Observation-based Simulations of Humidity and Temperature Using Quantile Regression

Andrew Poppick[*1] and Karen A. McKinnon[2]

[1]Department of Mathematics and Statistics, Carleton College
[2]Department of Statistics and Institute of the Environment and Sustainability, University of California, Los Angeles

This is a non peer-reviewed preprint.

Abstract

The human impacts of changes in heat events depend on changes in the joint behavior of temperature and humidity. Little is currently known about these complex joint changes, either in observations or projections from general circulation models (GCMs). Further, GCMs do not fully reproduce the observed joint distribution, implying a need for simulation methods that combine information from GCMs with observations for use in impact studies. We present an observation-based, conditional quantile mapping approach for the simulation of future temperature and humidity. A temperature simulation is first produced by transforming historical temperature observations to include projected changes in the mean and temporal covariance structure from a GCM. Next, a humidity simulation is produced by transforming humidity observations to account for projected changes in the conditional humidity distribution given temperature, using a quantile regression model. We use the Community Earth System Model Large Ensemble (CESM1-LE) to estimate future changes in summertime (June - August) temperature and humidity over the Continental United States (CONUS), and then use the proposed method to create future simulations of temperature and humidity at stations in the Global Summary of the Day dataset. We find that CESM1-LE projects decreases in summertime humidity across CONUS for a given deviation in temperature from the forced trend, but increases in the risk of high dew point on historically hot days. In comparison with CESM1-LE, our observation-based simulation largely projects smaller changes in the future risk of either high or low humidity on days with historically warm temperatures.

1   Introduction

Assessing the potential societal impacts of changes in future heat events requires an understanding of projected changes in both temperature and humidity. For example, high humidity is a contributor to human heat stress during heat events (Barreca, 2012) and some projections show substantial increases by the end of century in the risk of humid heat events that exceed theoretical limits on the human body's ability to self-regulate temperature through evaporative cooling (Pal and Eltahir, 2016; Coffel et al., 2017). In contrast, hot and dry conditions increase wildfire risk (Seager et al., 2015), among other impacts. Changes in hot and humid or hot and dry events can be affected by distributional changes beyond the means of each variable, such as changes in the underlying local relationship between temperature and humidity as well as changes in variability at multiple scales, implying a need for methods that are sensitive to these potentially complex changes.

There is limited existing work that addresses changes in local, joint changes in temperature and humidity either in observations or in general circulation models (GCMs). Many recent studies focus on

---

*Corresponding author: apoppick@carleton.edu

univariate summaries that may be useful for a particular impact of interest, such as wet bulb or wet bulb globe temperatures (as an indicator for human comfort) (Knutson and Ploshay, 2016; Pal and Eltahir, 2016; Coffel et al., 2017; Li et al., 2017; Lee and Min, 2018; Li et al., 2020) or vapor pressure deficit (as an indicator for crop health or wildfire risk) (Seager et al., 2015; Hsiao et al., 2019). While univariate summaries can be useful for studying particular impacts, a more general approach is desirable and requires multivariate methods. Studies that do consider changes in both temperature and humidity have tended to focus on limited quantities, such as changes in temperature and humidity on the 1% warmest days (Fischer and Knutti, 2013), for monthly averages (Simmons et al., 2010), in univariate quantities during specific definitions of hot and dry or hot and humid events (Schoof et al., 2017), or at a small number of spatial locations (Pryor and Schoof, 2016; Yuan et al., 2019). Fischer and Knutti (2013) recognized a need for more work in understanding joint changes in temperature and humidity, but a detailed understanding remains lacking in the literature.

Moreover, future impact studies may require not only an understanding of projected changes in temperature and humidity in GCMs, but instead realistic bivariate simulations of these variables. It is well understood that raw GCM output is insufficient for these purposes, because GCM output forced with historical forcings does not fully reproduce observed climate variable distributions; see John and Soden (2007), Brands et al. (2013), Tian et al. (2013), and Zhao et al. (2015) for examples specifically evaluating GCM simulations of temperature and humidity or heat stress (see also IPCC (2013), Chapter 9). This fact is not specific to temperature or humidity simulations, and a number of methods have been proposed to combine observations with model output to produce better calibrated future simulations, typically labeled "bias correction" methods (see, e.g., Ho et al. (2012); Hawkins et al. (2013); Cannon et al. (2020) for reviews of the main types of methods). Ho et al. (2012) separate popular methods into two classes: those that modify GCM output in an attempt to correct biases with observations, and those that modify observations in an attempt to account for GCM projected changes. Because the term "bias correction" connotes the former but not the latter type methods, we instead refer to these methods as "model-based" or "observation-based" simulation methods. While no simulation method resolves all potential defects of future simulations (Dixon et al., 2016; Lanzante et al., 2018), observation-based simulations have the attractive property that they easily preserve most of the higher-order behavior of observational distributions and generally require more statistical modeling of distributions in GCM output (where available data is typically abundant and signal to noise ratios are relatively high) than in observations (where data is typically limited and signal to noise ratios are lower). Recent work has extended observation-based simulation methods to incorporate more complex changes from GCM output, such as distinct variability changes at different timescales (Leeds et al., 2015; Poppick et al., 2016) or seasonally-varying changes to the full marginal distribution of a climate variable (Haugen et al., 2019).

Most of the aforementioned simulation methods are univariate, but there has been a recent increase in proposed methods for multivariate simulations. Indeed, Zscheischler et al. (2019) emphasizes the importance of multivariate simulation methods, using temperature and humidity as one example where there is such a need, because separately producing simulations of each variable fails to address inter-variable dependencies that can have relevant effects on impacts studies. To our knowledge, existing bivariate or multivariate methods are model-based (e.g., Piani et al. (2010); Vrac and Friederichs (2015); Mehrotra and Sharma (2015, 2016); Cannon (2018); Vrac (2018); Guo et al. (2019)) and so may struggle to realistically simulate higher-order distributional features that are not explicitly corrected in the simulation procedure. We therefore see a benefit in developing observation-based multivariate simulation methods.

This work makes two main contributions to the literature on joint changes in temperature and humidity. One is an analysis of projected changes in the relationship between summertime daily temperature and humidity over the continental United States (CONUS) in the Community Earth System Large Ensemble (CESM1-LE) (Kay et al., 2015). Our methodology is based on a quantile regression (Koenker and Bassett Jr, 1978) model for humidity given local and global temperature, allowing for a flexible and unifying approach to studying the relationship between humidity and temperature in different parts of the humidity distribution (i.e., dry vs. humid days), over the course of the summer, and changes thereof with increasing

global mean temperature. Quantile regression has been used previously in the climate literature, e.g., to study distributional changes in observed temperatures (Reich, 2012; Matiu et al., 2016; McKinnon et al., 2016; Rhines et al., 2017; Gao and Franzke, 2017), in temperatures from GCM output (Haugen et al., 2018), and recently, changes in the relationship between temperature and humidity in observations (McKinnon and Poppick, 2020) and for a limited number of CONUS gridcells within CESM1-LE but without an explicit model for changes (Yuan et al., 2019).

The second and primary contribution of this work is a proposed observation-based bivariate simulation method for daily summertime humidity and temperature. A temperature simulation is first produced by transforming historical temperature observations to account for GCM projected changes in mean and temporal covariance, using a method based on those proposed in Leeds et al. (2015) and Poppick et al. (2016). A simulation of humidity is then produced by transforming humidity observations to account for changes both in temperature itself as well as in the underlying relationship between humidity and temperature. Our methodology relies on the aforementioned quantile regression analysis and can be thought of as an observation-based "quantile-mapping" approach to simulation (as in Haugen et al. (2019)). To our knowledge, the proposed method is the first such observation-based method that can account for bivariate changes and is built specifically for simulating temperature and humidity.

The remainder of this paper is organized as follows. In Section 2, we define the variables in our analysis and describe both the GCM data used to study changes in the relationship between humidity and temperature and the observations we use for producing our observation-based simulation. In Section 3, we describe our method for producing an observation-based temperature simulation, which is an input into our humidity simulation, and illustrate how the simulation procedure works at one location. In Section 4, the primary contribution of this work, we then describe the proposed humidity simulation and illustrate the bivariate simulation at the same location. We provide an analysis of changes over CONUS, as well as summaries of our resulting observation-based simulations, in Section 5. In Section 6, we provide a discussion and concluding remarks. Additional technical details may be found in the Appendix and statistical model validation may be found in the Online Supplement.

2 Data

The simulation procedure that we propose in Sections 3 and 4 requires temperature and humidity data from both GCM output and observations. Unlike for temperature, many different and related variables are used as measures of humidity. In this paper, we use the *dew point*, which is the temperature to which an air parcel would need to be cooled in order for the moisture in it to condense. We use dew point for several reasons, including that it is directly measured by the weather stations in the observational data we use (see below), is a reasonable indicator of human comfort (Davis et al., 2016), and can be understood straightforwardly as a measure of humidity because it does not change with temperature if the moisture content of the air is fixed (unlike relative humidity). At some points in this paper, we also refer to the *dew point depression*, which is defined as the difference between (dry bulb) temperature and dew point. While our simulations and results are presented in terms of dew point, the dry bulb temperature and dew point determine the relative humidity value, so a relative humidity simulation is implicit in the proposed procedure.

2.1 Climate model data

Climate model output for temperature and humidity are from CESM1-LE, a 40-member initial condition ensemble run with the fully coupled 1-degree latitude-longitude version of CESM1. Historical forcings (Lamarque et al., 2010) are used for the years 1920-2005 and the RCP8.5 scenario (Meinshausen et al., 2011) is used for the years 2005-2100. In this study, we focus on boreal summer (JJA) daily average temperature and dew point over a subregion of the global grid that contains CONUS (latitudes between 25° N and 50° N, and longitudes between 66° W and 125° W, and gridcells at least 50% land).

Daily or sub-daily near-surface humidity data is only available (as reference height specific humidity, QREFHT) in CESM1-LE as six-hourly data during the time periods 1990-2005, 2026-2035, and 2071-2080. To calculate daily average dew point from the available data, we first use the QREFHT and surface pressure (PS) values to calculate six-hourly values of dew point using the equations found in Table 1 of Willett et al. (2014); we then average the four six-hourly values to compute a daily value. To be consistent with this calculation, daily temperatures are computed by averaging the four concurrent six-hourly reference height temperature (TREFHT) values. In cases where the calculated dew point exceeds the daily temperature value (0.03% of cases), it is set to the value equal to the temperature minus the minimum positive calculated dew point depression at the gridcell.

The statistical models we describe below for changes in the distribution of local temperature and dew point require an estimate of the forced trend in global mean temperature (GMT). An estimate from CESM1-LE is obtained by first averaging the annually-averaged TREFHT values across gridcells (weighting by geographic area) and across the first 35 ensemble members (for the full 1920-2100 ensemble output). We exclude the last five runs, completed at the University of Toronto, in this calculation due to concerns about the consistency of the results when considering trends in the GMT (CESM Project, 2020); however, they are included in the local temperature and dew point data described above, as we have found no evidence of a discrepancy in the distribution of these local variables over CONUS. The raw global annual-mean ensemble-mean temperature value is then further smoothed using a lowess smoother with a 5% span (approximately 9 years) and tricubic weighting to obtain our final estimate of the GMT forced trend. See the Supplement Figure S1 for the ensemble mean GMT anomalies and lowess smoothed estimated trend. While alternative methods for GMT trend estimation are available (e.g., Poppick et al. (2017)), we use lowess smoothing here for simplicity since studying the global trend is not of primary interest in this work.

2.2   Observational data

Observational temperature and dew point data are from the Global Summary of the Day (GSOD) database provided by the National Center for Environmental Information (NCEI). Data are typically collected on an hourly or three-hourly basis and are averaged to daily values. We restrict our analysis to stations within the United States and we use data from the years 1973-2018, where the record is most complete. We retain only stations that meet the relatively strict criterion that there are no missing temperature values during JJA over the time period studied; this corresponds to about 25% of CONUS stations in the database and appears to be geographically representative (see Figure 1). The requirement of no missing temperature values makes the proposed temperature simulation (relying on fast Fourier transforms) more convenient; however, this requirement could be relaxed by using interpolation methods to fill in missing values or by using an alternative temperature simulation method that does not require complete data, since the proposed dew point simulation can be implemented with any temperature simulation method deemed suitable by the user.

3   Univariate Simulation of Temperatures

The novel observation-based dew point simulation that we propose in Section 4 is conditional on a temperature simulation. Many methods exist for producing univariate simulations of temperature that combine information from GCM output with observations, as discussed in Section 1. The procedure used here is a modified version of those proposed in Leeds et al. (2015) and Poppick et al. (2016), which are extensions of the "delta method". The delta method simulates future temperatures by adding a GCM projected future mean trend to historical observations; however, temperature variability changes are also potentially important and unaccounted for by the delta method. Variability changes can be timescale-dependent (e.g. day-to-day vs. inter-annual variability changes can differ), implying full changes to the temporal

covariance structure. The proposed procedure accounts for these temporal covariance structure changes by modifying the observed temperatures' spectral density to account for GCM projected changes.

Before describing the simulation method, we introduce some notation. Throughout, we write $\tilde{x}$ for a quantity in GCM output that is analogous to the quantity x in reality. We write $\mathbf{T}_y^{(h)} = (T_{y,1}, ..., T_{y,n_{\mathrm{JJA}}})$ for the vector of the $n_{\mathrm{JJA}} = 92$ observed JJA daily temperatures in the y'th year of the observational record, and $\tilde{\mathbf{T}}_y^{(h)}$ and $\tilde{\mathbf{T}}_y^{(f)}$ for temperatures in historical and future portions of the GCM output (suppressing notation for the GCM ensemble member number for the sake of clarity).

We assume that temperatures can be separated into a mean forced response component and a residual component representing internal variability (also possibly changing in distribution in response to forcing):

$$\mathbf{T}_y^{(h)} = \boldsymbol{\mu}_y^{(h)} + \boldsymbol{\epsilon}_y^{(h)},$$

and similarly for $\tilde{\mathbf{T}}_y^{(h)}$ and $\tilde{\mathbf{T}}_y^{(f)}$, where $\boldsymbol{\mu}_y^{(h)}$ is the mean forced component in year y (not necessarily constant across days within year y) and $\boldsymbol{\epsilon}_y^{(h)}$ is the temperature deviation from that forced component. (In our terminology, the "mean forced component" can be thought of as the true average temperature over infinite realizations of internal variability, and would include both changes due to radiative forcing and the baseline average value.) Our temperature simulation procedure requires information about projected *changes* in both the mean forced component and internal temperature variability. The GCM change in mean forced component is denoted

$$\tilde{\boldsymbol{\Delta}}_y = \tilde{\boldsymbol{\mu}}_y^{(f)} - \tilde{\boldsymbol{\mu}}_y^{(h)}. \tag{1}$$

We represent changes in the temporal covariance structure of the temperature deviations, $\boldsymbol{\epsilon}_y$, in terms of changes in spectral densities. Under the reasonable approximation that JJA internal temperature variability, captured by $\boldsymbol{\epsilon}_y$, is statistically stationary within a year, it has a spectral density that is denoted $a_y^{(h)}(\omega)$ for frequency ω, interpreted as the variance attributable to fluctuations at frequency ω. We write $\mathbf{a}_y^{(h)} = \left(a_y^{(h)}(0/n_{\mathrm{JJA}}), a_y^{(h)}(1/n_{\mathrm{JJA}}), ..., a_y^{(h)}((n_{\mathrm{JJA}} - 1)/n_{\mathrm{JJA}}) \right)$ for the vector of spectral densities on the Fourier frequencies (noting that the spectral density must be periodic with $a(\omega) = a(1 - \omega)$ for $\omega \in (0, 1/2)$). Likewise, we write $\tilde{\mathbf{a}}_y^{(h)}$ and $\tilde{\mathbf{a}}_y^{(f)}$, and we write

$$\tilde{\boldsymbol{\rho}}_y = \left(\frac{\tilde{a}_y^{(f)}(0/n_{\mathrm{JJA}})}{\tilde{a}_y^{(h)}(0/n_{\mathrm{JJA}})}, ..., \frac{\tilde{a}_y^{(f)}((n_{\mathrm{JJA}} - 1)/n_{\mathrm{JJA}})}{\tilde{a}_y^{(h)}((n_{\mathrm{JJA}} - 1)/n_{\mathrm{JJA}})} \right) \tag{2}$$

for the projected change in spectral density, expressed as a ratio of variances.

In Section 3.1, we first summarize the proposed procedure that uses the above quantities to produce an observation-based future simulation that captures projected changes in mean and temporal covariance. Because the true forced mean component and distribution of internal variability is not fully known in either the observations or GCM, these quantities must be estimated using statistical methods. In Section 3.2, we therefore describe the statistical models that are used to estimate the quantities required for the proposed procedure. We then illustrate the method at an example location in Section 3.3.

3.1 Simulation method

With the notation established above, our proposed simulation of daily JJA temperatures, $\hat{\mathbf{T}}_y^{(f)}$, for the y'th year of the future period is

$$\hat{\mathbf{T}}_y^{(f)} = \boldsymbol{\mu}_y^{(h)} + \tilde{\boldsymbol{\Delta}}_y + \mathcal{F}^{-1}\mathrm{diag}(\sqrt{\tilde{\boldsymbol{\rho}}_y})\mathcal{F}(\mathbf{T}_y^{(h)} - \boldsymbol{\mu}_y^{(h)}), \tag{3}$$

where $\mathcal{F}$ is the discrete Fourier transform matrix, i.e., has entries $\mathcal{F}_{j,t} = e^{-2\pi i(j-1)(t-1)/n_{\mathrm{JJA}}}$. (While written as matrix multiplication, in practice this is calculated using a fast Fourier transform algorithm.) In words, starting with observed JJA temperatures from the y'th year of the historical record, we

1. subtract the estimated mean forced component from the observed temperatures to obtain the observed temperature deviation from the mean (herein simply the "temperature deviation"),

2. calculate the discrete Fourier transform of the observed temperature deviations,

3. multiply by frequency-dependent projected square root spectral density changes,

4. invert the Fourier transformation to obtain future temperature deviations reflecting GCM projected variability changes, and

5. add back the estimated forced response from the observations and the GCM projected change in mean to obtain the future temperature simulation.

This procedure produces a simulation that has the mean changes and (approximately) the covariance changes projected by the GCM (equations (1) and (2)), but otherwise retains many features of the observed temperatures (including, e.g., spatial coherences; see Poppick et al. (2016)). If $\tilde{\boldsymbol{\rho}}_y = 1$ (i.e., there are no projected variability changes), then the procedure reduces to the basic delta method. Other versions of the delta method extend the basic method to allow for *marginal* variability changes (i.e., same variability change at each frequency); step 3 is an extension of this idea, where instead we are separately multiplying by the change in standard deviation associated with fluctuations at each of the Fourier frequencies, producing changes to the full temporal covariance structure. As written, the simulation procedure is applied separately to temperatures in each year y, but in practice this can be done simultaneously in a vectorized fashion and the resulting simulation procedure can be implemented very quickly.

3.2   Statistical models for mean and variability changes

The quantities $\boldsymbol{\mu}_y^{(h)}$, $\tilde{\boldsymbol{\Delta}}_y$, and $\tilde{\boldsymbol{\rho}}_y$ in equation (3) are unknown and must be estimated from the available data using statistical methods.

One approach to estimation, although not ours, would be to rely on methods that attempt to avoid strong assumptions about the functional form of the above quantities; as a simple example, one might estimate $\tilde{\boldsymbol{\Delta}}_y$ as the difference in ensemble mean local temperatures between the two time points. However, this approach has a number of disadvantages. First, the ensemble mean local temperature (for example) is a very noisy estimate of the true model forced component, due to the presence of considerable internal variability in local temperatures and an ensemble of only 40 runs. This problem is even more pronounced when estimating more complex quantities (such as the changes in spectral densities, $\tilde{\boldsymbol{\rho}}_y$) or when estimating $\boldsymbol{\mu}_y^{(h)}$ from observations (where we have only a single observational record). Moreover, these kinds of methods do not readily provide an interpretable summary of how the relevant functions are changing across years within the forcing scenario of interest.

Instead, our approach is based on the fact that some aspects of the forced response of local atmospheric variables scale approximately with the GMT forced response (e.g., Santer et al. (1990); Dai et al. (2015); Li et al. (2020) and many others; see also IPCC (2013), Section 12.4.2). We use this idea to constrain the functional forms of the changes. The resulting parametric statistical models make stronger assumptions but give estimates that are both less noisy and easier to interpret. Validation of the statistical models presented below, and comparison with empirical estimates obtained only by averaging across ensemble members, may be found in supplement Sections S2.1 and S2.2.

We model the mean forced component of local temperature on the d'th day of the y'th year in the historical observational record as

$$\mu_{y,d}^{(h)} = \beta_0 + \beta_1 \tilde{G}_y^{(h)} + \gamma(d) + \eta(d)\tilde{G}_y^{(h)}, \tag{4}$$

where $\tilde{G}_y^{(h)}$ is the smoothed GMT anomaly value in year y of the historical period from the GCM, the β's are unknown parameters, and $\gamma(d)$ and $\eta(d)$ are unknown functions. We use $\tilde{G}_y^{(h)}$ as our estimate

of the GMT forced response in the historical period *for the observations* under the assumption that the GCM adequately captures the forced response of GMT. The function $\gamma(d)$ allows for seasonally varying mean temperatures and $\eta(d)$ allows for that seasonal cycle to change as the GMT changes; these functions are both parameterized using the first two seasonal harmonics. Since $\tilde{G}_y$ is smooth in time, the resulting estimates of forced mean changes in local temperature also vary smoothly in time, with additional flexibility to allow for seasonality in the mean changes. The GCM mean forced component is modeled similarly, so that the projected change in mean temperature on the d'th day of the y'th year (comparing future vs. historical time periods) is

$$\tilde{\Delta}_{y,d} \equiv \tilde{\mu}_{y,d}^{(f)} - \tilde{\mu}_{y,d}^{(h)} = \left(\tilde{\beta}_1 + \tilde{\eta}(d)\right)\left(\tilde{G}_y^{(f)} - \tilde{G}_y^{(h)}\right). \tag{5}$$

Modeling local mean changes as proportional to global mean changes results in well-fitting mean functions for the time periods, region, and climate change scenario analyzed here. More complex local mean emulators (e.g., Castruccio et al. (2014)) could be used if more a general simulation (e.g., for multiple climate change scenarios) were required.

Similarly, we model the changes in spectral densities at frequency ω as

$$\log(\tilde{\rho}_y(\omega)) = \tilde{\delta}(\omega)\left(\tilde{G}_y^{(f)} - \tilde{G}_y^{(h)}\right), \tag{6}$$

where $\tilde{\delta}(\omega)$ is a smooth function in frequency (smoothing enforced via a kernel smoother). If $\tilde{\delta}(\omega)$ is constant, then model (6) says that only the marginal standard deviation of temperatures changes over time (and scales with GMT); otherwise, the model allows different changes at different frequencies and therefore changes to the full temporal covariance structure. Positive values of $\delta(\omega)$ correspond to frequencies with variability increasing with increasing GMT, and similarly for negative values of $\delta(\omega)$. See the Sections A1.1 and A1.2 for additional modeling and estimation details for models (4) and (6), respectively.

3.3   Illustration of method

Here we show an illustration of the procedure described above for a location near Minneapolis, MN; the observational data we use is from a station located at Minneapolis-Saint Paul International Airport, and the changes from CESM1-LE are estimated using the nearest gridcell. For an analysis over all of CONUS, see Section 5. For illustration, the years 1996-2005 of the observations are used to simulate the years 2071-2080 of the RCP8.5 scenario. For reference, the average change in the GMT forced mean between these two time periods in CESM1-LE is approximately 3.3°C.

First, we show the CESM1-LE projected changes in mean and variability at this gridcell. Figure 2(a) shows changes in local mean temperature per degree warming of GMT. Specifically, we show the function $\tilde{\beta}_1 + \tilde{\eta}(d)$ vs. d from equation (5) along with approximate 95% pointwise confidence intervals; see Section A1.1 for a description of how the confidence intervals are obtained. Averaged over the summer, this gridcell shows warming of about 1.6°C per degree warming in GMT, but the later part of summer warms most quickly. Figure 2(b) shows the relative change in the square root spectral density at this gridcell per degree warming of GMT (i.e., we show the function $\exp\left\{\tilde{\delta}(\omega)/2\right\} - 1$, again along with approximate 95% pointwise confidence intervals; see Section A1.2), which can be interpreted as the relative change in standard deviation for fluctuations with frequency ω. In Figure 2(b), the value at $\omega = 0$ is labeled "JJA annual average"; this is because the discrete Fourier transform at $\omega = 0$ is proportional to the time average, so that $\sqrt{\tilde{\rho}_y(0)}$ is the change to the standard deviation of the JJA annual average in year y in equation (3). Summertime lower-frequency variability at this gridcell increases with GMT (e.g., the standard deviation of the JJA average temperature is estimated to increase by about 6% per degree increase in GMT), but the variability changes at higher frequencies are estimated to be negligible.

These changes in means and variability are then incorporated into the simulation procedure (equation (3)) to produce an observation-based simulation of future temperatures. Figures 2(c) - (f) show an illustration of observed temperatures and the resulting simulation. We show observed and simulated daily

values over the first year of the record (Figures 2(c) and (d)), as well as JJA average values across years (Figures 2(e) and (f)), to illustrate the differences in variability changes at short and longer timescales, respectively. The observed temperatures and simulated series are shown on top (Figures 2(c) and (e)), while deviations from the forced mean (i.e., $T_{y,d}^{(h)} - \mu_{y,d}^{(h)}$ and similar) are shown on bottom (Figures 2(d) and (f)) to illustrate how the simulation procedure affects temperature variability. Because high-frequency variability changes at this location are projected to be negligible, simulated daily temperature deviations look very similar to the observed values in the first year and the resulting simulation is driven by mean changes; however, lower-frequency variability is projected to increase at this location, resulting in JJA average values that are more variable in the simulation than in the original observations.

4 Observation-based Conditional Quantile Mapping of Dew Point Given Temperature

In this section, we propose an observation-based simulation method for dew point that depends on a temperature simulation such as that proposed in Section 3. The proposed dew point simulation falls into the general category of "quantile mapping" simulation methods, and works by transforming observed dew point quantile levels given observed temperatures to new dew points given simulated temperatures, accounting for GCM projected changes in the underlying relationship between dew point and temperature.

As in Section 3, we first summarize the proposed simulation procedure in Section 4.1, then describe the statistical models that are used to estimate the quantities required for the proposed procedure in Section 4.2, and finally provide an illustration of the simulation at one location in Section 4.3.

4.1 Simulation method

Conceptually, we want to simulate future dew point values, given a future temperature simulation, via a quantile mapping that transforms the observed distribution into a future distribution. Writing $\mathbf{D}^{(h)}$ for the observed dew point values and $\hat{\mathbf{D}}^{(f)}$ for the observation-based future simulation (and $\mathbf{T}^{(h)}, \hat{\mathbf{T}}^{(f)}, \tilde{\mathbf{G}}^{(h)}$, and $\tilde{\mathbf{G}}^{(f)}$ for the observed local temperature and future simulation, and historical and future smoothed GMT forced trend estimate, as above), our *ideal* simulation procedure would entail the following transformation:

$$
\begin{aligned}
\hat{\mathbf{D}}^{(f)} &= F_{\mathbf{D}^{(f)}}^{-1}\left(\boldsymbol{\tau}^* | \hat{\mathbf{T}}^{(f)}, \tilde{\mathbf{G}}^{(f)}\right), \text{ where} \\
\boldsymbol{\tau}^* &= F_{\mathbf{D}^{(h)}}\left(\mathbf{D}^{(h)} | \mathbf{T}^{(h)}, \tilde{\mathbf{G}}^{(h)}\right)
\end{aligned}
\tag{7}
$$

and where $F_X(\cdot|\cdot)$ denotes the (conditional) cumulative distribution function of the variable X and therefore $F_X^{-1}(\cdot|\cdot)$ is the conditional quantile function. In words, the ideal procedure would be to:

1. determine the quantile levels associated with observed dew point values, given observed local temperatures and the GMT forced trend, and then

2. simulate future dew points using those quantile levels, but conditionally on future simulated local and global temperatures and under the projected future relationship between dew point and temperature.

However, in practice we do not know $F_{\mathbf{D}^{(f)}}^{-1}(\cdot|\cdot)$, the conditional quantile function of true future dew points, so our proposed method will replace this function with one that incorporates changes from the GCM, analogous to other observation-based simulation procedures.

4.2 Statistical model for conditional quantile functions

We model the relationship between dew point and temperature (and changes thereof) through a quantile regression model of the log dew point depression, with effects for GMT, seasonality, and the local

8

temperature deviation from its estimated forced mean. Writing $\tilde{Y}_{y,d} = \log(\tilde{T}_{d,y} - \tilde{D}_{d,y})$ for the log dew point depression in the GCM (again suppressing notation for the run, and here not distinguishing between historical vs. future time periods), our quantile regression model for the GCM is

$$F_{\tilde{Y}_{d,y}}^{-1}\left(\tau | \tilde{T}_{d,y}, \tilde{G}_y\right) = \tilde{\alpha}_{0,\tau} + \tilde{g}_\tau(d) + \tilde{\alpha}_{1,\tau}\tilde{G}_y + \tilde{h}_\tau(\tilde{T}_{d,y} - \tilde{\mu}_{d,y}), \tag{8}$$

where the function $\tilde{g}_\tau(\cdot)$ characterizes the seasonal cycle in the τ'th quantile for a fixed GMT and local temperature deviation, $\tilde{\alpha}_{1,\tau}$ characterizes the change in dew point depression with GMT for a fixed local temperature deviation and day of the year, and the function $\tilde{h}_\tau(\cdot)$ characterizes the relationship between the dew point depression and local temperature deviation for a fixed day of year and GMT. We use the two leading seasonal harmonics to parameterize $\tilde{g}_\tau(\cdot)$ and use a natural spline to parameterize $\tilde{h}_\tau(\cdot)$; see Section A1.3 for details. Working on the scale of log dew point depression allows us to easily enforce the constraint that the dew point depression must be positive, and therefore that dew point must be less than the temperature value. However, model (8) implies an equivalent model on the scale of the dew point itself, $\tilde{D}_{y,d} = \tilde{T}_{d,y} - e^{\tilde{Y}_{d,y}}$, i.e.,

$$F_{\tilde{D}_{d,y}}^{-1}\left(\tau | \tilde{T}_{d,y}, \tilde{G}_y\right) = \tilde{T}_{d,y} - \exp\left\{\tilde{\alpha}_{0,1-\tau} + \tilde{g}_{1-\tau}(d) + \tilde{\alpha}_{1,1-\tau}\tilde{G}_y + \tilde{h}_{1-\tau}(\tilde{T}_{d,y} - \tilde{\mu}_{d,y})\right\}. \tag{9}$$

So while in practice, the model is estimated on the scale of equation (8) (where the model is linear), changes in the log dew point depression can be translated into changes in dew point (but the τ'th quantile of dew point depression corresponds to the $1-\tau$'th quantile of dew point). See supplementary materials Section S2.3 for information about the quality of fit of this quantile regression model in CESM1-LE.

In model (8), changes in the conditional quantile function associated with changes in GMT are captured by the $\tilde{\alpha}_{1,\tau}$ parameters at each quantile level τ. In the spirit of univariate observation-based simulation procedures, we assume that the GCM captures changes with GMT through this parameter, but that the terms describing the seasonal cycle and dependence on local temperature in observations may be different from those in the GCM. As such, we take the future quantile function for our simulation to be

$$F_{\hat{D}_{d,y}^{(f)}}^{-1}\left(\tau | \hat{T}_{d,y}^{(f)}, \tilde{G}_y^{(f)}\right) = \hat{T}_{d,y}^{(f)} - \exp\left\{\alpha_{0,1-\tau} + g_{1-\tau}(d) + \tilde{\alpha}_{1,1-\tau}\tilde{G}_y^{(f)} + h_{1-\tau}(\hat{T}_{d,y}^{(f)} - \hat{\mu}_{d,y}^{(f)})\right\}.$$

This assumption would imply that the ideal simulation procedure (7) may be rewritten as

$$\begin{aligned}
\hat{D}_{y,d}^{(f)} &= \hat{T}_{d,y} - \exp\left\{\alpha_{0,1-\tau_{y,d}^*} + g_{1-\tau_{y,d}^*}(d) + \tilde{\alpha}_{1,1-\tau_{y,d}^*}\tilde{G}_y^{(f)} + h_{1-\tau_{y,d}^*}(\hat{T}_{d,y}^{(f)} - \hat{\mu}_{d,y}^{(f)})\right\}, \text{ where} \\
\tau_{y,d}^* &= F_{D_{y,d}^{(h)}}\left(D_{y,d}^{(h)} | T_{y,d}^{(h)}, \tilde{G}_y^{(h)}\right),
\end{aligned} \tag{10}$$

which is therefore our proposed procedure. Note that we estimate the quantile function for the observations over the historical period using the $\tilde{\alpha}_{1,\tau}$'s estimated from the GCM; this is consistent with the overall simulation approach and the idea that the GCM will have much more information about this term than is contained in the observational output, because of both the larger amount of data in the GCM as well as the larger changes in GMT.

In order to better understand the procedure (10), it is instructive to consider first its behavior in the special case that $\tilde{\alpha}_{1,\tau} = 0$ for all τ (i.e., the relationship between dew point depression and the temperature deviation does not change with GMT) and $\tilde{\rho}(\omega) = 1$ for all ω (i.e., temperature variability does not change). In this setting, it is straightforward to show that the simulation resulting from (10) is $\hat{D}_{y,d}^{(f)} = D_{y,d}^{(h)} + \tilde{\Delta}_{y,d}$, where $\tilde{\Delta}_{y,d}$ is the projected change in local mean temperature (equation (5)). That is, in this simplest setting, the future dew point value "follows" the mean temperature change, and also no modeling of the observed relationship between temperature and dew point is explicitly needed for the simulation.

In the general case, the procedure (10) can therefore be understood as an adjustment to that simple setting to account for changes with GMT and in the simulated temperature deviation value. The reason

that an explicit quantile regression model for *observed* dew points is now needed is because the future simulated temperature deviation differs from the observed deviation value if $\tilde{\rho}(\omega) \neq 1$, therefore requiring an estimate of the observed relationship $h_\tau(\cdot)$ to carry out the simulation. This implies that, if projected temperature variability changes are small, the dew point simulation will not be very sensitive to errors in estimating $h_\tau(\cdot)$ from observations, and primarily depends on the estimated changes from the GCM. This is important because there is more information available from the GCM about changes than there is information in the observed record about the baseline distribution. Thus the simulation procedure (10) shares the important feature with other observation-based simulation methods that it is not as dependent on explicitly modeling the behavior of observations as it is on modeling changes in the GCM output.

4.3   Illustration of method

Figure 3 shows the bivariate distributions of temperature and dew point in CESM1-LE, as well as in the observations and observation-based simulations, for the same location near Minneapolis, MN used in Figure 2. We show observed and simulated values for the years 1996-2006 and 2071-2080 during the month of July, along with the fitted median and 5th/95th percentile curves on July 16 when the GMT is equal to the average value in the respective time interval. (The displayed quantile curves are those for a fixed day of year, so are not strictly comparable to the data values shown, which vary in distribution over the month and across years; however, since this variation is relatively small compared to the differences between the future and historical periods, we still find it helpful to display for reference.)

Focusing first on CESM1-LE, the bivariate distributions are complex, as are the projected changes. The estimated $\tilde{\alpha}_{1,\tau}$ values are positive (Figure 3, top right), and larger in magnitude for the smaller dew point depression quantiles (or equivalently, the larger dew point quantiles); for the three displayed quantile curves, the estimated values are $\tilde{\alpha}_{1,0.05} = 0.13°\mathrm{C}^{-1}$, $\tilde{\alpha}_{1,0.5} = 0.081°\mathrm{C}^{-1}$, and $\tilde{\alpha}_{1,0.95} = 0.067°\mathrm{C}^{-1}$. This means that for a fixed temperature *deviation* (i.e., $\tilde{T}_{d,y} - \tilde{\mu}_{d,y}$, the covariate in model (8)), the dew point depression tends to be larger in the future time period than in the historical period (Figure 3, top right), especially on days that were expected to be more humid; that is, for a fixed realization of temperature variability around the forced mean, we see relative drying especially on the most humid days. However, since the mean temperature increases from the historical to the future time periods, a fixed temperature deviation corresponds to a higher temperature value, and e.g. we still see combined heat and humidity events that are more severe in the future than in the historical period, albeit less so than if the $\tilde{\alpha}_{1,\tau}$ values were negative (Figure 3, top left).

When we compare the bivariate distributions in CESM1-LE to those in the observations and our observation-based simulation (Figure 3, bottom), we see clear discrepancies between the two. For example, dew points show less variability in the observations for a fixed temperature than is apparent from CESM1-LE, and the estimated median and 5th/95th percentile curves are nearly monotonic in the observations whereas not in CESM1-LE. The simulation procedure (10) transfers the projected changes with GMT to the observed bivariate relationship, generating a simulation whose changes look similar to those in CESM but that is consistent with the original observations' distribution (and better preserves higher order features from the observations, like spatiotemporal relationships). Because of this difference in the underlying relationship between dew point and temperature, the aforementioned increase in future risk of high humidity at a historically high temperature value is smaller in the observation-based simulation compared to in CESM1-LE.

5   Results

The preceding sections describe our simulation procedure and illustrate both the projected changes from CESM1-LE and resulting simulations at one location near Minneapolis, MN. Here we provide a summary of projected changes from CESM1-LE over a region containing CONUS, and discuss the resulting simulations using GSOD station data. Simulations are produced using projected changes from the gridcell that is

nearest to the GSOD station in question and at least 50% land, and the years 1996-2005 are used to produce a simulation of the years 2071-2080 in the RCP8.5 scenario.

Recall that the temperature simulations require estimates of both local mean changes and changes in spectral densities. Figure 4, top, shows projected local mean temperature changes per degree warming in GMT, at three time points throughout the summer. That is, we show $\tilde{\beta}_1 + \tilde{\eta}(d)$ from equation (5), for d corresponding to the dates in question. Across CONUS, summertime local temperatures in CESM increase more quickly than in the GMT, and the increases are stronger towards the end rather than the beginning of the summer (especially in the Northwest). That is, the changes over CONUS are broadly similar to those shown in Figure 2(a) at the gridcell near Minneapolis.

Figure 4, bottom, shows projected changes in local temperature spectral densities per degree increase in GMT, at three different frequencies. As in Figure 2(b), we show the values $\exp\{\delta(\omega)/2\} - 1$, where $\delta(\omega)$ is as in equation (6); this corresponds to the relative change, per degree increase in GMT, in the standard deviation attributable to fluctuations at frequency ω. Likewise as in 2(b), the changes at $\omega = 0$ are labeled "JJA average" because these can be interpreted as changes in the standard deviation of the JJA average temperature in the resulting simulation. Summertime low-frequency temperature variability is largely projected to increase with increasing GMT, although changes in the Northeast and Southwest are estimated as negligible. Higher-frequency variability changes are more spatially variable, with decreases in the Northeast and Northwest and increases in the Southwest and Plains states.

After estimating changes in temperature distributions for our temperature simulation, the dew point simulation requires information about changes in the conditional quantile function of dew point. Figure 5 shows the estimated values of $\tilde{\alpha}_{1,\tau}$ from equation (8) for three quantile levels; recall that these are the changes in log dew point depression quantiles per increase in GMT, fixing the day of the year and the temperature deviation from its forced mean, and that the τ'th quantile for dew point depression corresponds to the $1 - \tau$'th quantile for dew point. As in Figure 3, top right, we see increases in dew point depression with increasing GMT, particularly in the lower quantiles. That is, for a fixed temperature deviation, we see relative drying particularly on the most humid days. The increases are smaller in the West and Southwest.

The observation-based simulations inherit changes in dew point depression for a fixed temperature *deviation* projected from the GCM (i.e., the parameter $\tilde{\alpha}_{1,\tau}$), but the resulting changes in dew point for a fixed temperature *value* can differ between the GCM and observation-based simulation because the same temperature value corresponds to different temperature deviations in equations (9) vs. (10) and involve different relationships in the GCM and observations, $\tilde{h}(\cdot)$ vs. $h(\cdot)$. Here we consider changes in the risk of high or low dew point values for a fixed temperature value. Figure 6 shows changes in the risk of historically high humidity events at a fixed historically high temperature, for CESM1-LE and the observation-based simulation. We express changes in risk in terms of odds ratios, comparing the future (2071-2080) odds of exceeding the historical (1990-2005) 95th percentile of dew point, conditional on local temperature being at the historical 95th percentile, to the historical odds. Overall, the probability of historically high humidity on a historically warm day increases in both CESM1-LE and the observation-based simulation, particularly in the Southeast. However, the increase in risk is typically larger in CESM1-LE than in the observation-based simulation. This result may stand in apparent contrast with Figure 5, which shows dew point depression increases (so dew point decreases) for a fixed temperature deviation in the future; however, a historically high temperature is associated with a smaller (or possibly negative) temperature deviation in the future, explaining the difference (e.g., compare the left and right panels of Figure 3 for a fixed x-axis value). See the Appendix for more details on the quantities we are showing here.

Figure 7 is analogous to Figure 6 but instead shows changes in the risk of historically *low*-humidity events at a fixed historically high temperature. Here we compare the 2071-2080 odds of dew point less than the 1990-2005 5th percentile of dew point, fixing local temperature at its 1990-2005 95th percentile, to the 1990-2005 odds. Unlike for humid heat events, the risk of dry heat events appears to decrease over CONUS except in the Northeast (particularly in later summer). However, the decrease is stronger in CESM1-LE than in the observation-based simulation, indicating that future low humidity heat events

may be a greater risk than indicated by CESM1-LE.

6   Discussion

In this work, we propose an observation-based joint simulation of future dew point and temperature that accounts for estimated changes from CESM1-LE in (a) mean temperature, (b) temperature variability at multiple timescales, and (c) changes in the relationship between dew point and temperature. We believe that the proposed simulation method is preferable compared to those based on attempting to statistically correct GCM runs, both because observation-based simulation procedures retain higher-order distributional features of the observations and because they require less explicit statistical modeling of observations (where data is limited) than of GCMs (where more data is available).

That said, observation-based simulations are of course limited by the observational record. The observational record is of limited length, but it may be possible to produce longer or multiple simulations by using a resampling procedure or through developing a statistical model of the variability (as in, e.g., McKinnon et al. (2017); McKinnon and Deser (2018)). In situ station measurements are also spatially and temporally limited, can contain systematic errors (e.g., issues discussed in Brown and DeGaetano (2009); Durre et al. (2010); Dunn et al. (2014); Rhines et al. (2015) and many others), and are point measurements that cannot be straightforwardly compared to a GCM gridcell, presenting potential challenges for any simulation method that requires information from observations (including model-based simulation procedures). We do not address these issues here, but we believe that it is advisable to repeat the simulation procedure with multiple data sources (if available) to address the effects of observational uncertainties.

Simulation procedures that combine information from GCM output with observations inherently require users to choose either what changes from the GCM should be reflected in the simulation (for observation-based procedures) or what features of the GCM output should be corrected (for model-based procedures). These notions become arguably more ambiguous in more complex settings, e.g., for multivariate simulations and where the simulation involves nonlinear transformations. We believe that one additional advantage of the approach taken here, which involves parametric statistical models to characterize GCM projected changes, is that such models make transparent what changes are inherited in the resulting simulation and these estimated changes are on their own a relevant summary of the GCM's behavior.

Our understanding of changes in compound extreme events, such as humid or dry heat events, remains limited. We hope that the statistical methods developed in this paper aid in the development of coherent and interpretable comparisons across other regions, forcing scenarios, and climate models, to ultimately enrich our understanding of these complex but important projected changes.

Acknowledgments

Data Availability Statement

The CESM1-LE data is available for download from the Climate Data Gateway at NCAR (formally known as the Earth System Grid) (https://doi.org/10.5065/d6j101d1). The GSOD weather station data is available at https://www1.ncdc.noaa.gov/pub/data/gsod/.

A1   Details on Statistical Models

In this appendix, we provide more detail on the statistical models used in this paper, and on how the relevant parameters are estimated.

A1.1 Mean temperature model

The functions $\gamma(d)$ and $\eta(d)$ in the mean temperature model (4) control seasonality and changes thereof with GMT, respectively. These functions are parameterized using the first two seasonal harmonics, i.e.

$$\gamma(d) = \sum_{k=1}^{2} \{\gamma_{k,1} \cos(2\pi k d/365)) + \gamma_{k,2} \sin(2\pi k d/365)\} \tag{11}$$

and similarly for $\eta(d)$. We have found that using only the first two harmonics is sufficient for modeling seasonality within JJA (see Supplement Section S2.1 for model validation) since most seasonal structure is dealt with by only examining JJA temperatures.

The parameters in the mean model (4) are estimated via least squares (separately for CESM1-LE and GSOD, but using all data within each respective dataset). Uncertainties in CESM1-LE projected mean changes (i.e., as shown in Figure 2(a)) are assessed using a residual block bootstrap, blocking by run in the 40-member ensemble. That is, we

1. calculate the residuals from the mean model, $\tilde{\mathbf{e}} = \tilde{\mathbf{T}} - \tilde{\mu}$,

2. resample residuals by resampling 40 runs with replacement from the 40 ensemble members, and

3. add the resampled residuals to the originally estimated mean function to obtain a bootstrap dataset, and refit the mean model (4).

This process is repeated 1,000 times and we use a 95% bootstrap percentile interval to display uncertainties.

A1.2 Temperature spectrum model

The function $\tilde{\delta}(\omega)$ in the temperature variability change model (6) controls the change in spectral density at frequency ω with GMT changes. This function is estimated similarly to the methods described in Poppick et al. (2016). We briefly describe the fundamental elements of the procedure, but refer readers to Poppick et al. (2016), Sections 4.1.1 and S2-S3, for more detail.

A preliminary estimate of $\tilde{\delta}(\omega)$ is first obtained by maximizing a block composite Whittle likelihood function (Dahlhaus et al., 1997) based on the separate periodograms of each year of JJA temperatures. The estimator based on maximizing this approximate likelihood combines information from the yearly periodograms across runs and yearly blocks, but still produces a rather rough estimate across frequency ω. This is therefore further smoothed across frequencies using an Epanechnikov kernel to obtain the final estimate. The bandwidth of the kernel smoother is chosen to minimize the sum of squared errors from leave-one-out cross-validation. We use a fixed bandwidth for all nonzero frequencies, but do not smooth the estimate of $\tilde{\delta}(0)$ beyond the initial estimate to avoid contaminating information about inter-annual variability changes with intra-seasonal changes.

Approximate standard errors of the resulting estimate of $\tilde{\delta}(\omega)$ can then be derived through a tedious but straightforward calculation involving the second derivative of the composite likelihood function and the kernel smoother weights. The 95% intervals shown in Figure 2(b) are the $\pm 2SE$ intervals around the final estimate of $\tilde{\delta}(\omega)$. The determination of statistical significance shown in Figure 4, bottom, is based on first calculating a p-value from the Z-statistic associated with the estimate of $\tilde{\delta}(\omega)$ and its standard error, and then controlling the false discovery rate (FDR) at 5% using the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995; Wilks, 2016). The FDR control procedure accounts for the multiple-testing problem arising from the fact that we are assessing significance at each of the 1,344 gridcells (so would expect many small p-values even if there were no gridcells with nonzero true changes).

A1.3 Dew point quantile regression model

The function $g_\tau(\cdot)$ in quantile regression model (8) controls seasonality in dew point (after controlling for the local temperature deviation and GMT trend) and the function $h_\tau(\cdot)$ controls the relationship between dew point and the local temperature deviation for a fixed day and GMT. As for the mean temperature model, we parameterize $g_\tau(\cdot)$ using the first two seasonal harmonics (see equation (11)). The function $h_\tau(\cdot)$ is parameterized using a natural spline with 10 degrees of freedom (i.e., 8 knots). The knots are chosen as the empirical quantiles of the local temperature deviation, with boundary knots equal to the minimum and maximum observed value.

Model (8) is estimated for the quantile levels $\tau = 0.005, 0.01, 0.05, 0.1, 0.25, 0.5, 0.75, 0.9, 0.95,$ 0.99, and 0.995. If quantile functions cross, this is dealt with by reordering fitted values (Chernozhukov et al., 2010). Values for other quantile levels needed for the dew point simulation are obtained via linear interpolation. That is, the value $\tau_{y,d}^*$ in equation (10) is calculated by linearly interpolating between the two quantile levels yielding fitted values closest to the observed value $D_{y,d}^{(h)}$; the simulation $\hat{D}_{y,d}^{(f)}$ is then calculated by linearly interpolating between the two corresponding future simulation values in equation (10).

Finally, note the reported GSOD data are first rounded to the nearest tenth degree Fahrenheit, then converted to Celsius and re-rounded to the nearest tenth degree Celsius; to avoid numerical issues that arise when estimating quantile regression models with discrete response variables (Machado and Silva, 2005), we therefore add a random jitter to each observed dew point depression value when fitting the model to observations: dew point depression is produced by jittering both temperature and dew point separately by $U_F \times \frac{5}{9} + U_C$, where U_F and $U_C \overset{\text{indep}}{\sim} \text{Unif}(-0.05, 0.05)$.

A2 Details on Calculations of Changes in Risk Probabilities

Figures 6 and 7 show changes in the risk of historically high humidity and low humidity heat events, respectively, at a fixed historically high temperature. Here we provide more details on the quantities we are showing.

Write $\bar{G}^{(h)}$ and $\bar{G}^{(f)}$ for the average GMT forced trend over the years 1990-2005 and 2071-2080, respectively. Write $\tilde{\mu}_d^{(h)}$ and $\tilde{\mu}_d^{(f)}$ for the estimated local mean temperature on day d (from equation (4)) in CESM1-LE given those historical or future GMT forced trend values. Write $\tilde{T}_{95,m}^{(h)}$ for the empirical 95th percentile of local temperatures in month m over the years 1990-2005 in CESM1-LE (calculated across days in the month, years, and the 40 ensemble members). Finally, write $\tilde{q}_{95,d}^{(h)}$ for the estimated 95th percentile of dew points in CESM1-LE on day d from model (8), given local temperature $\tilde{T}_{95,m}^{(h)}$ and GMT anomaly $\bar{G}^{(h)}$, i.e.,

$$
\begin{aligned}
\tilde{q}_{95,d}^{(h)} &= F_{\tilde{D}_{d,y}}^{-1}\left(0.95 | \tilde{T}_{d,y} = \tilde{T}_{95,m}^{(h)}, \tilde{G}_y = \bar{G}^{(h)}\right) \\
&= \tilde{T}_{95,m}^{(h)} - \exp\left\{\tilde{\alpha}_{0,1-0.95} + \tilde{g}_{1-0.95}(d) + \tilde{\alpha}_{1,1-0.95}\bar{G}^{(h)} + \tilde{h}_{1-0.95}(\tilde{T}_{95,m}^{(h)} - \tilde{\mu}_d^{(h)})\right\},
\end{aligned}
$$

and write $q_{95,d}^{(h)}$ for the analogous quantity from observations (replacing all of the above relevant quantities with their analogues from the observations).

The *future risk* of a historically high humidity heat event in CESM1-LE is defined as the probability on day d of exceeding $\tilde{q}_{95,d}^{(h)}$ given a historically high local temperature, $\tilde{T}_{95,m}^{(h)}$, but future GMT anomaly $\bar{G}^{(f)}$. That is,

$$
\begin{aligned}
\tilde{p}_{95,d}^{(f)} &= \Pr\left(\tilde{D}_{d,y} > \tilde{q}_{95,d}^{(h)} | \tilde{T}_{d,y} = \tilde{T}_{95,m}^{(h)}, \tilde{G}_y = \bar{G}^{(f)}\right) \\
&= 1 - F_{\tilde{D}_{d,y}}\left(\tilde{q}_{95,d}^{(h)} | \tilde{T}_{d,y} = \tilde{T}_{95,m}^{(h)}, \tilde{G}_y = \bar{G}^{(f)}\right),
\end{aligned}
$$

By definition, the *historical* risk is 0.05 (because $\tilde{q}_{95,d}^{(h)}$ is the estimated historical 95th percentile). We measure the change in risk in terms of the odds ratio,

$$\tilde{\omega}_{95,d} = \frac{\tilde{p}_{95,d}^{(f)}/(1 - \tilde{p}_{95,d}^{(f)})}{0.05/(1 - 0.05)},$$

and similarly $\omega_{95,d}$ for the observation-based simulation. If the odds ratio is greater than 1, the future risk is greater than the historical risk. Figure 6 shows $\log_{10}(\tilde{\omega}_{95,d})$ (first row), $\log_{10}(\omega_{95,d})$ (second row), and $\log_{10}(\omega_{95,d}/\tilde{\omega}_{95,d})$ (bottom row) on the days June 1, July 1, and August 1.

Figure 7 is similar except shows changes in the risk of historically low humidity heat events, which we define as dew points falling below the historical 5th percentile (rather than above the 95th percentile) on days with local temperature equal to the historical 95th percentile value.

References

Barreca, A. I. (2012). Climate change, humidity, and mortality in the United States. *Journal of Environmental Economics and Management 63*(1), 19–34.

Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological) 57*(1), 289–300.

Brands, S., S. Herrera, J. Fernández, and J. M. Gutiérrez (2013). How well do CMIP5 earth system models simulate present climate conditions in Europe and Africa? *Climate dynamics 41*(3-4), 803–817.

Brown, P. J. and A. T. DeGaetano (2009). A method to detect inhomogeneities in historical dewpoint temperature series. *Journal of Applied Meteorology and Climatology 48*(11), 2362–2376.

Cannon, A. J. (2018). Multivariate quantile mapping bias correction: an n-dimensional probability density function transform for climate model simulations of multiple variables. *Climate Dynamics 50*(1-2), 31–49.

Cannon, A. J., C. Piani, and S. Sippel (2020). Bias correction of climate model output for impact models. In *Climate Extremes and Their Implications for Impact and Risk Assessment*, pp. 77–104. Elsevier.

Castruccio, S., D. J. McInerney, M. L. Stein, F. Liu Crouch, R. L. Jacob, and E. J. Moyer (2014). Statistical emulation of climate model projections based on precomputed GCM runs. *Journal of Climate 27*(5), 1829–1844.

CESM Project (2020). Known issues. [Online; accessed 6-February-2020].

Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica 78*(3), 1093–1125.

Coffel, E. D., R. M. Horton, and A. de Sherbinin (2017). Temperature and humidity based projections of a rapid rise in global heat stress exposure during the 21st century. *Environmental Research Letters 13*(1), 014001.

Dahlhaus, R. et al. (1997). Fitting time series models to nonstationary processes. *The Annals of Statistics 25*(1), 1–37.

Dai, A., J. C. Fyfe, S.-P. Xie, and X. Dai (2015). Decadal modulation of global surface temperature by internal climate variability. *Nature Climate Change 5*(6), 555.

Davis, R. E., G. R. McGregor, and K. B. Enfield (2016). Humidity: A review and primer on atmospheric moisture and human health. *Environmental research 144*, 106–116.

Dixon, K. W., J. R. Lanzante, M. J. Nath, K. Hayhoe, A. Stoner, A. Radhakrishnan, V. Balaji, and C. F. Gaitán (2016). Evaluating the stationarity assumption in statistically downscaled climate projections: is past performance an indicator of future results? *Climatic Change 135*(3-4), 395–408.

Dunn, R., K. Willett, C. Morice, and D. Parker (2014). Pairwise homogeneity assessment of hadisd. *Climate of the Past 10*(4), 1501.

Durre, I., M. J. Menne, B. E. Gleason, T. G. Houston, and R. S. Vose (2010). Comprehensive automated quality assurance of daily surface observations. *Journal of Applied Meteorology and Climatology 49*(8), 1615–1633.

Fischer, E. M. and R. Knutti (2013). Robust projections of combined humidity and temperature extremes. *Nature Climate Change 3*(2), 126.

Gao, M. and C. L. Franzke (2017). Quantile regression–based spatiotemporal analysis of extreme temperature change in China. *Journal of Climate 30*(24), 9897–9914.

Guo, Q., J. Chen, X. Zhang, M. Shen, H. Chen, and S. Guo (2019). A new two-stage multivariate quantile mapping method for bias correcting climate model outputs. *Climate Dynamics 53*(5-6), 3603–3623.

Haugen, M. A., M. L. Stein, E. J. Moyer, and R. L. Sriver (2018). Estimating changes in temperature distributions in a large ensemble of climate simulations using quantile regression. *Journal of Climate 31*(20), 8573–8588.

Haugen, M. A., M. L. Stein, R. L. Sriver, and E. J. Moyer (2019). Future climate emulations using quantile regressions on large ensembles. *Advances in Statistical Climatology, Meteorology and Oceanography 5*(1), 37–55.

Hawkins, E., T. M. Osborne, C. K. Ho, and A. J. Challinor (2013). Calibration and bias correction of climate projections for crop modelling: an idealised case study over Europe. *Agricultural and Forest Meteorology 170*, 19–31.

Ho, C. K., D. B. Stephenson, M. Collins, C. A. Ferro, and S. J. Brown (2012). Calibration strategies: a source of additional uncertainty in climate change projections. *Bulletin of the American Meteorological Society 93*(1), 21–26.

Hsiao, J., A. L. Swann, and S.-H. Kim (2019). Maize yield under a changing climate: The hidden role of vapor pressure deficit. *Agricultural and Forest Meteorology 279*, 107692.

IPCC (2013). *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.* New York, NY: Cambridge University Press. T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, P. M. Midgley (eds.).

John, V. and B. J. Soden (2007). Temperature and humidity biases in global climate models and their impact on climate feedbacks. *Geophysical Research Letters 34*(18).

Kay, J. E., C. Deser, A. Phillips, A. Mai, C. Hannay, G. Strand, J. M. Arblaster, S. Bates, G. Danabasoglu, J. Edwards, et al. (2015). The Community Earth System Model (CESM) large ensemble project: A community resource for studying climate change in the presence of internal climate variability. *Bulletin of the American Meteorological Society 96*(8), 1333–1349.

Knutson, T. R. and J. J. Ploshay (2016). Detection of anthropogenic influence on a summertime heat stress index. *Climatic Change 138*(1-2), 25–39.

Koenker, R. and G. Bassett Jr (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.

Lamarque, J.-F., T. C. Bond, V. Eyring, C. Granier, A. Heil, Z. Klimont, D. Lee, C. Liousse, A. Mieville, B. Owen, et al. (2010). Historical (1850–2000) gridded anthropogenic and biomass burning emissions of reactive gases and aerosols: methodology and application. *Atmospheric Chemistry and Physics 10*(15), 7017–7039.

Lanzante, J. R., K. W. Dixon, M. J. Nath, C. E. Whitlock, and D. Adams-Smith (2018). Some pitfalls in statistical downscaling of future climate. *Bulletin of the American Meteorological Society 99*(4), 791–803.

Lee, S.-M. and S.-K. Min (2018). Heat stress changes over east asia under 1.5° and 2.0°C global warming targets. *Journal of Climate 31*(7), 2819–2831.

Leeds, W. B., E. J. Moyer, and M. L. Stein (2015). Simulation of future climate under changing temporal covariance structures. *Advances in Statistical Climatology, Meteorology and Oceanography 1*(1), 1–14.

Li, C., X. Zhang, F. Zwiers, Y. Fang, and A. M. Michalak (2017). Recent very hot summers in northern hemispheric land areas measured by wet bulb globe temperature will be the norm within 20 years. *Earth's Future 5*(12), 1203–1216.

Li, D., J. Yuan, and R. B. Kopp (2020). Escalating global exposure to compound heat-humidity extremes with warming. *Environmental Research Letters*.

Machado, J. A. F. and J. S. Silva (2005). Quantiles for counts. *Journal of the American Statistical Association 100*(472), 1226–1237.

Matiu, M., D. P. Ankerst, and A. Menzel (2016). Asymmetric trends in seasonal temperature variability in instrumental records from ten stations in Switzerland, Germany and the UK from 1864 to 2012. *International Journal of Climatology 36*(1), 13–27.

McKinnon, K. A. and C. Deser (2018). Internal variability and regional climate trends in an observational large ensemble. *Journal of Climate 31*(17), 6783–6802.

McKinnon, K. A. and A. Poppick (2020). Estimating changes in the observed relationship between humidity and temperature using noncrossing quantile smoothing splines. *Journal of Agricultural, Biological and Environmental Statistics*.

McKinnon, K. A., A. Poppick, E. Dunn-Sigouin, and C. Deser (2017). An "observational large ensemble" to compare observed and modeled temperature trend uncertainty due to internal variability. *Journal of Climate 30*(19), 7585–7598.

McKinnon, K. A., A. Rhines, M. P. Tingley, and P. Huybers (2016). The changing shape of northern hemisphere summer temperature distributions. *Journal of Geophysical Research: Atmospheres 121*(15), 8849–8868.

Mehrotra, R. and A. Sharma (2015). Correcting for systematic biases in multiple raw GCM variables across a range of timescales. *Journal of Hydrology 520*, 214–223.

Mehrotra, R. and A. Sharma (2016). A multivariate quantile-matching bias correction approach with auto-and cross-dependence across multiple time scales: implications for downscaling. *Journal of Climate 29*(10), 3519–3539.

Meinshausen, M., S. J. Smith, K. Calvin, J. S. Daniel, M. Kainuma, J.-F. Lamarque, K. Matsumoto, S. Montzka, S. Raper, K. Riahi, et al. (2011). The RCP greenhouse gas concentrations and their extensions from 1765 to 2300. *Climatic Change 109*(1-2), 213.

Pal, J. S. and E. A. Eltahir (2016). Future temperature in southwest Asia projected to exceed a threshold for human adaptability. *Nature Climate Change 6*(2), 197.

Piani, C., G. Weedon, M. Best, S. Gomes, P. Viterbo, S. Hagemann, and J. Haerter (2010). Statistical bias correction of global simulated daily precipitation and temperature for the application of hydrological models. *Journal of Hydrology 395*(3-4), 199–215.

Poppick, A., D. J. McInerney, E. J. Moyer, and M. L. Stein (2016). Temperatures in transient climates: improved methods for simulations with evolving temporal covariances. *The Annals of Applied Statistics 10*(1), 477–505.

Poppick, A., E. J. Moyer, and M. L. Stein (2017). Estimating trends in the global mean temperature record. *Advances in Statistical Climatology, Meteorology and Oceanography 3*(1), 33–33.

Pryor, S. and J. Schoof (2016). Evaluation of near-surface temperature, humidity, and equivalent temperature from regional climate models applied in type II downscaling. *Journal of Geophysical Research: Atmospheres 121*(7), 3326–3338.

Reich, B. J. (2012). Spatiotemporal quantile regression for detecting distributional changes in environmental processes. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 61*(4), 535–553.

Rhines, A., K. A. McKinnon, M. P. Tingley, and P. Huybers (2017). Seasonally resolved distributional trends of north american temperatures show contraction of winter variability. *Journal of Climate 30*(3), 1139–1157.

Rhines, A., M. P. Tingley, K. A. McKinnon, and P. Huybers (2015). Decoding the precision of historical temperature observations. *Quarterly Journal of the Royal Meteorological Society 141*(693), 2923–2933.

Santer, B. D., T. M. Wigley, M. E. Schlesinger, and J. F. Mitchell (1990). Developing climate scenarios from equilibrium GCM results. Technical Report 47, Max-Planck-Institut für Meteorologie.

Schoof, J., T. Ford, and S. Pryor (2017). Recent changes in U.S. regional heat wave characteristics in observations and reanalyses. *Journal of Applied Meteorology and Climatology 56*(9), 2621–2636.

Seager, R., A. Hooks, A. P. Williams, B. Cook, J. Nakamura, and N. Henderson (2015). Climatology, variability, and trends in the us vapor pressure deficit, an important fire-related meteorological quantity. *Journal of Applied Meteorology and Climatology 54*(6), 1121–1141.

Simmons, A., K. Willett, P. Jones, P. Thorne, and D. Dee (2010). Low-frequency variations in surface atmospheric humidity, temperature, and precipitation: Inferences from reanalyses and monthly gridded observational data sets. *Journal of Geophysical Research: Atmospheres 115*, D01110.

Tian, B., E. J. Fetzer, B. H. Kahn, J. Teixeira, E. Manning, and T. Hearty (2013). Evaluating CMIP5 models using airs tropospheric air temperature and specific humidity climatology. *Journal of Geophysical Research: Atmospheres 118*(1), 114–134.

Vrac, M. (2018). Multivariate bias adjustment of high-dimensional climate simulations: the rank resampling for distributions and dependences (r2d2) bias correction. *Hydrol. Earth Syst. Sci 22*, 3175–3196.

Vrac, M. and P. Friederichs (2015). Multivariate – intervariable, spatial, and temporal – bias correction. *Journal of Climate 28*(1), 218–237.

Wilks, D. (2016). "The stippling shows statistically significant grid points": How research results are routinely overstated and overinterpreted, and what to do about it. *Bulletin of the American Meteorological Society 97*(12), 2263–2273.

Willett, K., R. Dunn, P. Thorne, S. Bell, M. De Podesta, D. Parker, P. Jones, and C. Williams Jr (2014). HadISDH land surface multi-variable humidity and temperature record for climate monitoring. *Climate of the Past 10*(6).

Yuan, J., M. L. Stein, and R. E. Kopp (2019). The evolving distribution of relative humidity conditional upon daily maximum temperature in a warming climate. *Earth and Space Science Open Archive*, 38.

Zhao, Y., A. Ducharne, B. Sultan, P. Braconnot, and R. Vautard (2015). Estimating heat stress from climate-based indicators: present-day biases and future spreads in the CMIP5 global climate model ensemble. *Environmental Research Letters 10*(8), 084013.

Zscheischler, J., E. M. Fischer, and S. Lange (2019). The effect of bias adjustment on impact modeling. *Earth Sys. Dyn. 10*, 31–43.
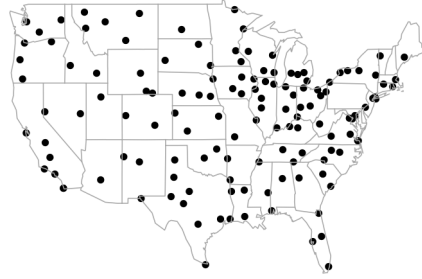
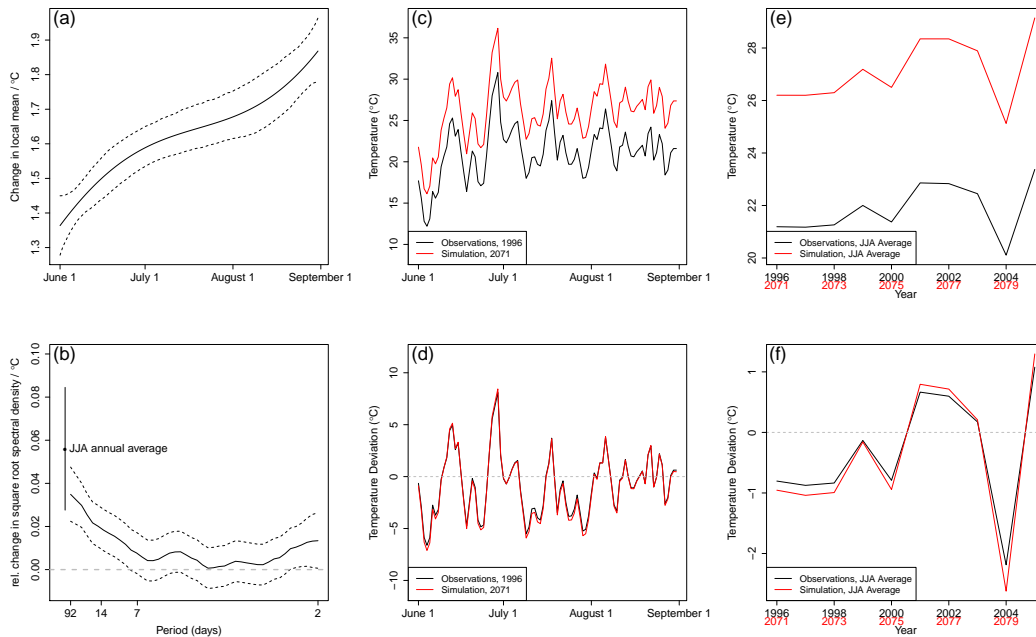Figure 1: Locations of GSOD stations used in our analysis.

Figure 2: (a) CESM1-LE projected change in mean temperature for a one-degree increase in GMT (with approximate 95% confidence interval, dashed), by day of the summer, for a gridcell near Minneapolis, MN. The summertime mean temperatures at this location increase more quickly than GMT, and the increase is greater towards the end rather than the beginning of the summer. (b) Estimated relative change in the square root spectral density for a one-degree increase in GMT at the same gridcell (with approximate 95% confidence interval, dashed). Variability increases at low frequencies but changes are negligible for higher frequencies. (c)-(f) Illustration of the proposed temperature simulation method using observations from Minneapolis-Saint Paul International Airport. Observations are from the years 1996-2005 and the simulation is for the years 2071-2080 of the RCP8.5 scenario. Top, temperatures; bottom, temperature deviations (i.e., temperature minus its estimated forced mean); left, first year of daily values in observations (black) and resulting simulation (red); right, JJA average values over the observed record and simulation period.
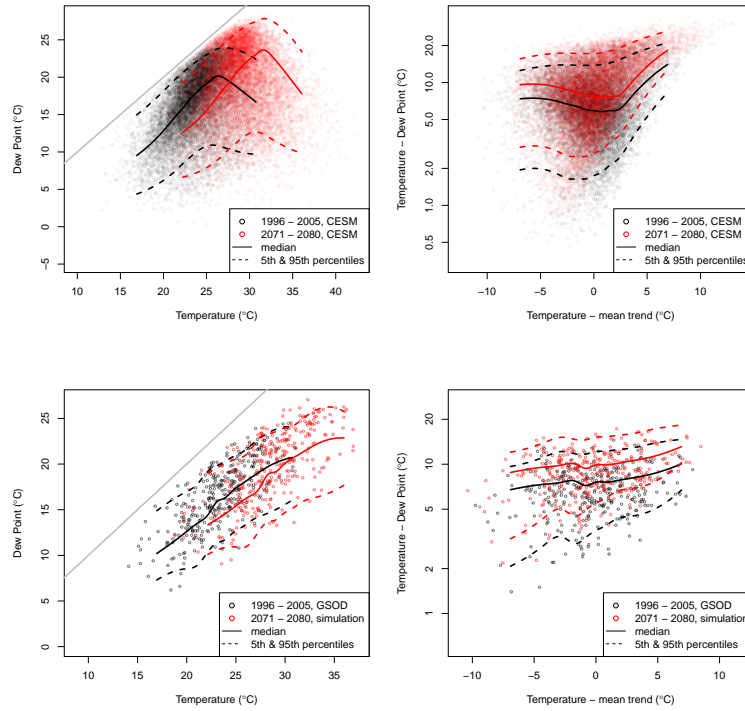
Figure 3: Comparison of historical (black) vs. future (red) bivariate distributions in CESM1-LE and observations vs. simulated values during the time periods 1996-2005 and 2071-2080 for the month of July. Top, CESM1-LE; bottom, observations and simulations; left, dew point vs. temperature; right, dew point depression (i.e., temperature minus dew point) vs. temperature deviation (i.e., temperature minus its estimated forced mean). Lines show the fitted median and 5th/95th percentile curves on July 16 when GMT is equal to its average value in the relevant time period.
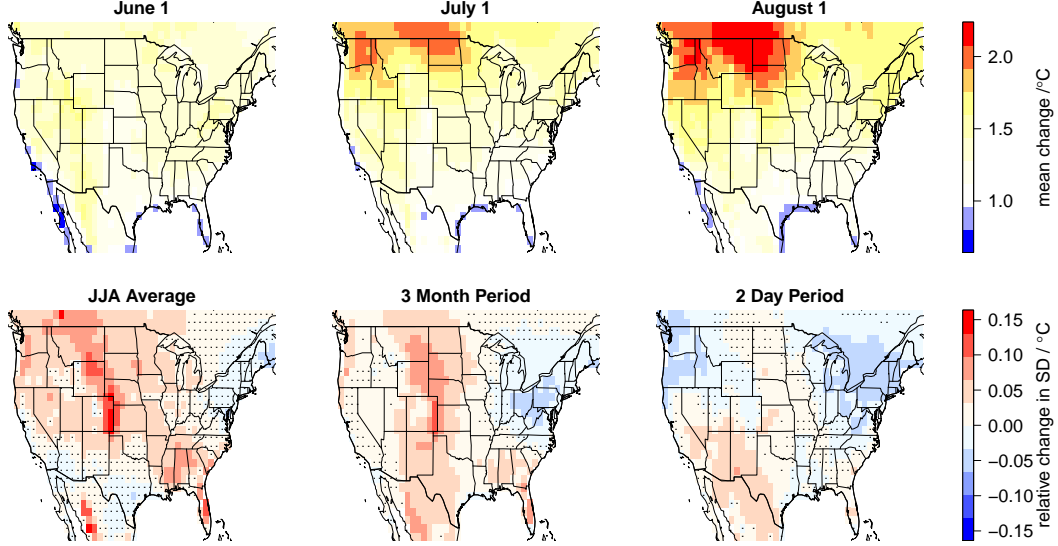
Figure 4: Top, projected mean temperature changes per degree increase in GMT on June 1, July 1, and August 1 (i.e., $\tilde{\beta}_1 + \tilde{\eta}(d)$ from equation (5), for d corresponding to the dates in question). Projected mean changes are stronger towards the end rather than the beginning of summer, as also illustrated for the example gridcell in Figure 2(a). Bottom, projected changes in temperature spectral densities per degree increase in GMT, for three different frequencies. The left plot is for $\omega = 0$, which corresponds to the change in the standard deviation of the JJA average temperature; the middle and right plots correspond to frequencies $\omega = 1/92$ and $\omega = 1/2$, respectively. We are displaying the values $\exp\{\delta(\omega)/2\} - 1$, where $\delta(\omega)$ is as in equation (6); this can be interpreted as the relative change, per degree increase in GMT, in the standard deviation attributable to fluctuations at the frequency ω. Gridcells are stippled if the estimated change is not significantly different from zero, controlling the false discovery rate at 5% (see Section A1.2 for details).
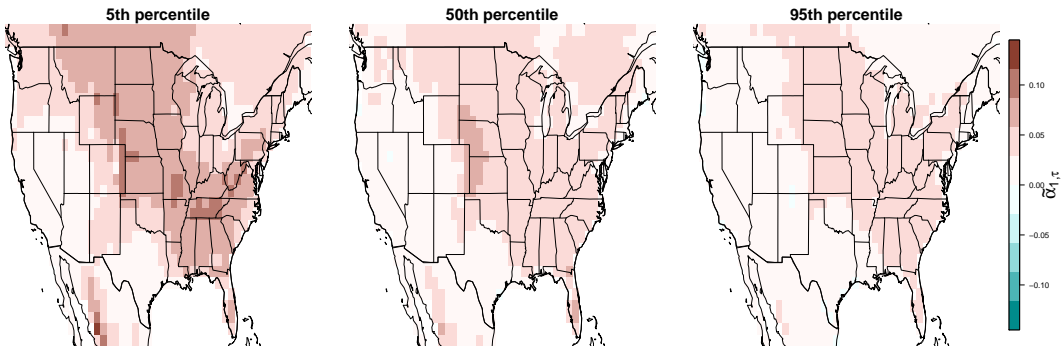

Figure 5: Projected changes in log dew point depression quantiles per degree increase in GMT, fixing the temperature deviation from the forced trend (i.e., $\tilde{\alpha}_{1,\tau}$ from equation (8)). Recall that the 5th percentile of dew point depression corresponds to the 95th percentile of dew point, and visa versa. Over land, summertime dew point depression increases with GMT, fixing the local temperature deviation, and more so on more humid days.
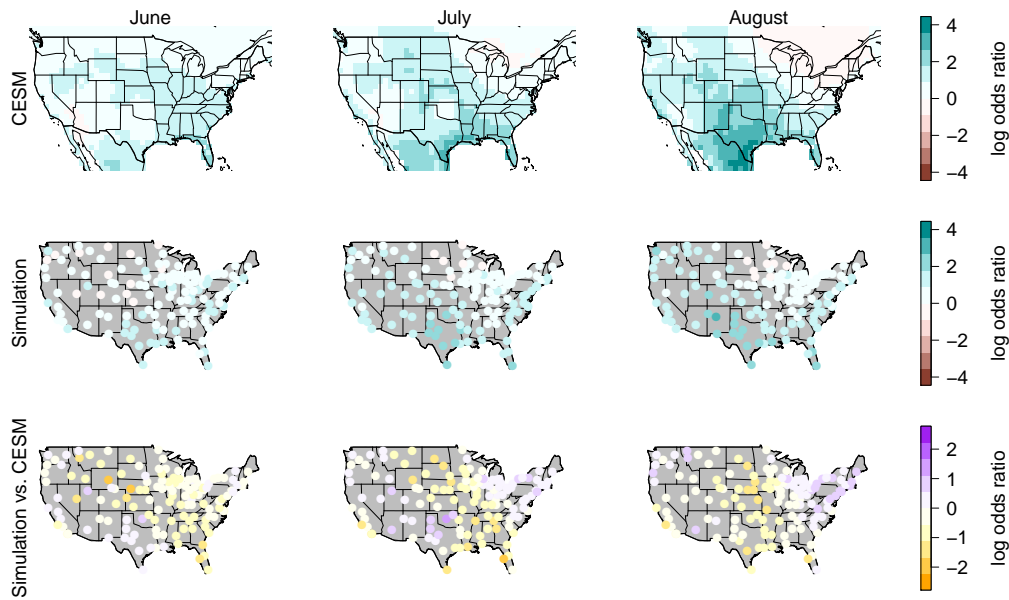
Figure 6: Change in risk of historically high (95th percentile) dew point on historically high temperature (95th percentile) days, by month, comparing the years 2071-2080 vs. 1990-2005 in CESM1-LE (top), in the observation-based simulation (middle), and the difference between the two (bottom). The top two rows show the log (base 10) odds ratio comparing the future vs. historical risk; positive values correspond to an increase in risk. The bottom row shows the difference between the middle and top rows; positive values indicate a larger future risk in the observation-based simulation than in CESM1-LE. See Appendix A2 for details.
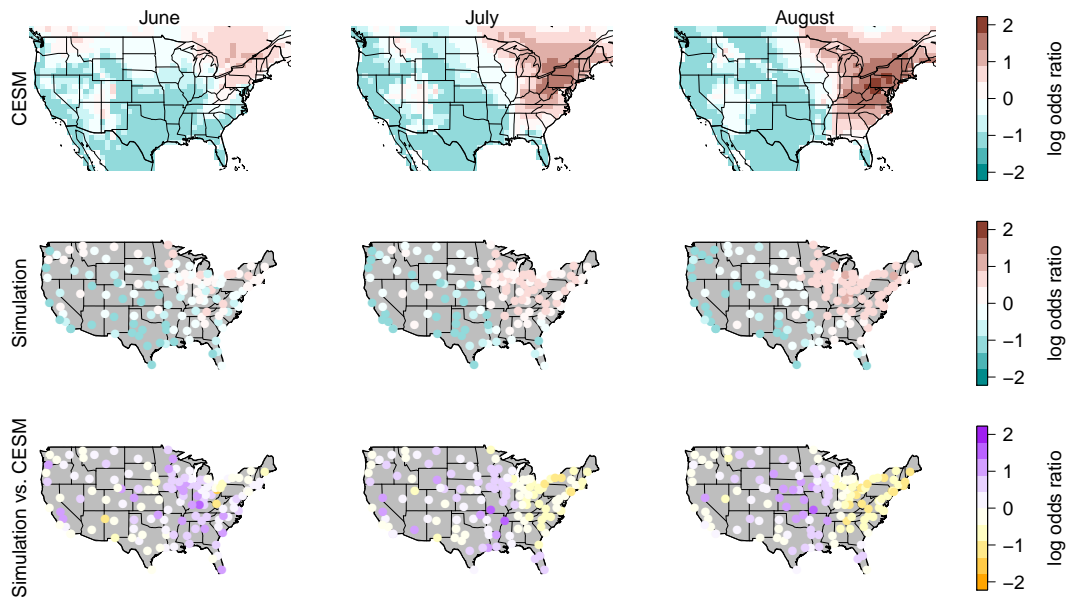
Figure 7: Change in risk of historically low (5th percentile) dew point on historically high temperature (95th percentile) days, by month, comparing the years 2071-2080 vs. 1990-2005 in CESM1-LE (top), in the observation-based simulation (middle), and the difference between the two (bottom). The top two rows show the log (base 10) odds ratio comparing the future vs. historical risk; positive values correspond to an increase in risk. The bottom row shows the difference between the middle and top rows; positive values indicate a larger future risk in the observation-based simulation than in CESM1-LE. See Appendix A2 for details.