Title

Estimating Changes in the Observed Relationship Between Humidity and Temperature Using Noncrossing Quantile Smoothing Splines

Journal

Journal of Agricultural, Biological and Environmental Statistics, 25(3)

ISSN

1085-7117

Authors

McKinnon, Karen A
Poppick, Andrew

Publication Date

2020-09-01

DOI

10.1007/s13253-020-00393-4

Peer reviewed

Estimating Changes in the Observed Relationship Between Humidity and Temperature Using Noncrossing Quantile Smoothing Splines

Karen A. McKinnon & Andrew Poppick

Journal of Agricultural, Biological and Environmental Statistics

ISSN 1085-7117

JABES
DOI 10.1007/s13253-020-00393-4

Springer

Supplementary materials for this article are available at https://doi.org/10.1007/s13253-020-00393-4.

Estimating Changes in the Observed Relationship Between Humidity and Temperature Using Noncrossing Quantile Smoothing Splines

Karen A. McKinnon and Andrew Poppick

The impacts of warm season heat extremes are dependent on both temperature and humidity, so it is critical to properly model their relationship, including how it may be changing. This presents statistical challenges because the bivariate temperature–humidity (measured here by dew point) distribution is complex and spatially variable. Here, we develop a flexible, semiparametric model based on quantile smoothing splines to summarize the distributional dependence of dew point on temperature, including how the dependence is changing with increasing global mean temperature. Noncrossing constraints enforce both the validity of the modeled distributions and the physical constraint that dew point cannot exceed temperature. The proposed method is first demonstrated with four synthetic, representative case studies. We then apply it to data from 2416 weather stations spanning the globe, with a focus on analyzing dew point trends during hot days. In general, dew point is increasing on both hot, humid and hot, dry days in the tropics and high latitudes, but decreasing in the subtropics, especially on hot, dry days. These changes appear to be mostly explained by changes in the temperature–dew point relationship, rather than by increases in temperature with a fixed temperature–dew point relationship.

Key Words: Dew point; Global Summary of the Day; High-dimensional regularization; Noncrossing quantiles; Quantile regression; Quantile smoothing splines.

1. INTRODUCTION

A large body of work has illustrated the role of anthropogenic climate change in increasing the probability of warm season heat waves (e.g., Meehl and Tebaldi 2004; Perkins et al. 2012), which remain the largest contributor among weather events to human mortality in

K. A. McKinnon (✉) Department of Statistics and Institute of the Environment and Sustainability, University of California, Los Angeles, Los Angeles, CA 90095, USA (E-mail: *kmckinnon@ucla.edu*).
A. Poppick Department of Mathematics and Statistics, Carleton College, Northfield, MN 55057, USA (E-mail: *apoppick@carleton.edu*).

Journal of Agricultural, Biological, and Environmental Statistics
https://doi.org/10.1007/s13253-020-00393-4

Published online: 15 May 2020

high-income countries (Wahlstrom and Guha-Sapir 2015). However, the impact of temperature on human comfort and mortality, as well as on plant health and productivity, is additionally dependent on the moisture content of the air. Physiologically, increased atmospheric moisture reduces the ability of the body to cool itself through sweating; epidemiological studies have found an increase in hospitalizations (Mastrangelo et al. 2007; Guirguis et al. 2014) and mortality (Barreca 2012) on hot and humid days that exceeds what would be expected on hot but less humid days. Conversely, summer-maturing crops such as corn, soy, and cotton are sensitive to low-humidity, high-heat conditions, which are associated with increased moisture demand from the atmosphere (Tanner and Sinclair 1983). Similarly, the occurrence and rate of spread of wildfires have been related to these hot, dry conditions, which desiccate the natural foliage, thus increasing the ease of combustion (Seager et al. 2015).

Depending on the impact of interest, it may be possible to collapse information about heat and humidity into a single variable and proceed with a traditional univariate analysis. For example, wet bulb temperature, defined as the temperature indicated by a thermometer covered in a water-soaked cloth, is a reasonable indicator of human comfort and has been used to identify regions in climate projections that may exceed theoretical limits for human tolerance within the next century (Pal and Eltahir 2016; Coffel et al. 2017). Vapor pressure deficit measures moisture demand from the atmosphere and plays a major role in crop health (Hsiao et al. 2019) and wildfire probability (Seager et al. 2015). However, a more general analysis of the relationship between temperature and moisture, and how it may be changing with climate change, requires a multivariate analysis.

The underlying dependence structure between climate variables such as temperature and humidity tends to be complex and spatially variable, and potentially inconsistent between the center and tails of the distribution. Misspecification of this dependence structure can lead to large errors in estimation of the probability of joint extremes (AghaKouchak et al. 2014; Zscheischler et al. 2018), so flexible statistical methods are required in order to model both the bivariate dependence structure and how it is changing.

One common approach to estimating multivariate distributions and the associated probability of compound extremes in climate data is the use of copulas, which allow for a separation of the marginal distributions from the dependence structure (Nelsen 2007). Copulas have been used to estimate the joint distribution of hot and dry events in California (AghaKouchak et al. 2014), storm surge and heavy precipitation along the United States coast (Wahl et al. 2015), and rainfall intensity and storm duration (De Michele and Salvadori 2003), among other applications. See Hao et al. (2018) for a review of the use of copulas in hydroclimate. While powerful, the copula approach requires choosing a copula family that encodes strong assumptions about the multivariate distribution (e.g., about the tail dependence structure) and typically is chosen for mathematical convenience. Copula models can further be difficult to interpret, especially if different copula families are chosen to model different regions due to spatial variability in dependence structures (as in, e.g., Ribeiro et al. 2019). To allow for greater flexibility, other researchers have developed nonparametric methods to estimate bivariate exceedance probabilities (Cooley et al. 2019).

Alternatively, it may be of interest to view one climate variable as a ‘driver’ (one controls or causes the other) and the other as the response for physical reasons, practical consider-

ations, or both. In the case of temperature and humidity, both of these motivations apply. Physically, the maximum (saturated) moisture content of the atmosphere is a function of temperature. Practically, many climate analyses focus primarily on temperature, so modeling other variables as a function of temperature can be advantageous. This framing opens up the use of flexible regression-based methods to model compound extreme events. We focus on quantile regression, which provides information about the sensitivity of a given quantile of a distribution to covariates (Koenker and Bassett 1978). Quantile regression has been used extensively to quantify time trends in the distribution of climate variables including daily temperature (Reich 2012; McKinnon et al. 2016; Rhines et al. 2017; Haugen et al. 2018), tropical cyclone intensity (Elsner et al. 2008), and drought (Chamaillé-Jammes et al. 2007), but the approach has typically not been used with multiple covariates as we do here. An appeal of quantile regression is the ability to nonparametrically model how the extremal percentiles of a conditional distribution are changing, as compared its center. The approach is complementary to fitting an extreme value distribution as a function of time-varying covariates (e.g., Katz et al. 2002; Ouarda and Charron 2019, for climate variables), which requires making parametric assumptions based on asymptotic approximations but may be used to extrapolate further into the tails of the distribution. We develop a quantile regression model to describe the humidity distribution conditional on local temperature, and identify changes in the relationship associated with increasing global temperatures.

Unlike dry-bulb temperature, which is the familiar temperature measurement made with a thermometer, there are many different but interrelated measures of humidity. We quantify humidity via dew point temperature, which is defined as the temperature to which an air parcel would need to be cooled in order for the moisture in it to condense. Dew point has the advantages that (1) it is measured directly by weather stations, (2) it can be interpreted straightforwardly as a metric for the amount of moisture in the air, (3) it has a minimal diurnal cycle so its daily average can be more straightforwardly calculated from a finite number of point-in-time measurements, and (4) it is a good indicator of human comfort (Davis et al. 2016). Dew point and (dry-bulb) temperature are connected in two key ways. First, dew point is physically constrained to not exceed temperature. Second, given sufficient moisture availability from the environment, dew point will increase with increasing temperature because a warmer atmosphere can hold more moisture. However, moisture availability is temporally and spatially variable, leading to substantial variability in the temperature–dew point relationship.

The characteristics of daily temperature and dew point data raise at least four modeling, estimation, and uncertainty quantification challenges. First, the bivariate dependence structure between temperature and dew point is spatially variable and complex and cannot generally be summarized by simple transformations to a bivariate normal distribution. Second, changes in the relationship between temperature and dew point are of particular interest to us and are likely to be a function of both local temperature and quantile, in addition to warming global temperatures. Third, the conditional quantiles of dew point should not exceed the associated values of local temperature. Finally, daily temperature and dew point exhibit temporal dependence due to climate processes that drive within-season correlation, which is important to account for in quantifying uncertainty in the estimated relationship between temperature and the quantiles of dew point. We address the first three issues through

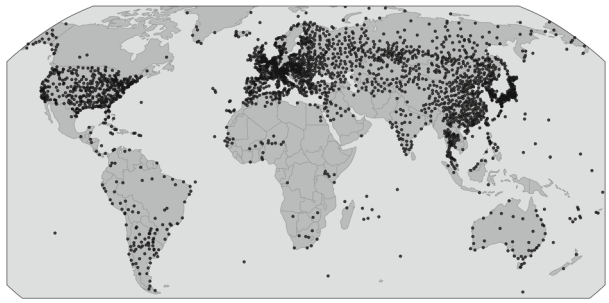

Figure 1. The spatial distribution of global summary of the day stations that have sufficient data to be included in our analysis (see main text).

our use of a regularized, noncrossing, semiparametric quantile regression model, and propose a nonparametric block bootstrap to estimate uncertainty accounting for within-season correlations. Our focus will be on changes in dew point distributions on the hottest days, motivated by their impacts, but the proposed methodology can also be used to study other aspects of the temperature–dew point relationship.

The remainder of the paper is organized as follows: We begin by introducing the data source for our application in Sect. 2. We then introduce our model based on noncrossing penalized quantile smoothing splines in Sect. 3. The choice of regularization parameter and estimation of uncertainty are discussed in Sects. 4 and 5. The model is first demonstrated on synthetic data in Sect. 6 and then applied to our global dataset of weather station data in Sect. 7. We discuss further model improvements and conclude in Sect. 8.

2. DATA

Temperature and dew point data are from the Global Summary of the Day (GSOD) database provided by the National Oceanic and Atmospheric Administration Climatic Data Center, which collates meteorological measurements from weather stations worldwide. Data are typically collected on an hourly or three-hourly basis, and measurements are averaged to create daily values within the GSOD database. We exclude any daily averages that rely on fewer than four values (where four observations would be associated with six-hourly data).

We restrict the analysis to 1973 to near-present (2018) because the dataset is most complete after 1973. As is common with weather data, stations in the GSOD database often have missing values. Data from a given station are only used if it has at least 80% coverage over 80% of years and has data during both the first and last three years of the period 1973–2018. Station locations that pass these criteria are not distributed equally worldwide (Fig. 1): Stations are relatively sparse outside of the USA and midlatitude Eurasia. Globally, stations tend to be denser in regions with higher population density with the major exceptions of equatorial Africa and much of the Indian subcontinent.

The focus of our analysis is on the warm season. We define the warm season on a station-by-station basis as the 60 consecutive days that are, across years, on average hottest at that

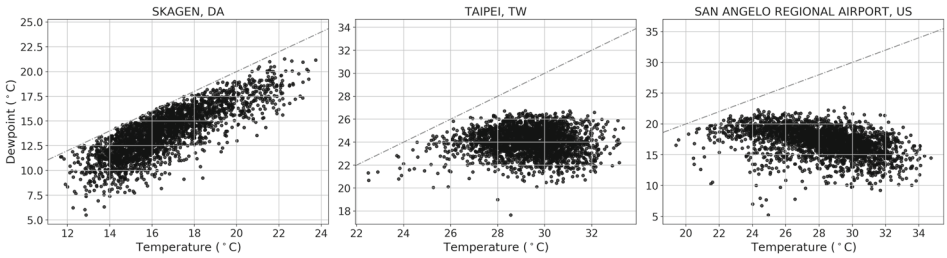

Figure 2. Three examples for warm season temperature–dew point relationships. In Skagen, Denmark, the climate is humid and maritime, so dew point tends to increase with increasing temperature. In Taipei, Taiwan, moisture becomes limited relative to the very high summer temperatures beyond approximately 26 °C, leading to a near-zero correlation of temperature and dew point during the warm season. Finally, at San Angelo Regional Airport in Texas, the summer climate is influenced by the North American monsoon, leading to a negative relationship. The one-to-one line is shown as a gray dot-dashed line and indicates the physical constraint that dew point cannot exceed temperature.

station, thereby allowing us to perform our analysis in a consistent manner globally Figure S1.

As discussed in the prior section, the relationship between warm season temperature and dew point is complex and exhibits substantial spatial variability. Three representative examples are shown in Fig. 2. In the first, from a weather station in Skagen, a coastal city in Denmark, temperature and dew point are positively correlated throughout the warm season, and it is not uncommon for the dew point to approach the dry-bulb temperature value; this is, however, less common on very warm days, indicating an important nonlinearity in the relationship for high dew point quantiles that is less evident at lower quantiles. The second example is from a weather station in Taipei, Taiwan, where there is almost no systematic relationship between temperature and dew point; when temperatures exceed about 26 °C, moisture is limited and dew point does not tend to increase. Finally, the third example is from central Texas, which is characterized by a monsoonal climate in which the atmospheric circulation can bring cool, moist air from the Gulf of Mexico, or warm, dry air from the interior continent, leading to a negative relationship between temperature and dew point.

3. METHODOLOGY

In order to flexibly model the dependence of dew point on temperature, as well as distributional changes in dew point, we develop a semiparametric model based on quantile smoothing splines and estimate the model parameters enforcing the constraint that quantile functions cannot cross. After first providing a brief overview of quantile regression, we describe the model, its interpretation and our estimation procedure.

Quantile regression is a nonparametric modeling tool originally introduced by Koenker and Bassett (1978). Rather than modeling the mean of a distribution as in ordinary linear regression, quantile regression provides a model for the τth quantile conditional on a set of covariates. In its basic formulation, the conditional quantiles q_τ are estimated through minimizing

$$\sum_{i=1}^{n} \rho_\tau(y_i - q_\tau(\mathbf{x}_i)) \tag{1}$$

where ρ_τ is the 'check function', $\rho_\tau(u) = u(\tau - I(u < 0))$, the y_i's are the data being modeled as a function of the covariates, $\mathbf{x}_i$, and $q_\tau(\cdot)$ is a generic function (in many cases chosen to be linear).

Given the dependence of dew point on temperature, as well as the expectation that the conditional distribution of dew point given temperature may be changing as the planet warms, we propose to model a given quantile of dew point, $q_\tau(t)$ on day t, as,

$$q_\tau(t) = \beta_{0,\tau} + \beta_{1,\tau} G'(t) + s_{0,\tau}(T(t)) + G'(t) s_{1,\tau}(T(t)) \tag{2}$$

where $T(t)$ is the local temperature measured at the same station as the dew point and $G'(t)$ is the smoothed global mean temperature anomaly (GMTA) based on the Berkeley Earth Surface Temperature dataset (Rohde et al. 2013). Anomalies in GMTA are calculated by removing the mean value across the available data, but neither $G'(t)$ nor $T(t)$ is detrended. To calculate the smoothed GMTA, we apply a low-pass filter to the raw GMTA data using a third-order Butterworth filter with a cutoff frequency of 1/10 years, although the results are largely insensitive to the choice of similar filters (Figure S3). We use $G'(t)$ rather than time itself as a proxy for the magnitude of anthropogenic influence on the climate to allow for nonlinear time trends (e.g., Smith et al. 2015), although we note that for the time period considered here, $G'(t)$ is dominated by a positive, largely linear trend (Figure S3). For a given local temperature, the dependence of the dew point quantiles on GMTA is specified as linear; this approximation is expected to hold over the period of study due to the relatively small trends compared to internal variability. Due to this linear dependence, as well as the smooth nature of $G'(t)$, the conditional quantiles of dew point evolve smoothly over time. Although both local temperature and the smoothed GMTA are generally increasing over time, the two are only weakly co-linear due to the large amount of variability at a local scale that is unrelated to global temperature increases; the correlation between the two time series has a median value of 0.14 and a 95% range of −0.09 to 0.41 across stations. The minimal co-linearity may lead to small increases in our standard errors, but will not bias our results. The functions $s_{k,\tau}, k \in \{0, 1\}$ are estimated via regularized noncrossing quantile smoothing splines, discussed below.

Each term in Eq. (2) summarizes different contributions to changes in dew point distributions. The term $s_{0,\tau}(T(t))$ captures the dependence of the dew point distribution on local temperature when the GMTA is zero. The term $G'(t)s_{1,\tau}(T(t))$ then characterizes the difference between the dew point-temperature relationship at the given GMTA value versus when GMTA is zero, therefore capturing time-varying changes in the dependence of dew point on temperature. The combined term $s_{1,\tau}(T(t)) + \beta_{1,\tau}$ is the change in q_τ for a 1 °C increase in GMTA, when the local temperature is equal to $T(t)$. Our analysis of the fitted model parameters in Sect. 7 will focus on the relative contributions of these terms.

Since the dependence of dew point on temperature varies widely as a function of local climatic conditions, it is difficult to prespecify a parametric form for the functions $s_{k,\tau}$. Instead, the $s_{k,\tau}$ are modeled as regularized quantile smoothing splines (Koenker et al. 1994).

The formulation of quantile smoothing splines in Koenker et al. (1994) adds a penalty on the total variation of the first derivative of $q_\tau(\cdot)$ in the loss function of Eq. (1); the minimizer of this penalized loss function is a linear spline with knots at each data value, which are the $T(t)$ values in our application. The penalty term provides a trade-off between fit to the data and overall smoothness of the spline; we penalize based on the total variation of the derivative of each $s_{k,\tau}$ independently. Thus, we minimize for each quantile,

$$\sum_{i=1}^{n} \rho_\tau\Big(y_i - q_\tau(t_i)\Big) + \lambda_{0,\tau} \sum_{j=1}^{n-1} |\mathbf{d}_j' \boldsymbol{\alpha}_{0,\tau}| + \lambda_{1,\tau} \sum_{j=1}^{n-1} |\mathbf{d}_j' \boldsymbol{\alpha}_{1,\tau}| \tag{3}$$

where the $\mathbf{d}_j'$'s are the second derivative operator, $\mathbf{d}_j' = (0, \ldots, 0, h_j^{-1}, -(h_{j+1}^{-1} + h_j^{-1}), h_{j+1}^{-1}, 0, \ldots, 0)$, and $h_j = T_{j+1} - T_j$ is the first difference of the ordered temperature values. All vectors are denoted in boldface. The $\boldsymbol{\alpha}_{k,\tau}$ are the values of the fitted splines $s_{k,\tau}$ at each knot (recalling that there is a knot at each value of temperature), which also correspond to the weights on each knot (see "Appendix"). The fully fit model for quantile τ is thus described by a parameter vector $\boldsymbol{\theta}_\tau = [\beta_{0,\tau}, \beta_{1,\tau}, \boldsymbol{\alpha}_{0,\tau}, \boldsymbol{\alpha}_{1,\tau}]$ of length $m = 2 + 2n$. The constants $\lambda_{k,\tau}$ are tuning parameters that control the smoothness of the resulting fitted splines. A small value of λ will encourage a noisier function with smaller residuals, whereas a larger value of λ will encourage a smoother function with larger residuals. In practice, when we apply our model to the representative example locations shown in Fig. 2, we find that the location of the active knots and the slope of the spline fits are effectively unchanged when using a single value of the regularization parameter rather than allowing for separate selection of the regularization parameter for each spline. As such, we reduce the parameter space for the model fit to a single $\lambda_\tau \equiv \lambda_{0,\tau} = \lambda_{1,\tau}$ for each quantile.

Quantile regression models with a small number of data points and/or multiple covariates tend to suffer from the 'embarrassing' problem of quantile crossing (He 1997), in which, for some values of the covariates, a lower quantile crosses a higher one, leading to the implication that the, e.g., 90th percentile of a distribution is greater than the 95th. To avoid this issue, it is necessary to place additional constraints on the model. We take a stepwise approach (Liu and Wu 2009) that begins with fitting the median quantile and then sequentially estimates the outer quantiles with the constraint that any $l + 1$ quantile cannot be smaller (and that any $l - 1$ quantile cannot be larger) than the lth quantile. Recalling that our application has the physical constraint that dew point cannot exceed temperature, we additionally add a noncrossing constraint for all quantiles that the estimated conditional quantile for dew point must be less than or equal to local temperature on the same day. The model is fit for quantiles spanning the 5th to 95th percentile in steps of five percent. The focus of our study is on the behavior of the outermost quantiles, although they are linked to the near-median quantiles via the noncrossing constraint. While the stepwise approach does lead to a dependence of the quantile fits on the choice and ordering of quantiles, we employ it rather than a synchronous estimate of all quantiles (e.g., Bondell et al. 2010) because, for the size of our dataset, the total number of parameters becomes computationally untenable (over 500,000 for 19 quantiles).

The fully constrained model can be solved by standard linear programming methods and is shown in matrix form in "Appendix." For a given value of τ and λ_τ, the estimated model coefficients $\boldsymbol{\theta}_\tau$ are solved for using an interior point method. We use the open-source solver ECOS (Domahidi et al. 2013) as implemented in the Python package CVXPY (Diamond and Boyd 2016). With approximately 3000 data points per station, the solver typically converges in 2–6 seconds on a single vCPU with an Intel Skylake processor and 3.75G memory, with the slower convergence rates occurring at the more extremal quantiles and for larger values of the regularization parameter λ_τ.

4. CHOICE OF REGULARIZATION PARAMETER

So far, we have not answered the question of how to choose the regularization parameters, λ_τ. Lacking prior knowledge about how regularization needs may vary across space and quantile, we take the empirical approach of selecting a best-fit λ_τ value separately for each station and quantile through minimizing the high-dimensional Bayesian information criterion (BIC, Lee et al. 2014). The use of a high-dimensional BIC, rather than a standard version, is necessary for our case because the number of potential parameters diverges with sample size. As shown by Lee et al. (2014) and also found in our own preliminary model-fitting process, this results in the standard BIC being too liberal for quantile regression, leading to an overly noisy fit of the data. The high-dimensional BIC, BIC^{H}, for each quantile-station pair is calculated as

$$\text{BIC}^{\text{H}}(\mathcal{S}) = \log\left(\frac{1}{n}\sum_{i=1}^{n}\rho_\tau\left(y_i - \hat{q}_\tau(t_i)\right)\right) + |\mathcal{S}|\frac{\log n}{2n}C_n. \tag{4}$$

The set $\mathcal{S}$ contains the d indices of the relevant predictors for a candidate model, such that $|\mathcal{S}| = d$. In our case, this is the total number of 'active knots', or knots where the slope of the fitted spline changes including both end points, for both splines, plus an additional two for the intercept and linear term. The addition of the multiplier C_n is the relevant modification from ordinary to high-dimensional BIC. There is not an established choice for the form of C_n, beyond that it is constrained so that, as $n \to \infty$, (1) $C_n \to \infty$ and (2) $C_n\frac{\log n}{n} \to 0$. Following the suggestion of Lee et al. (2014), we find that setting $C_n = \log m_n$, where m_n is the total number of model parameters (including nonactive knots) given a sample of size n, works well for our application based on simulations with synthetic datasets (Sect. 6). The 'likelihood' (first) term in $\text{BIC}^{\text{H}}(\mathcal{S})$ is calculated as if the data values were independent; despite this, we find Eq. (4) to be a reasonable mechanism for choosing λ, based on visual inspection and synthetic simulations. We return to uncertainty quantification under temporal dependence in Sect. 5.

In order to identify the optimal value of λ_τ for each spline fit, we first discretize λ_τ into 10 values in log space between 0.1 and 10. After identifying the two values of λ_τ that span the minimum in BIC^{H}, we further discretize between them in order to gain a more precise estimate. Performing this procedure, which tests 20 λ_τ values across the 19 quantiles that span the 5th to 95th percentiles takes 16–17 min on a single vCPU.

5. UNCERTAINTY ESTIMATION

Interpretation of the magnitude and even sign of trends in climate data is complicated by the large influence of internal variability. Fitting our model to the data as measured provides information about how the observed set of temperatures and dew points have changed over time, and with respect to increasing GMT. However, a different realization of internal climate variability given the same historical forcings would have resulted in a different set of observed data (e.g., Rind 1999; Deser et al. 2012; McKinnon and Deser 2018). To quantify the range of model fits possible if we had sampled a different sequence of internal variability, we propose to employ a nonparametric block bootstrap that retains the within-season correlation in the data through using a block size of one year.

Our approach uses case resampling, wherein each variable is resampled for a given yearly block. The inter-variable and within-season temporal correlation structure is retained since all three variables—GMTA, temperature, and dew point—are resampled as a block. In sampling designs where some of the predictor variables can be thought of as fixed (as with the GMTA trend in our setting), the case bootstrap will tend to produce conservative uncertainty estimates (He 2017); however, a residual bootstrap is not applicable here because the outer quantile functions are estimated in a stepwise fashion. It is possible that other resampling-based methods may be adapted to assess uncertainty, such as the generalized bootstrap and the Markov chain marginal bootstrap, as summarized in He (2017). However, we have chosen case resampling due to its simplicity combined with its success in simulations (Figure S6).

Our block bootstrapping approach proceeds as follows:

1. For a given station, fit Eq. (2) using the original data with the methodology described in Sect. 3 and the value of λ_τ identified in Sect. 4 for each quantile, τ.
2. Resample full years of the original temperature, dew point, and GMTA data with replacement.
3. Fit Eq. (2) to the resampled data using the same values of λ_τ as in Step 1.
4. Repeat steps 2 and 3 $N_{boot} = 1000$ times to estimate the uncertainty on all parameters.

The efficacy of this approach is demonstrated using synthetic simulations in the following section, by comparing our bootstrap intervals to those estimated from repeatedly generating synthetic data from the same generative process.

6. APPLICATION TO SYNTHETIC DATA

Before applying our proposed methodology to the GSOD station data, we test it on four synthetic examples in order to illustrate the properties of the proposed estimation and uncertainty quantification procedure. The parameters for the synthetic data are consistent with a typical midlatitude weather station, and the differences across cases are designed to represent different types of changes in dew point and its relationship to temperature under warming global temperatures.

6.1. Case 1

Perhaps the simplest model for changes in dew point distributions is one that depends entirely on local temperature. In this case, the dew point-temperature relationship remains the same as GMT increases (i.e., $\beta_{1,\tau}$ and $s_{1,\tau}$ are zero), but the distribution of dew point does change because the upper tail of local temperature is increasingly being sampled. In this case, given knowledge of the dew point dependence on temperature and the trends in temperature, it is possible to predict trends in the quantiles of dew point.

The synthetic data for this case are created through the following process:

1. Create a linearly increasing time series of daily GMTA, ranging from -0.5 to $0.5\,^{\circ}\mathrm{C}$ over 50 years.
2. Create an autocorrelated time series of daily temperature as $T(t) = \mu(t) + \epsilon(t)$ where (a) $\mu(t) = 15\,^{\circ}\mathrm{C} + 2G'(t)$ (the scaling with $G'(t)$ represents the fact that land typically warms faster than the global mean temperature), and (b) $\epsilon(t)$ is a Gaussian AR(1) time series with autocorrelation coefficient $\rho = 0.7$ and overall standard deviation $2\,^{\circ}\mathrm{C}$ (i.e., innovation standard deviation equal to $2\sqrt{1-\rho^2} \approx 1.52$).
3. Calculate dew point temperature as $T(t) - \Delta(t)$, where $\Delta(t)$ is the dew point depression (temperature minus dew point) and $\Delta(t) \sim \mathrm{Gamma}(4, 0.04T(t))$, with the second argument being the scale parameter in our parameterization.
4. Subselect the first 60 days from each year to represent the 'warm season.'

The autocorrelation coefficient for this and all cases is chosen to reflect the median value across weather stations in our dataset, but we also perform our synthetic model fits using an autocorrelation coefficient of $\rho = 0.95$, which is the maximum found across all stations. The synthetic model fits and uncertainties are largely unchanged (Figure S4).

In Fig. 3a, we compare the true conditional quantiles (gray lines) to the median (black lines) and spread (shading, pointwise 95% range) of the estimated conditional 5th, 50th, and 95th quantiles based on fits to 1000 realizations of the underlying generative process. The quantiles conditional on the GMTA being at the 25th (75th) percentile are shown as dashed (solid) lines; in this case, with no dependence on GMTA, they overlap entirely. The spread of the fitted splines increases at high and low temperatures because the data are more sparse, as well as for lower quantiles, which are further removed from the constraint that dew point cannot exceed temperature. Importantly, the fitted quantiles are essentially unbiased and do not show a dependence on GMTA, as expected.

6.2. Case 2

Case 2 is created with the same philosophy as Case 1, but allowing for a nonlinear dependence of dew point on temperature. The synthetic data are created with a process analogous to Case 1, except that $\Delta(t) \sim \mathrm{Gamma}(8, 0.002T(t)^2)$, providing curvature in the quantiles.

The fitted model is successful at capturing the nonlinear behavior through addition of more active knots. The fits are again largely unbiased and do not show a dependence on

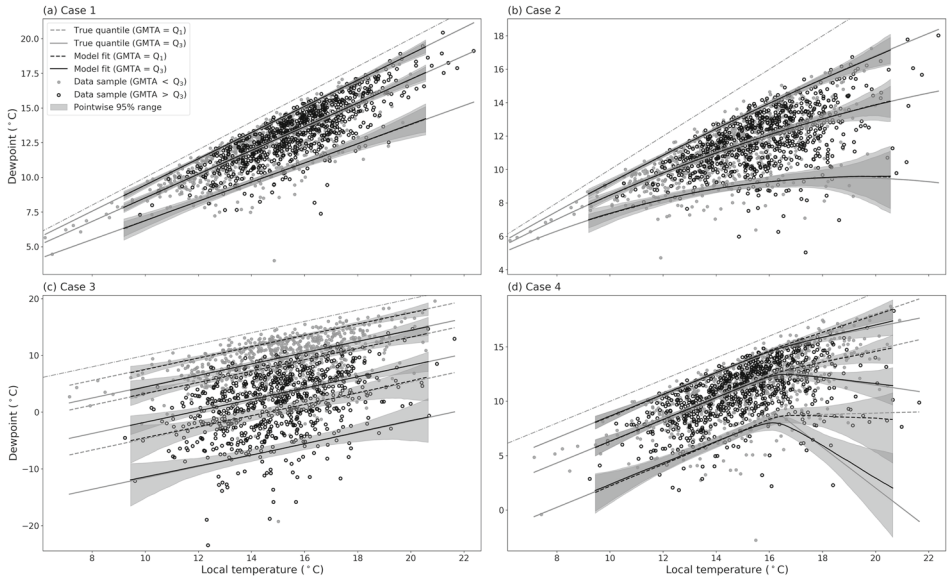

Figure 3. Four examples demonstrating model fit for synthetic data. See text for description of cases. In all panels, one realization of temperature and dew point data from the generative process is shown; the open dots indicate data when the GMTA is in the lower quartile, whereas the gray dots indicate that GMTA is in the upper quartile. The gray shading shows the 95% pointwise spread across the 500 fitted conditional quantiles. The median conditional on GMTA being at the 25th (75th) percentile is shown as a black dashed (solid) line. The true conditional quantiles for the 25th (75th) percentile of GMTA are shown in gray dashed (solid) lines. The one-to-one line is shown as a gray dot-dashed line and indicates the physical constraint that dew point cannot exceed temperature .

GMTA, although we note some evidence that the curvature in the 5th quantile may be slightly underestimated by the model. Variance in the conditional quantile fits again increases with decreased density of data at high and low temperatures, and low quantiles.

6.3. CASE 3

Another end member case would occur when increases in GMTA lead to moisture limitations at a given location, but there are no changes in local temperature. In other words, the temperature distribution is unchanged as GMTA increases, but the dew point decreases.

The synthetic data are created through the following process:

1. Create a linearly increasing time series of daily GMTA, ranging from $-0.5\,^{\circ}\mathrm{C}$ to $0.5\,^{\circ}\mathrm{C}$ over 50 years.

2. Create an autocorrelated time series of daily temperature as $T(t) = \mu_0 + \epsilon(t)$ where $\mu_0 = 15\,^{\circ}\mathrm{C}$, and $\epsilon(t)$ is a Gaussian AR(1) time series with autocorrelation coefficient $\rho = 0.7$ and overall standard deviation $2\,^{\circ}\mathrm{C}$ (as in Case 1).

3. Calculate dew point temperature as $T(t) - \Delta(t)$, where $\Delta(t) \sim \mathrm{Gamma}(5(G'(t) + 1), 2)$.

4. Subselect the first 60 days from each year to represent the 'warm season.'

In this case, the slope of the true conditional quantiles does not change with GMTA, but the difference between temperature and dew point for a given temperature increases with GMTA. Now that we have introduced a dependence on GMTA for the true quantiles, the fits conditional on the 25th (dashed) and 75th (solid) percentile of GMTA no longer overlap and effectively capture the dependence on GMTA.

6.4. Case 4

Finally, we have the case where moisture limitations increase with increasing GMTA, as in Case 3, but only at high temperatures due to, e.g., desiccation of the land surface.

The synthetic data are created in an analogous way to Case 3, but with $\Delta(t) \sim \text{Gamma}(4, 0.5(G'(t) + 0.5)T_{75}(t) + 1)$, where $T_{75}(t)$ is the difference between $T(t)$ and its 75th percentile if $T(t)$ is in the upper quartile, and otherwise zero.

Due to the large changes in the temperature–dew point relationship in only a subset of the dataset, we see the greatest differences between the fitted and true conditional quantiles in this example. Nevertheless, the biases are small compared to the trends and the 95% range. The variance in the estimator is also the largest among the case studies, but shifts in the temperature–dew point relationship would still typically be detectable given the magnitude of the changes in this simulation.

6.5. Parameter Uncertainty

In our subsequent observational analysis, we do not have access to an underlying generative process. Thus, as discussed in Sect. 5, we propose to use a block bootstrap to estimate uncertainty in our parameters. Here, we test this approach by comparing uncertainty intervals estimated using the bootstrap procedure to those calculated above from generating new datasets. In particular, for each synthetic case study, we generate a single realization from the associated generative process, perform our block bootstrap procedure 1000 times, and then compare the simulation-based intervals to the bootstrap-based intervals.

Estimation of conditional quantiles can fail when data are bootstrapped because the data distribution is no longer smooth, leading to objective functions that are not differentiable (Machado and Silva 2005; Rhines et al. 2015). To address this issue, we add a small amount of 'jitter' that is uniformly distributed around zero with a half-width of 0.005 °C to the bootstrapped datasets. A similar issue arises in our analysis of the observations, but because all data are rounded to 0.1 °F. In that case, we increase the half width to 0.05 °F, or 0.03 °C, to account for the uncertainty in the true value.

As discussed above, the case bootstrap is likely to lead to conservative uncertainty estimates because one of our covariates, GMTA, can be viewed as fixed. In contrast, since we are using a block size of one season in order to retain within-season autocorrelation, we only have 50 synthetic 'years' from which to sample from; if this number of years is insufficient to sample the true variability in the process, the bootstrap would underestimate the uncertainty in the parameters.

In general, we find that the conservative nature of the case bootstrap dominates, and the bootstrapped confidence intervals are larger than the intervals estimated from sampling

from the underlying generative process. We find a median (mean) overestimation of the width of the 95% pointwise intervals by 11% (13%), although there is substantial variability across case, dew point quantile, and temperature value (Figure S6), with lower quantiles more likely to show under- rather than overestimation of the 95% interval.

7. APPLICATION TO WEATHER STATION DATA

We now turn our attention to the daily warm season temperature and dew point data that motivated our methodology. We first analyze the three stations of Fig. 2 in order to provide intuition for the model parameters. Due to the high-dimensional nature of the model, we then focus on a small set of summary statistics to examine across all 2416 stations worldwide.

7.1. INDIVIDUAL STATION ANALYSIS

During the warm season in Skagen, Denmark (Fig. 4a), the model captures the apparent increase in dew point with temperature as well as an inflection point around 17 °C where dew point increases more slowly with temperature. In the upper half of the temperature distribution, the conditional quantiles exhibit minimal dependence on the GMTA, as indicated by comparing the black dashed and solid lines, which show the fitted quantiles conditional on GMTA being at the 25th (−0.22 °C) and 75th (0.21 °C) percentile, respectively. Thus, dew point trends are largely connected to trends in local temperature. In the remaining analysis, we primarily focus on dew point trends conditional on high temperatures, as measured by the 95th percentile, since impacts of humidity tend to be greatest at high temperatures in the warm season.

We now quantitatively assess the contribution to changes in dew point of the terms linearly dependent on GMTA, $\beta_{1,\tau} + s_{1,\tau}(T(t))$, versus the term capturing the dew point dependence on local temperature for a GMTA of zero, $s_{0,\tau}(T(t))$. The change in dew point due to the GMTA terms is calculated by subtracting the fit conditioned on a GMTA being at the 25th percentile from it being at the 75th percentile, with local temperature fixed at the 95th percentile over the full period. The change in dew point due to the local temperature term is calculated as the difference between the value of $s_{0,\tau}$ at the 95th percentile of local temperature in the first versus second half of the record. The uncertainty in the two terms is shown as the analog of one standard error (68%) range across the bootstrap samples, although we calculate it as the difference between the 84th and 16th percentiles of the bootstrap distribution because of the asymmetry present in many of the intervals. We focus on the 68% range for the observational analysis because uncertainty due to internal variability remains large compared to the signal of trends in dew point since 1973. As expected from our visual inspection of the spline fits, the changes in dew point conditional on the 95th percentile of temperature are largely explained by trends in local temperature, with the contributions linearly related to GMTA summing to near-zero (Fig. 4b). Furthermore, changes in dew point on hot days are relatively consistent across quantiles, with the dew point during hot, dry days increasing at the same rate as that during hot, humid days. While not the focus of

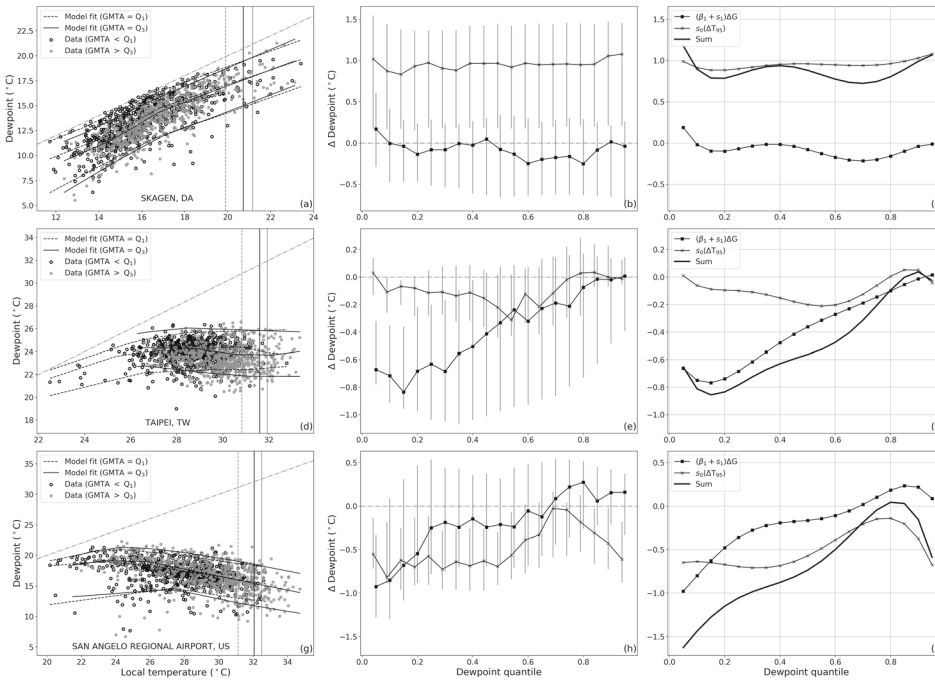

Figure 4. Model fits to the same warm season data shown in Fig. 2. The left column is analogous to the scatter plots for the synthetic data in Fig. 3: The open dots indicate data when the GMTA is in the lower quartile, whereas the gray dots indicate that GMTA is in the upper quartile. The spline fit conditional on GMTA being at the 25th (75th) percentile is shown as a black dashed (solid) line; splines fits are only shown spanning the range of temperature for the lower (upper) quartile of GMTA. The one-to-one line is shown as a gray dot-dashed line and indicates the physical constraint that dew point cannot exceed temperature. The black vertical line is the 95th percentile of temperature over the full record, whereas the dashed (solid) gray lines are the 95th percentiles from the lower (upper) 50% of GMTA. The middle column shows the empirical contribution of the two types of model terms to changes in dew point between the first and second half of the record, conditional on temperature being at the 95th percentile; vertical bars span the 16th-84th percentile range estimated through bootstrapping. The right column shows a smoothed version, as well as the total change, calculated through projection onto the first six Legendre polynomials. The legend in the right column also applies to the middle column, where it is omitted for clarity.

this analysis, the spline fits do indicate changes in the temperature–dew point relationship for low quantiles of temperature, with dew points decreasing for a given value of temperature.

In contrast, both Taipei, Taiwan (Fig. 4d) and San Angelo, Texas (Fig. 4g) exhibit more complex behavior. In Taipei, because dew point does not exhibit a dependence on temperature at temperatures over approximately 26 °C, changes in local temperature alone cannot explain changes in dew point (if such changes exist). At low quantiles of dew point, there is evidence of drying with increasing GMTA, but the changes at high quantiles of dew point are minimal. Finally, the underlying relationship between temperature and dew point is negative in San Angelo, so increasing local temperature generally leads to decreases in dew point, while the sensitivity to GMT switches from negative for low dew point quantiles to slightly positive at high quantiles.

Because there is no cross-quantile smoothness constraint enforced during the fitting process, yet we expect relatively smooth behavior across quantiles, we perform a final

post-processing step to smooth the quantile fits through projecting them onto the first six Legendre polynomials (Fig. 4c, f, i). The Legendre polynomials are a natural basis for smooth one-dimensional functions (Gibson et al. 1992; McKinnon et al. 2016).

7.2. Large-Scale Spatial Patterns: Hot, Dry Days

Having analyzed three stations, we would like to zoom out and understand the large-scale spatial patterns associated with trends in dew point, as well as the relative contribution from each model term. To do so, we summarize the fits at each station with two sets of values: the change in the 5th and 95th percentiles of dew point conditional on the 95th percentile in temperature. In other words, we inspect the changes in dew point on the hottest/driest days and hottest/most humid days. The contributions to the changes in dew point are calculated in the manner described in Sect. 7.1, and the values presented are from the smoothed quantile fits (i.e., the leftmost and rightmost points for each line in Fig. 4c, f, i). The spread in parameter values that could occur due to different sampling of internal variability is calculated separately for each station using case resampling; however, unlike the results presented for the three example stations, we only calculate 200 bootstrap samples for each station due to the computational demands of our model and the large number of stations. Through resampling our 1000 bootstrap samples calculated for our three station examples in Fig. 4, we estimate that this results in a median underestimation of width of our confidence intervals of 2% (varying by case and percentile level, with 95% range of 14% under-estimation to 10% overestimation). We assess significance against the null hypothesis of no trend in dew point for each component of the model: the trends dependent on changes in GMTA, versus those dependent on local temperature. To do so, we estimate a two-sided p value as $2\min[\Pr(\hat{\xi}^* < 0), \Pr(\hat{\xi}^* > 0)]$, where $\hat{\xi}^*$ is the bootstrap estimate of the parameter of interest. This corresponds to inverting the bootstrap percentile interval, i.e., if the $1 - \alpha$ interval has zero as an endpoint, then the p value is α. In order to account for multiple hypothesis testing across our stations, quantiles, and model terms, we then limit the false discovery rate (Benjamini and Hochberg 1995) to 0.1 for our assessment of significance, which is associated with a local significance cutoff at approximately the 0.01 level.

Since 1973, hot, dry days have generally gotten more humid in the tropics and high latitudes, and slightly drier in the already dry subtropics (Fig. 5e). This pattern is evident in both the component of the trends that are linearly dependent on changes in GMTA (Fig. 5c), and those that are a function of changing local temperature (Fig. 5a).

The majority of the spatial structure in dew point trends emerges due to changes in the dependence of dew point on temperature with increasing GMTA (like synthetic cases 3 and 4). There are some regions for which trends in dew point are driven by trends in local temperature (like synthetic cases 1 and 2), without a change in their relationship, particularly around Finland, as well as cases where the two types of changes have effects of opposite signs, particularly on the coasts around the Mediterranean Sea.

There is a weak, negative relationship on a station-by-station basis between warm season average dew point depression (temperature minus dew point, shown in Figure S2), and trends in the dew point on hot, dry days (shown in Fig. 5e, Pearson's correlation coefficient of -0.2), indicating that moisture trends on hot, dry days are, to a limited extent, enforcing

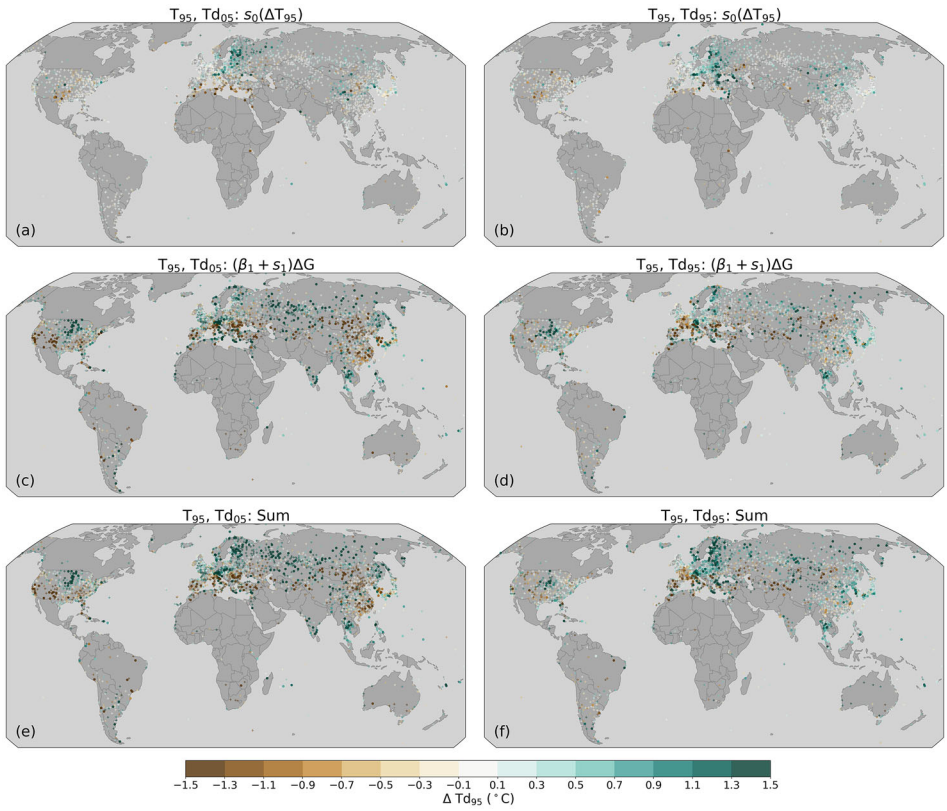

Figure 5. Changes in dew point on hot, dry (left) and hot, humid (right) days. **a**, **b** The change in dew point that can be explained by increasing local temperatures with no change in GMTA. **c**, **d** The change in dew point due to a change in the temperature–dew point relationship as a function of GMTA. **e**, **f** The sum of the two contributions, illustrating total change in dew point. Changes in dew point for each component identified as significant through limiting the false discovery rate to 0.1 are shown as circles, otherwise they are shown as plus signs.

climatological patterns. This 'drier get drier' effect is clear in regions like the Interior Western United States, the Iberian Peninsula, and interior Australia. However, there are notable deviations such as the substantial drying in eastern China and the southeastern United States, regions that are climatologically humid.

7.3. Large-Scale Spatial Patterns: Hot, Humid Days

The latitudinal pattern of trends in hot, humid days is similar to that for hot, dry days, with the tropics and high latitudes showing the greatest increases, although the trends in the subtropics are more likely to be slightly positive than negative (Fig. 5f). Analogous to trends in the 5th percentile, most of the spatial structure emerges due to changes in the temperature–dew point relationship (Fig. 5d) rather than sampling warmer temperatures with the same underlying relationship (Fig. 5b). The magnitude of the overall trends tend to be muted compared to those for the 5th percentile in dew point, in part because the upper quantiles of dew point are additionally constrained by not exceeding local temperature.

Some of the regions of greatest moistening are Finland and the Baltic states, where we do see a dominant contribution from changes in local temperature, and the northern Great Plains in the USA, where the trends are largely related to increases in GMTA. It is difficult to discern large-scale patterns in regions with sparse measurements, but the meridional pattern of moistening in the tropics and higher latitudes does appear evident in eastern Siberia, western Africa, western Australia, and Patagonia, with contributions primarily from changes in GMTA.

Drying of the most hot, humid days may help ameliorate some of the human health impacts of warming, especially in places that are climatologically humid. There is some evidence of this in western Europe, where recent heat waves have typically been dry; however, because drier soils typically have allowed for hotter (if drier) heat waves (Miralles et al. 2014; Whan et al. 2015), the extent to which drying reduces negative impacts is not immediately clear.

8. SUMMARY AND DISCUSSION

Some of the greatest impacts of climate change may be felt through the interaction of two or more variables, such as temperature and humidity. It is thus imperative to develop flexible statistical methods to better describe the observed nature of these relationships, with particular interest in the probability of joint, or compound, extremes.

Here, we have proposed to model the relationship between dew point, local temperature, and GMTA using nonparametric quantile smoothing splines. Through noncrossing constraints, we also encode the physical property that dew point cannot exceed local temperature. The method is tested using synthetic examples that isolate different types of changes in dew point–temperature relationships and appears to perform well except perhaps for an extreme case of large changes at high percentiles of temperature (Case Study 4). We also test and employ a nonparametric block bootstrap to quantify the uncertainty in our parameter estimates related to random sampling.

The application of our model to a near-global dataset of daily temperature and dew point over land provides important insights into changes in humidity on the hottest days. With the exception of small regions of western Eurasia, trends in low and high dew points on hot days appear to largely reflect changes in the temperature–dew point relationship, rather than simply following changes in local temperature. This further underscores the need for additional work developing models to describe changes in their relationship, since we may not be able to make proper predictions for future dew points using trends in local temperature alone.

There are a number of avenues for promising future work. First, our current approach does not take advantage of the spatial correlation structure that is inherent in climate data. This has at least two important implications. One is that there is a likely unrealistic lack of spatial smoothness in parameter estimates in some geographic locations (Fig. 5), and we are unable to provide parameter estimates for unobserved locations. Prior work using quantile regression has addressed this challenge through, e.g., fitting a spatial model to the parameter estimates estimated at each data location (Reich et al. 2011). Spatial information

has also been incorporated to ensure spatial smoothness in the underlying quantile function through the use of spatial basis functions (Franco-Villoria et al. 2019) or the application of a spatially correlated errors model (Lum and Gelfand 2012). Each of these approaches involves borrowing strength from nearby spatial locations and allows for inferences about locations where data are unavailable. However, it is not straightforward to apply these methods to our application for a number of reasons. First, the problem is high-dimensional with over 2000 station locations and thousands of parameters estimated at each station (although we note the latter could be reduced through summarizing the location of the active knots and slope of the intervening linear spline), leading to computational challenges. Second, the spatial covariance structure is likely to be nonstationary and anisotropic due to factors including continental geometry, the prevailing wind directions, and topography, so would itself require a complex model. Finally, we note that the extreme sparseness of data in some regions would limit the usefulness of interpolated parameter estimates, which will be highly variable. The second important implication is that, because the underlying data are themselves spatially correlated, the uncertainty in our parameter estimates should also exhibit spatial correlation. This is also a challenging problem to fully account for in the multivariate case with complex spatial covariance, but might be partially addressed without a full spatial model through performing our bootstrap resampling such that the full spatial domain is resampled in the same way in each iteration. These implications are best explored on a regional scale with a smaller dataset and will be investigated in future work.

As a second avenue for future work, it may be possible to enforce additional structure on the regularization parameter, λ. Recall that the regularization parameter controls the relative importance of the residuals versus the smoothness of the fitted quantiles in the minimization. Stations with greater variability in dew point will tend to have larger residuals from the fitted quantiles, so a larger regularization parameter will be required to produce the same smoothness in the conditional quantiles. Indeed, we find that the fitted values of λ across stations show a positive relationship to the magnitude of both temperature and dew point variability (Figure S5). Rather than selecting the value of λ for each station-quantile pair through minimizing an information criterion, as we do here, one could enforce a dependence on dew point variability. Interestingly, we do not find meaningful structure in the values of λ as a function of quantile, and some exploratory data analysis indicates that using the same λ at all quantiles for a given station has minimal impact on the result.

Finally, while analysis of the observations provides us with key information about current trends, they cannot alone inform expected future changes. For this application, it is advantageous to turn to climate models and their future projections. However, climate models typically exhibit biases in their variability, including the covariability between temperature and humidity (Schoof et al. 2019). To produce better calibrated bivariate simulations of projected temperature and dew point, it may be possible to combine the observed temperature–dew point relationships from this study with climate model-based estimates of changes in their relationship.

Our understanding and modeling of compound climate extremes remains nascent. The model proposed here provides a tool with which to parse changes in the temperature–dew point relationship from trends in dew point that can be directly explained by increasing

temperatures. Similar approaches may be advantageous for other sets of variables with strong underlying covariance, such as temperature and air quality.

9. DATA AND CODE AVAILABILITY

GSOD weather station data are publicly available at https://www1.ncdc.noaa.gov/pub/data/gsod/. Code is publicly available at https://github.com/karenamckinnon/humidity_variability.

ACKNOWLEDGEMENTS

AP received support from the Hewlett Mellon Fellowship from Carleton College.

[Received November 2019. Accepted April 2020.]

A. MATRIX FORM OF NONCROSSING QUANTILE SMOOTHING SPLINES

In order to solve for the parameters in our regularized quantile smoothing splines model with noncrossing constraints, we write it in standard form as follows:

$$\begin{aligned} &\min c^{\mathsf{T}} z \; s.t. \\ &\mathbf{A}z = b \\ &\mathbf{G}z \leq h \end{aligned} \tag{5}$$

and describe each term in this optimization problem below.

Throughout this appendix, bold face implies a two-dimensional matrix, whereas normal script indicates a vector or, in the case of τ and λ, a constant. Note that this stands in contrast to the main text, where vectors are bolded, in order to more clearly indicate the relevant matrix operations.

The parameter vector z and cost vector c are,

$$\begin{aligned} z &= \left[\theta^+ \; \theta^- \; \epsilon^+ \; \epsilon^- \; u^+ \; u^- \; v^+ \; v^-\right] \\ c &= \left[0_m \; 0_m \; \tau 1_n \; -\tau 1_n \; \lambda 1_{n-1} \; \lambda 1_{n-1} \; \lambda 1_{n-1} \; \lambda 1_{n-1}\right], \end{aligned} \tag{6}$$

where the m model parameters are contained in θ and the remaining components of z are the values of the residuals (ϵ's) and the second derivatives of the α_0 (u's) and α_1 values (v's), which are the fitted splines at each knot. The superscripts $+$ and $-$ denote the magnitude of the positive and negative components, respectively, such that, e.g., $\theta = \theta^+ - \theta^-$. The cost

vector contains the minimization of Eq. (3), with no penalty on the values of θ, a quantile-weighted penalty on the residuals, and the λ-weighted penalty on the second derivatives of the splines.

The equality term, $\mathbf{A}z = b$, is written as,

$$\mathbf{A} = \begin{bmatrix} \mathbf{X} & -\mathbf{X} & \mathbf{1}_{n\times n} & -\mathbf{1}_{n\times n} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} \\ \mathbf{D_0} & -\mathbf{D_0} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & -\mathbf{1}_{n\times(n-1)} & \mathbf{1}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} \\ \mathbf{D_1} & -\mathbf{D_1} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & -\mathbf{1}_{n\times(n-1)} & \mathbf{1}_{n\times(n-1)} \end{bmatrix}$$
$$b = \begin{bmatrix} y & 0_{n-1} & 0_{n-1} \end{bmatrix}, \tau \tag{7}$$

where the matrix $\mathbf{X}$ is the design matrix for Eq. (2), the $\mathbf{D}_k$ are the second derivative operators for the spline coefficients, α_k, and the length (dimensions) of the vectors (matrices) of zeros and ones are indicated by the subscripts. Thus, the first row of $\mathbf{A}$ contains the constraint that the data minus the conditional quantile fit equals the residuals. The second and third rows define that the second derivate of the spline coefficients equal u and v, respectively.

Finally, all noncrossing constraints, as well as the standard constraint that all are in the inequality term, $\mathbf{G}z \leq h$. Define $c = 2m + 2n + 4 \times (n-1)$ as the length of z. Then, for the median quantile, $\tau = 0.5$,

$$\mathbf{G} = \begin{bmatrix} & & & & -\mathbf{1}_{c\times c} & & & \\ \mathbf{X} & -\mathbf{X} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} \end{bmatrix}$$
$$h = \begin{bmatrix} 0_c & T \end{bmatrix}, \tag{8}$$

where T is the vector of sorted local temperatures. The first row of $\mathbf{G}$ constrains all parameter values in z to be positive, whereas the second row ensures that the conditional quantile of dew point does not exceed temperature.

For quantile levels above the median, $\tau_i = 0.5 + \delta_i$, where δ_i is positive and increasing, define q as the best fit conditional quantile for the τ_{i-1} quantile. Then the inequality constraint is slightly modified to:

$$\mathbf{G} = \begin{bmatrix} & & & & -\mathbf{1}_{c\times c} & & & \\ \mathbf{X} & -\mathbf{X} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} \\ -\mathbf{X} & \mathbf{X} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} \end{bmatrix}$$
$$h = \begin{bmatrix} 0_c & T & -q \end{bmatrix}, \tag{9}$$

where the last row in $\mathbf{G}$ is the noncrossing constraint for consecutive quantiles.

Finally, for quantile levels below the median, $\tau_i = 0.5 + \delta_i$, where δ_i is negative and decreasing, define q as the best fit conditional quantile for the τ_{i-1} quantile. In this case, we can remove the constraint that the conditional dew point quantile is less than the local temperature because it will already be contained within the noncrossing constraint. Thus,

the inequality constraint is:

$$\mathbf{G} = \begin{bmatrix} & & & & & & -\mathbf{1}_{c\times c} \\ \mathbf{X} & -\mathbf{X} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times n} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} & \mathbf{0}_{n\times(n-1)} \end{bmatrix} \tag{10}$$
$$h = \begin{bmatrix} 0_c & q \end{bmatrix}$$

REFERENCES

AghaKouchak, A., L. Cheng, O. Mazdiyasni, and A. Farahmand (2014), Global warming and changes in risk of concurrent climate extremes: Insights from the 2014 California drought, *Geophysical Research Letters*, *41*(24), 8847–8852.

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.

Bondell, H. D., B. J. Reich, and H. Wang (2010), Noncrossing quantile regression curve estimation, *Biometrika*, *97*(4), 825–838.

Chamaillé-Jammes, S., H. Fritz, and F. Murindagomo (2007), Detecting climate changes of concern in highly variable environments: Quantile regressions reveal that droughts worsen in Hwange National Park, Zimbabwe, *Journal of Arid Environments*, *71*(3), 321–326.

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), 014,001.

Cooley, D., E. Thibaud, F. Castillo, and M. F. Wehner (2019), A nonparametric method for producing isolines of bivariate exceedance probabilities, *Extremes*, pp. 1–18.

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.

De Michele, C., and G. Salvadori (2003), A generalized Pareto intensity-duration model of storm rainfall exploiting 2-copulas, *Journal of Geophysical Research: Atmospheres*, *108*(D2).

Deser, C., R. Knutti, S. Solomon, and A. S. Phillips (2012), Communication of the role of natural variability in future North American climate, *Nature Climate Change*, *2*(11), 775–779.

Diamond, S., and S. Boyd (2016), CVXPY: A Python-embedded modeling language for convex optimization, *The Journal of Machine Learning Research*, *17*(1), 2909–2913.

Domahidi, A., E. Chu, and S. Boyd (2013), ECOS: An SOCP solver for embedded systems, in *2013 European Control Conference (ECC)*, pp. 3071–3076, IEEE.

Elsner, J. B., J. P. Kossin, and T. H. Jagger (2008), The increasing intensity of the strongest tropical cyclones, *Nature*, *455*(7209), 92.

Franco-Villoria, M., M. Scott, and T. Hoey (2019), Spatiotemporal modeling of hydrological return levels: A quantile regression approach, *Environmetrics*, *30*(2), e2522.

Gibson, J. F., J. D. Farmer, M. Casdagli, and S. Eubank (1992), An analytic approach to practical state space reconstruction, *Physica D: Nonlinear Phenomena*, *57*(1-2), 1–30.

Guirguis, K., A. Gershunov, A. Tardy, and R. Basu (2014), The impact of recent heat waves on human health in California, *Journal of Applied Meteorology and Climatology*, *53*(1), 3–19.

Hao, Z., V. Singh, and F. Hao (2018), Compound extremes in hydroclimatology: a review, *Water*, *10*(6), 718.

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.

He, X. (1997), Quantile curves without crossing, *The American Statistician*, *51*(2), 186–192.

He, X. (2017), Resampling methods, in *Handbook of quantile regression*, edited by R. Koenker, V. Chernozhukov, X. He, and L. Peng, pp. 7–19, Chapman and Hall/CRC.

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*, 107,692.

Katz, R. W., M. B. Parlange, and P. Naveau (2002), Statistics of extremes in hydrology, *Advances in Water Resources*, *25*(8-12), 1287–1304.

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

Koenker, R., P. Ng, and S. Portnoy (1994), Quantile smoothing splines, *Biometrika*, *81*(4), 673–680.

Lee, E. R., H. Noh, and B. U. Park (2014), Model selection via Bayesian information criterion for quantile regression models, *Journal of the American Statistical Association*, *109*(505), 216–229.

Liu, Y., and Y. Wu (2009), Stepwise multiple quantile regression estimation using non-crossing constraints, *Statistics and its Interface*, *2*(3), 299–310.

Lum, K., and A. E. Gelfand (2012), Spatial quantile multiple regression using the asymmetric Laplace process, *Bayesian Analysis*, *7*(2), 235–258.

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

Mastrangelo, G., U. Fedeli, C. Visentin, G. Milan, E. Fadda, and P. Spolaore (2007), Pattern and determinants of hospitalization during heat waves: an ecologic study, *BMC Public Health*, *7*(1), 200.

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., 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.

Meehl, G. A., and C. Tebaldi (2004), More intense, more frequent, and longer lasting heat waves in the 21st century, *Science*, *305*(5686), 994–997.

Miralles, D. G., A. J. Teuling, C. C. Van Heerwaarden, and J. V.-G. De Arellano (2014), Mega-heatwave temperatures due to combined soil desiccation and atmospheric heat accumulation, *Nature geoscience*, *7*(5), 345.

Nelsen, R. B. (2007), *An introduction to copulas*, Springer Science & Business Media.

Ouarda, T. B., and C. Charron (2019), Changes in the distribution of hydro-climatic extremes in a non-stationary framework, *Scientific Reports*, *9*(1), 1–8.

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.

Perkins, S., L. Alexander, and J. Nairn (2012), Increasing frequency, intensity and duration of observed global heatwaves and warm spells, *Geophysical Research Letters*, *39*(20).

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.

Reich, B. J., M. Fuentes, and D. B. Dunson (2011), Bayesian spatial quantile regression, *Journal of the American Statistical Association*, *106*(493), 6–20.

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.

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.

Ribeiro, A. F., A. Russo, C. M. Gouveia, and P. Páscoa (2019), Copula-based agricultural drought risk of rainfed cropping systems, *Agricultural Water Management*, *223*, 105,689.

Rind, D. (1999), Complexity and climate, *Science*, *284*(5411), 105–107.

Rohde, R., R. Muller, R. Jacobsen, S. Perlmutter, A. Rosenfeld, J. Wurtele, J. Curry, C. Wickham, and S. Mosher (2013), Berkeley Earth temperature averaging process, *Geoinformatics & Geostatistics: An Overview*, *1*(2), 1–13.

Schoof, J., S. Pryor, and T. Ford (2019), Projected changes in United States regional extreme heat days derived from bivariate quantile mapping of CMIP5 simulations, *Journal of Geophysical Research: Atmospheres*.

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.

Smith, S. J., J. Edmonds, C. A. Hartin, A. Mundra, and K. Calvin (2015), Near-term acceleration in the rate of temperature change, *Nature Climate Change*, *5*(4), 333.

Tanner, C. B., and T. R. Sinclair (1983), Efficient water use in crop production: research or re-search?, in *Limitations to efficient water use in crop production*, edited by H. M. Taylor, W. R. Jordan, and T. R. Sinclair, pp. 1–27, American Society of Agronomy, Madison, WI.

Wahl, T., S. Jain, J. Bender, S. D. Meyers, and M. E. Luther (2015), Increasing risk of compound flooding from storm surge and rainfall for major US cities, *Nature Climate Change*, *5*(12), 1093.

Wahlstrom, M., and D. Guha-Sapir (2015), The human cost of weather-related disasters, 1995–2015, *Tech. rep.*, Geneva, Switzerland.

Whan, K., J. Zscheischler, R. Orth, M. Shongwe, M. Rahimi, E. O. Asare, and S. I. Seneviratne (2015), Impact of soil moisture on extreme maximum temperatures in Europe, *Weather and Climate Extremes*, *9*, 57–67.

Zscheischler, J., S. Westra, B. J. Van Den Hurk, S. I. Seneviratne, P. J. Ward, A. Pitman, A. AghaKouchak, D. N. Bresch, M. Leonard, T. Wahl, et al. (2018), Future climate risk from compound events, *Nature Climate Change*, *8*(6), 469.