Lawrence Berkeley National Laboratory

Recent Work

Title

Quantifying the effect of interannual ocean variability on the attribution of extreme climate events to human influence

Permalink

<https://escholarship.org/uc/item/5jf0x9r9>

Journal

Climate Dynamics, 49(9-10)

ISSN

0930-7575

Authors

Risser, MD
Stone, DA
Paciorek, CJ
et al.

Publication Date

2017-11-01

DOI

10.1007/s00382-016-3492-x

Peer reviewed

Quantifying the effect of interannual ocean variability on the attribution of extreme climate events to human influence

**Mark D. Risser, Dáithí A. Stone,
Christopher J. Paciorek, Michael
F. Wehner & Oliver Angélil**

Climate Dynamics

Observational, Theoretical and
Computational Research on the Climate
System

ISSN 0930-7575

Volume 49

Combined 9-10

Clim Dyn (2017) 49:3051-3073

DOI 10.1007/s00382-016-3492-x

Your article is protected by copyright and all rights are held exclusively by Springer-Verlag Berlin Heidelberg. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Quantifying the effect of interannual ocean variability on the attribution of extreme climate events to human influence

Mark D. Risser¹ · Dáithí A. Stone² · Christopher J. Paciorek³ · Michael F. Wehner² · Oliver Angéllil⁴

Received: 28 June 2016 / Accepted: 4 December 2016 / Published online: 12 January 2017
© Springer-Verlag Berlin Heidelberg 2017

Abstract In recent years, the climate change research community has become highly interested in describing the anthropogenic influence on extreme weather events, commonly termed “event attribution.” Limitations in the observational record and in computational resources motivate the use of uncoupled, atmosphere/land-only climate models with prescribed ocean conditions run over a short period, leading up to and including an event of interest. In this approach, large ensembles of high-resolution simulations can be generated under factual observed conditions and counterfactual conditions that might have been observed in the absence of human interference; these can be used to estimate the change in probability of the given event due to anthropogenic influence. However, using a prescribed ocean state ignores the possibility that estimates of attributable risk might be a function of the ocean state. Thus, the uncertainty in attributable risk is likely underestimated, implying an over-confidence in anthropogenic influence. In this work, we estimate the year-to-year variability in calculations of the anthropogenic contribution to extreme

weather based on large ensembles of atmospheric model simulations. Our results both quantify the magnitude of year-to-year variability and categorize the degree to which conclusions of attributable risk are qualitatively affected. The methodology is illustrated by exploring extreme temperature and precipitation events for the northwest coast of South America and northern-central Siberia; we also provides results for regions around the globe. While it remains preferable to perform a full multi-year analysis, the results presented here can serve as an indication of where and when attribution researchers should be concerned about the use of atmosphere-only simulations.

Keywords Climate change · Anthropogenic · Event attribution · Extreme weather · Risk ratio

1 Introduction

In recent years, public interest has motivated the climate change research community to become highly interested in describing the influence of anthropogenic emissions on extreme weather events, commonly termed “event attribution” (National Academies of Sciences, Engineering, and Medicine 2016). Limitations in the observational record motivate the use of climate models to estimate anthropogenic influence and risk: first, extremes are (by definition) rare, meaning their occurrence in the observational record is sparse; second, no observations at all are available for a counterfactual world without human influence. The use of observations from a century ago has issues both in terms of representing anthropogenic influence and of data quality. Instead, climate models can be used to simulate the weather conditions in both a factual world driven by observed boundary conditions

Electronic supplementary material The online version of this article (doi:10.1007/s00382-016-3492-x) contains supplementary material, which is available to authorized users.

✉ Mark D. Risser
mdrisser@lbl.gov

¹ Climate and Ecosystem Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

² Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

³ Department of Statistics, University of California Berkeley, Berkeley, CA, USA

⁴ Climate Change Research Centre, University of New South Wales, Kensington, USA

(e.g., greenhouse gas concentrations, solar luminosity) and a counterfactual world driven by what those boundary conditions might have been in the absence of historical human influence on the climate. However, extreme weather is often localized, which requires these models to be run at a resolution fine enough to resolve the physics in the model that generate extreme weather. And, for a probabilistic approach examining rare events, large ensembles of simulations are needed for both scenarios.

This leads to a computational problem: fully-coupled, high resolution models that collectively model land, ocean, and atmospheric processes come with an extremely high computational cost. Fully-coupled models have long memory, meaning that it can take decades to centuries of model time to move away from biases introduced by the initial conditions. An additional issue is that significant biases can arise in full models of the climate system, wherein biases in ocean conditions can induce biases in the atmospheric state, which in turn may reinforce the ocean biases.

To navigate these difficulties when considering atmospheric extremes, an alternative approach involves using uncoupled atmosphere/land-only models with prescribed ocean conditions. The boundary conditions for both the factual and counterfactual worlds now include the observed ocean state, with the counterfactual values adjusted according to a spatially and seasonally evolving estimate of the warming due to anthropogenic emissions. These models have less spin-up time, and large ensembles can be obtained much more quickly. Additionally, fixing the ocean conditions can provide some benefit in that it removes dependence on the quality of the ocean model's simulation of low frequency processes like El Niño and on biases in the ocean model's ocean surface state.

Of course, this approach involves some important assumptions relating to the ocean surface state and the way the atmosphere interacts with the ocean. While researchers may be interested in the conditional case (e.g. during El Niño events), when the results are interpreted as a general overall assessment of the role of anthropogenic emissions, these conditions become important assumptions. All attribution analyses involve some elements of uncertainty (Jeon et al. 2016; Hansen et al. 2014; Pall et al. 2011), but particular types of uncertainty are specific to the atmosphere/land-only modeling approach. First and foremost, the strategy of fixing the sea surface temperatures (SSTs) does not account for the internal variability of the ocean system. This involves three important assumptions.

1. **Invariance to ocean state:** The anthropogenic influence on extreme weather is assumed to be independent of the ocean state, so for instance anthropogenic influence is considered identical during El Niño events as during La Niña events, ignoring the possibility that the

atmospheric response to these two types of events may have different sensitivities to anthropogenic forcing.

2. **No change in ocean variability:** The nature of ocean variability, for instance the frequency and spatial structure of El Niño events, is assumed to be unaffected by anthropogenic emissions, beyond any spatial or seasonal gradients imposed by the attributable ocean warming estimate.
3. **Unimportance of atmosphere–ocean interaction:** Short time-scale atmosphere–ocean interactions, for instance that might occur under tropical cyclones, are assumed to be unimportant.

This paper specifically examines the accuracy of the first assumption.

The impact of ocean variability on atmosphere-model-based event attribution assessments has yet to be quantified. Identifying the effect of oceanic variability boils down to investigating how the experimental design of a single-year, atmosphere-only attribution framework impacts attribution statements. For example, attribution statements may be highly sensitive to the specific temporal period used for the study, over periods short enough that a long-term trend is unimportant (Otto et al. 2015; Wolski et al. 2014). These sensitivities might vary for different events and different regions in the world, and thus it would be helpful to identify events and regions of the world for which atmospheric model designs are adequate for reasonably accurate event attribution assessment.

In this work, we develop and implement a hierarchical statistical model that allows us to address the uncertainty introduced in attribution studies from the use of uncoupled, atmosphere/land-only model simulations (in addition to sampling uncertainty from the limited number of simulations). Specifically, this framework allows us to quantify the effect of ocean variability on statements of risk, while incorporating and adjusting for long-term trends. Furthermore, the model is fit using a Bayesian framework, which allows for the development of a diagnostic tool for identifying event types and regions of the globe for which single-year, atmosphere-only climate simulations may be sufficient for assessing the anthropogenic contribution to risk. Finally, we present a framework within which future single-year, atmosphere-only event attribution studies can assess the robustness of their results to the effect of ocean variability.

The paper is organized as follows. In Sect. 2, we introduce our event attribution framework and describe the specific climate model simulations used in this analysis. In Sect. 3, we outline the hierarchical statistical model, including a discussion of the implied model for the risk ratio. Section 4 introduces several diagnostic tools for applying estimates of the effect of ocean variability in a general

event attribution study. In Sect. 5, we apply the statistical model to two case studies, and in Sect. 6 we describe and apply the diagnostic tool generally to several dozen regions around the world. Section 7 concludes the paper.

2 Attribution framework

2.1 Climate model simulations

We examine the frequency of extreme monthly values of temperature and precipitation in two ensembles of simulations of the CAM5.1 global atmosphere/land climate model. The model is run in its conventional $\sim 1^\circ$ longitude/latitude configuration (Neale et al. 2012). Simulations have been run under the experiment protocols of the C20C+ Detection and Attribution Project (Stone and Pall 2016), following two historical scenarios (Angélil et al. 2016b). The first set of simulations (ALL) is driven by observed boundary conditions of atmospheric chemistry (greenhouse gases, tropospheric and stratospheric aerosols, ozone), solar luminosity, land use/cover, and the ocean surface (temperature and ice coverage). The second set of simulations (NAT) is driven by what observed boundary conditions might have been in the absence of historical anthropogenic emissions: the anthropogenic component of atmospheric chemistry is set to year-1855 values, ocean temperatures are cooled by a seasonally- and spatially-varying estimate of the warming attributable to anthropogenic emissions (that is the same across all simulations), and sea ice concentrations are adjusted for consistency with the ocean temperatures (Stone and Pall 2016). Simulations within a scenario differ only in the starting conditions. Here we examine data during the 1982–2013 period; while simulations are available for some earlier years, we want to ensure that incorporation of satellite data into the sea surface temperature product used by the simulations in 1981 does not influence our conclusions (Hurrell et al. 2008). Furthermore, this provides a period in which risk ratios should be of the same order of magnitude across years (even if exhibiting a substantial long-term trend), which would not be the case as one goes further back in time. In this analysis, we make this assumption for second order terms but not for first order terms. For each scenario we have 50 simulations during 1982–1996, with an additional 50 (hence 100 total) during 1997–2010, and an additional 300 (400 total) during 2011–2013 (Table 1). The data and further details on the simulations are available at <http://portal.nersc.gov/c20c>.

2.2 Using the risk ratio for attributional statements

In order to make attribution statements for the impact of anthropogenic influences on the occurrence of extreme

Table 1 Ensemble sizes for the CAM5.1 simulations from 1982–2013

Year range(<i>t</i>)	Ensemble size (<i>n_t</i>)
1982–1996	50
1997–2010	100
2011–2013	400

events, we follow a probabilistic extreme event attribution approach (Allen 2003; Stone and Allen 2005; Hansen et al. 2014), by considering occurrence probabilities p_A and p_N . These quantities represent the probability that a predefined event will occur over a predefined spatial and temporal domain, with p_A representing the observed world or “world as it is” (henceforth the so-called ALL forcings scenario), including both natural and anthropogenic influences on the climate, and p_N representing the counterfactual scenario of the “world as it may have been” (henceforth the so-called NAT forcings scenario) without anthropogenic effects. We consider the ratio of these probabilities, or the risk ratio

$$RR = \frac{p_A}{p_N},$$

which mirrors the measurement of risk used in epidemiology and environmental law (Stone and Allen 2005) and has a number of desirable properties (National Academies of Sciences, Engineering, and Medicine 2016). A risk ratio of larger than one indicates that anthropogenic influences have increased the likelihood of an event’s occurrence, while a risk ratio of less than one indicates a decreased likelihood of occurrence.

In this work, we are interested in estimating the risk ratio over time, at a total of T time increments. In what follows, the time increments represent years, but this framework could similarly be used for months, seasons, or decades, for example. Therefore, we must consider

$$RR_t = \frac{p_{At}}{p_{Nt}}, \quad t = 1, \dots, T,$$

where p_{At} and p_{Nt} are the occurrence probabilities for a predefined time increment t .

2.3 Modeling probabilities and event definition

In general, while there are a variety of approaches for estimating p_A and p_N , two popular methods are the nonparametric (or binomial) approach and extreme value analysis. The binomial approach is desirable in that it makes no assumptions about the underlying distribution of the variable of interest (e.g., precipitation or temperature), however it becomes essentially useless when the event of interest

is extremely rare or unobserved. Extreme value analysis, on the other hand, requires additional assumptions on the behavior of the variable of interest but is able to estimate the probability of extremely rare events, even those that do not occur frequently or at all in an ensemble or the observational record.

In what follows, we considered events of a pre-specified magnitude, and simply used an empirical quantile to define a threshold for determining whether or not a particular event is “extreme.” Furthermore, for simplicity, we chose a quantile that is extreme but not too extreme: namely, the one in ten year event. In this case, this nonparametric binomial approach is sufficient for modeling probabilities.

3 Hierarchical statistical model for ALL and NAT probabilities

Unlike other approaches that have explored estimation and uncertainty quantification for individual time points or years (Pall et al. 2011), recall that our goal is to make systematic statements about the behavior of the risk ratio for multiple time points. In doing so, we can make statements about the effect of various explanatory variables on event probabilities and the risk ratio in addition to quantifying interannual variability. A hierarchical statistical modeling framework allows us to address each of these questions by borrowing information over time.

Using a nonparametric approach for modeling the probabilities, the first level of the hierarchical model relates the climate model simulations to the occurrence probabilities. While we are interested in modeling the risk ratio across different years, in order to account for seasonality we define monthly random variables

$$\begin{aligned} Z_{Aij} &= \text{number of ALL simulations in month } j \text{ of year } t \\ &\quad \text{in which the event occurred} \\ Z_{Ntj} &= \text{number of NAT simulations in month } j \text{ of year } t \\ &\quad \text{in which the event occurred;} \end{aligned} \quad (1)$$

each of these random variables is modeled in a statistical framework as arising from a binomial experiment with n_t total trials and corresponding success probabilities p_{ktj} , for $k \in \{A, N\}$, $j = 1, \dots, 12$, and $t = 1, \dots, T$. Furthermore, conditional on the probabilities $\mathbf{p} = \{p_{ktj}\}$, the random variables $\mathbf{Z} = \{Z_{ktj}\}$, are independent. In other words, the joint probability mass function (pmf) of $\mathbf{Z}$ conditional on $\mathbf{p}$ is

$$p(\mathbf{Z}|\mathbf{p}) = \prod_{k \in \{A, N\}} \prod_{t=1}^T \prod_{j=1}^{12} p(Z_{ktj}; p_{ktj}), \quad (2)$$

where the marginal pmfs $p(Z_{ktj}; p_{ktj})$ are binomial.

The second level of the hierarchy consists of a model that allows the unknown parameters to vary based on

covariate information, while also tying together parameters that correspond to the same time period (i.e., p_{Atj} and p_{Ntj} for common t and j). For this level, we used

$$\text{logit } p_{ktj} \equiv \log \left(\frac{p_{ktj}}{1 - p_{ktj}} \right) = \mathbf{x}_{kt}^\top \boldsymbol{\beta}_k + \alpha_t + \delta_t 1_{\{k=A\}} + \gamma_j, \quad (3)$$

for $k \in \{A, N\}$, $t = 1, \dots, T$, and $j = 1, \dots, 12$. Here, $\mathbf{x}_{kt} = (1, x_{kt1}, \dots, x_{ktp})$ is a vector of covariates for scenario $k \in \{A, N\}$, $\boldsymbol{\beta}_k = (\beta_{k0}, \dots, \beta_{kp})$ is a scenario-specific vector of unknown regression coefficients, and $1_{\{\cdot\}}$ is an indicator function. The strategy in (3) is a mixed effects logistic regression, a special case of the more general statistical approach of generalized linear mixed models (see, e.g., McCulloch and Searle 2004). Intuitively, including the scenario-specific covariates $\mathbf{x}_{kt}$ accounts for any long-term trends present in p_N and p_A .

The final level of the hierarchy ties together the year-specific effects $\{\alpha_t, \delta_t; t = 1, \dots, T\}$ and month-specific effects $\{\gamma_j; j = 1, \dots, 12\}$ in order to allow for borrowing of information across years and months (called “partial pooling”). This level of the model specifies

$$\alpha_t \stackrel{\text{iid}}{\sim} N(0, \tau^2), \quad \delta_t \stackrel{\text{iid}}{\sim} N(0, \sigma^2), \quad \gamma_j \stackrel{\text{iid}}{\sim} N(0, \omega^2) \quad (4)$$

where “iid” stands for “independent and identically distributed” and $N(a, b)$ denotes a univariate Gaussian (normal) distribution with mean a and variance b . In terms of the probabilities, the α_t terms represent yearly deviations in event probabilities (above and beyond deviations due to the long-term trend) that are common to both ALL and NAT scenarios (e.g., El Niño or La Niña events), while the δ_t terms represent deviations above and beyond the α_t that are specific to the ALL scenario (for example, the magnitude or probability of an event for ALL in an El Niño or La Niña year may be different than for NAT). These δ_t terms have a specific importance for how the risk ratio is modeled, and will be explored further in Sect. 3.1. Finally, the γ_j terms represent within-year deviations from the average yearly probabilities, which account for seasonality: for example, in the extratropics of the northern hemisphere, the probability of a hot extreme is essentially zero from October to March, while a hot extreme might be relatively common in July. Therefore, γ_{12} (December) would likely be a large negative number (corresponding to smaller probabilities on the logit scale) while γ_7 (July) would likely be a large positive number (corresponding to larger probabilities on the logit scale).

Several notes should be made regarding the statistical modeling of the monthly effects. First, because the population for the γ_j terms represents only 12 months (and because this sample of 12 represents all possible months), we impose the restriction that $\sum_{j=1}^{12} \gamma_j = 0$ to ensure that the sample

average is equal to the population average; also, this constraint ensures we can interpret the δ_t terms directly with respect to the yearly risk ratio (see Sect. 3.1). Second, note that the γ_j terms depend on neither k (the scenario) nor t (the year); the implication is that we impose the same seasonality on the simulations in both ALL and NAT and a constant seasonality over time. While this may be a restrictive assumption, we make this choice for several reasons. First, it simplifies the parameter interpretations and allows us to use a single parameter to capture the variability in the risk ratio over time (again see Sect. 3.1). Also, where the annual cycle in probability has a strong seasonal peak, it is only really the values for that season that matter. Finally, note that the exchangeable model for the γ_j terms in (4) does not require that the seasonality be modeled smoothly in time. We acknowledge that our approach is somewhat simplistic, but in doing so we avoid having to customize the model to account for the differing seasonalities of each region and each event type. Instead, the simplified approach of (4) is highly flexible and can automatically adjust to the type of seasonality present for a particular event type and region (accounting for both the length of seasonality and the magnitude of differences across seasons).

The partial pooling of (4) is in contrast to *complete pooling*, which would fix each $\alpha_t \equiv \alpha$, $\delta_t \equiv \delta$, and $\gamma_j \equiv \gamma$, and *no pooling*, which would allow each α_t , δ_t , and γ_j to be estimated independently of all others. Partial pooling, a popular strategy in random effects modeling (Gelman et al. 2013) and meta analysis (DerSimonian and Laird 1986), offers a compromise between these two by allowing each year (or month) to have its own effect (i.e., unique α_t , δ_t , and γ_j) while borrowing strength across years (or months) by requiring these effects to come from a common distribution. The partial pooling done here is a standard form of statistical shrinkage that seeks to increase the signal to noise ratio by using the data from all years to help inform the time-point specific effects, with the degree to which information is shared across time points controlled by the data itself.

Up until this point, the model has been presented in general form; in this application, the model was fit within a Bayesian paradigm. In a Bayesian analysis, the unknown parameters are assumed to be random variables, and hence must be assigned a prior distribution. This distribution summarizes all knowledge about the unknown parameters *a priori*, or prior to observing data. For notational simplicity, set $\theta = (\alpha, \delta, \gamma, \beta_A, \beta_N, \tau^2, \sigma^2, \omega^2)$; denote the prior as $p(\theta)$, and rewrite the likelihood (2) in terms of these parameters, i.e., $p(\mathbf{Z}|\theta)$. Bayes' Theorem supplies the required machinery to update the prior distribution with observed data (here, $\mathbf{Z}$) to arrive at the posterior distribution:

$$p(\theta|\mathbf{Z}) = \frac{p(\mathbf{Z}|\theta)p(\theta)}{\int_{\theta} p(\mathbf{Z}|\theta)p(\theta)d\theta}, \quad (5)$$

i.e., the updated knowledge about θ after observing $\mathbf{Z}$. While the posterior (5) is not available in closed form (due to the intractable integral in the denominator of (5); this is often the case), Markov chain Monte Carlo (MCMC) methods can be used to obtain joint samples from the posterior distribution of θ , upon which all subsequent inference is based.

While one advantage of using Bayesian methods (in general) is that any known prior information relating to the parameters can be incorporated by way of the prior distribution $p(\theta)$, we instead chose to use a “noninformative” prior distribution, which specifies essentially no information about the parameters of the model (in order to avoid any prior biases).

Additional documentation, including more information on the prior specification and computational details for fitting the hierarchical model, is available in the supplemental materials section online. Furthermore, software for fitting the hierarchical model using the R programming language and the data sets used in Sect. 5 are available for download at http://bitbucket.org/markdrisser/timerr_package.

3.1 Risk ratio modeling

Now, consider what the three-level model specified by (2), (3), and (4) imply about a model for the risk ratio. While the occurrence probabilities have been defined monthly, we are actually interested in averaging over the seasonality to arrive at yearly probabilities, defined as $p_{kt} = \frac{1}{12} \sum_{j=1}^{12} p_{ktj}$. The reason for doing so is that we are less interested in describing how the risk ratio changes monthly (e.g., how the risk ratio in July, 1995 is different from that of August, 1995) and more interested in how the risk ratio changes over longer time scales (years). The fact that the occurrence probabilities are defined for each month (i.e., p_{ktj}) is motivated by the need to account for seasonality.

Therefore, the risk ratio for an individual year, averaged over all months (from (3)), is

$$RR_t = \frac{p_{At}}{p_{Nt}} = \frac{\frac{1}{12} \sum_{j=1}^{12} \text{logit}^{-1}(\mathbf{x}_{At}^T \boldsymbol{\beta}_A + \alpha_t + \delta_t + \gamma_j)}{\frac{1}{12} \sum_{j=1}^{12} \text{logit}^{-1}(\mathbf{x}_{Nt}^T \boldsymbol{\beta}_N + \alpha_t + \gamma_j)}. \quad (6)$$

To understand the implications of (6) more clearly, note that in dealing with extreme events, the probabilities are often small (i.e., near zero), in which case $\text{logit}(p) = \log(p) - \log(1-p) \approx \log(p)$ and $\text{logit}^{-1}(x) \approx \exp\{x\}$. Therefore, we can approximate the risk ratio (6) as

$$\begin{aligned} RR_t &\approx \frac{\frac{1}{12} \sum_{j=1}^{12} \exp\{\beta_{A0} + \beta_{A1}x_{At1} + \alpha_t + \delta_t + \gamma_j\}}{\frac{1}{12} \sum_{j=1}^{12} \exp\{\beta_{N0} + \beta_{N1}x_{Nt1} + \alpha_t + \gamma_j\}} \\ &= \frac{\exp\{\beta_{A0} + \beta_{A1}x_{At1} + \alpha_t + \delta_t\} \times \frac{1}{12} \sum_{j=1}^{12} \exp\{\gamma_j\}}{\exp\{\beta_{N0} + \beta_{N1}x_{Nt1} + \alpha_t\} \times \frac{1}{12} \sum_{j=1}^{12} \exp\{\gamma_j\}} \\ &= \exp\{\beta_{A0} + \beta_{A1}x_{At1} + \delta_t\} / \exp\{\beta_{N0} + \beta_{N1}x_{Nt1}\} \\ &= RR_0 \times \exp\{\beta_{A1}x_{At1} - \beta_{N1}x_{Nt1}\} \times \exp\{\delta_t\} \end{aligned} \quad (7)$$

where, for simplicity, we assume that there is only a single covariate of interest along with an intercept (i.e., $\mathbf{x}_{kt}^\top \boldsymbol{\beta}_k = \beta_{k0} + x_{kt1} \beta_{k1}$). From (7), the yearly risk ratio is approximated by three pieces: first, $RR_0 = \exp\{\beta_{A0} - \beta_{N0}\}$, or the “baseline” risk ratio for the entire time interval; second, $\exp\{\beta_{A1}x_{At1} - \beta_{N1}x_{Nt1}\}$, a multiplicative scaling due to the covariates; and finally $\exp\{\delta_t\}$, a scaling for the risk ratio in a particular year. Note that the α_t term, which appears in both numerator and denominator and represents yearly deviations in event probabilities for both scenarios, cancels out. Furthermore, because the γ_j are constant across scenario, the impact of seasonality also cancels out. Thus, we can interpret $\sigma^2 = \text{Var}(\delta_t) \approx \text{Var}(\log RR_t)$ as the effect of oceanic internal variability on the risk ratio or, alternatively, general variability in the p_{At} above and beyond variability in the p_{Nt} . This single parameter represents the effect of prescribed ocean variability on the risk ratio after adjusting for a long-term trend by way of relevant covariate information.

4 A diagnostic tool for event attribution sensitivity to oceanic variability

In addition to the sensitivity of the quantitative results, a natural question to ask in light of quantifying the effect of ocean variability is: do the qualitative results of an event attribution study change based on the effect of oceanic internal variability? Put another way, for the years considered in the analysis (1982–2013), for what proportion of years would the qualitative conclusion of an attribution analysis have changed based on the effect of ocean internal variability, assuming a present-day level of anthropogenic influence? In this framework, a “change” in an attribution analysis means concluding the presence of an anthropogenic effect (in one direction or another) instead of concluding the absence of an anthropogenic effect (or vice versa), and is relative to the event type and how strong the evidence must be to establish anthropogenic influence. It may also include changes in qualitative descriptions of the degree of anthropogenic influence, for instance “moderate” to “small”. The answer to this question may not be obvious from the numerical value of σ alone, as it also depends on the proximity of RR_0 to the boundaries between qualitative categories. Note that the statistical modeling approach allows us to impose a “stationary” climate (with present-day anthropogenic influence) by adjusting the probability estimates (and therefore the risk ratio estimates) by fixing each year to have the same covariate values, allowing comparison across each of the years. (For an illustration, see Sect. 6).

We formalize the question above by estimating π , the proportion of years for which the qualitative inference of

a 1-year, atmosphere-only attribution analysis remains the same over many different years (all arising from a stationary climate) in spite of internal variability. In other words, π is the proportion of all years under consideration for which the risk ratio significantly exceeds (or does not exceed) a particular cutoff. Intuitively, large π (near 1) indicates that the interval estimate of the risk ratio exceeds (or does not exceed) a threshold for most years; small π (near 0) indicates that the interval estimate for the risk ratio exceeds (or does not exceed) the threshold for only a few years. For each region and event type, π can be estimated using the MCMC samples (more details will be provided in Sect. 6).

We set our criterion as follows: if the interval estimate of π for a particular region/event type and specified cutoff, say (π_L, π_U) , is entirely included in the set $[0, 0.05] \cup (0.95, 1]$, then the risk ratios are homogeneous (relative to the threshold) and the qualitative inference of an attribution analysis is insensitive to the specific year chosen for analysis. On the other hand, if the interval estimate for π falls completely inside $[0.05, 0.95]$, then there is evidence of heterogeneity in both the risk ratio (relative to the threshold) and the corresponding qualitative statement about the risk. If the estimate of π includes parts of both $[0, 0.05] \cup (0.95, 1]$ and $[0.05, 0.95]$, then while there is some evidence of heterogeneity, we are unable to make a definitive statement one way or the other.

Therefore, we can use this tool to classify each region and event type into one of the following: in answer to the question “Does the effect of ocean variability impact the consistency of attribution statements?”, either

1. Yes, the conclusions of an attribution study vary over time; i.e., $(\pi_L, \pi_U) \subset [0.05, 0.95]$.
2. Most likely yes, but the results are inconclusive; i.e., $(\pi_L, \pi_U) \not\subset [0.05, 0.95]$ but $(\pi_L, \pi_U) \cap [0.05, 0.95] \neq \{\emptyset\}$.
3. No, the conclusions of an attribution study do not vary over time (relative to a threshold); i.e., either $(\pi_L, \pi_U) \subset [0, 0.05]$ or $(\pi_L, \pi_U) \subset (0.95, 1]$.

This strategy will be illustrated in Sect. 6.

5 Case studies using climate model ensembles

The methods outlined in Sect. 3 are illustrated with two case studies, both using the ensembles of the CAM5.1 atmosphere/land model from the C20C+ D&A Project described in Sect. 2. The ensembles are used to estimate risk ratio for monthly (calendar month) hot, cold, and wet extremes; dry extremes are omitted here because the zero-bound is encountered regularly for many regions on a monthly basis.

Temperature and precipitation are aggregated spatially into 58 geopolitical land-regions, which correspond to the regions used for the Weather Risk Attribution Forecast (WRAF; <http://www.csag.uct.ac.za/~daithi/forecast>), shown in Fig. 1. Within each region, temperature and precipitation are further aggregated by calendar month. For year t and month j , we then have n_{tj} measurements, where n_t is the number of simulations in year t . The one-in-ten year event for monthly data corresponds to the $100(1-0.1/12)$ th = 99.17th percentile for large extremes (hot, wet) and the $100(0.1/12)$ th = 0.83th percentile for small extremes (cold). In order to avoid an inappropriate bias from the more recent years with larger ensemble sizes, the empirical threshold percentiles were calculated from the 50-member ensemble that covers the entire time period of the study, from 1982–2013. For cold and wet events, the empirical percentiles were calculated from the ALL simulations; for hot events, the empirical percentiles were calculated from the NAT simulations (in order to ensure sufficient sampling in the ALL simulations, as one in ten year hot events in ALL typically do not occur in NAT). Our interest here is in the probability (and risk ratio) of extreme monthly events, irrespective of when they occur during the year; hence we average over the monthly probabilities to obtain the yearly probabilities as defined in Sect. 3.1.

The case studies explore the risk ratio from year to year, and consider two WRAF regions (labelled in Fig. 1): (a) the southern Andean Community (comprising Peru and Bolivia) and (b) Krasnoyarsk (a federal subject in central Russia). As will be illustrated, these two regions were chosen as case studies because they exhibit very different behavior in monthly extremes and their relation to sea

surface conditions. The empirical probabilities of each event type for the southern Andean Community and Krasnoyarsk are provided in Appendix 2.

Because our interest is in understanding possible limitations of the atmospheric modeling approach arising from the specifics of the experiment design, rather than other possible limits arising from factors like the quality of the climate model, we do not perform any evaluation of the CAM5.1 model as part of our analysis. However, as part of their reassessments of recent event attribution studies, Angélil et al. (2016b) included comparisons against a range of observational products as part of their criteria for pursuing with the reassessments using the CAM5.1 simulations used here; the events covered regions of similar size to those examined here, for durations ranging from one to several months. Angélil et al. (2016a) have also compared the frequency of daily temperature and precipitation extremes against those estimated from reanalysis products for the same 58 regions studied here, and concluded that in general any disagreement lies within the uncertainty encapsulated by the range of reanalysis products.

Recall that the δ_t terms in (3) capture variation above and beyond an overall trend in the RR_t . As a way to describe the overall trend in the extreme event probabilities, a scenario-specific global mean temperature covariate was used. More specifically, the covariate used is the average surface air temperature over land and ocean between 50°S and 50°N, calculated as an ensemble average from the 50 members that covered the entire time period from 1982 to 2013 (the average was smoothed using the low-pass filter described in Appendix 3.A of Trenberth et al. 2007). The scenario-specific values are plotted in Fig. 2. The use

Fig. 1 Weather risk attribution forecast (WRAF) regions (<http://www.csag.uct.ac.za/~daithi/forecast>). The two regions used as case studies are labelled as: **a** the southern Andean Community and **b** Krasnoyarsk

Fig. 2 Global mean temperature measurements between 50°S and 50°N, calculated as an land-ocean ensemble average from the 50 members that covered the entire time period from 1980 to 2013, using the low-pass filter described in Appendix 3.A of Trenberth et al. (2007)

of actual values, rather than say a linear trend fit, means that we are making no assumptions of the temporal shape of the climatic response to external forcings; the latitude range prevents strong influence from small variations in the sea ice edge.

5.1 Case study 1: Southern Andean community

The temporal estimates of the event probabilities and risk ratio for each event type over the southern Andean Community are shown in Fig. 3. For hot extremes note that high temperature events are extremely common in the ALL forcings scenario (the p_{At} are as large as 0.6) but very unlikely in the NAT scenario, yielding median risk ratio estimates and also lower bounds on those estimates that are extremely large (with $p_N \approx 0$, $RR \approx \infty$). Even though they occur on average 0.83% of the time (by definition), hot events for NAT are essentially only possible during large El Niño years (e.g., 1983 and 1998), but even then the risk ratio estimates are significantly larger than 1. Even after adjusting for the long term trend imposed by the changes in global temperature, the posterior distribution for σ for hot events in Fig. 5 shows that ocean variability has a very large impact on the log risk ratio, the largest for any event type in this region. This is not surprising, as the southern Andean Community is known to be strongly affected by ocean conditions in the eastern tropical Pacific.

In terms of the probability and risk ratio estimates, Fig. 3 shows that cold extremes have essentially the opposite behavior of hot extremes for this region: extremely cold events are common in the NAT scenario and very uncommon (on average 0.83% by definition) in the ALL scenario. As a result, the risk ratio estimates are extremely small and significantly less than one for all years, particularly for the

most recent years. From Fig. 5 we see that there is still a large effect of the ocean variability on log risk for cold events, albeit somewhat less so than for hot events.

The results for precipitation are perhaps the most interesting for the southern Andean Community. For wet extremes, the event probabilities in each scenario are much more comparable, leading to risk ratio estimates much closer to 1. For the ALL scenario, it is quite easy to pick out the biggest El Niño events, as these years yield the largest event probabilities (1983, 1989, 1998, and 2011). However, note that these years also correspond to large probabilities in the NAT scenario, with the result that the median risk ratio estimates are in fact comparable to neighboring years. Because the ratio of the uncertainty in both p_A and p_N to their median value is smaller for higher probabilities (i.e., El Niño years), the corresponding uncertainty in the risk ratio estimates is also smaller. Finally, while the effect of ocean variability on the log risk ratio is smaller for wet events than hot or cold (Fig. 5), the effect is decidedly non-zero; this is visible in the estimates of risk ratio in Fig. 3 with respect to a number of non-overlapping interval estimates (e.g., 1983 vs. 2011).

5.2 Case study 2: Krasnoyarsk, Russia

Similar plots of the estimated probabilities, estimated risk ratio, and posteriors for σ in Krasnoyarsk are shown in Figs. 4 and 5.

For each event type, the effect of ocean variability is greatly reduced for Krasnoyarsk relative to the southern Andean community, which is not at all surprising given the high latitude continental nature of the region. Hence, individual El Niño/La Niña years do not stand out in the same way as for the southern Andean community. This is in spite

Fig. 3 Probabilities and risk ratio estimates for extremes in the southern Andean Community

Fig. 4 Probabilities and risk ratio estimates for extremes in Krasnoyarsk, Russia

Fig. 5 Posterior distributions for σ , the effect of ocean variability on log risk ratio, for each event type in the southern Andean Community (*top*) and Krasnoyarsk (*bottom*)

of the estimates of the risk ratio having visible trends over time, the implication being that this trend is almost completely described by the long-term trend in global mean temperature (i.e., the covariate). There is a large difference between scenarios for the probabilities of cold extremes for Krasnoyarsk, leading to risk estimates that are significantly less than 1. The fact that this is the case even for the 1980s may be related to the effect of Soviet-era aerosol pollution on winter cold spells (but which is outside the consideration of this paper).

6 Global results summarizing the effect of ocean variability

To summarize the effect of ocean variability more broadly for all of the WRAF regions, we provide one-dimensional “maps” (sorted by central latitude) of the posterior distributions of σ for each event type, which show the median

estimate as well as the 95% probability interval (called a “credible interval” in a Bayesian framework, as opposed to confidence interval); see Figs. 10, 11, and 12.

However, before we describe the color schemes, we first discuss how the shadings are determined. As mentioned in Sect. 4, for each region and event type, we can determine whether or not the effect of ocean variability causes different conclusions to be drawn from attribution studies in different years. In order to do this, we need to adjust the probability and risk ratio estimates such that the different years behave as if they occurred under a stationary climate, which is done as follows. From (6), with a single covariate (global mean temperature), the (unadjusted) risk ratio for a particular year t is calculated as

$$RR_t = \frac{\frac{1}{12} \sum_{j=1}^{12} \text{logit}^{-1}(\beta_{A0} + \beta_{A1}x_{At} + \alpha_t + \delta_t + \gamma_j)}{\frac{1}{12} \sum_{j=1}^{12} \text{logit}^{-1}(\beta_{N0} + \beta_{N1}x_{Nt} + \alpha_t + \gamma_j)}.$$

To transform this estimate to behave as though it arose from a stationary climate, we can substitute common covariate values x_A^* and x_N^* for the year-specific values x_{At} and x_{Nt} , while maintaining the year-specific effects α_t and δ_t to allow for year-to-year differences within the stationary climate state. That is, we can instead use the adjusted risk ratio $\widetilde{RR}_t$, where

$$\widetilde{RR}_t = \frac{\frac{1}{12} \sum_{j=1}^{12} \text{logit}^{-1}(\beta_{A0} + \beta_{A1}x_A^* + \alpha_t + \delta_t + \gamma_j)}{\frac{1}{12} \sum_{j=1}^{12} \text{logit}^{-1}(\beta_{N0} + \beta_{N1}x_N^* + \alpha_t + \gamma_j)} \quad (8)$$

for $t = 1982, \dots, 2013$ to get a sense of what the risk ratio would look like over the time period had there been no changes to the climate state. Specifically, we chose to set x_k^* equal to the average global mean temperature in scenario $k \in \{A, N\}$ from the five most recent years (2009–2013).

As an illustration of how this changes the risk ratio estimates, consider the plot in Fig. 6, which shows the raw (unadjusted) risk ratio estimates and the adjusted risk ratio estimates for hot events in the northern European Economic Area (EEA), shown as the blue region in northern Europe in Fig. 1. This region is unusual in that it has experienced strong anthropogenic influence (via land cover change and production of various types of aerosols) for a much longer time, with the balance of these effects resulting in a summer cooling effect prior to the 1990s in the CAM5.1 simulations. This late crossover exists in climate models generally, but with different timings, and means that risk ratio calculations for European events can

be sensitive to the definition of the counterfactual “natural” climate.

Using the adjusted risk ratios, we are now ready to estimate the π probabilities for each region and event type. As mentioned in Sect. 4, π is estimated using each of the MCMC joint posterior samples. Each MCMC sample yields a point-estimate time series of the estimated risk; aggregated over all posterior samples, these time series define the 95% credible intervals shown in, e.g., Fig. 6. For each sample, we then calculate the proportion of years (out of the total $T = 32$) for which the risk ratio exceeds (or does not exceed) a particular cutoff. Then, aggregating over all of the MCMC samples, we can obtain a posterior distribution on π for each region and event type. Once we have the posteriors, we are prepared to classify each region as outlined in Sect. 4.

The interval estimates of π across all regions for each event type are shown in Figs. 7, 8, and 9 and are shaded according to the classification described in Sect. 4. Then, using the same shading, estimates of the effect of ocean variability on the results of a one-year, atmosphere-model-only attribution study (i.e., σ) are shown in Figs. 10, 11, and 12. The shaded plots of σ highlight the somewhat counterintuitive relationship between the magnitude of the ocean variability’s impact and the corresponding impact on qualitative statements over time. The effect of anthropogenic emissions is now strong enough that the conclusions that $RR > 1$ for hot events, and $RR < 1$ for cold events, are rarely challenged

Fig. 6 Unadjusted (RR_t ; left) and adjusted ($\widetilde{RR}_t$; right) risk ratio for hot events in the northern European Economic Area (EEA), shown as the blue region in northern Europe in Fig. 1. The adjusted risk is calculated using a fixed covariate value (global mean temperature) for all years

Fig. 7 95% interval estimates of π for an exceedance threshold of 1 for hot events, with the horizontal black line denoting the 0.95 cutoff for π . The intervals represent an estimate of the proportion of years for which the risk ratio is greater than 1, assuming a stationary climate (i.e., using $\bar{RR}_t$). Black intervals represent regions for which oceanic

variability potentially impacts the consistency of attribution statements over time. The varying shades of red, on the other hand, are used to denote regions for which oceanic variability does *not* affect consistency of attribution statements, for several different thresholds (i.e., $RR > 1$, $RR > 2$, and $RR > 10$)

Fig. 8 95% interval estimates of π for *not* exceeding a threshold of 1 for cold events, with the horizontal black line denoting the 0.95 cutoff for π . The intervals represent an estimate of the proportion of years for which the risk ratio is less than 1, assuming a stationary climate (i.e., using $\bar{RR}_t$). Black intervals again represent regions for which

oceanic variability potentially impacts the consistency of attribution statements over time. Alternatively, the varying shades of blue are used to denote regions for which oceanic variability does *not* affect consistency of attribution statements, for several different thresholds (i.e., $RR < 1$, $RR < 1/2$, and $RR < 1/10$)

by considering the effect of ocean variability; even the $RR > 2$ or $RR < \frac{1}{2}$ conclusions still hold for most regions. What is perhaps surprising is that the regions most affected lie in the mid- and high-latitudes, where the All-Hist probabilities (p_A) are largely unaffected by SSTs (as reflected in the low skill of seasonal weather forecasts over these regions) and the σ values are thus quite low. In contrast, the regions with the most robust conclusions of an anthropogenic influence, with $RR > 10$ or $RR < \frac{1}{10}$ being robust in many cases, lie in the tropics, where the effect of SSTs on p_A is high (as reflected in the relatively high skill of seasonal forecasts) and the σ values are much higher. In other words, conclusions based on classification into the qualitative influence of anthropogenic emissions are strongly anticorrelated to conclusions based on numerical sensitivity (as measured by σ) in terms of their sensitivity to the effect of interannual ocean variability.

radiative boundary conditions, whereas temperatures in extratropical areas are also heavily influenced by dynamical processes. This means that interannual variability is lower over tropical land, and thus that for the same amount of mean warming the risk ratio is much higher there than over the extratropics, all other things being equal. For instance, the hot and cold risk ratios for the southern Andean Community are well beyond 100 and $\frac{1}{100}$ respectively, in recent years, whereas they are within 10 and $\frac{1}{10}$ for Krasnoyarsk and the northern European Economic Area. Hence, tropical regions have much more leeway to accommodate the moderately larger (as in a factor of three rather than several orders of magnitude) uncertainty described by σ , whereas many extratropical regions do not. This result may not hold if the hot categories were defined by, for instance, whether $RR > 10^4$, but we argue that such a definition would be pedantic in the light of any practical interpretation of the role of anthropogenic emissions.

The comparison of the quantitative and qualitative conclusions for wet months differs somewhat from the comparisons for hot and cold months. In this case there is no

Fig. 10 95% interval estimates of σ , the effect of ocean variability, for hot events across WRAF regions. The colors correspond to those in Fig. 7

discernible geographic pattern in the influence of SST variability, at least in part resulting from the fact that the estimated risk ratios are close to—and consistent with—unity in the first place. However, there is also a similarity in that σ values are higher in the tropics; while there is no anti-correlation in this case, the σ values themselves are poor indicators of the categories. Given the lack of relationship between latitude and the color categories in Figs. 9 and 12, we provide a map of the regions shaded by their color group (see Fig. 15).

7 Discussion

This study was conducted to understand, both quantitatively and qualitatively, how many recent and current studies assessing the role of anthropogenic emissions on individual extreme atmospheric events may be biased because their adoption of an experimental design that conditions on the observed ocean state. While quantitative estimates for events in the tropics are found, as expected (Otto et al. 2015), to be more sensitive to the conditioning on the ocean

state, in fact for hot and cold events the qualitative effect on conclusions is larger in the extratropics. While tests were performed only for events over a certain set of regions, all approximately 2 million km² in size, and covering a calendar month, results should be broadly transferrable to events covering overlapping regions of varying sizes and lasting for different durations in most cases (Angélil et al. 2014b), with the possible exception of small-scale, short-duration wet events (Angélil et al. 2014a).

In Sect. 1 we noted that the atmospheric modeling approach to diagnosing the role of anthropogenic emissions in extreme weather, as introduced by Pall et al. (2011), rests on three major assumptions in addition to uncertainties common to other approaches: that results are invariant to ocean state; that there is no change in ocean variability; and that coupled atmosphere–ocean interactions are unimportant. At first glance, these three assumptions may be expected to be more valid for the extratropics than the tropics. Concerning the third assumption, winds can more easily alter the SSTs in the tropics due to the shallower thermocline there, meaning that the third assumption should be more questionable in the tropics. It is not clear

In order to assess the effect of using atmosphere-only climate models on attribution statements for extremes, using a Bayesian hierarchical statistical approach we developed a corresponding decision model for estimating the probability of exceedance of a threshold given a large set of related ensembles of simulations, as well as for estimating the

Fig. 12 95% interval estimates of σ , the effect of ocean variability, for wet events across WRAF regions. The colors correspond to those in Fig. 9

Table 2 Comparison between the importance of the three major assumptions underlying the atmospheric modelling approach to diagnosing the role of anthropogenic emissions in extreme weather, as distinguished between the tropics and extratropics

Event type	Tropics	Extratropics
<i>1. Results are invariant to ocean state</i>		
Hot and cold months	Technically invalid, but not relevant in practice	Validity unclear, but can be quite important in practice
Wet months	Technically invalid, could be important factor	Validity unclear, may be important in practice
<i>2. No change in ocean variability</i>		
Hot and cold months	Validity questionable due to some control from ocean variability, possibly dominant factor	Weak control from ocean variability, so likely valid
Wet months	Not valid due to strong control from ocean variability, possibly dominant factor	Weak control from ocean variability, so likely valid
<i>3. Coupled atmosphere–ocean interactions are unimportant</i>		
Hot, cold, and wet months	Stronger interactions, importance not clear, perhaps invalid for some events (e.g. tropical cyclones)	Weak interactions, so likely valid

risk ratio estimated from pairs of such probabilities. This model should have broader application beyond the analysis performed in this paper. For instance, if climate model simulations are available for multiple years, as is the case in

the C20C+ D&A Project, then more accurate estimates of the posterior probabilities and the risk ratio are possible by incorporating data from multiple years. The use of a covariate not only provides the possibility for formal detection

of long-term trends in extremes, but also for detecting the influence of other factors such as volcanic eruptions and El Niño events. In addition to using data from multiple years and including covariates, it may also be possible to borrow information from across multiple estimates of the ocean warming attributable to anthropogenic emissions and from across various atmospheric models; however, in practice the selection of the prior distribution would be less obvious in these cases than it is across multiple years.

Finally, a more direct application of the results presented in this paper could be in the assessment of the robustness of calculated attribution conclusions to the effect of ocean variability. In other words, the estimates provided in this paper of the effect of oceanic variability could potentially be combined with estimates of the risk ratio made by future single-year, atmosphere-only attribution studies to make a statement about the “true” risk ratio, that is, the risk ratio informed by all possible years. The spatio-temporal characteristics of observed events are unlikely to exactly coincide with the monthly-region definitions used in this paper. However, event attribution results appear to be transferrable in a predictable way for both the spatial and temporal definitions for hot and cold events and to some degree in the spatial definition for wet events (Angélil et al. 2014b), although perhaps not at the local scale for wet events (Angélil et al. 2014a).

As a concrete example, suppose a hypothetical research team conducts a one-year, atmosphere-only attribution analysis that involves estimating the risk ratio for a certain event type in a certain region (for example in, Pall et al. (2011)). Furthermore, suppose that their new region can be associated with one of the regions used in our analysis (see Sect. 5), and that the study concerns one of the atmospheric variables used in Sect. 5. The research team might want to know: given the sampling variability in the new study, how much concern needs to be taken over the fact that the variability introduced by the ocean was ignored? In other words, what statements can be made about the overall risk ratio across all years based on the results from a single year?

In order to answer this question, consider the following framework. Suppose the new study plans to use output from an ensemble of climate models, say $X_{t1}, \dots, X_{tm}$, in order to estimate the risk ratio. Define the new study's estimate of the (log) risk ratio to be $\hat{\xi}_t \equiv \hat{\xi}(X_{t1}, \dots, X_{tm})$; the “ t ” subscript denotes that the estimator is for the risk ratio in a specific year or time period t . Whatever the functional form of the estimator, we might assume that its sampling distribution is something like $\hat{\xi}_t \sim N(\xi_t, v^2/n)$; that is, the sampling distribution of the estimator $\hat{\xi}_t$ is Gaussian and centered on ξ_t , the true log risk ratio for year t , with the uncertainty introduced by finite sampling represented by the variance v^2/n . For large enough n , the central limit

theorem ensures that the Gaussian assumption provides a good approximation.

As the notation implies, a confidence interval for ξ_t using only $\hat{\xi}_t$ and v^2/n is in fact only an interval estimate for the (log) risk ratio in a single year. In fact, the true log risk ratio for year t (ξ_t) itself comes from a population distribution which involves variability, in this case arising from the unaccounted-for effect of the ocean's internal variability. That is, we might have something like $\xi_t \sim N(\mu, \sigma^2)$, where now μ represents a “population” mean (the population being all possible years) and σ^2 represents the magnitude of oceanic internal variability (for events of a certain type in the specified region). Note that the σ^2 here is the same quantity as the one introduced in Sect. 3.

Recall that we wish to make a statement about the overall risk ratio across all years based on results from a single year: i.e., use $\hat{\xi}_t$ to estimate the population distribution $N(\mu, \sigma^2)$. Statistical theory allows us to derive the sampling distribution of $\hat{\xi}_t$ conditional on μ and the variance components, which is

$$\hat{\xi}_t | \mu \sim N(\mu, v^2/n + \sigma^2). \quad (9)$$

(for details, see Appendix 1). In what follows, we assume that the sampling variance v^2 and additional variance component σ^2 are known. In fact, these components are of course unknown and must be estimated: v^2 based on the form of the estimator, and σ^2 from the statistical model we propose in this paper. This assumption ignores any possible uncertainty in these estimates and may be inappropriate: we emphasize that this is a very rough approximation. Other assumptions might further compromise the following approach, for example assuming that the population distribution for ξ_t is also Gaussian.

In any case, using this sampling distribution, we are equipped to estimate features of the population distribution, based on $\hat{\xi}_t$. Specifically, we might be interested in estimating a particular percentile of (9), namely

$$\phi_p = \mu + c_p \sqrt{v^2/n + \sigma^2},$$

where the critical value c_p is such that $\mathbb{P}(Y \leq c_p) = p$, where Y is a standard Gaussian random variable (with mean 0 and variance 1). Note that $\phi_{0.5}$ corresponds to a confidence interval for μ . A $100(1 - \alpha)\%$ confidence interval for ϕ_p is then

$$\left(\hat{\xi}_t + (c_p - z_{\alpha/2}) \sqrt{v^2/n + \sigma^2}, \hat{\xi}_t + (c_p + z_{\alpha/2}) \sqrt{v^2/n + \sigma^2} \right), \quad (10)$$

where $z_{\alpha/2}$ the critical value for a $100(1 - \alpha)\%$ confidence interval. (Details for the derivation of (10) are also provided in Appendix 1. For example, a 95% confidence interval for

the lower 5th percentile of the population distribution, i.e., $\phi_{0.05} = \mu - 1.645\sqrt{v^2/n + \sigma^2}$, is

$$\left(\hat{\xi}_t + (-1.645 - 1.96)\sqrt{v^2/n + \sigma^2}, \hat{\xi}_t + (-1.645 + 1.96)\sqrt{v^2/n + \sigma^2}\right)$$

or

$$\left(\hat{\xi}_t - 3.605\sqrt{v^2/n + \sigma^2}, \hat{\xi}_t + 0.315\sqrt{v^2/n + \sigma^2}\right). \quad (11)$$

The main idea is that this formula could be used by the hypothetical research team to make a statement about the population distribution of risk values over all possible years, based on a single year. The research team would provide values of $\hat{\xi}_t$ and v^2/n , but could then refer to this paper for an estimate of σ^2 that corresponds to the correct geographic region and event type. If the research team finds that the (log) risk ratio for a particular year is significantly larger than $\log 1 = 0$, then they might wish to estimate whether a lower percentile of the population distribution, e.g. $\phi_{0.05}$, is also larger than $\log 1 = 0$. If so (i.e., if the interval in (11) is entirely above zero), then the team would have evidence that their results are robust to the fact that they only considered a single year in their study.

Acknowledgements This work was supported by the Regional and Global Climate Modeling Program of the Office of Biological and Environmental Research in the Department of Energy Office of Science under contract number DE-AC02-05CH11231. This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favor-

Appendix 1: Details for the derivation of the confidence interval of ϕ_p

Recall the setting introduced in Sect. 7, the goal being to make a confidence statement regarding the population distribution of risk from all possible years, using only a risk estimate from a single year.

The sampling distribution of $\hat{\xi}_t$ conditional on μ is derived as follows. Again, recall that the sampling distribution of $\hat{\xi}_t$ conditional on ξ_t is $N(\xi_t, v^2/n)$: also $\xi_t \sim N(\mu, \sigma^2)$. The sampling distribution of interest is calculated by averaging over ξ_t :

$$p(\hat{\xi}_t|\mu) = \int_{\xi_t} p(\hat{\xi}_t|\xi_t)p(\xi_t|\mu)d\xi_t,$$

where the implicit conditioning on v^2 and σ^2 is suppressed in the notation. Given that $p(\hat{\xi}_t|\xi_t) = N(\xi_t, v^2/n)$ and $p(\xi_t|\mu) = N(\mu, \sigma^2)$, the closed-form solution is well-known (this setup is identical to the derivation for the marginal distribution of the data in a Normal-Normal Bayesian posterior calculation). The result is that

$$p(\hat{\xi}_t|\mu) = N(\mu, v^2/n + \sigma^2). \quad (12)$$

Next, to derive a confidence interval for ϕ_p , first note that using (12) we can obtain a $100(1 - \alpha)\%$ confidence interval for μ as

$$\left(\hat{\xi}_t - z_{\alpha/2}\sqrt{v^2/n + \sigma^2}, \hat{\xi}_t + z_{\alpha/2}\sqrt{v^2/n + \sigma^2}\right).$$

Because $\phi_p = f(\mu) = \mu + c_p\sqrt{v^2/n + \sigma^2}$ is a linear function of μ , statistical theory says that $f(\hat{\xi}_t) \sim N(f(\mu), \text{Var}[f(\hat{\xi}_t)])$, so that a $100(1 - \alpha)\%$ confidence interval for $\phi_p = f(\mu)$ is

$$\left(f(\hat{\xi}_t) - z_{\alpha/2}\sqrt{\text{Var}[f(\hat{\xi}_t)]}, f(\hat{\xi}_t) + z_{\alpha/2}\sqrt{\text{Var}[f(\hat{\xi}_t)]}\right).$$

Since $\text{Var}[f(\hat{\xi}_t)] = \text{Var}\hat{\xi}_t = v^2/n + \sigma^2$, the confidence interval for ϕ_p is

$$\left(\left[\hat{\xi}_t - c_p\sqrt{v^2/n + \sigma^2}\right] - z_{\alpha/2}\sqrt{v^2/n + \sigma^2}, \left[\hat{\xi}_t - c_p\sqrt{v^2/n + \sigma^2}\right] + z_{\alpha/2}\sqrt{v^2/n + \sigma^2}\right)$$

ing by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

or

$$\left(\hat{\xi}_t + (c_p - z_{\alpha/2})\sqrt{v^2/n + \sigma^2}, \hat{\xi}_t + (c_p + z_{\alpha/2})\sqrt{v^2/n + \sigma^2}\right),$$

which is what is given in (10).

Appendix 2: Supplemental Figures

See Figs. 13, 14 and 15.

Fig. 13 Empirical event probabilities averaged across each year for the southern Andean Community

Fig. 14 Empirical event probabilities averaged across each year for Krasnoyarsk, Russia

Fig. 15 Map representation of the color codings for wet events in Figs. 9 and 12. While there is no relationship between latitude and the colors, there is noticeable clustering

References

- Allen M (2003) Liability for climate change. *Nature* 421:891–892
- Angélil O, Stone DA, Pall P (2014a) Attributing the probability of South African weather extremes to anthropogenic greenhouse gas emissions: spatial characteristics. *Geophys Res Lett* 41:3238–3243. doi:10.1002/2014GL059760
- Angélil O, Stone DA, Tadross M, Tummon F, Wehner M, Knutti R (2014b) Attribution of extreme weather to anthropogenic greenhouse gas emissions: sensitivity to spatial and temporal scales. *Geophys Res Lett* 41:2150–2155. doi:10.1002/2014GL059234
- Angélil O, Perkins-Kirkpatrick S, Alexander LV, Stone D, Donat MG, Wehner M, Shiogama H, Ciavarella A, Christidis N (2016a) Comparing regional precipitation and temperature extremes in climate model and reanalysis products. *Weather Clim Extremes* 13:35–43. doi:10.1016/j.wace.2016.07.001. <http://www.sciencedirect.com/science/article/pii/S2212094715300220>
- Angélil O, Stone D, Wehner M, Paciorek CJ, Krishnan H, Collins W (2016b) An independent assessment of anthropogenic attribution statements for recent extreme temperature and rainfall events. *J Clim*. doi:10.1175/JCLI-D-16-0077.1
- DerSimonian R, Laird N (1986) Meta-analysis in clinical trials. *Controlled Clin Trials* 7(3):177–188. doi:10.1016/0197-2456(86)90046-2. <http://www.sciencedirect.com/science/article/pii/0197245686900462>
- Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB (2013) *Bayesian Data Analysis*, 3rd edn. Chapman and Hall/CRC
- Hansen G, Auffhammer M, Solow AR (2014) On the attribution of a single event to climate change. *J Clim* 27:8297–8301
- Hurrell JW, Hack JJ, Shea D, Caron JM, Rosinski J (2008) A new sea surface temperature and sea ice boundary dataset for the Community Atmosphere Model. *J Clim* 21:5145–5153
- Jeon S, Paciorek CJ, Wehner MF (2016) Quantile-based bias correction and uncertainty quantification of extreme event attribution statements. *Weather Clim Extremes*. doi:10.1016/j.wace.2016.02.001. <http://www.sciencedirect.com/science/article/pii/S2212094715300220>
- McCulloch CE, Searle SR (2004) *Generalized linear mixed models*. In: *Wiley series in probability and statistics: applied probability and statistics*. Wiley. <https://books.google.com/books?id=XpL38GzFqL0C>
- National Academies of Sciences, Engineering, and Medicine (2016) *Attribution of extreme weather events in the context of climate change*. The National Academies Press, Washington, DC. doi:10.17226/21852
- Neale RB, Chen CC, Gettelman A, Lauritzen PH, Park S, Williamson DL, Conley AJ, Garcia R, Kinnison JF, Lamarque, Marsh D, Mills M, Smith AK, Tilmes F, Vitt, Morrison H, Cameron-Smith P, Collins WD, Iacono MJ, Easter RC, Ghan SJ, Liu X, Rasch PJ, Taylor MA (2012) Description of the NCAR community atmosphere model (CAM 5.0). Tech. rep., NCAR Technical Note NCAR/TN-486+STR
- Otto FEL, Boyd E, Jones RG, Cornforth RJ, James R, Parker HR, Allen MR (2015) Attribution of extreme weather events in Africa: a preliminary exploration of the science and policy implications. *Clim Change* 132:531–543. doi:10.1007/s10584-015-1432-0
- Pall P, Aina T, Stone DA, Stott PA, Nozawa T, Hilberts AGJ, Lohmann D, Allen MR (2011) Anthropogenic greenhouse gas contribution to flood risk in England and Wales in Autumn 2000. *Nature* 470:382–385

- Stone DA, Allen MR (2005) The end-to-end attribution problem: from emissions to impacts. *Clim Change* 71:303–318
- Stone DA, Pall P (2016) A benchmark estimate of the effect of anthropogenic emissions on the ocean surface (**In preparation**)
- Trenberth K, Jones P, Ambenje P, Bojariu R, Easterling D, Tank AK, Parker D, Rahimzadeh F, Renwick J, Rusticucci M, Soden B, Zhai P (2007) Observations: Surface and Atmospheric Climate Change. In: *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge
- Wolski P, Stone D, Tadross M, Wehner M, Hewitson B (2014) Attribution of floods in the Okavango Basin, Southern Africa. *J Hydrol* 511:350–358