noted that in a realistic scenario, α is affected by both anthropogenic climate change and internal variability. The scenario of internal variability alone analyzed here is artificial, of course, and does not capture that, in the current climate, GMST is a reasonable proxy for anthropogenic climate change. Analyzing simulations that only include internal variability nevertheless provides insight into the effects of internal variability in more complete scenarios.

To examine this potential role of internal variability, we start with the case of the hot weather associated with the bushfires in southeastern Australia and perform a GEV analysis using a 1800-yr CESM preindustrial run. This run holds anthropogenic CO₂, other anthropogenic greenhouse emissions, volcanic (natural) emissions, land use changes, and anthropogenic aerosols at values estimated for the year 1850. We divide the run into 18 data segments of length 94 years, equal to the length of the observed record and analogous to having 18 simulations of historical climate without ACC. The distribution of α values for the GEV preindustrial analysis is shown by the green bars in Fig. 1a. We find a large and positive α calculated from essentially all model output segments centered around a median value of 4.9°C °C^{−1}, or that extreme summer temperatures increase by 4.9°C for every 1°C GMST warming. Note that, following WWA, we are using 4-yr running average GMST values for this computation.




FIG. 1. Effect of internal climate variability on α distribution. PDFs of α for (a) the Australian 7-day moving mean of the maximum temperature related to the analysis of the 2019–20 bushfires; (b) the Siberian maximum daily temperatures that were used in the analysis of the 2020 heat wave; and (c) the Madagascar 2-yr precipitation used in the analysis of the drought of 2019–21. The analyses are based on the 1800-yr CESM preindustrial run with 94-, 98-, and 70-yr segments, respectively (green), and the 35 LE members (blue). The lengths of these segments correspond to the extent of the observed records used in CC2021, OK2021, and HW2022. The median values for the CESM preindustrial and LE for each of the cases are indicated in the boxes. The median values are Australia: $\alpha_{\text{CTRL}} = 4.9$, $\alpha_{\text{LE}} = 0.4$; Siberia: $\alpha_{\text{CTRL}} = 1.1$, $\alpha_{\text{LE}} = 2.5$; and Madagascar: $\alpha_{\text{CTRL}} = -0.02$, $\alpha_{\text{LE}} = 0.09$.

It is difficult to imagine a physical causal mechanism that would lead to such an enormous amplification of the GMST signal in Australian heat waves. A more plausible interpretation is that internal variability—for example, ENSO—could increase both extreme temperatures over Australia and, to a smaller degree, GMST, generating covariance between extreme events and GMST. In this scenario, GMST is not an indicator of forcing of local extremes; rather, both local extremes and GMST are influenced by a third factor. It is known that several modes of internal variability affect extreme temperatures in Australia, including ENSO, the Indian Ocean dipole, and the southern annular mode (Hendon et al. 2007; Ummenhofer et al. 2009; King et al. 2020). Min et al. (2013) find that ENSO and the Indian Ocean dipole can lead to warming of extreme temperatures in southeastern Australia by around 1°C, indicating that

internally generated variability can have an effect that is similar in magnitude to ACC.

WWA recognizes that α may be influenced by internal variability and, as noted, uses a 4-yr running mean GMST to filter out the effects of short-term internal variability. This approach, however, may not fully remove the ENSO signal. Furthermore, there is also the potential for nonlinear interactions between ENSO and GMST as well as the presence of slower modes of variability such as the Pacific decadal oscillation (PDO).

A similar analysis of Australian heat extremes made using the CESM Large Ensemble, which includes anthropogenic forcing, gives a distribution of α that is closer to zero (blue bars in Fig. 1) than in the unforced simulations (green bars). This indicates that both ACC and internal variability influence regional extremes and GMST, such that estimates of α will reflect multiple processes. In simulations with time-varying anthropogenic forcing, however, there is no way to disentangle the effect of internal variability on α using the WWA approach as described in the published literature. We conclude that the magnitude of α may be a biased indicator of anthropogenic influence. Including indices of internal variability as covariates (by adding parameters that represent the dependence of the distribution on regional climate variability modes such as ENSO) may help to disentangle their respective contributions. Yet, the projection of internal variability onto GMST can still confound the results, and adding additional parameters to the distribution to represent the dependence of the distribution on regional climate variability modes such as ENSO may not be justified given that below we find that even adding the single parameter α can be questioned due to the short observed record.

These results highlight a potential major difficulty with analyses based on a fit of distribution parameters to GMST in a more realistic scenario that includes both internal variability and ACC: the calculated α , and therefore the shift in the probability of extreme events, may be due to internal variability rather than only ACC. When analyzing the observational record in section 3c, we find $\alpha \approx 2^\circ\text{C } ^\circ\text{C}^{-1}$. This suggests a 2° warming of extreme temperatures for a 1° warming of the GMST, which is at the upper end of the range of values estimated from CESM Large Ensemble (LE). Based on our results from the CESM preindustrial run with a time-invariant anthropogenic forcing, we therefore suspect that the estimate obtained for α from the observed record is biased high by internal variability.

Analysis of the CESM preindustrial run for the case of Siberian heat waves (green bars in Fig. 1b) shows that the value for α deduced from a record similar in length to that of the observations ranges between 1 and 4, although there is again no time-varying anthropogenic forcing in the corresponding model run. This large and variable response implies that a nonzero α does not necessarily imply an effect of ACC. It may suggest that natural variability affects the value of α for Siberian heat waves during some decades but not others. This analysis shows that α equal to zero is an unsuitable null hypothesis for the purposes of evaluating the presence of an anthropogenic influence on extreme events (for context, we

note again that WWA does not explicitly define a null hypothesis).

For the case of Madagascar drought (Fig. 1c), we find a large spread in α for both the preindustrial run as well as the CESM Large Ensemble, with both centered near zero. The fact that α distribution from the preindustrial run is centered around zero suggests that a null hypothesis of α equal to zero may be appropriate, but the fact that α from the Large Ensemble is also centered around zero indicates that there is little potential for this test to discern an ACC influence on extremes. That is, low extremes of precipitation show no systematic increase or decrease related to GMST in either the forced or unforced runs such that this test has low statistical power. We further consider this case below using observations.

The analysis in this paper shows that internal variability can lead to a large α . We conclude that in a more typical scenario analyzed using the WWA methodology and applied to observations under varying anthropogenic GHG and aerosols, the value of α may still be affected by internal variability and, therefore, may bias the estimates of the effect of ACC on the return time of extreme events. One could justify the WWA use of the 4-yr averaged GMST as being a proxy for known greenhouse gas concentrations. If CO_2 were indeed used, the analysis of the preindustrial run would, of course, show no dependence, eliminating this problem, though then introducing issues associated with lags between GHG concentrations and warming (Proistosescu and Huybers 2017). Because WWA studies use GMST, we follow that methodology here.

The result that internal variability may strongly affect α when the 4-yr average GMST is used to link extreme events to ACC suggests that it might be useful to reevaluate some of the WWA results based on the fit to extreme distributions. It is possible, of course, that a strong anthropogenic signal in a few decades could significantly affect extreme events and overwhelm the effects of internal variability in cases such as the Siberian heat wave. The attribution of extreme temperatures may be more difficult in cases like the Australian example, where the effects of internal variability are large, as indicated by the large α detected in the CESM preindustrial simulation. In such cases, it would be important to correctly identify all relevant covariates—such as the possible effect of ENSO in the Australian case—and note that misleading results for α may be obtained if one relies only on GMST.

b. Siberian heat wave

We now explore other facets of the attribution of the Siberian heat wave in the summer of 2020 to anthropogenic change, notwithstanding the foregoing finding that the interpretation of α is complicated by the presence of internal variability. We focus specifically on the CC2021 analysis of the June daily maximum temperatures because it allows us to address the question of whether the data justify the use of nonstationary (GMST dependent) extreme value distributions.

While it is understood based on theoretical considerations that a GEV is the appropriate representation of block maxima of data (Coles et al. 2001), it is not obvious for a given problem what minimum block size justifies using a GEV and




FIG. 2. Siberian heat wave: (a) examining the fit of a GEV and normal distribution to Verkhoyansk station data. The red curve indicates the empirical CDF of the highest daily maximum temperature for each year for the Verkhoyansk station data. The purple and blue lines denote fits by a GEV (with a constant location parameter) and a normal distribution, respectively. The light and dark green shadings correspond to the 90th and 95th percentile confidence intervals (section 2c). (b) Justification for using a GMST-dependent mean. Boxplots of the deviance statistic between a GEV distribution without a GMST-dependent mean relative to a GEV distribution with a GMST-dependent mean using the Siberian heat wave data with 5000 bootstrap resamples. The edges of the boxes correspond to the first and third quartiles. The deviance statistic for the station observations is shown in green as the leftmost box. The deviance statistic was also computed for each of the 35 members from the CESM LE (blue). The 95th percentile significance level from $\chi^2_{k=1}$ distribution with one degree of freedom is shown by the horizontal red line, corresponding to the addition of 1 degree of freedom by including α in the fitted model.

specifically if a year of daily maximum data suffices. This is especially an issue given that the temperature is autocorrelated and, therefore, more data are needed per averaging block to justify the use of a GEV. It is, therefore, worthwhile to check how well other standard distributions perform in fitting the data.

A cumulative distribution estimate of the highest June daily maximum temperature observed each year between 1926 and 2020 at the Verkhoyansk station is shown in Fig. 2a along with a 95% confidence interval of the CDF obtained by bootstrapping. In the first step of our analysis, we consider cumulative distribution function fits to the observed distribution, assuming the data follow either a GEV or normal distribution. The fitted distributions are specified here to be stationary [by setting $\alpha = 0$ and $\mu = \mu_0$ in Eqs. (1)–(4)]. Fits are obtained by maximizing the likelihood of the data, and the results are shown in Fig. 2a. Both distributions generally fit the observations well, with the normal fit completely within the 95% range of the observed distribution, and with the GEV fit being good except at the highest temperature, where it slightly underestimates the probability of obtaining the highest temperatures. There is only one data point above 35°C, so the empirical CDF is poorly constrained there in any case. Interestingly, the normal distribution and GEV CDFs look similar and lie within the confidence intervals defined for the empirical CDF, despite the normal distribution having fewer parameters than the GEV. It is worth noting that a GEV and a normal distribution may assign very different probabilities to extreme events, even when the two distributions provide an adequate fit to the observed

CDF. These modeling decisions are therefore of critical importance for attribution studies, which usually focus as much or more on the estimation of the exceedance probability and return period of an event—and, consequently, on the changing likelihood and fraction of attributable risk of an event—as they do on the change in intensity. It is therefore useful to consider alternate distributional forms.

The quality of fits indicates that an extreme value distribution is not empirically a better choice than a normal distribution for fitting the data. Similarly, a nonstationary distribution with $\alpha \neq 0$ may not be necessary given that the stationary distributions seem to provide a satisfactory fit to the empirical cumulative distributions, and we further examine this next. CDFs are commonly used to differentiate between distribution functions, and the results of a Kolmogorov–Smirnov test (Supplemental Fig. 1 in the online supplemental material for normal, GEV, and empirical distributions) also suggest that the normal distribution is consistent with the observed data. An application of Occam’s razor principle suggests that the simpler normal distribution should be used rather than a GEV if there is not a significant improvement in the distributional fit with the GEV. This is also discussed by Philip et al. (2020) who mention explicitly that the extreme value distributions may not outperform a simpler (normal) distribution in some cases. Otto et al. (2018) opted to use a normal distribution in lieu of a GPD for attribution analysis of a cold spell in Peru because the normal “uses all data, rather than the GPD, which samples only the cold tail.” CC2021 found that a normal distribution fits the data well but proceeded to use a




FIG. 3. Siberian heat wave: uncertainty in α . (a) Scatterplot of the daily max temperature at the Verkhoyansk station against GMST. The red line indicates a linear regression fit to the data, and the blue line indicates a linear fit using the location parameter and α derived from the GEV maximum likelihood estimate. [(b); red] PDF of the GEV-fit α distributions from the station data, using 5000 bootstrap resamplings of the observed time series. The α value from the observations (no resampling) is shown by the vertical red line. [(b); blue] A histogram of the GEV-fit α PDFs derived from the CESM LE, where the α values are calculated from each ensemble member with no resampling. The average α value across the LE is indicated with the blue line. There are 94 annual data points for each ensemble member.

GEV for the bulk of the analysis. These results question the need for extreme value distribution functions when it comes to attribution analysis. One might want to use proper selection criteria [the Akaike information criterion (AIC) or the deviance statistic also known as (AKA) the log-likelihood ratio test which we use in this work, or similar] to rigorously decide which of the models performs best.

We now proceed to examine the role of GMST in the fit to a GEV distribution [Eq. (4)] in this particular application, following CC2021. A linear regression of GMST anomalies against daily max temperatures gives a slope of $1.1^{\circ}\text{C }^{\circ}\text{C}^{-1}$ (Fig. 3a), where this value indicates that the annual maximum daily temperatures at Verkhoyansk increase slightly more quickly than GMST. The squared Pearson correlation, however, is only $r^2 = 0.01$, with a p value of 0.2, indicating that GMST is only a weakly significant indicator of trends in Siberian temperature maxima.

Figure 3b shows the maximum likelihood fit of a nonstationary GEV, including α , for multiple bootstrapping samples as specified in the methods section. The value of α estimated from the observations is $1.8^{\circ}\text{C }^{\circ}\text{C}^{-1}$, larger than the slope of 1.1 obtained from the linear regression. This discrepancy between α and the slope suggests that the conclusions from an

attribution analysis can be highly sensitive to the choice of the fitted distribution, again underscoring the importance of model selection. The [2.5%, 97.5%] confidence interval is $[-1.2, 4.8]$, while the 5% value is -0.7 (Fig. 3b; red plots). Both indicate that α is not significantly different from zero. It is noteworthy that the red distribution in Fig. 3b is bimodal, presumably because some of the bootstrapped samples included the 2020 event and some did not. This suggests that the fitted distribution is very sensitive to the inclusion or exclusion of the event (again consistent with Miralles and Davison 2023); that is, α would be significantly nonzero if the event is included but is not significant when it is omitted. While CC2021 did not provide an uncertainty range for α , they note that the change in intensity of Siberian heat waves at a fixed probability is estimated from their inferred PDF as 1.04°C (with a 95% confidence interval from 0.35 to 3.4). The α is directly related to the change in intensity: The change in intensity is $k \times a$, where k is the observed GMST change between the two periods of interest, usually 1.2°C . Assuming a GMST change of 1.2°C , we infer the value of α in CC2021 is 1.3 with a 95% confidence interval from 0.42 to 4.1. The central estimate of α is similar to our results but entails a smaller uncertainty range, likely due to our decision to include the 38°C event in our bootstrap resampling, but which was not done in CC2021. The bootstrapped distribution for all GEV parameters for the observations is given in supplemental Fig. 2.

The deviance statistics D (section 2d) indicate that the addition of a model parameter α , representing the GMST dependence, is justified in the case of the Verkhoyansk observations (Fig. 2b, leftmost bar is above the threshold indicated by the red horizontal line). Yet the deviance statistic uncertainty range includes a large range that is below the significance line, which seems appropriate given that the uncertainty range for α includes negative values.

To provide context for the interpretation of the Verkhoyansk observational results, we also examine output from the CESM Large Ensemble, which includes time-varying anthropogenic forcing, at the grid point containing the Verkhoyansk station. We find that all ensemble members within the CESM Large Ensemble indicate a positive value of α , with a mean value of $2.60^{\circ}\text{C }^{\circ}\text{C}^{-1}$ (Fig. 3b; blue plots). Similarly, 33 of 35 ensemble members give a deviance statistic that supports the use of a nonstationary component (i.e., a GMST dependence) in the GEV statistic (Fig. 2b).

The fact that the LE results indicate a stronger link between Siberian extremes and GMST (i.e., a larger α) is consistent with the LE, on average, simulating an Arctic-wide summer warming rate of $1.8^{\circ}\text{C }^{\circ}\text{C}^{-1}$ warming in GMST, whereas an observational analysis only indicates a rate of 1.3°C Arctic JJA warming per degree Celsius in GMST (using GISTEMP for both; Figs. 4a,b). This greater amplification of Arctic warming may be related to the fact that CESM1 has a relatively high climate sensitivity of 4.1 K (Gettelman et al. 2019). That said, one expects heat waves to be strongly influenced by the JJA climatology, and Fig. 4c shows that different ensemble members, corresponding to different realizations of internal variability, show very different dependencies of the Arctic JJA climatology on GMST. It is difficult, therefore, to estimate from the single realization corresponding




FIG. 4. Internal variability and the attribution of Siberian heat waves. (a) A scatterplot of the annual average JJA Arctic (north of 70°N) surface temperatures against the 4-yr running mean of GISTEMP annual GMST. (b) As in (a), but for the 35 members of the CESM LE over the period 1926–2019. (c) PDF of the slopes of annual average JJA Arctic surface temperatures against the 4-yr running mean of annual GMST within each member of the CESM LE with 5000 bootstrap resamples (red). The same distribution is indicated for GISTEMP observations in blue. The vertical red line indicates the mean slope across the CESM LE, and the vertical blue line indicates the slope determined from observations.

to the observed record to what degree the Arctic JJA climatology is affected by ACC versus internal variability, adding to the uncertainty in attributing Siberian heat waves. To address this issue, the WWA protocol includes a step where the same analysis is applied to climate models, leveraging a range of climate models so that results are not dependent entirely on one GCM, and results are reported from both sources.

Unsurprisingly, if data from multiple ensemble members are analyzed simultaneously as a single dataset, which mimics the availability of more data, the distribution of the GEV parameters narrows, generally converging around the mean value of the parameters across individual ensemble member fits (Fig. 5). The 95% confidence range for α (red asterisks in

Fig. 5a) decreases from a large range of [1, 4] (°C °C⁻¹) with one ensemble member to a smaller [2, 3] range with 10 ensemble members. Similar behavior is seen for the other GEV parameters. This implies that with a sufficiently long record, one can reduce the uncertainty in the GEV parameters, as expected, but that might require much more data than is expected to be available in the historical record anytime soon.

The overall result for the GMST dependence for the Siberian heat wave analysis, which is the basis of the attribution to ACC, is largely consistent with the results found from CC2021 but with three important caveats. First, we found in the previous subsection that internal variability can lead to GMST dependence that biases the value of α and hence the attribution results




FIG. 5. Siberian heat wave: uncertainty in GEV parameters as a function of the number of ensemble members used. PDF contour plots for each of the GEV parameters as a function of the number of ensemble members used from the CESM LE and the parameter value. One ensemble member is equivalent in length to the historical record. The PDFs for (a) α , (b) ξ , (c) σ , and (d) μ_0 are shown. The PDFs were derived from the highest daily maximum Siberian temperature in June, as in CC2021, using the CESM LE. Contours are drawn every 0.3 for (a), 2 for (b), 1.5 for (c), and 0.7 for (d), all starting at 0.05. We denote the 95% confidence range for each number of ensemble members by the red asterisks.

(Fig. 1a). Second, whereas a GEV fit a priori makes sense to fit to block maxima of daily temperature extremes, a normal distribution with $\alpha = 0$ does not appear to be ruled out by the observations (Fig. 2a). This may be a result of both the relatively short observational record and the autocorrelated nature of the temperature extremes. Still, a nonzero α is found to meaningfully improve the statistical fit to the CESM Large Ensemble Siberian extremes based on deviance statistics. Finally, the uncertainty range for α , which was not explicitly

evaluated in the original WWA analysis, includes negative values (Fig. 3b).

c. Australian bushfire heat wave

We now consider the analysis of possible connections between the hot weather that may have contributed to the Australian bushfires of the summer of 2020–21 and ACC, using both observations and the CESM Large Ensemble, following the observational analysis of OK2021. We again find that a




FIG. 6. Australian bushfire heat wave: (a) testing the justification for a GEV. The red curve indicates the empirical CDF of the highest annual 7-day running mean of daily maximum temperature over the area defined in OK2021 for observations. The purple and blue lines denote GEV (with a constant mean) and normal distribution fits, respectively. The light and dark green shadings correspond to the 90th and 95th percentile confidence intervals. (b) Justification for using a GMST-dependent mean. Boxplots of the deviance statistic between a GEV without a GMST component and a GEV with a GMST component using the Australia bushfire heat wave example with 5000 bootstrap resamples. The edges of the boxes correspond to the first and third quartiles. The deviance statistic was computed for each of the 35 members from the CESM LE (blue) and compared with the 95th percentile significance level from χ^2 distribution with 1 degree of freedom (red line). A boxplot for the deviance statistic for observations is shown in green.

normal distribution and a GEV distribution (both with a constant mean/location parameter) fit the 7-day moving mean of daily temperature maxima in the observations without a GMST dependence well (Fig. 6a). In this case, a normal distribution fit to the empirical distribution has a mean absolute error (MAE) of 0.01, equal to the MAE for the GEV fit. Supplemental Fig. 3 shows results for normal, GEV, and empirical distributions that again suggest that the normal distribution aligns well with the observed data. The alignment may reflect that the small number of available data do not constrain the PDF very well. Given that the normal distribution has one less parameter, it seems to offer a simpler approach to representing the data. The deviance statistic for the observations does suggest a significant GMST component (leftmost bar in Fig. 6b). However, given the possibility that this component is driven by internal variability rather than by anthropogenic change (section 3a, Fig. 1a), one must conclude again that even if α significantly differs from zero, this might reflect a bias introduced by the effects of internal variability. This underscores the need to include covariates representing relevant internal variability based on our physical understanding. The deviance statistic test for the CESM Large Ensemble shows that in only five of the 35 ensemble members is the extra complexity of adding GMST dependence to the distribution justified (Fig. 6b). This raises the question of whether the influence of GMST, and therefore of ACC in the WWA formulation, is a necessary part of the statistical description of extreme temperatures in Australia as represented in the CESM Large Ensemble. It should be noted that our results relate only to one climate

model that, of course, does not necessarily represent all the relevant physical processes at work in the real world. The use of multiple climate models in the WWA analysis helps ensure a wider range of process representations are included.

The regression analysis of Australian extreme temperatures and GMST (Fig. 7a) as well as the range of α values estimated from observations (Fig. 7b; red), both show a GMST dependence near a 2°C change in extreme temperatures per degree increase in GMST. As mentioned above, a sensible null hypothesis for heat waves in a warming climate is that the extreme values shift with the mean warming (Tziperman 2022). The GMST dependence calculated here suggests that Australian heat waves increase at nearly twice the rate of the GMST. Such a strong response requires a physical mechanism. Some candidates are 1) the fact that land heats up more quickly than the ocean, so GMST always tends to lag behind land surface temperatures; 2) soil drying can lead to further increases in warming over land as less water is available for evaporative cooling, etc. It is not clear that these can lead to a doubling of the GMST effect, but further exploration of this amplification is beyond the scope. The alternative is that this is a result of internal variability—possibly, ENSO—affecting both Australian heat extremes and the GMST, consistent with Fig. 1a. Again, this suggests that using this value of α for attribution may lead to biased results. Finally, Fig. 7b (blue) shows a large range of GMST dependencies for different ensemble members in the CESM Large Ensemble, again highlighting that a large component of this dependence most likely represents internal variability




FIG. 7. Australian bushfire heat wave: distribution of α values. (a) Scatterplot of the highest annual 7-day running mean of daily maximum temperatures against GMST (following OK2021). The red line indicates a linear regression fit to the data, and the blue line indicates a linear fit using the location parameter and α derived from the GEV maximum likelihood estimate. [(b); red] PDF of the GEV-fit α distributions from observations, using bootstrap resampling. The α value from the observations without resampling is shown by the red line (at 1.7). [(b); blue] Effect of the internal variability and number of data from a model LE on the uncertainty in α for the Australian Bushfire heat wave example. Shown is a histogram of the GEV-fit α PDFs derived from the α value from each ensemble member with no resampling. The mean α value across the LE is indicated with the blue line.

rather than anthropogenic change. The distribution of all GEV parameters for the observations is given in supplemental Fig. 4. How the GEV parameters scale with the number of CESM Large Ensemble members used is indicated in supplemental Fig. 5.

d. Madagascar drought

Finally, we consider the attribution analysis for the 2019–21 drought in Madagascar, following the lead of HW2022. This case differs from the two examined above in that it analyzes precipitation extremes that are positive by definition. HW2022, therefore, used a different extreme value distribution function, the GPD, which is meant to represent the high tail of the distribution of extreme events over a specified threshold (Coles et al. 2001). As discussed in section 2, HW2022 analyzed the driest 20% of the 2-yr averages of precipitation data from 1951 to 2020 multiplied by minus one to turn minima (droughts) into maxima. We follow their example, with the important caveat made in section 2b that the validity of using a GPD this way for precipitation data is not guaranteed, given that there is a hard maximum at zero precipitation that seems inconsistent

with the definition of the GPD. We also note that taking the lowest 20% of 2-yr precipitation data points amounts to only 14 data points, likely too low a number for a reliable attribution.

Figure 8d shows the MRL plot (section 2e) for the Madagascar drought case. As a reminder, this plot needs to scale linearly within the range of the chosen threshold value (section 2e) in order for the GPD analysis to be justified and self-consistent. For small and large values of precipitation, the MRL curve appears to curve nonlinearly, indicating that a GPD fit is likely invalid using those thresholds. For precipitation thresholds ranging from roughly 1.6 to 2 mm day⁻¹, the MRL curve plateaus. The orange line shows the expected MRL line, $[\sigma + u\xi/(1 - \xi)]$; that slope does not align with the MRL curve around the value of 1.9 corresponding to the 20th percentile precipitation threshold chosen by HW2022. In general, we suggest that this MRL analysis should be helpful in examining the self-consistency of applying a GPD to precipitation data, though the small number of data points appears an obstacle for a reliable MRL analysis in this particular case. Specifically, it is not obvious if the required linear range exists to a degree that justifies using a GPD distribution and allows for selecting a threshold.

We find again that an extreme value distribution is not required by the data distribution. The GPD fit to the bottom 20th percentile precipitation data is no better than a lognormal distribution fitted to all of the precipitation data (Fig. 8a). Furthermore, both the lognormal and GPD fits indicate no need of a GMST dependence, in keeping with the small value of α found by HW2022. A lognormal distribution seems a reasonable null hypothesis for the positive precipitation data. In fact, the empirical CDF is practically unconstrained at its lower tail because there are no data points below a value of 1.25 mm day⁻¹. Supplemental Fig. 6 for GPD, lognormal, and empirical distributions also suggests the lognormal distribution looks reasonable, albeit acknowledging the highly limited sample size in this case.

The deviance statistic for the Madagascar precipitation observations is marginally significant at the $X\%$ level (leftmost green bar is above the red significance line, Fig. 8b), indicating that the inclusion of the additional model parameter α is justified. Uncertainty plots for the GPD parameters fit to these precipitation observations are indicated in supplemental Fig. 7. The bootstrap estimates for α parameter (supplemental Fig. 8a) range over both negative and positive values, indicating that it does not significantly differ from the null. None of the 35 ensemble members exceed the deviance statistic significance level. Figure 1c shows that there is no relation between these extreme drought events and GMST. One interpretation is that this is due to the absence of any strong effect of internal variability modes on both GMST and Madagascar drought, in contrast with the analysis of Australian temperatures in Fig. 1a. Also, Fig. 8b may suggest that the CESM Large Ensemble does not adequately capture the effect of ACC on droughts in Madagascar.

As a reminder, the threshold parameter of the GPD, corresponding to the maximum precipitation value considered as an extreme drought year in this attribution case, was assumed by HW2022 to vary exponentially with the GMST $\bar{T}$ as




FIG. 8. Madagascar drought: (a) evaluating GPD and lognormal fits to rainfall data. An empirical CDF of rainfall is shown (red line, with 90th and 95th percentile confidence intervals shown by red shadings), based on 24-month running means of daily precipitation over the Madagascar area defined in HW2022 using ERA5 precipitation. A GPD fit to the bottom 20th percentile precipitation (black) and a lognormal fit to the full distribution (blue). (b) The deviance statistic between a GPD with and without a GMST component in the CESM LE (blue) and observations (green). (c) Scatterplot of the 2-yr running mean of ERA5 daily precipitation over the Madagascar area defined in HW2022 against the 4-yr running mean of annual GMST for the lowest 20% precipitation data. The dashed blue line shows the exponential GPD log-likelihood fit for the threshold parameter for the observations without resampling [Eq. (7)] following HW2022. (d) MRL plot (section 2e) for different precipitation thresholds. The blue shading defines a region of confidence calculated as ± 1 standard deviation. The red line denotes the threshold used in the GPD fit of HW2022 (i.e., the 20th percentile). The orange line represents the full linear equation from Eq. (9) [$\sigma + u\xi/(1 - \xi)$].

$u = u_0 \exp(\alpha \bar{T}/u_0)$ with $u_0 < 0$. If droughts become more severe with increasing GMST, we expect the threshold parameter for the distribution of the negative precipitation to get larger (less negative) with increasing GMST, and therefore α/u_0 to be negative and α to be positive. The value of α_{GPD} calculated from observations is small and indistinguishable from zero (Fig. 8c). The exponential fit was assumed by HW2022 to ensure that the scaling factor $\exp(\alpha \bar{T}/u_0)$ applied

to both u_0 and σ_0 is strictly positive, so the distribution has fixed dispersion (Philip et al. 2020). A scatterplot of ERA5 precipitation against GMST anomalies looks notably linear (see the dashed blue line in Fig. 8c), consistent with the Taylor expansion of the exponential function, which is as follows:

$$u = u_0 \exp(\alpha \bar{T}/u_0) \approx u_0 (1 + \alpha \bar{T}/u_0) = u_0 + \alpha \bar{T}.$$

The small value of α found for observations indicates that the increase in GMST has had little effect on the statistics of 2-yr-average precipitation in Madagascar. Whereas this is also the conclusion reached by HW2022, we found that the data can be fitted with a lognormal distribution rather than the GPD and without GMST dependence (Fig. 8a). Further, the deviance statistic only marginally justifies adding GMST dependence for observations and does not justify adding such dependence for model output (Fig. 8b). The fact that only 14 data points are involved in the fit seems much too small to draw any meaningful conclusions on an extreme value distribution. Supplemental Fig. 9a indicates that even if using the equivalent of 35 times as much data, the uncertainty range for α includes both positive and negative values. One could, therefore, conclude that there is no clear relationship between GMST and low rainfall in this region.

4. Conclusions

We were inspired in this work by a series of attribution studies by the “World Weather Attribution (WWA) initiative” and evaluated the part of their methodology that involves a fit of extreme value distribution functions to observations. The full protocol used by this effort (Philip et al. 2020) involves additional tools and model analyses, including selecting and defining extreme events, analyzing trends in observational data, using multimodel ensembles to simulate event scenarios with and without human-induced climate change, performing quality checks on model accuracy, and synthesizing results to determine climate change’s influence on the event’s intensity and likelihood. A central WWA method, nevertheless, involves fitting an extreme value distribution function to an observed record, with the distribution parameters depending on global mean surface temperature (GMST), whose variation is taken to represent ACC. For example, the location parameter of an extreme value distribution (GEV) is sometimes made a linear function of GMST $\bar{T}$ as $\mu = \mu_0 + \alpha \times \bar{T}$. The GMST dependence included in the function is then used to calculate the change in return time of extreme events. This method was applied to many extreme events such as heat waves, cold spells, floods, droughts, and more (e.g., van Oldenborgh et al. 2016; van der Wiel et al. 2017; van Oldenborgh et al. 2017; CC2021; OK2021; HW2022). We examined the robustness of the results of three example WWA studies, demonstrated when the results may be unreliable due to the effects of internal variability, and suggested several additional statistical tests, reviewed below, that may increase confidence in the results of such an analysis.

We began by examining the GMST dependence of extreme heat wave events in Siberia (following CC2021) and Australia (following OK2021), as well as drought in Madagascar (following OK2021), using data from a CESM preindustrial simulation. We found a significant GMST dependence in the first two cases, indicating that such dependence can result from internal variability rather than only represent the effects of ACC. The GMST dependence, in these cases, is a result of regional internal variability such as ENSO or other climate variability modes, which affects the extreme events being examined

and also weakly affects GMST. This points out that the value of α calculated in the more realistic case where both effects are present includes the effects of internal variability in addition to those of ACC. The value of α in the Australian case is especially large and could naively be interpreted as indicating that every degree increase in GMST is associated with a 5° increase in Australian heat waves. Heuristically, the large values of α can be thought of as a result of internal variability (possibly ENSO) affecting GMST by 0.2° for every 1° effect it has on Australian heat waves. As noted earlier, this issue could be mitigated in part by including indices for modes of internal variability as covariates in the model fitting.

We tested if an extreme value distribution is indeed needed by plotting the empirical cumulative distribution function (CDF) with bootstrapping error bars (Figs. 2a, 6a, and 8a). We found that fits using a standard distribution (normal or lognormal) fall within the error bars of the empirical distribution, indicating that the observations do not empirically allow for distinguishing between the characteristics of extreme value and normal distributions. The differing tails of these distributions, however, lead to different extreme event return times. Similarly, the standard distributions seem to fit the data well without a GMST dependence, indicating that such dependence may not be justified by the data. We further tested if the addition of a GMST dependence of the distribution parameters is justified statistically using the log-likelihood ratio test, also known as the deviance statistics. When adding a parameter—in this case, α —one expects the fit to improve, but the question is whether the improvement is beyond what is expected simply due to the addition of one more model parameter, so that one is convinced that the additional parameter represents a genuine effect in the data. We found that in the Australia and Madagascar cases, there is no justification for the addition of α .

The Madagascar drought attribution (HW2022) was based on only fourteen of the lowest 2-yr-averaged precipitation rates, hardly sufficient to construct a meaningful empirical CDF or to constrain the four parameters of the GPD distribution used for that purpose. In addition, computing 2-yr averages of precipitation and applying an extreme value distribution may lead to inconsistencies a priori given that averages converge to a Gaussian distribution (under the central limit theorem). The GPD fit to the record, in this case, is meant for high values above a threshold, and the expected mean needs to vary linearly with the threshold chosen. We follow Coles et al. (2001) in suggesting that a mean residual life (MRL) plot can generally be used to test the self-consistency of using a GPD. In this case, it was not clear that these drought events are well represented by a GPD, but with the aforementioned overarching concern regarding the number of data points. To deal with the loss of data associated with this approach and difficulties in choice of an optimal precipitation threshold, Naveau et al. (2016) suggest instead using a statistical model fit to the full distribution of precipitation data, which has the added benefit of following the rules of extreme value theory for both high and low precipitation events. While our results—that attribution of 2-yr-averaged precipitation is not possible—agree with those of HW2022, we emphasized our different perspective: It is unlikely that attribution would be possible

based on such a small number of data points, as well as the other difficulties pointed out above. In contrast, other approaches evaluating changes in the seasonal cycle of precipitation have led to the attribution of shifts in seasonal rainfall patterns in southern Madagascar to climate change (Rigden et al. 2024).

An important issue raised by this work is the discrepancy between the modeled and observed findings, as seen, for example, in the deviance statistic results for the observations versus CESM Large Ensemble (Figs. 2b, 6b, and 8b). This may be due to a particular sequence of internal variability in the observations, measurement error in the observational record, or issues with the simulation of extreme weather in the CESM Large Ensemble. It would be helpful to test if these discrepancies also occur with other models. The use of multiple models by WWA (Philip et al. 2020; OK2021) may help mitigate against this kind of problem.

Overall, our results suggest uncertainty in the interpretation of the dependence of the parameters of extreme event distributions on the global mean surface temperature as reflecting ACC. The proposed and demonstrated deviance statistic, MRL plot, and a careful analysis of uncertainty ranges, especially for α , may help avoid misguided confidence in attribution results. As noted in the introduction, the WWA does not exclusively rely upon the empirical methodology evaluated here to reach conclusions, and also evaluates other historical data, model simulations, and the dynamical context. Philip et al. (2020) also highlight some caveats, including the necessity of testing for other probability laws (e.g., Gaussian and gamma) and covariates (ENSO, PDO, etc.), among other issues. Nevertheless, the observations-based component of the WWA methodology examined here still figures prominently in their analyses, and we have highlighted several caveats with its interpretation, foremost that internal variability can give rise to the appearance of strong covariance between GMST and local extremes. It is important to study when the ACC signal emerges distinctly relative to internal variability on a case-by-case basis. The ACC signal is known to have emerged in most regions of the world for temperature extremes, while this may not be the case for variables like precipitation that have much higher variability. Time of emergence is a factor that should always be considered when discussing the results of an attribution study. Our results do not provide evidence against a relationship between GMST and local extremes but, rather, that internal variability needs to be taken into account when fitting a distribution function to the record. As currently implemented, the WWA approach can lead to first-order biases in some cases. It, therefore, seems that it may be best to use a multimethod approach to ensure a robust attribution approach (e.g., Shepherd et al. 2018; Leach et al. 2024; Robin and Ribes 2020; Sippel et al. 2024).

We suggest two key takeaways. First, controlling for internal variability, possibly by making the distribution parameters depend on additional covariates such as ENSO, appears important to accurately estimate the mean ACC trend. This improvement upon the current method should be further developed. The inclusion of additional covariates is vital as a way to avoid biases in the estimate of ACC effects on extreme event statistics.

Second, empirical attribution techniques should be examined by applying the approach to simulations with and without changes to anthropogenic emissions. Application to such simulations would make it much easier to identify the likely influence of other factors on the local response in the absence of ACC, and so to ensure that all relevant modes are included in the statistical model, assuming the models are right in this respect.

Acknowledgments. P. S. is supported by the Harvard Global Institute. E. T. is supported by DOE Grant DE-SC0023134 and thanks the Weizmann Institute for its hospitality during parts of this work. P. H. is supported by NSF Award 2123295. The CESM project is supported primarily by the National Science Foundation (NSF). We acknowledge the CESM Large Ensemble Community Project and supercomputing resources provided by NSF/CISL/Yellowstone. We also acknowledge instructive conversations with Angela Rigden.

Data availability statement. All data used in this paper are available in public repositories. The climatological output from the CESM Large Ensemble and CESM preindustrial runs used in this paper is publicly available using the <https://intake-esm.readthedocs.io/en/stable/> Python package. All codes and data used here are available under the public open science foundation depository at <https://osf.io/fycu2/> and <https://doi.org/10.17605/OSF.IO/FYCU2>.

REFERENCES

- Allen, M., 2003: Liability for climate change. *Nature*, **421**, 891–892, <https://doi.org/10.1038/421891a>.
- Christidis, N., P. A. Stott, and F. W. Zwiers, 2015: Fast-track attribution assessments based on pre-computed estimates of changes in the odds of warm extremes. *Climate Dyn.*, **45**, 1547–1564, <https://doi.org/10.1007/s00382-014-2408-x>.
- Ciavarella, A., and Coauthors, 2021: Prolonged Siberian heat of 2020 almost impossible without human influence. *Climate Change*, **166**, 9, <https://doi.org/10.1007/s10584-021-03052-w>.
- Clarke, B., F. Otto, R. Stuart-Smith, and L. Harrington, 2022: Extreme weather impacts of climate change: An attribution perspective. *Environ. Res. Climate*, **1**, 012001, <https://doi.org/10.1088/2752-5295/ac6e7d>.
- Coles, S., J. Bawa, L. Trenner, and P. Dorazio, 2001: *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics, Vol. 208. Springer, 224 pp.s.
- De Paola, F., M. Giugni, F. Pugliese, A. Annis, and F. Nardi, 2018: GEV parameter estimation and stationary vs. non-stationary analysis of extreme rainfall in African test cities. *Hydrology*, **5**, 28, <https://doi.org/10.3390/hydrology5020028>.
- Eyring, V., and Coauthors, 2021: Human influence on the climate system. *Climate Change 2021: The Physical Science Basis*, V. Masson-Delmotte et al., Eds., Cambridge University Press, 423–552, <https://doi.org/10.1017/9781009157896.005>.
- Faranda, D., S. Bourdin, M. Ginesta, M. Krouma, R. Noyelle, F. Pons, P. Yiou, and G. Messori, 2022: A climate-change attribution retrospective of some impactful weather extremes of 2021. *Wea. Climate Dyn.*, **3**, 1311–1340, <https://doi.org/10.5194/wcd-3-1311-2022>.

- Gottelman, A., and Coauthors, 2019: High climate sensitivity in the Community Earth System Model version 2 (CESM2). *Geophys. Res. Lett.*, **46**, 8329–8337, <https://doi.org/10.1029/2019GL083978>.
- Gillett, N. P., and Coauthors, 2016: The detection and attribution model intercomparison project (DAMIP v1.0) contribution to CMIP6. *Geosci. Model Dev.*, **9**, 3685–3697, <https://doi.org/10.5194/gmd-9-3685-2016>.
- Harrington, L. J., and Coauthors, 2022: Limited role of climate change in extreme low rainfall associated with southern Madagascar food insecurity, 2019–21. *Environ. Res. Climate*, **1**, 021003, <https://doi.org/10.1088/2752-5295/aca695>.
- Hegerl, G., and F. Zwiers, 2011: Use of models in detection and attribution of climate change. *Wiley Interdiscip. Rev.: Climate Change*, **2**, 570–591, <https://doi.org/10.1002/wcc.121>.
- Hendon, H. H., D. W. J. Thompson, and M. C. Wheeler, 2007: Australian rainfall and surface temperature variations associated with the Southern Hemisphere annular mode. *J. Climate*, **20**, 2452–2467, <https://doi.org/10.1175/JCLI4134.1>.
- Hersbach, H., and Coauthors, 2020: The ERA5 global reanalysis. *Quart. J. Roy. Meteor. Soc.*, **146**, 1999–2049, <https://doi.org/10.1002/qj.3803>.
- IPCC, 2021: Summary for policymakers. *Climate Change 2021: The Physical Science Basis*, V. Masson-Delmotte et al., Eds., Cambridge University Press, 3–32, <https://doi.org/10.1017/9781009157896.001>.
- , 2023: Current status and trends. *Climate Change 2023: Synthesis Report*, H. Lee and J. Romero, Eds., Cambridge University Press, 35–115, <https://doi.org/10.59327/IPCC/AR6-9789291691647>.
- Kay, J. E., and Coauthors, 2015: The Community Earth System Model (CESM) large ensemble project: A community resource for studying climate change in the presence of internal climate variability. *Bull. Amer. Meteor. Soc.*, **96**, 1333–1349, <https://doi.org/10.1175/BAMS-D-13-00255.1>.
- King, A. D., A. J. Pitman, B. J. Henley, A. M. Ukkola, and J. R. Brown, 2020: The role of climate variability in Australian drought. *Nat. Climate Change*, **10**, 177–179, <https://doi.org/10.1038/s41558-020-0718-z>.
- Knutson, T., J. Kossin, C. Mears, J. Perlwitz, and M. Wehner, 2017: Detection and attribution of climate change. *Climate Science Special Report: Fourth National Climate Assessment, Volume I*, U.S. Global Change Research Program, 114–132.
- Leach, N. J., and Coauthors, 2024: Heatwave attribution based on reliable operational weather forecasts. *Nat. Commun.*, **15**, 4530, <https://doi.org/10.1038/s41467-024-48280-7>.
- Lenssen, N. J. L., G. A. Schmidt, J. E. Hansen, M. J. Menne, A. Persin, R. Ruedy, and D. Zyss, 2019: Improvements in the GISTEMP uncertainty model. *J. Geophys. Res. Atmos.*, **124**, 6307–6326, <https://doi.org/10.1029/2018JD029522>.
- Lobeto, H., M. Menendez, and I. J. Losada, 2021: Future behavior of wind wave extremes due to climate change. *Sci. Rep.*, **11**, 7869, <https://doi.org/10.1038/s41598-021-86524-4>.
- Lorenz, E. N., 1969: Atmospheric predictability as revealed by naturally occurring analogues. *J. Atmos. Sci.*, **26**, 636–646, [https://doi.org/10.1175/1520-0469\(1969\)26<636:APARBN>2.0.CO;2](https://doi.org/10.1175/1520-0469(1969)26<636:APARBN>2.0.CO;2).
- Lott, F. C., N. Christidis, and P. A. Stott, 2013: Can the 2011 East African drought be attributed to human-induced climate change? *Geophys. Res. Lett.*, **40**, 1177–1181, <https://doi.org/10.1002/grl.50235>.
- Min, S.-K., W. Cai, and P. Whetton, 2013: Influence of climate variability on seasonal extremes over Australia. *J. Geophys. Res. Atmos.*, **118**, 643–654, <https://doi.org/10.1002/jgrd.50164>.
- Miralles, O., and A. C. Davison, 2023: Timing and spatial selection bias in rapid extreme event attribution. *Wea. Climate Extremes*, **41**, 100584, <https://doi.org/10.1016/j.wace.2023.100584>.
- National Academies of Sciences, 2016: *Attribution of Extreme Weather Events in the Context of Climate Change*. The National Academies Press, 186 pp.
- National Center for Atmospheric Research, 2024: Known issues—Large Ensemble Community Project (LENS). Community Earth System Model, accessed 29 August 2024, <https://www.cesm.ucar.edu/community-projects/lens/known-issues>.
- Naveau, P., R. Huser, P. Ribereau, and A. Hannart, 2016: Modeling jointly low, moderate, and heavy rainfall intensities without a threshold selection. *Water Resour. Res.*, **52**, 2753–2769, <https://doi.org/10.1002/2015WR018552>.
- Otto, F. E. L., S. Philip, S. Kew, S. Li, A. King, and H. Cullen, 2018: Attributing high-impact extreme events across timescales—A case study of four different types of events. *Climatic Change*, **149**, 399–412, <https://doi.org/10.1007/s10584-018-2258-3>.
- , and Coauthors, 2020: Challenges to understanding extreme weather changes in lower income countries. *Bull. Amer. Meteor. Soc.*, **101**, E1851–E1860, <https://doi.org/10.1175/BAMS-D-19-0317.1>.
- Pascale, S., S. B. Kapnick, T. L. Delworth, and W. F. Cooke, 2020: Increasing risk of another Cape Town “day zero” drought in the 21st century. *Proc. Natl. Acad. Sci. USA*, **117**, 29 495–29 503, <https://doi.org/10.1073/pnas.2009144117>.
- Pendergrass, A. G., 2018: What precipitation is extreme? *Science*, **360**, 1072–1073, <https://doi.org/10.1126/science.aat1871>.
- Perkins, S. E., A. J. Pitman, N. J. Holbrook, and J. McAneney, 2007: Evaluation of the AR4 climate models’ simulated daily maximum temperature, minimum temperature, and precipitation over Australia using probability density functions. *J. Climate*, **20**, 4356–4376, <https://doi.org/10.1175/JCLI4253.1>.
- Philip, S., and Coauthors, 2020: A protocol for probabilistic extreme event attribution analyses. *Adv. Stat. Climate Meteor. Oceanogr.*, **6**, 177–203, <https://doi.org/10.5194/ascmo-6-177-2020>.
- Pinto, I., R. Barimalala, S. Philip, M. Zachariah, S. Sengupta, and M. Vahlberg, 2023: Extreme poverty rendering Madagascar highly vulnerable to underreported extreme heat that would not have occurred without human-induced climate change. Grantham Institute-Climate Change and the Environment Tech. Rep., 32 pp., <https://spiral.imperial.ac.uk/entities/publication/21b00871-bb0e-4140-87ee-3ea8404444a8>.
- Proistosescu, C., and P. J. Huybers, 2017: Slow climate mode reconciles historical and model-based estimates of climate sensitivity. *Sci. Adv.*, **3**, e1602821, <https://doi.org/10.1126/sciadv.1602821>.
- Rigden, A., C. Golden, D. Chan, and P. Huybers, 2024: Climate change linked to drought in southern Madagascar. *npj Climate Atmos. Sci.*, **7**, 41, <https://doi.org/10.1038/s41612-024-00583-8>.
- Robin, Y., and A. Ribes, 2020: Nonstationary extreme value analysis for event attribution combining climate models and observations. *Adv. Stat. Climate Meteor. Oceanogr.*, **6**, 205–221, <https://doi.org/10.5194/ascmo-6-205-2020>.
- Seneviratne, S., and Coauthors, 2021: Weather and climate extreme events in a changing climate. *Climate Change 2021: The Physical Science Basis*, V. Masson-Delmotte et al., Eds.,

- Cambridge University Press, 1513–1766, <https://doi.org/10.1017/9781009157896.013>.
- Shepherd, T. G., and Coauthors, 2018: Storylines: An alternative approach to representing uncertainty in physical aspects of climate change. *Climate Change*, **151**, 555–571, <https://doi.org/10.1007/s10584-018-2317-9>.
- Sippel, S., and Coauthors, 2024: Could an extremely cold central European winter such as 1963 happen again despite climate change? *Wea. Climate Dyn.*, **5**, 943–957, <https://doi.org/10.5194/wcd-5-943-2024>.
- Stott, P. A., D. A. Stone, and M. R. Allen, 2004: Human contribution to the European heatwave of 2003. *Nature*, **432**, 610–614, <https://doi.org/10.1038/nature03089>.
- , N. P. Gillett, G. C. Hegerl, D. J. Karoly, D. A. Stone, X. Zhang, and F. Zwiers, 2010: Detection and attribution of climate change: A regional perspective. *Wiley Interdiscip. Rev.: Climate Change*, **1**, 192–211, <https://doi.org/10.1002/wcc.34>.
- , and Coauthors, 2016: Attribution of extreme weather and climate-related events. *Wiley Interdiscip. Rev.: Climate Change*, **7**, 23–41, <https://doi.org/10.1002/wcc.380>.
- The World Weather Attribution Project, 2024: World Weather Attribution. Accessed 29 August 2024, <https://www.worldweatherattribution.org/>.
- Trevino, A. M., K. A. McKinnon, and P. J. Huybers, 2020: Extremely wrong: When limited data leads to biases in general extreme values distribution estimates. *2020 AGU Fall Meeting*, Online, Amer. Geophys. Union, Abstract A147-0010, <https://ui.adsabs.harvard.edu/abs/2020AGUFMA147.0010T/abstract>.
- Tziperman, E., 2022: *Global Warming Science: A Quantitative Introduction to Climate Change and Its Consequences*. Princeton University Press, 336 pp.
- Ummenhofer, C. C., M. H. England, P. C. McIntosh, G. A. Meyers, M. J. Pook, J. S. Risbey, A. S. Gupta, and A. S. Taschetto, 2009: What causes southeast Australia's worst droughts? *Geophys. Res. Lett.*, **36**, L04706, <https://doi.org/10.1029/2008GL036801>.
- van der Wiel, K., and Coauthors, 2017: Rapid attribution of the August 2016 flood-inducing extreme precipitation in South Louisiana to climate change. *Hydrol. Earth Syst. Sci.*, **21**, 897–921, <https://doi.org/10.5194/hess-21-897-2017>.
- van Oldenborgh, G. J., A. Van Urk, and M. Allen, 2012: The absence of a role of climate change in the 2011 Thailand floods. *Bull. Amer. Meteor. Soc.*, **93**, 1047–1049.
- , F. E. L. Otto, K. Haustein, and K. AchutaRao, 2016: The heavy precipitation event of December 2015 in Chennai, India. *Bull. Amer. Meteor. Soc.*, **97**, S87–S91, <https://doi.org/10.1175/BAMS-D-16-0129.1>.
- , and Coauthors, 2017: Attribution of extreme rainfall from Hurricane Harvey, August 2017. *Environ. Res. Lett.*, **12**, 124009, <https://doi.org/10.1088/1748-9326/aa9ef2>.
- , and Coauthors, 2021: Attribution of the Australian bushfire risk to anthropogenic climate change. *Nat. Hazards Earth Syst. Sci.*, **21**, 941–960, <https://doi.org/10.5194/nhess-21-941-2021>.
- Vautard, R., and P. Yiou, 2009: Control of recent European surface climate change by atmospheric flow. *Geophys. Res. Lett.*, **36**, L22702, <https://doi.org/10.1029/2009GL040480>.
- Yamaguchi, M., J. C. L. Chan, I.-J. Moon, K. Yoshida, and R. Mizuta, 2020: Global warming changes tropical cyclone translation speed. *Nat. Commun.*, **11**, 47, <https://doi.org/10.1038/s41467-019-13902-y>.
- Yiou, P., R. Vautard, P. Naveau, and C. Cassou, 2007: Inconsistency between atmospheric dynamics and temperatures during the exceptional 2006/2007 fall/winter and recent warming in Europe. *Geophys. Res. Lett.*, **34**, L21808, <https://doi.org/10.1029/2007GL031981>.
- Yoshida, K., M. Sugi, R. Mizuta, H. Murakami, and M. Ishii, 2017: Future changes in tropical cyclone activity in high-resolution large-ensemble simulations. *Geophys. Res. Lett.*, **44**, 9910–9917, <https://doi.org/10.1002/2017GL075058>.
- Zwiers, F. W., X. Zhang, and Y. Feng, 2011: Anthropogenic influence on long return period daily temperature extremes at regional scales. *J. Climate*, **24**, 881–892, <https://doi.org/10.1175/2010JCLI3908.1>.