

A MARKOV-SWITCHING MODEL FOR HEAT WAVES

BY BENJAMIN A. SHABY*, BRIAN J. REICH†, DANIEL COOLEY‡
AND CARI G. KAUFMAN§

*Pennsylvania State University**, *North Carolina State University†*,
Colorado State University‡ and *University of California, Berkeley§*

Heat waves merit careful study because they inflict severe economic and societal damage. We use an intuitive, informal working definition of a heat wave—a persistent event in the tail of the temperature distribution—to motivate an interpretable latent state extreme value model. A latent variable with dependence in time indicates membership in the heat wave state. The strength of the temporal dependence of the latent variable controls the frequency and persistence of heat waves. Within each heat wave, temperatures are modeled using extreme value distributions, with extremal dependence across time accomplished through an extreme value Markov model. One important virtue of interpretability is that model parameters directly translate into quantities of interest for risk management, so that questions like whether heat waves are becoming longer, more severe or more frequent are easily answered by querying an appropriate fitted model. We demonstrate the latent state model on two recent, calamitous, examples: the European heat wave of 2003 and the Russian heat wave of 2010.

1. Introduction. When widespread heat waves occur, they dominate news reports and inspire passionate discussions about climate change and public policy. The European heat wave of 2003 was estimated to have caused up to an estimated 70,000 additional deaths [Robine et al. (2008)] and cost the 2011 equivalent of \$16 billion [Munich Re (2003), Parry et al. (2007)]. The Russian heat wave of 2010 was responsible for an estimated 55,000 excess deaths, a 25% reduction in agriculture and \$15 billion in economic loss [Barriopedro et al. (2011)]. Perhaps because of their high public visibility and disastrous public health and economic consequences, heat waves are the subject of a great deal of scientific research [e.g., Amengual et al. (2014), Clark, Brown and Murphy (2006), Easterling et al. (2000), Fischer and Schär (2010), Frich et al. (2002), Hanlon, Morak and Hegerl (2013), Huth, Kyselý and Pokorná (2000), Meehl and Tebaldi (2004), Otto et al. (2012), Schär et al. (2004)]. Here we build a model, based on an informal notion of what a heat wave is, that may be used for studying such events. Our chief objective for building such a model is for it to be highly interpretable while still realistically characterizing the upper tail of the temperature distribution.

Received February 2015; revised August 2015.

Key words and phrases. Latent state, extremes, generalized Pareto distribution, extremal dependence.

We are aware of very few studies that have applied extreme value theory to the analysis of heat waves. [Furrer et al. \(2010\)](#) applied a conditional points over threshold model to daily temperatures to make inferences about the frequency, intensity and duration of heat waves. A more recent example is [Reich, Shaby and Cooley \(2014\)](#), who modeled serially dependent points above a high threshold using a transformed max-stable process. In addition to their temporal structure, heat waves have potentially important spatial features, which neither these works nor ours attempt to analyze.

Part of the difficulty in analyzing heat waves might be that there is little agreement on exactly what a heat wave is [[Huth, Kyselý and Pokorná \(2000\)](#), [Karl and Knight \(1997\)](#), [Khaliq et al. \(2005\)](#), [Palecki, Changnon and Kunkel \(2001\)](#)]. For example, [Huth, Kyselý and Pokorná \(2000\)](#) defined a heat wave as “the longest continuous period (i) during which the maximum daily temperature was at least T_1 in at least three days, (ii) whose mean maximum daily temperature was at least T_1 , and (iii) during which the maximum daily temperature did not drop below T_2 ” for some specified temperatures T_1 and T_2 [this definition was also used by [Meehl and Tebaldi \(2004\)](#) and [Peng et al. \(2011\)](#)]. [Reich, Shaby and Cooley \(2014\)](#) defined a heat wave as a run of consecutive days above some threshold. [Furrer et al. \(2010\)](#) avoided explicit definitions by pairing their statistical model with a stochastic weather generator, producing draws from which the characteristics of any desired definition of a heat wave can be inferred [the model in [Reich, Shaby and Cooley \(2014\)](#) also has this potential]. Our model uses an implicit definition of a heat wave according to membership in a latent state. Once it is fit using MCMC, it can then function as a weather generator, so it is capable of accommodating any definition of a heat wave that is germane to a given application.

We use a Bayesian hierarchical model with latent state variables that control whether the temperature for each day is assigned the heat wave state or the nonheat wave state. Temporal dependence in the latent state variables is modeled through a simple two-state Markov chain, with one parameter in the transition matrix controlling the frequency of heat waves and the other controlling the persistence of heat waves. For each day, the posterior probability of the state variable represents the degree of confidence with which it is classified as being part of a heat wave.

By employing two states, our model allows the temporal dependence of temperatures that occur in the heat wave state to differ from the dependence structure when temperatures are behaving “typically.” The heat wave state is modeled with a Markovian extreme-value threshold-exceedance model that allows temperatures to exhibit extremal (asymptotic) dependence [[Coles \(2001\)](#), Section 8.4], thereby capturing the persistence of heat waves. The nonheat wave state is modeled with Gaussian dependence structure. Importantly, Gaussian dependence cannot exhibit extremal dependence [[Sibuya \(1959\)](#)]. A Markov model similar to our within-heat-wave component was used by [Smith, Tawn and Coles \(1997\)](#), who studied daily minimum temperature exceedances in Wooster, Ohio. However, whereas [Smith, Tawn and Coles \(1997\)](#) fixed a high threshold and fit the Markov model to the

exceedances, treating all other observations as censored, here we assign the extreme value Markov model to those temperatures that are in the heat wave state, and membership in the state is estimated from the data.

2. A Markov-switching model for threshold exceedances. We begin building our model by informally defining a heat wave as a period of persistent extremely high temperatures. This simple notion leads naturally to a model with two states, one representing days that are part of a heat wave and one representing all other days. Persistence implies positive temporal dependence in the state variables, and extremeness implies temperatures that lie in the upper tail of the distribution. Because our primary focus is the behavior of the upper tail, it is important to appropriately capture tail dependence, and we rely on models suggested by extreme value theory to do so. We seek a parsimonious model that represents both persistence and extremeness in the most interpretable way possible, while still providing a realistic fit to the data.

To define the latent two-state model, let $S_1, \dots, S_T \in \{0, 1\}$ denote the state of the temperature process on each day. The state variable S_t takes a value of 1 if day t is in the heat wave state, and a value of 0 otherwise. The state variables S_1, \dots, S_T are dependent in time according to a Markov chain structure with transition matrix

$$\mathbf{A} = \begin{bmatrix} 1 - a_0 & a_0 \\ 1 - a_1 & a_1 \end{bmatrix}.$$

The parameter $a_0 = P(S_t = 1 | S_{t-1} = 0)$ determines the probability of entering a heat wave, and the parameter $a_1 = P(S_t = 1 | S_{t-1} = 1)$ determines the probability of remaining in a heat wave.

Let the time series $\mathbf{Y} = (Y_1, \dots, Y_T)^T$ denote the observed temperature on days $1, \dots, T$. The distribution of each Y_t will depend on whether or not the corresponding S_t positively indicates membership in the heat wave state. Furthermore, because daily temperature data exhibits strong temporal dependence, we specify a dependence structure for \mathbf{Y} , even conditional on \mathbf{S} . Perhaps the simplest way to model temporal dependence in \mathbf{Y} is through a Markov process.

Since we are assuming a Markov structure for $\mathbf{Y}|\mathbf{S}$, the likelihood of $\mathbf{Y}|\mathbf{S}$ may be written as the product of conditional densities

$$L(\mathbf{y}|\mathbf{s}) = f(y_1|s_1) \prod_{t=2}^T f(y_t|y_{t-1}, s_t, s_{t-1}; \boldsymbol{\theta}),$$

where $\mathbf{y} = (y_1, \dots, y_T)^T$ is the vector of observed temperatures and $\boldsymbol{\theta}$ is a vector of parameters that indexes the set of conditional distributions. Therefore, the conditional likelihood of $\mathbf{Y}|\mathbf{S}$ may be completely specified by four families of conditional distributions $Y_t|Y_{t-1}, S_t = i, S_{t-1} = j$ for $i, j \in \{0, 1\}$. This type of model, depicted graphically in Figure 1, is sometimes referred to as a Markov-switching model [Frühwirth-Schnatter (2006)]. Markov-switching models resemble hidden

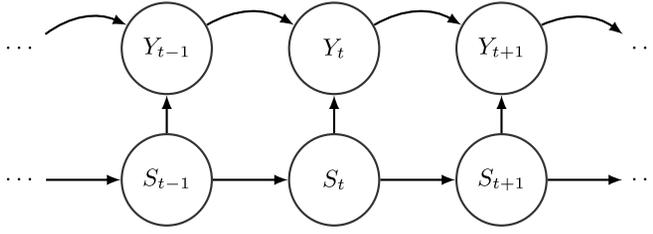


FIG. 1. Graphical representation of the heat wave model. The state variables S_1, \dots, S_T are modeled as a two-state Markov chain. The distribution of each Y_t depends on its corresponding state variable S_t , for $t = 1, \dots, T$. Finally, conditionally on S_1, \dots, S_T , the observations Y_1, \dots, Y_T are also modeled as a Markov process. This structure is sometimes referred to as a Markov-switching model. In a hidden Markov model, there are no arrows directly connecting the observations Y_1, \dots, Y_T .

Markov models, the latter differing in that Y_1, \dots, Y_T are conditionally independent given S_1, \dots, S_T .

Conditioning on the state variables to separate the likelihood into two main components, one arising from the heat wave state and one arising from the nonheat wave state, endows each component with an immediate interpretation: a dedicated tail model for the heat wave state and a model for the bulk of the distribution for the nonheat wave state. Building a separate model for the tail of the distribution is common practice in extreme value analysis. The key concern is that any distributional assumptions designed to fit the bulk of the distribution well may be insufficiently flexible to accommodate the behavior of the tail, and attempts to fit the entire distribution, including the tail, nonparametrically are frustrated by the dearth of data in the tail. A related consideration when assuming the entire distribution comes from a single parametric model is that any fitting procedure will encourage fidelity to the main part of the distribution at the expense of the tail for the simple reason that, by definition, there is much more data in the bulk than in the tail. This is especially undesirable when one is primarily interested in learning about the tail, as we are here. The state variables provide a convenient construct for building separate models for the tail and the main part of the distribution into the likelihood.

We now turn our attention to the conditional likelihoods $f(y_t | y_{t-1}, s_t, s_{t-1}; \theta)$. In the heat wave state, we assume that the temperature is in the far right tail of the distribution. Extreme value theory says that the marginal distribution of values in the upper tail is well approximated by a generalized Pareto distribution (GPD) [Coles (2001)], which has survivor function

$$P(Y > y | Y > u) = \left(1 + \frac{\xi}{\sigma}(y - u)\right)_+^{-1/\xi},$$

where u is a high threshold, σ is a scale parameter, and ξ is a shape parameter that controls the thickness of the tail. Thus, conditionally on $S_t = 1$, we want Y_t

to follow a GPD. Furthermore, we want consecutive observations Y_{t-1} and Y_t , given $S_{t-1} = 1$ and $S_t = 1$ (i.e., given both are in the heat wave state), to exhibit *extremal dependence* because it is clearly seen in the data (see Figures 4, 5). Extremal dependence between Y_{t-1} and Y_t exists if there is positive probability that both observations lie in the asymptotic tail of their bivariate distribution, that is, if $\lim_{c \rightarrow \infty} P(Y_t > c | Y_{t-1} > c) > 0$ for Y_{t-1} and Y_t having the same marginal distribution.

To build the conditional likelihoods for the case where both Y_{t-1} and Y_t are in the heat wave state, $f(y_t | y_{t-1}, s_t = 1, s_{t-1} = 1; \theta)$, with the desired extreme value properties, we follow [Smith, Tawn and Coles \(1997\)](#) and construct

$$(1) \quad f(y_t | y_{t-1}, s_{t-1} = 1, s_t = 1; \theta) = \frac{f(y_{t-1}, y_t | s_{t-1} = 1, s_t = 1; \theta)}{f(y_{t-1} | s_{t-1} = 1; \theta)},$$

where the joint density $f(y_{t-1}, y_t | s_{t-1} = 1, s_t = 1; \theta)$ is a parametric family with GPD margins and extremal dependence, and $f(y_{t-1} | s_{t-1} = 1; \theta)$ is the density of the GPD. This definition affords some flexibility in that any valid joint density with GPD margins and extremal dependence may be used, and several choices for bivariate parametric families are known [[Coles \(2001\)](#), Chapter 8]. Here, we choose the simplest such family, the logistic family with parameter $\alpha \in (0, 1]$. The bivariate logistic model may be defined through its cumulative distribution function (CDF),

$$(2) \quad G(z_{t-1}, z_t) = \exp\left\{-\left(z_{t-1}^{-1/\alpha} + z_t^{-1/\alpha}\right)^\alpha\right\},$$

where z_{t-1} and z_t are derived from y_{t-1} and y_t by applying the transformation from GPD to unit Fréchet, $z = -\log(F_{\text{GPD}}(y))^{-1}$, where F_{GPD} denotes the CDF of the GPD. The bivariate likelihoods $f(y_{t-1}, y_t | s_{t-1} = 1, s_t = 1; \theta)$ are obtained by differentiating (2) with respect to y_{t-1} and y_t . For the logistic model, smaller values of α indicate stronger dependence, with $\alpha \rightarrow 0$ representing complete dependence and $\alpha = 1$ representing complete independence.

The next case that we consider is when $S_{t-1} = S_t = 0$, indicating that times $t-1$ and t are both members of the nonheat wave state. This case is modeled simply as an AR(1) process with mean μ , variance σ_N^2 and autocorrelation parameter $\phi \in (0, 1)$ (negative autocorrelation is physically implausible). That is, the conditional densities for days outside of heat waves are modeled directly as

$$f(y_t | y_{t-1}, s_t = 0, s_{t-1} = 0) = N(\mu + \phi(y_{t-1} - \mu), \sigma_N^2),$$

where, unlike in the case of the logistic family, larger values of the dependence parameter ϕ indicate stronger dependence. The Gaussian AR(1) process is appropriate for the bulk of the temperature distribution (see Figures 2, 3), but, unlike the logistic Markov process defined by (1) and (2), the AR(1) process is asymptotically independent and therefore inadequate as a model for the tail behavior. More

elaborate models for the bulk of the temperature distribution are possible, but because our focus is on the tail, we use the simplest available structure that seems to fit the data, the Gaussian AR(1).

Finally, we must specify the heterogeneous cases $\{S_{t-1} = 0, S_t = 1\}$ and $\{S_{t-1} = 1, S_t = 0\}$. These represent the transitions into and out of heat waves. The approach we take here is again similar to that of [Smith, Tawn and Coles \(1997\)](#) in that we define the conditional densities through corresponding bivariate densities as in (1), which again have logistic dependence defined through (2). The difference from the $\{S_{t-1} = S_t = 1\}$ case is that here one of the marginal distributions is Gaussian rather than GPD. This necessitates two modifications. The first is that for $\{S_{t-1} = 0, S_t = 1\}$ (the transition into a heat wave), the density in the denominator of (1) is normal. The second is that in both heterogeneous cases, one of the z variables in (2) is the result of a transformation from normal to unit Fréchet, $z = -\log(\Phi[(y - \mu)/\sigma_N])^{-1}$, rather than both being the result of the transformation from GPD to unit Fréchet. Here, we use a single dependence parameter α_{01} to characterize the temporal dependence between the first day of a heat wave and the day before, and between the last day of a heat wave and the day after.

Explicit formulas for the bivariate likelihoods for the days corresponding to transitions into and out of heat waves, that is, for $y_{t-1}, y_t | s_{t-1} = 0, s_t = 1$ and $y_{t-1}, y_t | s_{t-1} = 1, s_t = 0$, are constructed as follows. First, transform both margins to $U(0, 1)$ by taking, for $j = t - 1, t$,

$$u_j = \begin{cases} \Phi\left(\frac{y_j - \mu}{\sigma_N}\right), & \text{when } s_j = 0, \\ 1 - \left[1 + \frac{\xi(y_j - u)}{\sigma}\right]^{-1/\xi}, & \text{when } s_j = 1, \end{cases}$$

where $\Phi(\cdot)$ is the standard normal CDF. Next, transform both margins to unit Fréchet using $z_j = -\log(u_j)^{-1}$, for $j = t - 1, t$. The bivariate likelihood is then

$$f(y_{t-1}, y_t | s_{t-1}, s_t; \theta) = K_{t-1} K_t (V_{t-1} V_t - V_{t-1,t}) e^V,$$

where $V = (z_{t-1}^{-1/\alpha_{01}} + z_t^{-1/\alpha_{01}})^{\alpha_{01}}$, $V_{t-1,t} = (1 - 1/\alpha_{01})(z_{t-1} z_t)^{-1/\alpha_{01} - 1} V^{1-2/\alpha_{01}}$ and, for $j = t - 1, t$, $V_j = z_j^{-1/\alpha_{01} - 1} V^{1-1/\alpha_{01}}$ and

$$K_j = \begin{cases} \varphi\left(\frac{y_j - \mu}{\sigma_N}\right) z_j^2 \exp(1/z_j), & \text{when } s_j = 0, \\ \sigma^{-1} u_j^{1+\xi} z_j^2 \exp(1/z_j), & \text{when } s_j = 1, \end{cases}$$

where $\varphi(\cdot)$ is the standard normal probability density function (p.d.f.).

The bivariate likelihoods for days within heat waves, that is, for $y_{t-1}, y_t | s_{t-1} = 1, s_t = 1$, are exactly the same, only the dependence parameter α is substituted for α_{01} .

Most of the model parameters have direct physical interpretations, so posterior inference on them immediately tells us something about the nature of the observed heat waves. First and foremost, the state variables S_1, \dots, S_N indicate whether or not each day is classified as being in a heat wave. By looking at the posterior state probabilities for each day, we can easily retrospectively identify when and for how long heat waves occurred, according to the model. The Markov transition probability a_0 represents the propensity of the system to enter into heat waves, and a_1 represents the propensity of heat waves to persist once they get started. Hence, together a_0 and a_1 describe the expected number and duration of heat waves. The GPD parameters u, σ and ξ characterize the severity of the heat waves, with u representing the minimum temperature needed to attain heat wave status. The dependence parameters α, α_{01} and ϕ together control the strength of the temporal dependence in the temperature series. Finally, μ and σ_N^2 describe the marginal behavior of the temperature on days that are not in heat waves.

An interesting feature of the model is that short-lived extremely high temperatures are not necessarily classified as heat waves. The Markov chain structure of the state variables encourages the model to consider duration in its classification criteria, so single very hot days, for example, will tend to have low posterior probability of being in the heat wave state. This phenomenon is illustrated in the results of the case study found in the next section.

3. Case studies. As case studies, we select temperature time series from Paris and Moscow, which both recently suffered through high-profile heat waves. To simplify the analysis, we extract daily maximum temperatures from the summer months (JJA) from the years 1990–2011 (92 days per year over 22 years), a time period that is short enough that a stationarity assumption is plausible. Temperature data is available through the European Climate Assessment Database (<http://eca.knmi.nl/>). For the time period under consideration, the Paris time series is complete, and the Moscow time series contains just two missing values. To remove seasonal effects in the JJA data, we de-seasonalize using the following procedure. First, we fit a penalized spline to the JJA temperatures, using absolute error, rather than squared error, as a loss function (using the `qsreg` function in the R package `fields` [Nychka, Furrer and Sain (2014)], with the smoothing parameter chosen by the default generalized cross-validation criterion). In this way, we do not allow the magnitude of the extremes to unduly influence the calculation of the climatological average. Next, we subtract the fitted spline function from the raw data. Finally, to aid interpretation, we add back the (constant in time) overall median temperature so that the magnitude remains on the same scale as the original data, but without the seasonal cycle. Even after removing the seasonally-varying median in this way, it is possible that some seasonality remains, for example, a seasonally-varying variance, although we find no evidence of this. With the (first-order) seasonal cycle removed, now assume that all model parameters are constant in time.

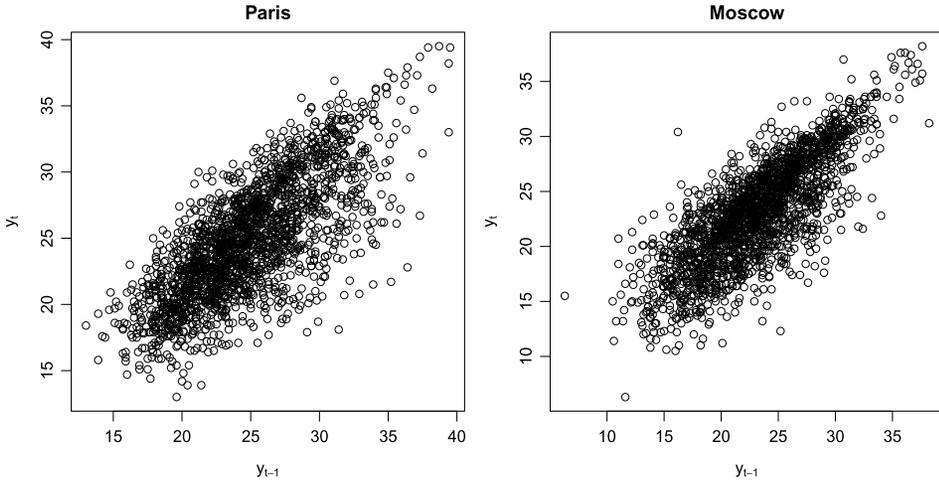


FIG. 2. JJA temperatures y_t (i.e., at day t) plotted against the temperature of the previous day, y_{t-1} (in $^{\circ}\text{C}$). Paris is shown in the left panel, Moscow on the right panel. There is strong autocorrelation in both cases.

3.1. *Exploratory analysis.* To check the validity of the model, we run a variety of diagnostics. Since most of the data does not lie in heat waves, the AR(1) portion of the model is the easiest to check. A scatterplot of y_{t-1} vs. y_t is shown in Figure 2, which indicates strong autocorrelation at lag 1 in both Paris and Moscow.

Next, we plot empirical partial autocorrelation functions for each city in Figure 3. In both cities, we see a large value at lag 1, quickly decaying to near zero by lag 2. This pattern is consistent with an AR(1) model. To further check the validity of the AR(1) assumption, we fit AR(p) models to each year of data separately and choose p using AIC. In the majority of years, AIC chooses $p = 1$. We conclude that the simple structure we have specified for the nonheat wave days is adequate.

To check for extremal dependence, we again look at a scatterplot of y_t against y_{t-1} , but this time we first transform the data to the Fréchet scale using a rank transformation (Figure 4). If asymptotic dependence were not present, Figure 4 would show points lining up along the y_{t-1} - and y_t -axes. What we see instead is that points lie in the interior of the plot, a pattern that indicates asymptotic dependence [Coles (2001)].

As an additional check for extremal dependence, we examine estimates of the quantity $\chi = \lim_{u \rightarrow \infty} \hat{\chi}(u)$, where $\chi(u) = P\{Y_t > u | Y_{t-1} > u\}$ [Coles, Heffernan and Tawn (1999)], for the two cities. A value of $\chi = 0$ indicates asymptotic independence, while any $0 < \chi < 1$ indicates asymptotic dependence. In practice, we estimate $\hat{\chi}(u)$ for many values of u and examine its behavior as u gets large. Plots of $\hat{\chi}(u)$ are shown in Figure 5, with the x -axis transformed to the quantile scale for clarity. In both cases, the curves remain comfortably away from zero, except at the

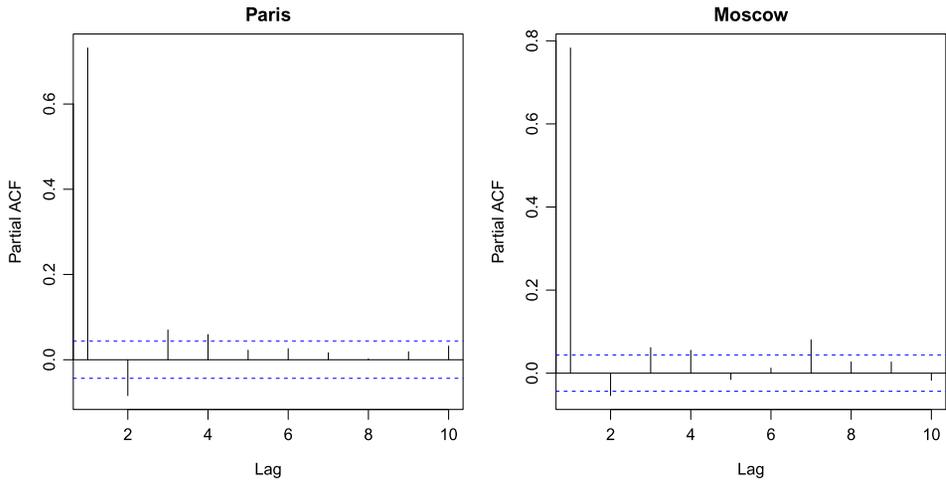


FIG. 3. *Partial autocorrelation functions for the temperature data. Paris is shown in the left panel, Moscow on the right panel. The large value at lag 1 and the small values at all other lags are consistent with the AR(1) assumption. Values between the dashed lines are not significantly different from zero at $\alpha = 0.05$.*

far right-hand edge of the Moscow plot where there is almost no data, again suggesting that asymptotic dependence is present in the data at lag 1. (For additional exploratory analysis, see the Supplemental Materials [Shaby et al. (2015)].)

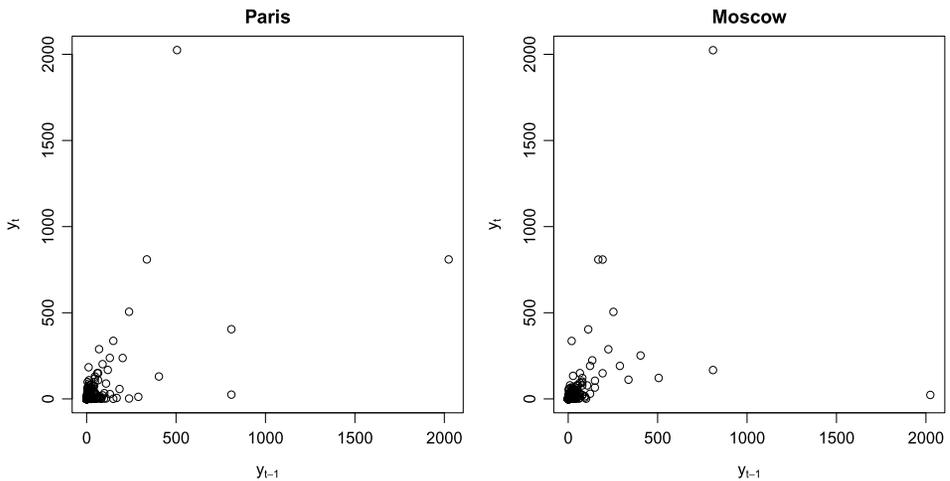


FIG. 4. *JJA daily temperatures y_t plotted against the temperature of previous day, y_{t-1} , on the Fréchet scale. Paris is shown in the left panel, Moscow on the right panel. The presence of many points lying in the interior of the plot (i.e., away from the axes) suggests strong asymptotic dependence at lag 1.*

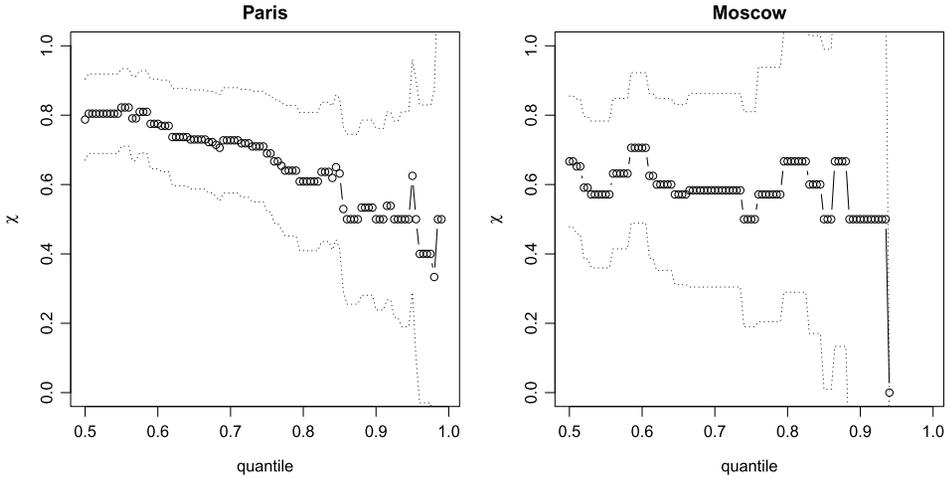


FIG. 5. Estimates of $\chi(u)$, for increasing values of u (with the u -axes on the quantile scale). Paris is shown in the left panel, Moscow on the right panel. Since $\chi = 0$ indicates extremal independence, these plots indicate the presence of asymptotic dependence.

3.2. Prior specification and computing. Since these exploratory checks are consistent with the proposed model, we move ahead. The next step is to specify prior distributions on the model parameters. For the GPD marginal parameters, we choose a vague normal for $\log \sigma$, a uniform on $(-0.5, 0.5)$ for ξ , and a normal with a small variance centered on the 0.98 quantile (33° for Moscow and 35° for Paris) for the threshold u . The priors on ξ and u are informative, but, we believe, justified. Previous studies of summer high temperatures routinely estimate ξ at around -0.22 , so the chosen uniform prior will have little effect other than ensuring that the posterior have no support on $(-\infty, -0.5]$, the region for which the GPD is not regular (i.e., standard likelihood results do not apply [Smith (1985)]). The tight normal prior on u encourages the GPD to be applied only to the tail of the distribution, but not so far into the tail as to be irrelevant. The informative prior on u is necessary to achieve good convergence, and less restrictive than the standard practice in extreme value analysis of fixing u at a prespecified value. A sensitivity analysis for the prior on u is reported in the Supplementary Materials.

For the Gaussian marginal parameters, we choose vague normal priors for both μ and $\log \sigma_N^2$. The dependence parameters α , α_{01} and ϕ are given uninformative uniform $(0, 1)$ priors. Finally, conjugate beta priors are specified for the Markov transition probabilities a_0 and a_1 .

Posterior simulation is carried out using a block Gibbs sampler, with conjugate updates for a_0 and a_1 , and Metropolis updates for all other model parameters. The state variables S_1, \dots, S_T are updated jointly using a forward-filtering backward-sampling algorithm [Frühwirth-Schnatter (2006)] within the Gibbs sampler. Missing values are handled seamlessly by treating them as unknown parameters and drawing from their predictive distributions at each MCMC iteration.

3.3. *Results.* For each day in the study period, the sampler outputs the posterior probability of being in a heat wave. A useful place to start examining the results is by looking at these posterior probabilities for time periods that include the famous heat waves that motivated this study. Figure 6 shows the temperature time series from the summer of 2003 in Paris and the summer of 2010 in Moscow.

The y-axis in Figure 6 is the observed temperature, and the color of each dot is proportional to the posterior probability of membership in the heat wave state. These plots show that the model correctly identifies these well-known events. Furthermore, it locates fairly clear beginning and end points of each event, the locations of which are not so obvious from just looking at the time series in the case of Moscow in 2010.

The bottom panel of Figure 6 also demonstrates the interesting feature described at the end of Section 2, where a very hot day in Moscow on August 17 is not classified as a heat wave, even though it was noticeably warmer than other days that were classified as heat waves. This is because the value of a_1 that the model estimates defines a heat wave as having rather strong persistence with high probability, and August 17 stands apart from its closest neighbors, whereas the cooler days toward

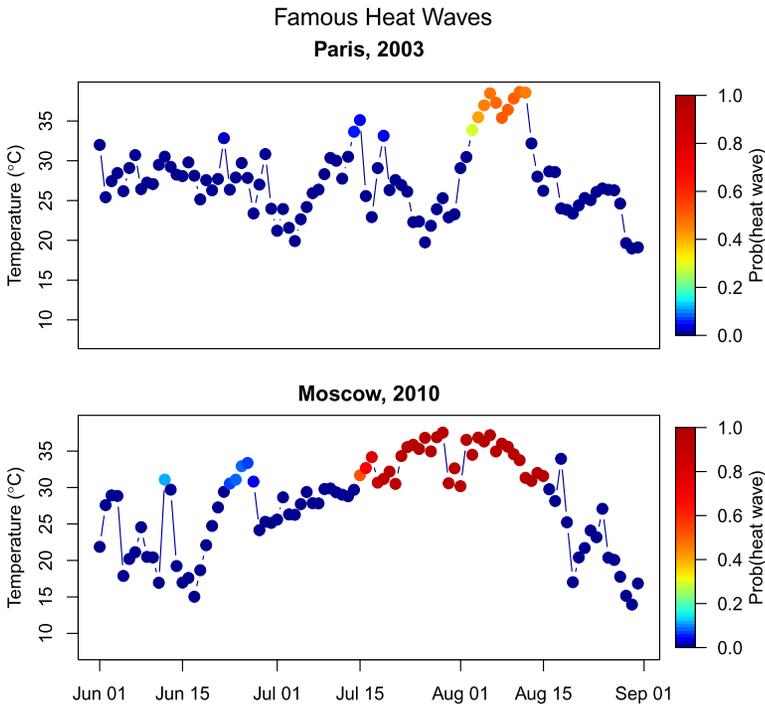


FIG. 6. Temperature data from the European heat wave of 2003, which hit France especially hard, and the Russian heat wave of 2010. Color coding of the dots indicates the posterior probability of being in State 1. Note that in Moscow, August 17 was hotter than, say, August 14, but the model did not classify it as a heat wave because it was a single hot day. Similarly in Paris in the middle of July.

the end of the 2010 heat wave were members of a contiguous mass. This feature of the model conforms to our notion of what a heat wave is: it must be both very hot and persistent—just being hot is not enough. In this way, the fitted model contains an implicit definition of a heat wave, inclusion in State 1 given the data.

Figure 7 shows two summer high temperature time series from Moscow, 1996 and 2001. The model output suggests that there might have been a heat wave in 1996. However, even though the annual maximum in 1996 was higher than the annual maximum in 2001, the hot period in 2001 was classified more confidently as being a heat wave. This is again because of persistence; the hot period in 1996 lasted a short time and showed weak temporal dependence, while the hot period in 2001, though cooler, lasted longer and showed stronger temporal dependence consistent with the posterior estimate of α .

Figure 8 shows kernel density estimates of the posterior densities of the parameters a_1 and α . The Markov transition probability a_1 is the probability of remaining in a heat wave, given that one has already started. Comparing the two curves, it appears that the model is more confident that Moscow (solid curve) tends to have persistent heat waves than it is about Paris (dashed curve), although for both cities

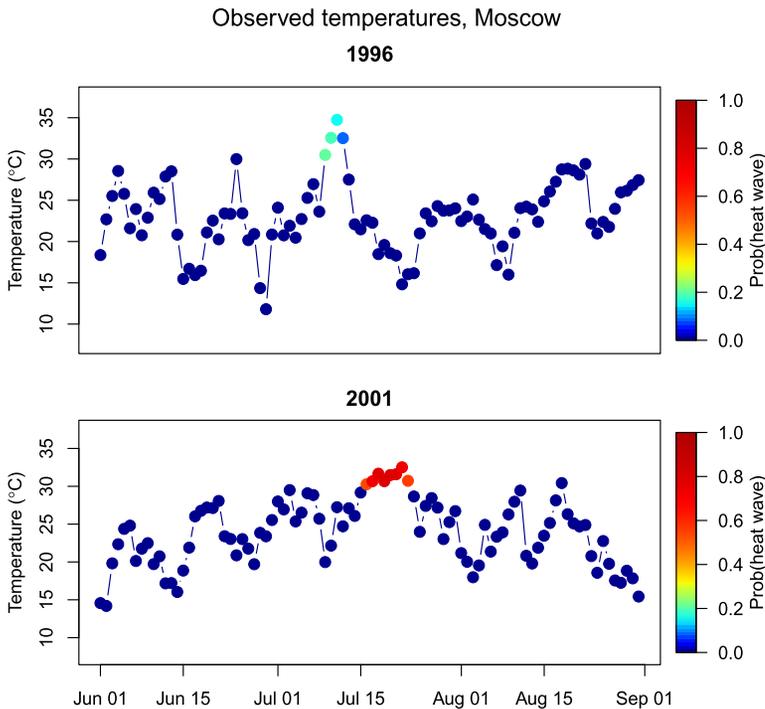


FIG. 7. *Even though the annual maximum in 1996 was higher than that in 2001, the model classifies the hot period in 2001 more confidently as a heat wave, probably because of the prolonged period and strong temporal dependence in 2001.*

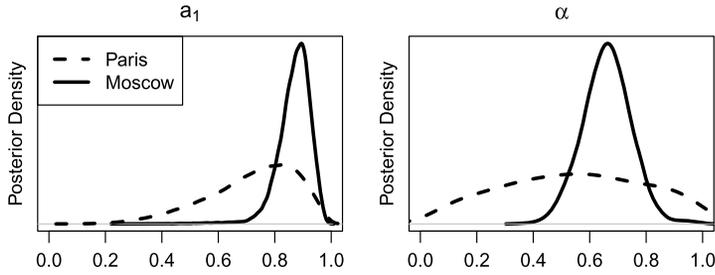


FIG. 8. Kernel density estimates of the marginal posterior densities of a_1 and α .

almost all of the mass lies well to the right of 0.5. The logistic dependence parameter α controls the temporal dependence of temperatures within heat waves. From Figure 8, we see that while posterior means for the two cities are similar, the model again allows posterior mass to concentrate more for Moscow at around 0.6, well away from the extreme cases of independence and complete dependence. The relatively high posterior precision in Moscow probably reflects the larger number of heat waves that occurred there during the study period [Figure 9(b)]. The combined interpretation of the two parameters shown in Figure 8 is that with high probability, Moscow, when it does experience heat waves, tends to experience longer heat waves with more stable temperatures.

The expected length and frequency of the heat waves is directly calculable from a_0 and a_1 , and posterior distributions of these expectations are straightforward to estimate from the MCMC sample. This type of exercise is useful for making predictions in a stationary world. In addition to looking ahead using expectation-type calculations, it is interesting to do retrospective analysis of actual heat waves during the study period.

Figure 9 demonstrates a retrospective analysis. The left-hand panel (a) shows the posterior probability mass function (p.m.f.) of the length of heat waves that occurred in Paris and Moscow from 1990–2011. The most prominent feature of Figure 9(a) is the large amount of probability mass for Moscow at large durations. This massive right tail mostly reflects the extremely long 2010 heat wave. Figure 9(b) shows that heat waves in Moscow tended to be more numerous than those in Paris. However, we see in Figure 9(c) that heat waves in Paris tended to be much hotter. Putting together these three characteristics, it appears that heat waves in Paris from 1990–2011 were hotter, though shorter and less frequent, than those in Moscow.

An anonymous referee points out that this behavior is as expected. Moscow has a more continental climate, enabling stable anticyclonic conditions (or *blocking* episodes), associated with clear skies and excesses of downward solar radiation, to persist for long periods. In contrast, Western Europe is under the influence of the jet stream and its westerly winds directly coming from the Western Atlantic ocean, which inhibits the maintenance of blocking situations. The higher temperatures observed in Paris are likely due to the effect of latitude.

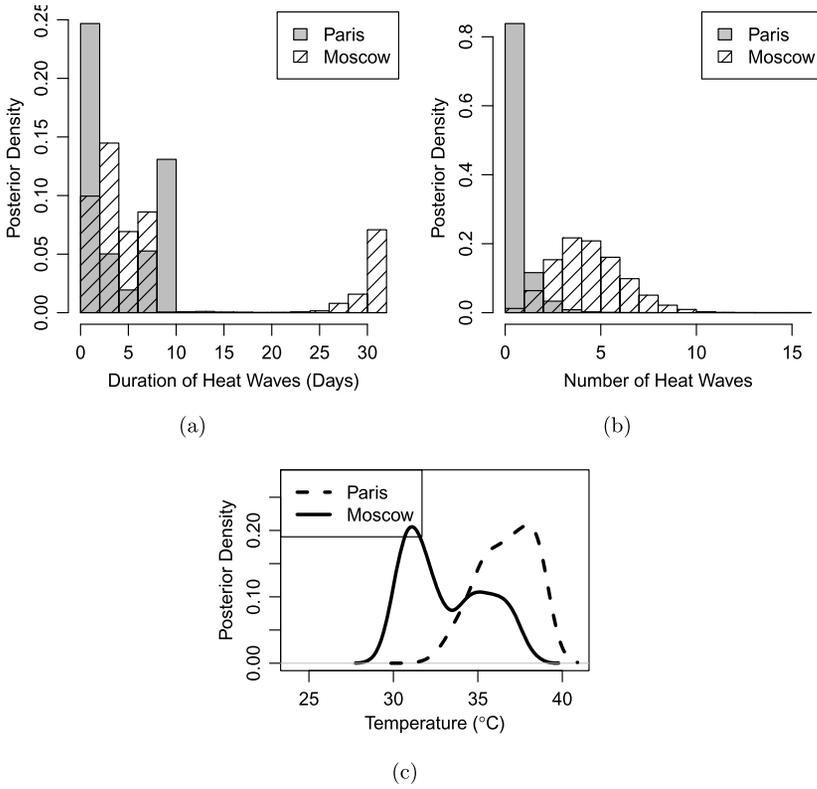


FIG. 9. A retrospective analysis of heat waves during the study period. Panel (a) shows the posterior p.m.f. of the length of the heat waves that occurred in Paris (shaded) and Moscow (crosshatched), and panel (b) shows the posterior p.m.f. of the number of heat waves that occurred. Panel (c) shows that the temperatures that occurred during heat waves were much higher in Paris than in Moscow.

3.4. *Alternative definitions of heat waves.* To explore the behavior of the model under alternative definitions of heat waves, we use it as a stochastic weather generator and compute the posterior distribution of frequency, duration and mean temperature of heat waves, where heat waves are defined according to criteria found in the literature. For each MCMC iteration, we simulate 500 summers worth of random draws from the model, conditional on the model parameters at that iteration. This results in a posterior sample of summers, from which we can apply any definitions of heat waves that we choose. Following Meehl and Tebaldi (2004), we use two common criteria. The first defines a heat wave as the three-day period in any given year with the highest average low temperature, based on the idea that stretches without relief from extreme heat may have large health impacts [Karl and Knight (1997)]. Since we are working with daily high temperatures rather than daily lows, we modify this “worst annual event” definition accordingly. The second considers two thresholds T_1 and T_2 and defines a heat wave as the

longest contiguous period during which the daily high temperature exceeds T_1 at least three times, the daily high temperature is always above T_2 , and the average daily high temperature is greater than T_1 [Huth, Kyselý and Pokorná (2000)]. The thresholds T_1 and T_2 are set, respectively, at the 0.975 and 0.81 empirical quantiles [Meehl and Tebaldi (2004)].

Figures 10 (Paris) and 11 (Moscow) show features of our simulated heat waves, under the three definitions of heat waves (the definition implicit in our latent state model—a contiguous block of days for which $S = 1$, the threshold-based definition, and the “worst annual event” definition). We plot the posterior density of the duration of heat waves [Figures 10(a) and 11(a)], the frequency of heat waves [Figures 10(b) and 11(b)] and the mean temperature during heat waves [Figures 10(c)

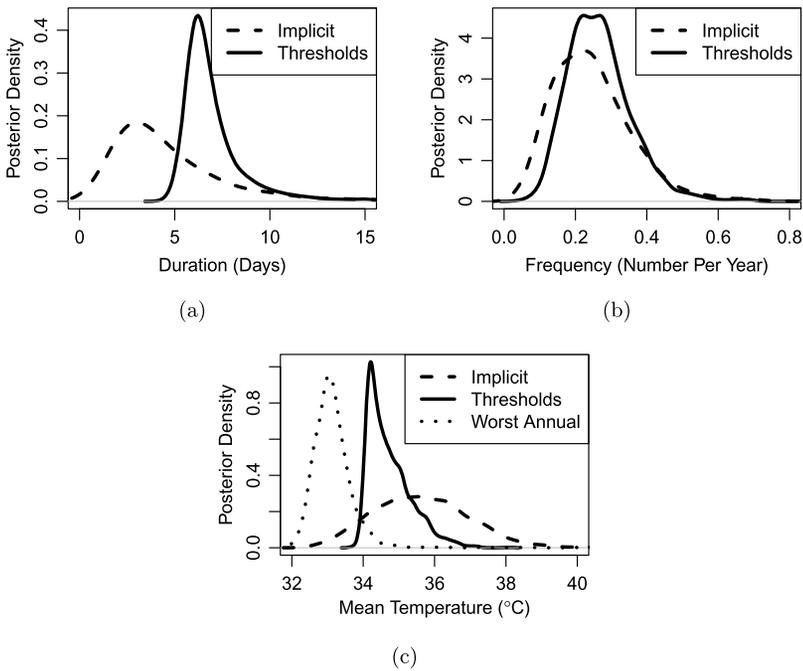


FIG. 10. Comparisons of Paris heat waves under different definitions. Panel (a) shows kernel density estimates of the posterior distribution of the duration of heat waves, where heat waves are defined implicitly by the latent state model, as well as using a threshold definition. The threshold definition produces heat waves that are longer than the implicit definition. Panel (b) shows the posterior densities of the frequency of heat waves under the same two definitions of heat waves. The two definitions (coincidentally) produce heat waves at similar frequencies. Panel (c) shows the posterior densities of the mean daily high temperatures during heat waves, under the two previous definitions, plus the “worst annual event” definition. This latter definition is less restrictive, and hence produces heat waves that are cooler than the other two. All posterior distributions are sampled by drawing from the latent state model, conditional on model parameters at each iteration of the MCMC sampler.

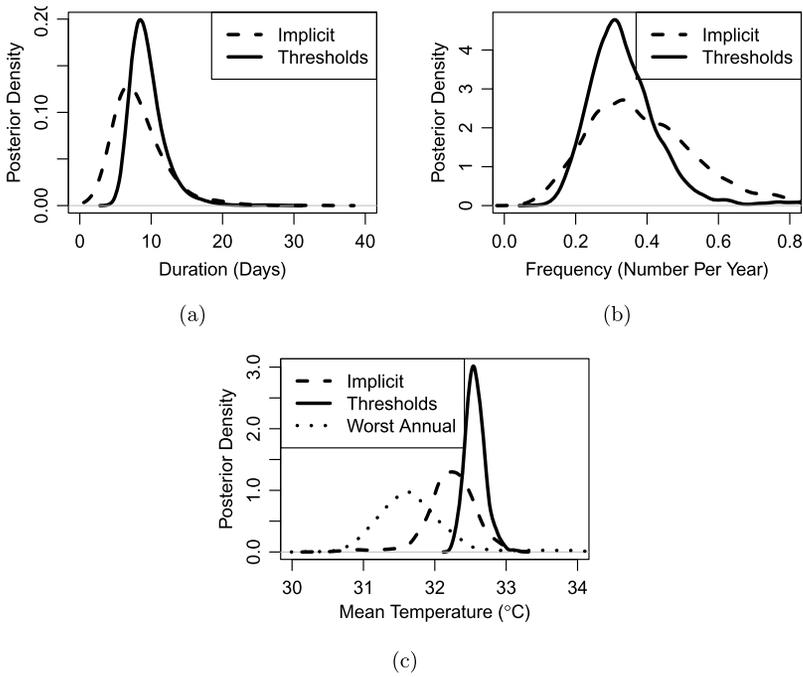


FIG. 11. Comparisons of Moscow heat waves under different definitions. Panels (a), (b) and (c) show the same basic patterns seen in Figure 10. In comparison with Figure 10, the most noticeable difference is that posterior densities for the implicit definition and the threshold definition coincide more closely for Moscow than for Paris. Just as for Paris, Panel (c) shows that the “worst annual event” definition produces the least severe heat waves, as expected.

and 11(c)]. By definition, the “worst annual event” type of heat wave occurs exactly once per year for three days, making duration and frequency trivial.

The overall patterns in heat wave characteristics are similar across cities. The implicit latent state definition produces shorter heat waves than the threshold definition. The frequency of implicitly-defined and threshold-based heat waves is similar, but the posterior distribution is slightly more diffuse for the implicit definition. For the mean daily high temperature during heat waves, the “worst annual event” heat waves are cooler than the other two, which is expected because the implicit and threshold definitions find heat waves less frequently than once annually, and hence exclude the less extreme annual events. The posterior distribution of mean temperatures is noticeably more peaked for the threshold than for the implicit definition.

3.5. Assessing model fit. Popular tools for assessing the fit of a Bayesian model to a given data set (e.g., deviance information criterion [Spiegelhalter et al. (2002)], proper scoring rules [Gneiting and Raftery (2007)]) work by comparing the fits of competing models and choosing the one with the highest score. Here,

our goal is interpretability, not achieving the best possible fit to the data, so these tools are not ideal. However, we still need to determine whether the fit is adequate. To check compatibility with the data without making comparisons among competing models, we use the posterior predictive checks of [Gelman, Meng and Stern \(1996\)](#). The idea is to make posterior predictive draws from the model and see whether those draws resemble the observed data set according to a suite of relevant summary statistics. If the summaries of the observed data set fall within an acceptable range of the summaries of the simulated data sets, then the model is deemed adequate.

We have already made many draws from the posterior predictive distribution from the analysis in Section 3.4, so we can use those for model assessment. Since we are interested in heat waves, we choose summary statistics that describe the extremes of the temperature distribution. To assess the marginal fit, we compute the 0.99 and 0.999 empirical quantiles. To assess the fit of the dependence, we compute an estimate of the extremal index [[Ferro and Segers \(2003\)](#)] at the 0.975 observed quantile (denoted as $\hat{\nu}$ in Table 1), which can be interpreted as the inverse of the mean size of clusters of observations above the chosen threshold, as well as $\hat{\chi}(u)$ at time lags of 1 and 5 days and at several values of u . We compute each summary statistic for each posterior predictive draw and report the 0.025 and 0.975 quantiles. If the observed statistics fall within their corresponding predictive intervals, we declare that the model fits the data satisfactorily.

Table 1 shows the results of the posterior predictive checks. For both Paris and Moscow, the observed summary statistics fall within the posterior 95% intervals of the chosen statistics. The only hint of a problem is for $\hat{\chi}(32)$ at lag 1 for Moscow, where the observed statistic falls on the endpoint of the predictive interval. Other than that, the predictive diagnostic indicates that the model provides a suitable fit to the data.

TABLE 1

Posterior predictive intervals for summary statistics. Each column corresponds to a summary statistic. The statistics $q_{0.99}$ and $q_{0.999}$ (the empirical 0.99 and 0.999 quantiles) describe the extremes of the marginal predictive distributions, and the statistics $\hat{\nu}$ (the extremal index) and $\hat{\chi}_h(u)$ (extremal dependence at time lag h and threshold u) describe the strength of the asymptotic dependence. For both Paris and Moscow, the top and bottom rows correspond to the 0.025 and 0.975 quantiles of the posterior draws, and the middle row corresponds to the observed quantities

		$q_{0.99}$	$q_{0.999}$	$\hat{\nu}$	$\hat{\chi}_1(28)$	$\hat{\chi}_5(28)$	$\hat{\chi}_1(32)$	$\hat{\chi}_5(32)$	$\hat{\chi}_1(36)$	$\hat{\chi}_5(36)$
Paris	$q_{0.025}$	33.62	33.62	0.205	0.529	0.232	0.303	0.030	0.000	0.000
	obs.	35.33	38.48	0.561	0.607	0.297	0.389	0.102	0.462	0.231
	$q_{0.975}$	38.32	43.17	0.867	1.000	1.000	1.000	1.000	0.686	0.386
Moscow	$q_{0.025}$	32.50	34.83	0.225	0.494	0.165	0.244	0.000	0.000	0.000
	obs.	34.13	36.89	0.245	0.662	0.356	0.600	0.280	0.375	0.125
	$q_{0.975}$	35.82	39.25	0.754	0.758	0.517	0.600	0.333	0.571	0.167

4. Discussion. We have presented a simple Bayesian latent state model for studying heat waves. The chief virtue of this model is its interpretability; the latent state vector \mathbf{S} directly indicates which days are part of heat waves, and the Markov transition probabilities represent the frequency and duration of heat waves. In addition to the easy interpretability, the latent state model has the advantage that, unlike other extreme value models, there is no need to prespecify a threshold over which the GPD applies or to use censored likelihoods for data below the threshold, giving it a certain elegance.

One feature that is extremely useful is the ability to sample from the fitted model, allowing it to act as a weather generator (Section 3.4). In this way, our model is adaptable to any operational definition of heat waves that is germane to the application at hand. An additional feature of our approach is that missing values, endemic to meteorological data, are easily handled by treating them as unknown parameters.

A limitation of our approach is that it only considers maximum daily temperatures, but other aspects of heat waves might be of interest to analysts. A more complete picture of heat waves might include, for example, the daily minimum temperature or measures of heat stress.

In the case studies, we have assumed stationary that might not be realistic for other data sets, but extensions are straightforward. A more intricate analysis would allow, for example, a seasonal trend in u and μ . In addition, it would be a simple matter to include an inter-annual trend for model parameters, which would allow for the investigation of long-term changes in the behavior of heat waves.

Another improvement would be to borrow strength across space using Gaussian process priors, which would provide improved statistical efficiency at the expense of a more complicated forward-filtering backward-sampling algorithm. Candidates for spatial priors are the Markov transition probabilities a_0 and a_1 , the GPD parameters u , σ and ξ , and possibly even the dependence parameters α , α_{01} and ϕ . A related extension is to replace the two-state Markov chain on the states with a logistic regression-type model, where the state probabilities would depend on a spatially and temporally dependent random process, and possibly on covariates as well. This type of model is appealing, but it sacrifices much of the interpretability that is so desirable.

Acknowledgments. We would like to thank the three anonymous referees, as well as the Associate Editor, for their insightful comments.

SUPPLEMENTARY MATERIAL

Additional analysis (DOI: [10.1214/00-10.1214/15-AOAS873SUPP](https://doi.org/10.1214/00-10.1214/15-AOAS873SUPP); .pdf). The Supplement contains additional exploratory analysis related to the case study, a sensitivity analysis for the prior distribution of the GPD threshold u , and results of the model run on 2003 temperatures at several additional sites throughout Western Europe.

REFERENCES

- AMENGUAL, A., HOMAR, V., ROMERO, R., BROOKS, H. E., RAMIS, C., GORDALIZA, M. and ALONSO, S. (2014). Projections of heat waves with high impact on human health in Europe. *Glob. Planet. Change* **119** 71–84.
- BARRIOPEDRO, D., FISCHER, E. M., LUTERBACHER, J., TRIGO, R. M. and GARCÍA-HERRERA, R. (2011). The hot summer of 2010: Redrawing the temperature record map of Europe. *Science* **332** 220.
- CLARK, R. T., BROWN, S. J. and MURPHY, J. M. (2006). Modeling northern hemisphere summer heat extreme changes and their uncertainties using a physics ensemble of climate sensitivity experiments. *J. Climate* **19** 4418–4435.
- COLES, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer London, Ltd., London. [MR1932132](#)
- COLES, S., HEFFERNAN, J. and TAWN, J. (1999). Dependence measures for extreme value analyses. *Extremes* **2** 339–365.
- EASTERLING, D. R., EVANS, J. L., GROISMAN, P. Y., KARL, T. R., KUNKEL, K. E. and AMBENJE, P. (2000). Observed variability and trends in extreme climate events: A brief review. *Bull. Am. Meteorol. Soc.* **81** 417–425.
- FERRO, C. A. T. and SEGERS, J. (2003). Inference for clusters of extreme values. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **65** 545–556. [MR1983763](#)
- FISCHER, E. M. and SCHÄR, C. (2010). Consistent geographical patterns of changes in high-impact European heatwaves. *Nat. Geosci.* **3** 398–403.
- FRICH, P., ALEXANDER, L. V., DELLA-MARTA, P., GLEASON, B., HAYLOCK, M., KLEIN TANK, A. M. G. and PETERSON, T. (2002). Observed coherent changes in climatic extremes during the second half of the twentieth century. *Clim. Res.* **19** 193–212.
- FRÜHWIRTH-SCHNATTER, S. (2006). *Finite Mixture and Markov Switching Models*. Springer, New York. [MR2265601](#)
- FURRER, E. M., KATZ, R. W., WALTER, M. D. and FURRER, R. (2010). Statistical modeling of hot spells and heat waves. *Clim. Res.* **43** 191–205.
- GELMAN, A., MENG, X.-L. and STERN, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statist. Sinica* **6** 733–807. [MR1422404](#)
- GNEITING, T. and RAFTERY, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* **102** 359–378. [MR2345548](#)
- HANLON, H. M., MORAK, S. and HEGERL, G. C. (2013). Detection and prediction of mean and extreme European summer temperatures with a multimodel ensemble. *J. Geophys. Res.* **118** 9631–9641.
- HUTH, R., KYSELÛ, J. and POKORNÁ, L. (2000). A GCM simulation of heat waves, dry spells, and their relationships to circulation. *Clim. Change* **46** 29–60.
- KARL, T. R. and KNIGHT, R. W. (1997). The 1995 Chicago heat wave: How likely is a recurrence? *Bull. Am. Meteorol. Soc.* **78** 1107–1120.
- KHALIQ, M. N., ST-HILAIRE, A., OUARDA, T. and BOBEE, B. (2005). Frequency analysis and temporal pattern of occurrences of southern quebec heatwaves. *Int. J. Climatol.* **25** 485–504.
- MEEHL, G. A. and TEBALDI, C. (2004). More intense, more frequent, and longer lasting heat waves in the 21st century. *Science* **305** 994–997.
- MUNICH RE (2003). Annual review: Natural catastrophes 2003. Available at www.munichre.com.
- NYCHKA, D., FURRER, R. and SAIN, S. (2014). fields: Tools for spatial data. R package version 7.1.
- OTTO, F. E. L., MASSEY, N., OLDENBORGH, G. J., JONES, R. G. and ALLEN, M. R. (2012). Reconciling two approaches to attribution of the 2010 Russian heat wave. *Geophys. Res. Lett.* **39** L04702–L04707.

- PALECKI, M. A., CHANGNON, S. A. and KUNKEL, K. E. (2001). The nature and impacts of the July 1999 heat wave in the midwestern United States: Learning from the lessons of 1995. *Bull. Am. Meteorol. Soc.* **82** 1353–1367.
- PARRY, M. L., CANZIANI, O. F., PALUTIKOF, J. P., VAN DER LINDEN, P. J. and HANSON, C. E., eds. (2007). *Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the IPCC*. Cambridge Univ. Press, Cambridge.
- PENG, R. D., BOBB, J. F., TEBALDI, C., MCDANIEL, L., BELL, M. L. and DOMINICI, F. (2011). Toward a quantitative estimate of future heat wave mortality under global climate change. *Environ. Health Perspect.* **119** 701–706.
- REICH, B. J., SHABY, B. A. and COOLEY, D. (2014). A hierarchical model for serially-dependent extremes: A study of heat waves in the western US. *J. Agric. Biol. Environ. Stat.* **19** 119–135. [MR3257905](#)
- ROBINE, J. M., CHEUNG, S. L. K., LE ROY, S., VAN OYEN, H., GRIFFITHS, C., MICHEL, J. P. and HERRMANN, F. R. (2008). Death toll exceeded 70,000 in Europe during the summer of 2003. *C. R., Biol.* **331** 171–178.
- SCHÄR, C., VIDALE, P. L., LÜTHI, D., FREI, C., HÄBERLI, C., LINIGER, M. A. and APPENZELLER, C. (2004). The role of increasing temperature variability in European summer heat-waves. *Nature* **427** 332–336.
- SHABY, B. A., REICH, B. J., COOLEY, D. and KAUFMAN, C. G. (2015). Supplement to “A Markov-switching model for heat waves.” DOI:[10.1214/15-AOAS873SUPP](#).
- SIBUYA, M. (1959). Bivariate extreme statistics. *Ann. Inst. Statist. Math.* **11** 195–210.
- SMITH, R. L. (1985). Maximum likelihood estimation in a class of nonregular cases. *Biometrika* **72** 67–90. [MR0790201](#)
- SMITH, R. L., TAWN, J. A. and COLES, S. G. (1997). Markov chain models for threshold exceedances. *Biometrika* **84** 249–268. [MR1467045](#)
- SPIEGELHALTER, D. J., BEST, N. G., CARLIN, B. P. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 583–639. [MR1979380](#)

B. A. SHABY
DEPARTMENT OF STATISTICS
PENNSYLVANIA STATE UNIVERSITY
313 THOMAS BUILDING
UNIVERSITY PARK, PENNSYLVANIA 16802
USA
E-MAIL: bshaby@psu.edu

D. COOLEY
DEPARTMENT OF STATISTICS
COLORADO STATE UNIVERSITY
STATISTICS 217
FT. COLLINS, COLORADO 80523
USA
E-MAIL: cooley@stat.colostate.edu

B. J. REICH
DEPARTMENT OF STATISTICS
NORTH CAROLINA STATE UNIVERSITY
5146 SAS HALL, BOX 8203
RALEIGH, NORTH CAROLINA 27695
USA
E-MAIL: bjreich@ncsu.edu

C. G. KAUFMAN
DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA, BERKELEY
315 EVANS HALL
BERKELEY, CALIFORNIA 94720
USA
E-MAIL: cgk@stat.berkeley.edu