

Technical Note: Modeling Spatial Fields of Extreme Precipitation – A Hierarchical Bayesian Approach

Bianca Rahill-Marier¹, Naresh Devineni^{2*} and Upmanu Lall³

¹NCX, New York, NY.

5 ²Department of Civil Engineering, City University of New York (City College), New York, NY 10031.

³Department of Earth and Environmental Engineering, Columbia Water Center, Columbia University, New York, NY 10027.

Correspondence to: Naresh Devineni (ndeiveneni@ccny.cuny.edu)

Abstract. We introduce a hierarchical Bayesian model for the spatial distribution of rainfall corresponding to an extreme event of a specified duration that could be used with regional hydrologic models to perform a regional hydrologic risk analysis. An extreme event is defined if any gaging site in the watershed experiences an annual maximum rainfall event, and the spatial field of rainfall at all sites corresponding to that occurrence is modeled. Applications to data from New York City demonstrate the effectiveness of the model for providing spatial scenarios that could be used for simulating loadings into the urban drainage system. Insights as to the homogeneity in spatial rainfall and its implications for modeling are provided by considering partial pooling in the Hierarchical Bayesian framework.

15 1 Introduction

For an existing urban drainage network, a proper consideration of the spatial structure of extreme rainfall events is important for an assessment of the effectiveness of the network for handling urban flooding subsequent to rainfall events of varying duration, especially as concerns emerge as to the resilience of the system under a changing climate. Often, investigators focus on return period analysis of extreme rainfall at a site considering annual maxima or peaks over threshold for a specific rainfall duration. In a regional context, spatial models of annual maximum rainfall are sometimes considered (Renard et al., 2006; Renard and Lang, 2007; Dyrddal et al., 2015). However, since the annual maximum is unlikely to occur for a given event at all sites, these models do not represent the actual structure of potential extreme rainfall events. Thus, existing models for spatial rainfall extremes cannot be used to provide forcing for the performance of an existing drainage network (natural or constructed) under an extreme rainfall event. We address this situation in this note by considering that the rainfall events of interest for a specified duration are ones where any one of the sites in the region experiences an annual maximum event, and the spatial field of rainfall of interest is then the field associated with each such event.

In the exploratory analyses performed for New York City we noted that the structure of storms that lead to annual maximum events at different gages in the region may not be the same and that the basic statistics of rainfall vary across sites. We assemble a dataset of the annual maximum rainfall for each specified duration at each station. Let's denote this as A_{djt} , for

30 duration d , site j and year t . These events may not occur on the same day of each year across the stations. Second, we consider the rainfall field at all stations associated with the annual maximum at any one station and call this the “spatial field” (SF), R_{djki} , where i is identified as an event such that $R_{dkki} = A_{dkt}$ for site $k = j$ for year t . R_{djki} then has the rainfall at all sites j , for the event where site k has an annual maximum. As a result, the total number of events, i , may be much larger than the total number of years of data, N . A spatial event field of rainfall is thus conditional on the occurrence of an annual maximum rainfall for any station. Of interest is $f(R_{djki}|A_{dkt})f(A_{dkt})$, where $f(\cdot|.)$ and $f(\cdot)$ refer to a conditional and marginal probability distribution, respectively. For example, in 1979, the annual maximum 12-hour rainfall for Central Park rainfall gage is 3.29 inches. At this time of the day, eight other rainfall stations around New York City had 12-hour rainfall accumulations of 1.70 inches (Essex Fells), 2 inches (JF Kennedy International Airport), 2.73 inches (LaGuardia Airport), 0.79 inches (Newark International Airport), 1.68 inches (New Brunswick), 0.67 inches (New Milford), 0.98 inches (Staten Island), and 0.87 inches (Watchung), respectively. However, not all of those other accumulations are annual maximum events in their respective sites. Hence, $A_{12-CP-1979} = 3.29$ inches, and $R_{12-j-CP-i} = [3.29, 1.70, 2.00, 2.73, 0.79, 1.68, 0.67, 0.98, 0.87]$. This is similar to the issue noted by Asquith and Famigletti (2000). See Table A.1 in the appendix for a full example of the spatial rainfall fields in 1979. The hierarchical Bayesian models developed consider the spatial field SF with the goal of providing an approach for stochastically generating representative spatial fields of rainfall for a specified duration, such that at least one site in the region experiences an annual maximum event. This is fundamentally different from the traditional rainfall frequency analysis which models annual maximum data at each site independently or using a covariance structure from the annual maximum data at all sites. In this paper, we formulated and tested a simple model that could directly explore whether or not, and to what extent there was opportunity to pool regional information on extreme rainfall events to describe plausible spatial fields of extreme rainfall.

45

50 In Section 2, we present the data and the context for the application to the Greater New York area. In Section 3, we describe the details of the multivariate hierarchical Bayesian models. The results are discussed in Section 4. Finally, in Section 5, we present a summary and conclusions.

2 Data Description

55 2.1 Greater New York Area Context

The Greater New York City area has high density of man-made infrastructure and hence a complex hydrological landscape. Like many older cities, sanitary and industrial wastewater, and rainwater and street runoff are collected in the same sewers and conveyed together to treatment plants. Approximately 60 percent of NYC’s drainage area is served by these combined sewers. Flooding and combined sewer overflows (CSOs) are a concern, and innovative solutions for using the sewer system itself as flood storage by pumping water to different areas during a storm has been suggested. Such hydrologic system upgrades need to be informed by the spatial variability of extreme precipitation.

60

There are very few rain gage records that are longer than twenty or thirty years. Persistent data quality issues further reduce the available data. This challenge of reconciling sparse data with spatially variable hydrological networks and meteorological phenomena is common to many urban areas. Widely accepted design standards are derived from a set of intensity-duration-frequency curve developed using a daily rainfall record from 1903 to 1951 (NYCDEP 2008). An analysis of precipitation extremes for the region is offered in Wilks and Cember (1993), using daily rainfall data, and McKay and Wilks (1995), using hourly rainfall data. None of these analyses consider the spatial correlation of rainfall.

2.2 Precipitation Data

The precipitation data was obtained from the National Climatic Data Center (NCDC 2013). Rain gages were selected based on the proximity to New York City, data quality, and length of the historical record. Twenty-nine stations were initially identified as lying within a 100-mile radius of Central Park with over 25 years of continuous hourly rainfall records. Sixteen stations were excluded since the resulting data quality was too poor. The final dataset consists of the remaining nine stations (Table 1); abbreviations for each station used in the figures throughout are provided in the first column.

[Table 1 - in here]

75

2.3 Diagnostics and Spatial Concordance

The Shapiro-Wilks test (Shapiro and Wilk, 1965; Royston, 1992) was applied to the log-transformed annual maxima series at each station for each storm duration from 1-hour to 24-hours. For these 216 time series, the null hypothesis of the appropriateness of the Log Normal distribution was not rejected for 98% of the sites at the 1% level, 88% at the 5% level, and 81% at the 10% level. Though other distributions, such as the Generalized Extreme Value (GEV), Pearson, Log Pearson Type III (LP3) are popular for extreme precipitation modeling, in our application to the spatial field of rainfall, only one of the sites experiences the annual maximum event and others may not be extreme values. In such a setting, the Log Normal distribution can often be an appropriate representation (Raiford et al. 2007), as seen here. Consequently, to illustrate the idea, we consider a lognormal distribution with spatial correlation across the sites for the NYC example. Other choices could very well be made.

A heat map showing the fraction of annual maximums that occur simultaneously is provided in Figure 1. For these plots, we define simultaneous storms to be those beginning within $\pm n$ hours of each other (where n is a multiple of the event duration) to allow for the movement of a storm event over the area, and to identify distinct, independent rainfall events. We see that precipitation extremes, even within a relatively local area, are frequently not simultaneous. As expected, the simultaneous fraction of concurrent area increases as the storm duration increases. However, even for a 24-hour duration less than 60 percent of the events are concurrent. This is true even for the four closest stations – JFK, LGA, Central Park and Staten Island that are typically used to inform hydrologic design in New York City.

90

[Figure 1 - in here]

This diagnostic analysis highlights the importance of considering the spatial structure of extreme rainfall for an event with a
95 specified duration.

3 Methodology

A hierarchical Bayesian approach that provides the ability to partially pool model parameters across the rain gage sites was developed. Full pooling would imply that a parameter (e.g., the mean, variance etc. of the distribution) was homogeneous across the sites. No pooling would imply that each site is independent. Partial pooling is an intermediate step that allows
100 information to be shared across sites at a level informed by the data. This results in a multi-level model, where model parameters are estimated at each site, but are assumed to be drawn from the parameters of distributions that are specified at the regional level for each parameter (Gelman and Hill 2007). Such an approach has been implemented for hydrometeorological extremes in Lima and Lall (2010) and Kwon et al. (2008), and for paleoclimate reconstructions by Devineni et al (2013).

105 3.1 Spatial Fields Hierarchical Model conditioned on the site experiencing an annual maximum

In this model, we consider a conditional process, where site k has experienced an annual maximum event, and the corresponding rainfall amounts, $R_{d|k}$ at all sites are observed. The logarithm of rainfall is considered to be Normally distributed, and a multivariate Normal distribution is specified for each site k , where an annual maximum has occurred. For each such condition, we consider partial pooling of the mean rainfall across all sites, and consider the spatial covariance
110 across sites. We consider that the spatial field of rainfall may actually be different depending on which site experiences an annual maximum. The hierarchical model is described as below.

$$\begin{aligned} \mathbf{Y}_k &\sim MVN(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \\ \mu_{kj} &\sim N(\omega_k, \sigma_k^2) \\ &\text{Priors} \end{aligned}$$

$$\begin{aligned} 115 \quad \boldsymbol{\Sigma}_k &\sim \text{Inv} - \text{Wishart}(\Lambda, \mathbf{v}) \\ \omega_k &\sim N(0, 1000) \\ \sigma_k^2 &\sim U(0, 100) \end{aligned} \tag{1}$$

\mathbf{Y}_k is the log of the rainfall field $R_{d|k}$ across all sites corresponding to when station k has an annual maximum. For the New York City application, it is a matrix of 64 (years) by 9 (stations) for a given duration, and station k . \mathbf{Y}_k is assumed to follow a
120 multivariate normal distribution with a vector of station means $\boldsymbol{\mu}_k$ and covariance across stations specified by a 9-by-9

matrix Σ_k . At the second level of the model, the station-specific means μ_j are assumed to be Normally distributed with a common mean ω_k and variance σ_k^2 . This is a partial pooling approach with no covariates, as outlined in Gelman and Hill (2007). A non-informative conjugate prior, the inverse-Wishart distribution, is assumed for Σ_k where Λ is the scale matrix and \mathbf{v} is the degrees of freedom (Gelman et al. 2004). If Σ_k is a j -by- j matrix, we assume \mathbf{v} equivalent to $(j + 1)$ and Λ equal to the j -by- j identity matrix (\mathbf{I}). This is equivalent to a uniform prior on each variance element of the correlation matrix (Gelman and Hill 2007). We give σ_k^2 a non-informative uniform prior, and ω_k a non-informative conjugate normal prior, for computational convenience (Gelman and Hill 2007; Gelman et al. 2004).

There are nine stations, and therefore there are nine distinct datasets \mathbf{Y}_k and nine distinct models for each storm duration. For extreme rainfall events, i.e., those that exceed a nominal design return period, we outline a simulation strategy from these models that pools simulated fields together that represent regional extreme events.

3.2 Spatial Fields Single-Level Model

We consider a subset of the previous model where the assumption that the mean log-rainfall is drawn from a common spatial mean is relaxed. This leads to the simpler, no-pooling model represented below.

$$\begin{aligned} \mathbf{Y}_k &\sim MVN(\boldsymbol{\mu}_{ks}, \boldsymbol{\Sigma}_{ks}) \\ \mu_{kjs} &\sim N(0, 1000) \end{aligned} \quad (2)$$

$$\boldsymbol{\Sigma}_{ks} \sim \text{Inv} - \text{Wishart}(\Lambda, \mathbf{v})$$

As in the hierarchical model, \mathbf{Y}_k is the log of the SF when station k is at an annual maximum. The vector of precipitation means across j stations (including station k) is $\boldsymbol{\mu}_{ks}$, with a subscript s to indicate single-level model.

3.3 Spatial Fields Simulation for a regional T-year return period

The Spatial Fields model can be used to simulate rainfall fields corresponding to an annual maximum occurring at any one of the sites, k . Next, if we are interested in design rainfall fields represented by the T -year return period across the domain we follow a two-step process. First, based on the model that is fit, we identify the T year return period annual maximum rainfall event for each site, k . Then from simulations of the multivariate rainfall fields using the model we identify all cases where the rainfall at site k exceeds the T -year event for that site, and take the corresponding simulated rainfall field across all sites, j . The process is outlined below:

- i. *Threshold Calculation:* For each return period (T) and rainfall duration (D) a precipitation threshold is computed for each station using the posterior mean and variance from the station k 's hierarchical Bayesian model. The threshold was computed using frequency factors K for the normal distribution (Guo 2006) and the equation below.

150
$$\log(Y_{T,dk}) = \widetilde{\mu}_k + K(T) * \widetilde{\Sigma}_{kk} \quad (3)$$

For example, letting $k = 1$ for Central Park, we compute the SF model from \mathbf{Y}_1 . We extract μ_{11} , the Central Park mean and Σ_{11} , the Central Park variance and use them in equation 3 above.

- 155 ii. *Simulate Multivariate field for \mathbf{Y}_k* : From the hierarchical Bayesian model defined in (1) simulate a large number of realizations M (e.g., equal to 10,000), of the rainfall fields \mathbf{Y}_k corresponding to the case when site k has an annual maximum. These are based on draws from the posterior distributions of the parameters, and hence incorporate a consideration of parameter uncertainty.
- iii. *Extract Subset of Simulations that exceed the T -Year event at site k* : Retain a subfield \mathbf{Z}_k from \mathbf{Y}_k , such that $Y_{dkkm} > Y_{T,dk}$, and $m = 1 \dots M$, is the index of the simulation.
- 160 iv. *Pool the T -year return period fields*: The \mathbf{Z}_k are subsamples of rainfall fields from each of the nine models, such that an equal number of draws from each of the k fields is selected. For $T=100$ years, on average 100 such samples will be generated for $M=10000$, from each station, and 900 total fields are then available for our application to the New York City data for design or reliability analyses. For our illustrations here, we sampled the same number of fields that are obtained from applying steps (i) to (iv) on the observed rainfall fields data. In essence, the spatial fields corresponding to a T -year return period event are first derived from observed rainfall fields data; these T -year return period spatial rainfall fields then form the baseline to which the simulated T -year return period spatial rainfall fields derived from the posterior runs of the hierarchical model are compared to. Note that since there may be multiple sites with annual maxima per event i in the original R_{djk_i} data, and that these are contained in each random field indexed by k , and we modeled this spatial field, the concurrence of high rainfall at those sites will also be reproduced in the simulations. Similarly, the incidence of high rainfall at multiple stations will also be correctly reproduced across the pooled data across the K simulations.
- 165
- 170

3.4 Model Fitting and Convergence

Two hundred and sixteen models (one for each duration and each site) were fit using JAGS (Just Another Gibbs Sampler) (Plummer et al., 2003; Denwood et al., 2016). It uses a Markov chain Monte Carlo (MCMC) simulation algorithm (a Gibbs Sampler for the current example), to simulate the posterior probability distribution of parameters. A random normal distribution was used for vector of station means ($\boldsymbol{\mu}, \boldsymbol{\mu}_k, \boldsymbol{\mu}_{ks}$) and a random Wishart distribution was used for the precision. Similarly, a non-informative conjugate normal prior for ω_k and a non-informative uniform prior for σ_k^2 are assumed. In JAGS, the normal distribution is parameterized in terms of precision instead of covariance ($\boldsymbol{\Sigma}, \boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}_{ks}$) as is noted by convention in the model formulas above. We simulated four chains, ran the model for 20,000 iterations and the first half of

175

180 the simulations were discarded as burn-in.

4 Results and Analysis

4.1 Bayesian Model Checking

For each model, the convergence of the posterior distribution of each parameter was checked using the shrink factor proposed by Gelman and Rubin (1992) - values under 1.1 for all parameters suggest that the model has converged. 185 Convergence plots (showing the mixing of the four chains) were visually checked for all cases. All models converged appropriately with each parameter attaining a shrink factor between 1.0 and 1.1, and the large majority reaching 1.0.

We compare the performance of the two SF models using the Deviance Information Criterion (DIC) and pD, recommended *in* Gelman et al. (2004). The scores were virtually identical for the two types of SF models for each rainfall duration. Next, we considered whether the common mean in the hierarchical model converged as successfully as other parameters; it did. 190 Gelman and Hill (2007) suggest that when there are only a small number of groups and the group-level standard deviation is large, multi-level modeling may not add much information. The resulting model will not necessarily perform worse and will likely resemble the model without pooling (as it does here). The posterior parameters for the resulting simulations are essentially identical (Figure 2a) and the posterior for the mean only very slightly shrunk (Figure 2b).

[Figure 2 - in here]

195 Next, we consider how the return period spatial rainfall fields identified in the SF hierarchical models compare to observed return period spatial fields. We do so by plotting the empirical cumulative distribution function of the return period event field estimated from the hierarchical model and comparing it with the empirical cumulative distribution function of the return period event field derived from the observed data. The 99% credible interval obtained from the posterior simulations of the hierarchical models is also presented to represent the uncertainty. Though the Bayesian models can easily simulate any 200 return period, the reliability of the empirical estimate is dependent on the length of record so it is only reliable as a goodness of fit measure for shorter return periods. Results from all nine stations for the 10-year 12 hour accumulated rainfall are presented in Figure 3. The plots for all nine stations at the 1-hr and 24-hr durations and ten-year return period are provided in Figure A1 and Figure A2 respectively, of the appendix.

[Figure 3 - in here]

205 The observed 10-year 12 hour return period fields for all stations are within the 99% credible interval of the simulations, for the most part. The SF model simulated tails that are much greater in magnitude than the observed, as expected. The uncertainty band is also wider at the tail end of the distribution. New Milford is an exception with slight underprediction of the left tail. These differences are amplified at the left tail for higher durations (i.e., the 24 hour storms). The shorter durations (1 hour storms) seem to be modeled well, albeit with wider credible intervals.

210 Figure 4 presents a slightly different view of the results with a focus on how well various durations for a shorter return
period event (5-year events) are simulated for a specific station – Staten Island. While 1-hour 5-year return period storm
fields are well simulated, there seems to be a bias (with the observed field’s left tail falling outside the 99% credible interval
of the simulated) with increasing durations. It’s difficult to identify exactly why this might be the case without significant
additional exploratory analysis of the Staten Island data. However, we note that the simulations reflect additional information
215 provided by the other stations in the model. In this application the Staten Island site record is twenty years shorter than the
record at the other sites, and the shift in the simulations reflects the shift in the rainfall in the other sites over this period.
Thus, the pooling strategy is instrumental in using the joint distribution across sites from the common period of record to
provide simulations that cover the shift in the period. Of course, if the correlation across sites has also changed over the
period of missing data, the algorithm is incapable of replicating such a change.

220 **[Figure 4 - in here]**

5 Summary and Conclusions

For larger cities, a consideration of the drainage network, and the spatial dependence in rainfall at different durations is
important to consider, at least from the perspective of assessing the performance and resilience of the network, and perhaps
also for design considerations. We were interested in formulating and testing a simple model that could directly explore
225 whether or not, and to what extent there was opportunity to pool regional information on extreme rainfall events to describe
plausible spatial fields of extreme rainfall. This led to postulating and testing a Bayesian model that considers the spatial
field of rainfall associated with an annual maximum occurrence at any site. We considered the application of model to
relatively long rainfall time series from the New York City region. Initial exploratory analyses suggested that the rainfall
characteristics and storm tracks varied by event and by season across the region, such that distinct clusters could be
230 identified, suggesting that the region had a heterogeneous spatial structure with respect to extreme rainfall (Hamidi et al.
2017). Our applications further clarified the nature of this heterogeneity. It is interesting to also note from the New York City
analysis that there is support for pooling the spatial covariance of rainfall across all sites (irrespective of which one
experienced an annual maximum rain event for a given duration), even though often the exceedance probability distributions
of rainfall for a given duration may differ across sites, even after partial pooling. The hierarchical Bayesian framework
235 permits a consideration of the uncertainty in parameter and model structure and helps us identify the level of homogeneity
that may be appropriate for representing the processes underlying a particular data set.

Rain gages are the preferred data source for extreme event modeling because of their long-record, but incorporating radar in
addition to rain gages could provide the spatial density needed to explore how event rainfall characteristics relate to specific
meteorological phenomena or to provide comparable simulations to existing stochastic models. The radar information would

240 contain considerably more spatial detail necessary for building the type of model exemplified here. However, radar rainfall records are much shorter, and consequently, one needs to develop a methodology to appropriately blend the shorter but spatially richer radar data with the longer but spatially sparse gage data. Our algorithm can be readily applied to a mix of radar and rain gage data. However, some extensions need to be pursued to address the very different record lengths of each data source.

245 We used a Log Normal distribution applied to rainfall for each duration, to illustrate our approach. The goodness of fit tests supports this assumption, and this permits some confidence in the kind of conclusions we drew from the applications to the New York City data. However, other models such as the GEV or Generalized Pareto or other choices for the distribution could very well be considered. The point here was to highlight the need to consider spatial covariance and an appropriate blending of local and regional data sources through partial pooling.

250 **Appendix**

Auxiliary Figures

This appendix includes Table A1, Figure A1, and Figure A2. Table A1 provides an example derivation of extreme rainfall fields. Figure A1 and A2 provide the results and baseline comparisons of the 1-hour and 24-hour ten year return period events from the SF models for the nine stations in and around New York City.

255 **Code Availability**

The code for conducting the analysis presented in this paper can be made available upon request.

Data Availability

260 Rainfall data for the analysis of the nine sites in and around New York City can be obtained from the public source provided in the references, National Climate Data Center, <https://www.ncdc.noaa.gov/data-access/quick-links>. The authors can be contacted for any details on the methodology.

Author Contributions

ND and UL designed the study and wrote and edited the manuscript. BRM and ND conducted the analysis. BRM wrote the original draft of the paper.

Competing Interests

265 The authors declare that they have no competing interests.

Financial Support

This research is supported by the Department of Energy Early CAREER Award No. DE-SC0018124 for the second author and National Science Foundation, Water Sustainability and Climate (WSC) program – award number: 1360446. Partial support was also provided by the American International Group (AIG) under the “Climate Informed Global Flood Risk
270 Assessment” project.

References

- Asquith, W.H., and J. S. Famiglietti.: Precipitation areal-reduction factor estimation using an annual-maxima centered approach, *Journal of Hydrology*, 230, 55-69, 2000.
- Denwood, M. J., et al. (2016). runjags: An R package providing interface utilities, model templates, parallel computing
275 methods and additional distributions for MCMC models in JAGS. *Journal of Statistical Software*, 71(9), 1–25.
- Devineni, N., Lall, U., Pederson, N., & Cook, E.: A tree-ring-based reconstruction of Delaware River basin streamflow using hierarchical Bayesian regression. *Journal of Climate*, 26(12), 4357-4374, 2013.
- Dyrddal, A. V., Lenkoski, A., Thorarinsdottir, T. L., and Stordal, F.: Bayesian hierarchical modeling of extreme hourly precipitation in Norway, *Environmetrics*, 26, pages 89-106, 2015.
- 280 Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin.: *Bayesian Data Analysis*, Chapman & Hall/CRC, 2004.
- Gelman, A., and J. Hill.: *Data Analysis Using Regression and Multilevel/Hierarchical Models*, Cambridge University Press, 2007.
- Gelman, A., and D. B. Rubin.: Inference from iterative simulation using multiple sequences (with discussion), *Statistical Science*, 7, 457-511, 1992.
- 285 Guo, J. C. Y.: *Urban Hydrology and Hydraulic Design*. Chapter 4: Frequency Analysis, Water Resources Publication, LLC, 2006.
- Hamidi, A., N. Devineni, J.F. Booth, A. Hosten, R.R. Ferraro, and R. Khanbilvardi.: Classifying Urban Rainfall Extremes Using Weather Radar Data: An Application to the Greater New York Area, *J. Hydrometeor.*, 18, 611–623, 2017.
- Kwon, H.H., C. Brown, and U. Lall.: Climate Informed flood frequency analysis and prediction in Montana using
290 hierarchical Bayesian modeling, *Geophysics Research Letters*, 35, 1-6, 2008.
- Lima, C.H.R, and U. Lall.: Spatial scaling in a changing climate: A hierarchical bayesian model for non-stationary multi-site annual maximum and monthly streamflow, *Journal of Hydrology*, 383, 307-318, 2010.

- Lunn, D. J., A. Thomas, N. Best, and D. Spiegelhalter.: WinBUGS—A Bayesian modelling framework: Concepts, structure, and extensibility. *Stat. Comput.*, 10, 325–337, 2000.
- 295 McKay, M., and D. S. Wilks.: Atlas of Short-Duration Precipitation Extremes for the Northeastern United States and Southeastern Canada, Northeast Regional Climate Center Research Publication RR 95, 26p, 1995.
- National Climate Data Center, <https://www.ncdc.noaa.gov/data-access/quick-links>. Accessed July 2013.
- Plummer, M., et al. (2003). Jags: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing (Vol. 124)*.
- 300 Raiford, J.P., N. M. Aziz, A. A. Khan, and D. N. Powell.: Rainfall Depth-Duration-Frequency Relationships for South Carolina, North Carolina, and Georgia, *American Journal of Environmental Science*, 3, 78-84, 2007.
- Renard, B., V. Garreta, and M. Lang.: An application of Bayesian analysis and Markov Chain Monte Carlo methods to the estimation of a regional trend in annual maxima, *Water Resources Research*, 42, 17p, 2006.
- Renard, B., M. Lang.: Use of a Gaussian copula for multivariate extreme value analysis: Some case studies in hydrology, *Advances in Water Resources*, Volume 30, Issue 4, Pages 897-912, 2007.
- 305 Royston, J.P.: Approximating the Shapiro-Wilk W-Test for Non-Normality. *Statistics and Computing* 2, 117-119, 1992.
- Shapiro, S.S., and M.B. Wilk.: An Analysis of Variance Test for Normality (Complete Samples). *Biometrika* 52, 591-611, 1965.
- Spiegelhalter, D., A. Thomas, N. Best, and W. Gilks.: BUGS 0.5: Bayesian inference using Gibbs sampling manual (version ii). Medical Research Council Biostatistics Unit Manual, 59 pp, 1996.
- 310 The New York City Department of Environmental Protection (NYC DEP): The NYC DEP Climate Change Program Assessment and Action Plan: Report 1. Chapter 2: Potential Climate Change Impacts on the DEP, 32-46, 2008.
- Wilks, D.S., and R. P. Cember.: Atlas of Precipitation Extremes for the Northeastern United States and Southeastern Canada, Northeast Regional Climate Center Research Publication, RR 93-5, 40p, 1993.
- 315

320

Table 1: Rain gage stations in New York City and surroundings

Abb.	Location	Latitude	Longitude	Elevation (ft)	Start	End
CP	New York Central Park Observation Belvedere Tower, NY	40.66889	-73.9602	39.6	5/1/1948	7/29/2012
EF	Essex Fells Service Building, NJ	40.8314	-74.2858	106.7	7/4/1949	8/1/2012
JFK	New York JF Kennedy International Airport, NY	40.63861	-73.7622	3.4	1/1/1949	7/29/2012
LGA	New York LaGuardia Airport	40.77944	-73.8803	3.4	5/1/1948	7/29/2012
NB	New Brunswick 3 SE, NJ	40.4719	-74.4365	26.2	6/1/1968	2/1/2006
NW	Newark International Airport, NJ	40.6825	-74.1694	2.1	5/1/1948	7/29/2012
NM	New Milford, NJ	40.961	-74.015	3.7	5/31/1946	6/30/1980
SI	New York Westerleigh, NY (Staten Island)	40.63333	-74.1167	24.4	5/1/1948	9/1/1992
WT	Watchung, NJ	40.66222	-74.4164	79.2	6/1/1948	8/1/2012

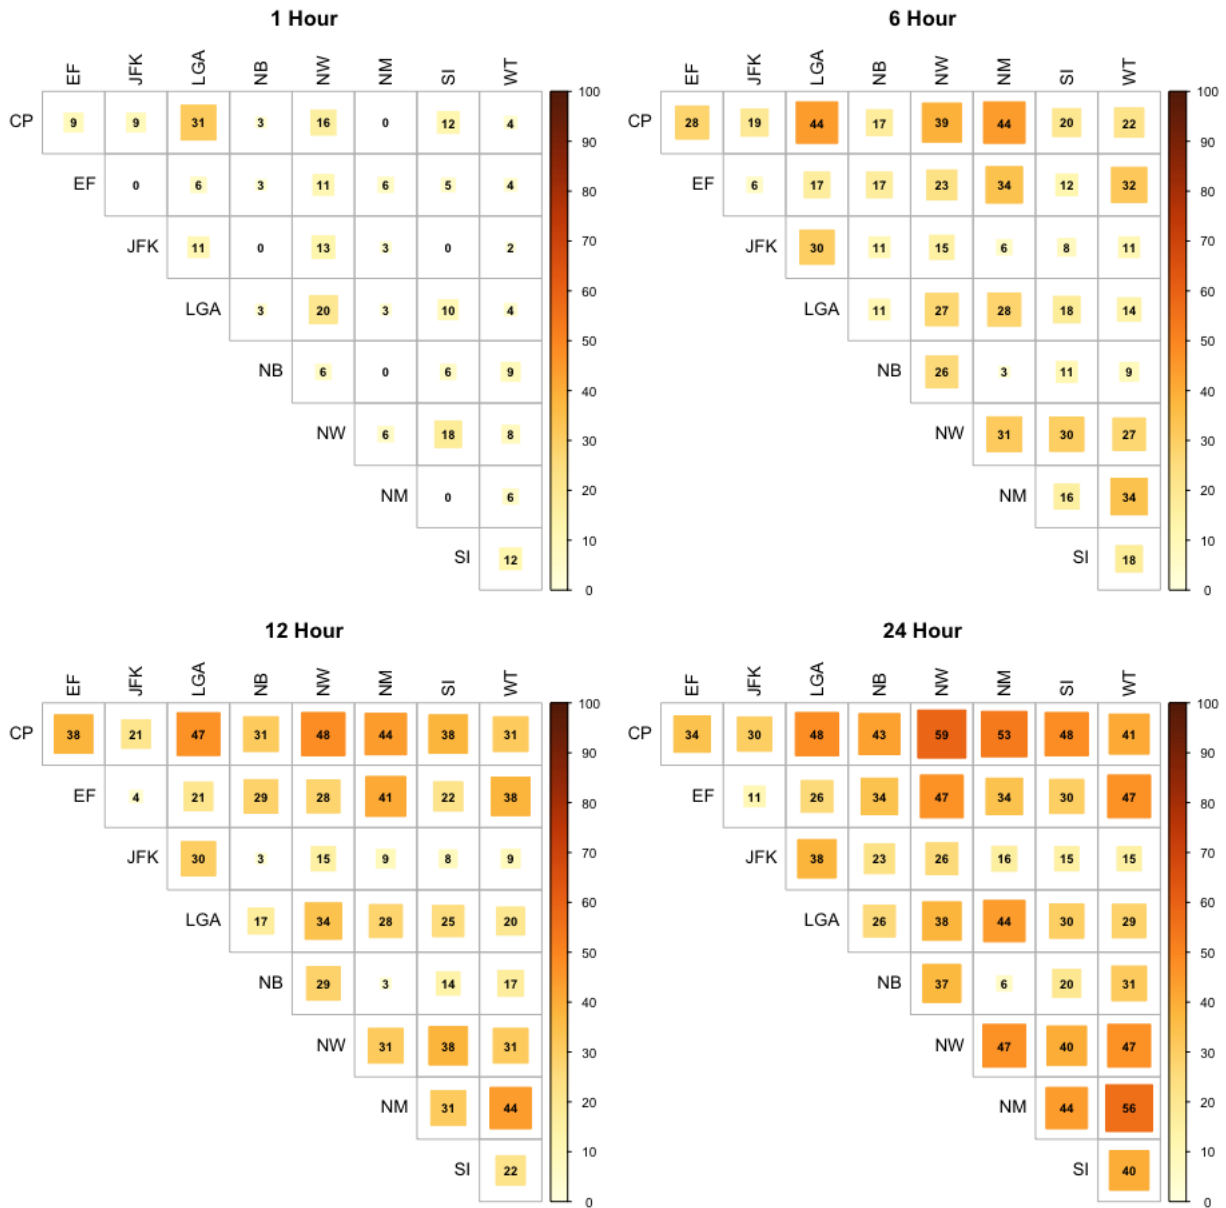


Figure 1: Percent of simultaneous or near-simultaneous annual maxima events shown for the site-by-site comparison for nine sites and 1-hr, 6-hr, 12-hr, and 24-hr storms.

330

335

340

Figure 2: Empirical cumulative distribution function of the posterior distribution of (a) simulations and (b) mean parameters for Central Park 12-hour hierarchical and non-hierarchical SF model. Posterior means and simulations are shown on an untransformed scale (i.e., the mean is log mean). The Empirical cumulative distribution function derived from the observed data is also shown in 2(a) along with the band that depicts the 99% credible interval of the posterior simulations.

345

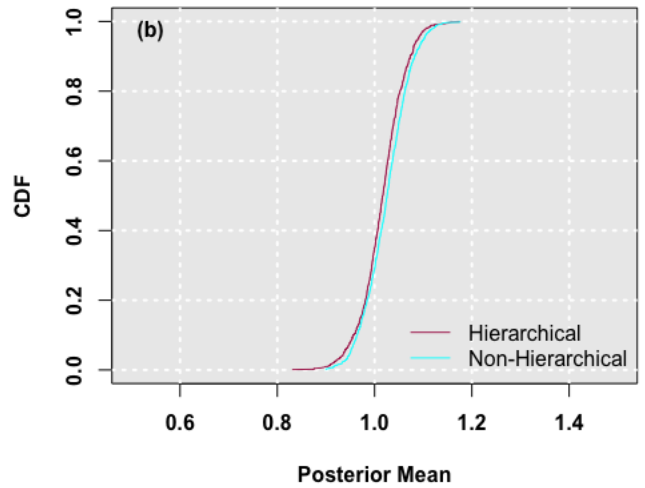
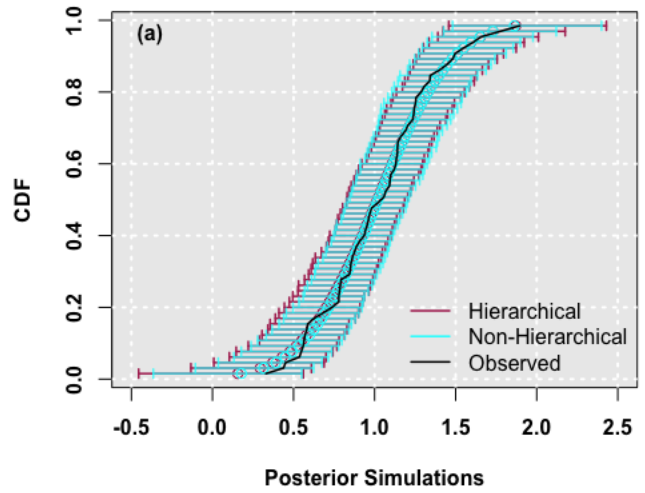
350

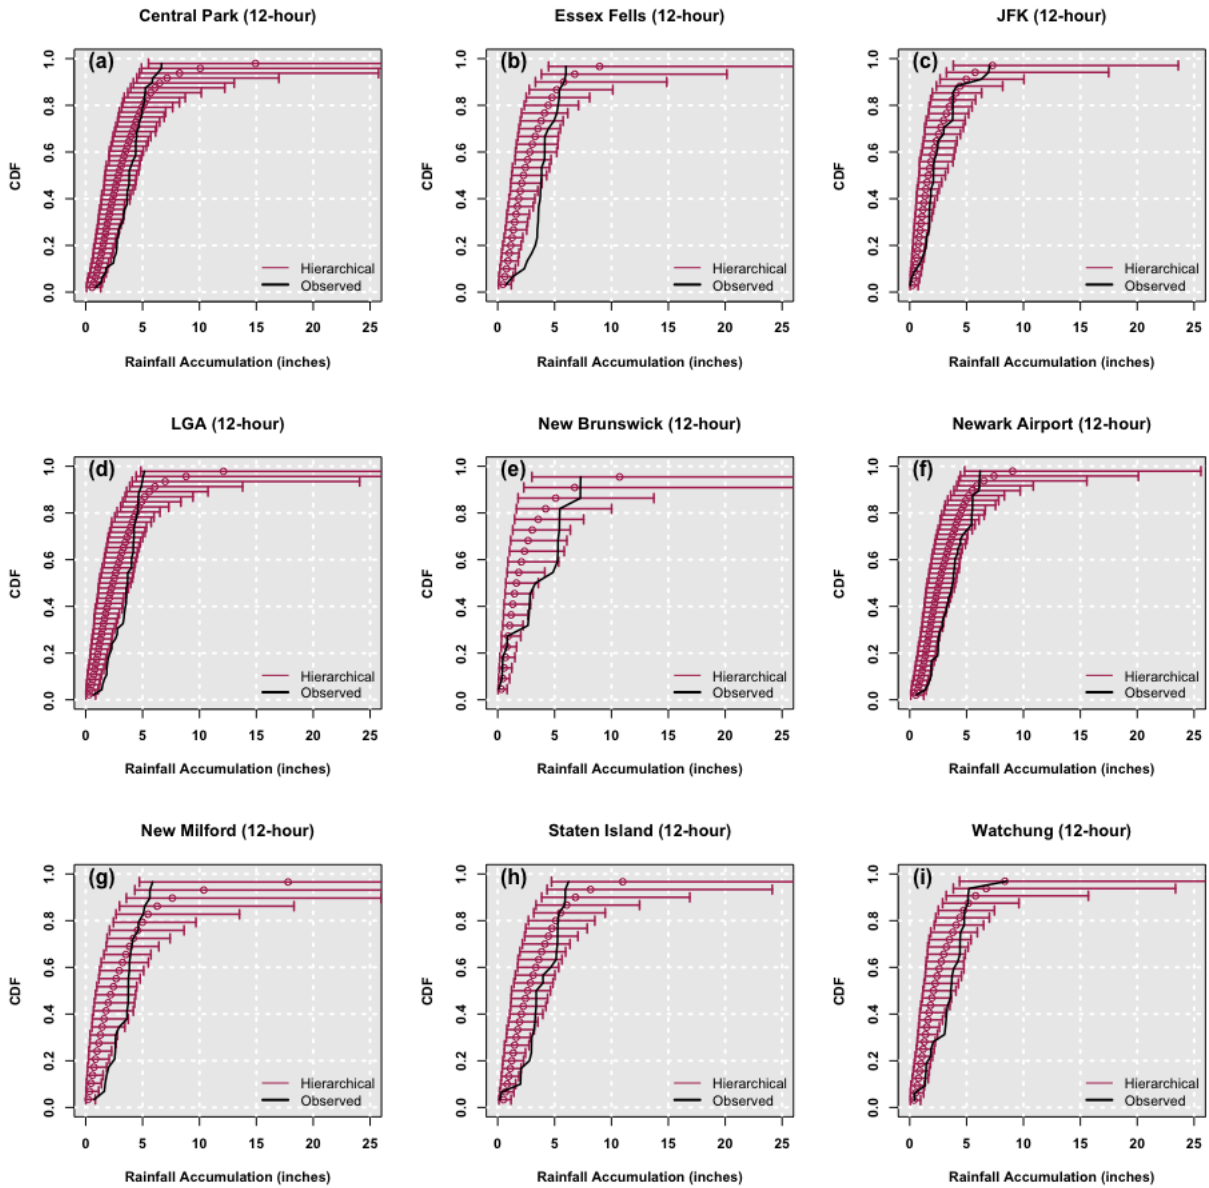
355

360

365

370

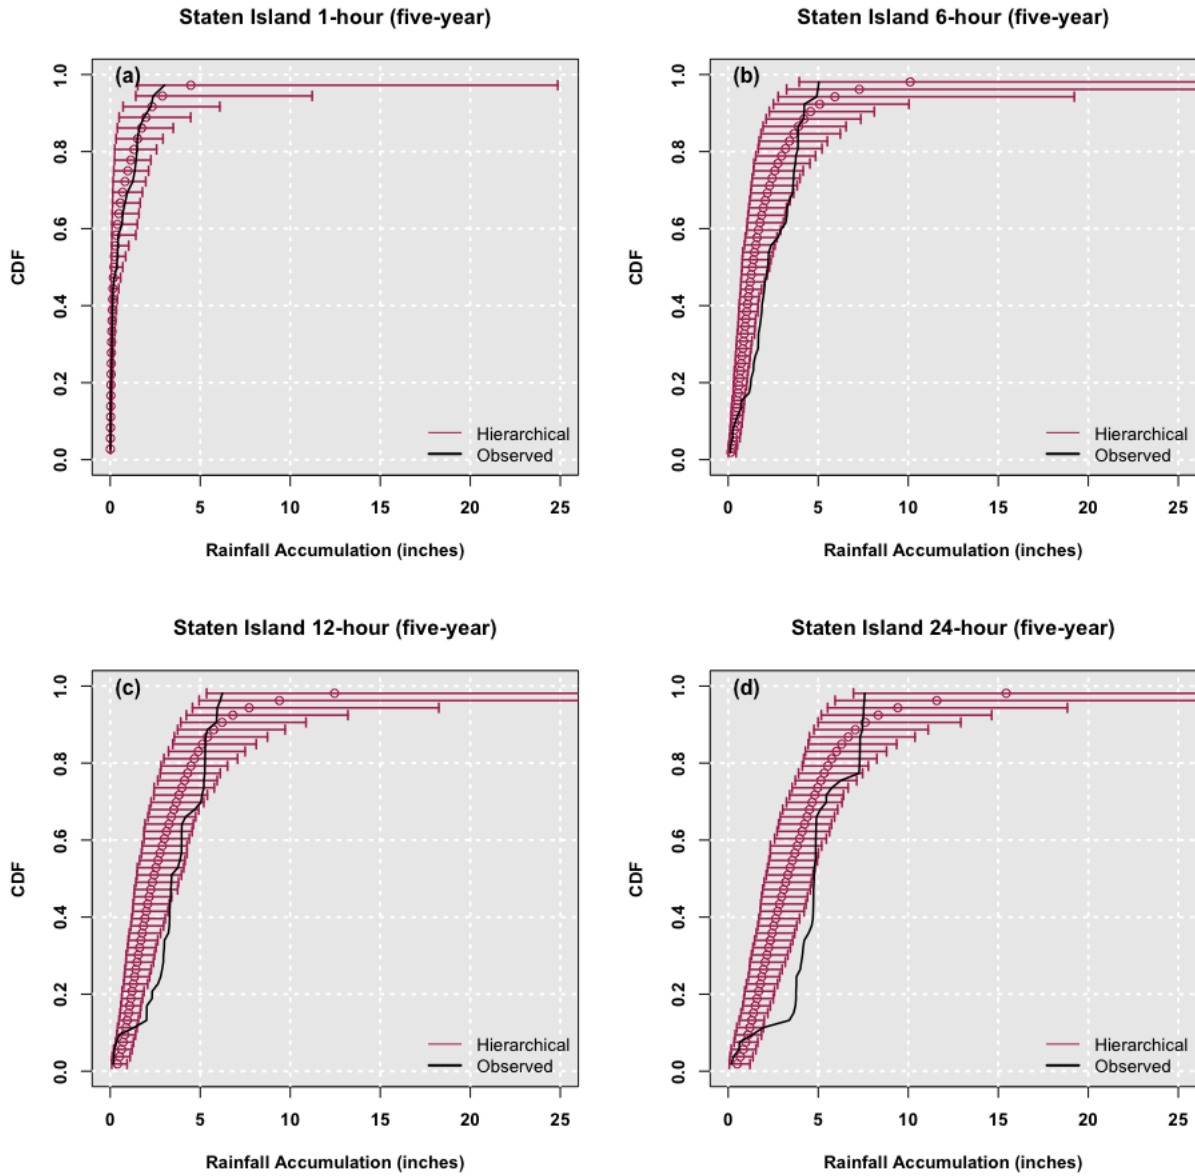




375

Figure 3: Empirical cumulative distribution function of the 12-hr ten-year return period event from SF hierarchical models for the nine stations (Central Park, LGA, New Milford – top to bottom left panel; Essex Fells, New Brunswick, Staten Island - top to bottom, middle panel; JFK, Newark, Watchung - top to bottom, right panel). The Empirical cumulative distribution function derived from the observed 12-hr ten-year return period field is also shown. The band depicts the 99% credible interval of the posterior simulations.

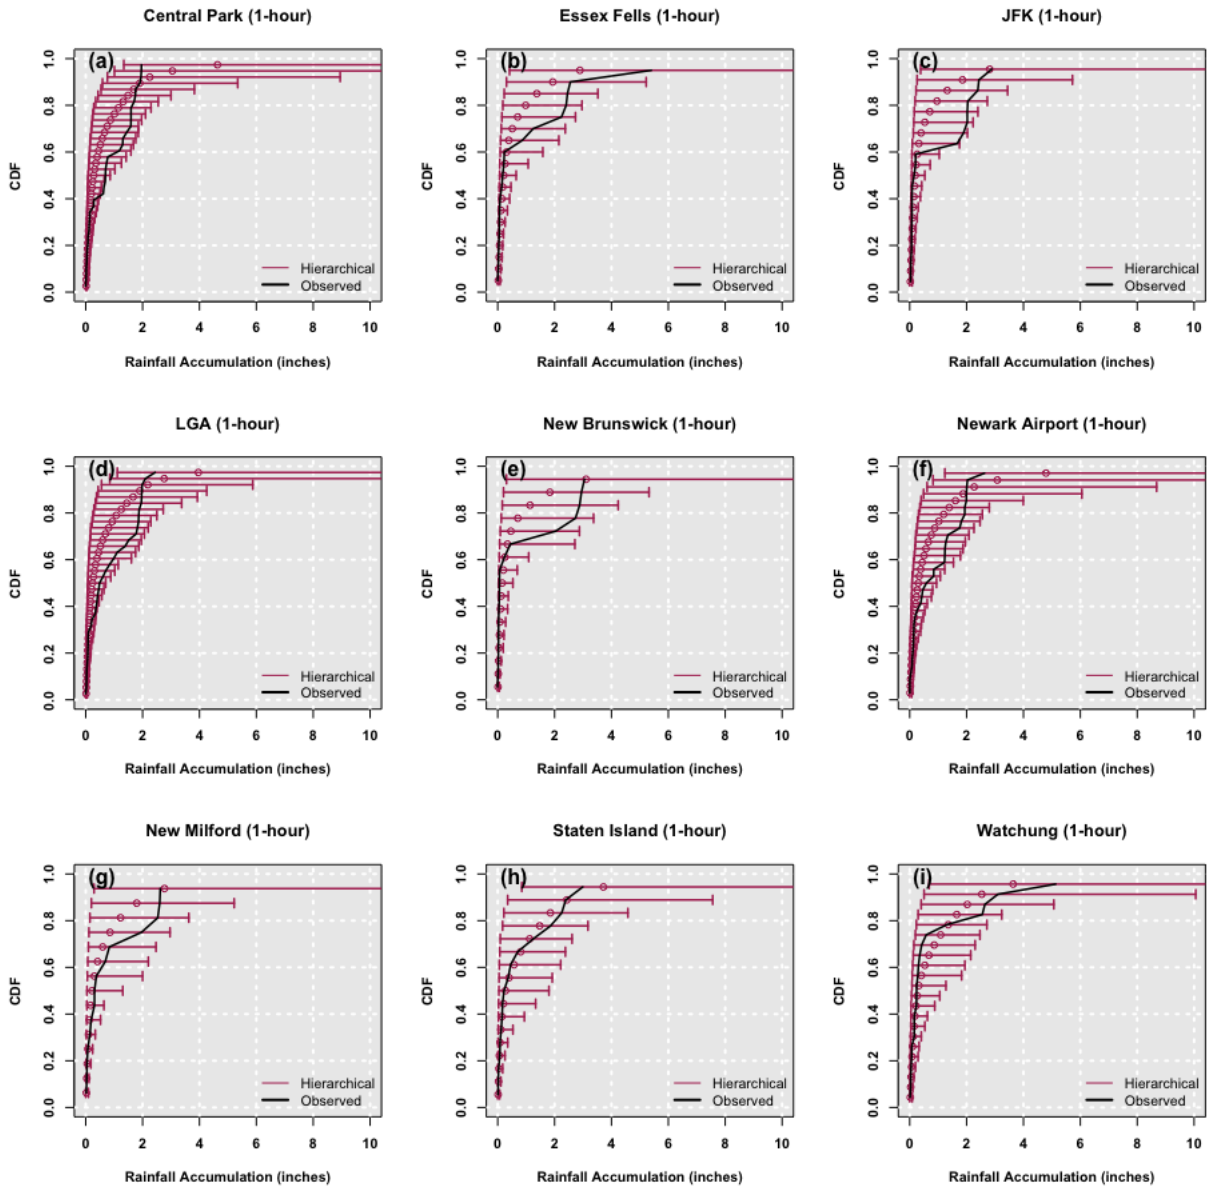
380



385 **Figure 4:** Empirical cumulative distribution function of 5-year return period events for 1-hour (top left), 6-hour (top right), 12-hour (bottom left), and 24-hour (bottom right) storms from Staten Island SF hierarchical model simulations. The empirical cumulative distribution function of the observed 5-year return period event fields at the respective durations are also shown. The band depicts the 99% credible interval of the posterior simulations.

390 **Table A.1: Illustration of spatial rainfall fields for an example year, 1979 for 12-hour accumulated rainfall. Anchor**
station column is the condition where annual maximum rainfall is identified (A_{dkt}). The corresponding row presents
the rainfall for other stations when the anchor station is experiencing an annual maximum event ($R_{djk|A_{dkt}}$). For
instance, the third row shows the spatial rainfall field when JF Kennedy International Airport has an annual
maximum event in 1979 (2.03 inches of 12-hour rainfall is shown in bold font to indicate the annual maximum rainfall
395 **at the conditioning station). The 12-hour rainfall recorded at the other stations simultaneously are presented across,**
in the row. Notice that none of those other stations' recorded rainfall is the at-site annual maximum event.

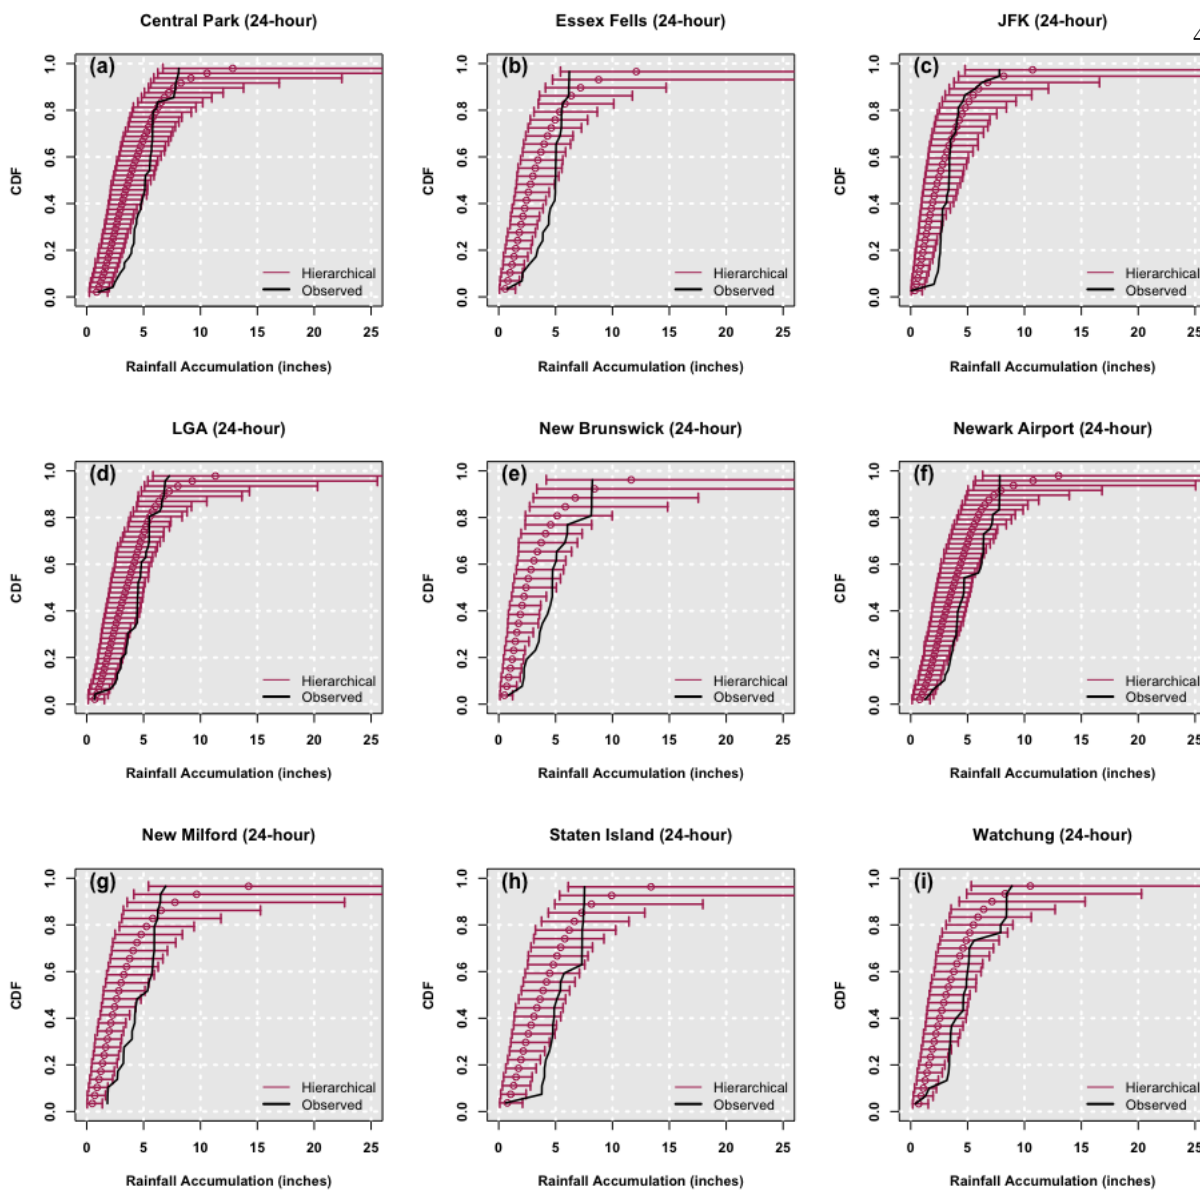
		Other Stations that form the Spatial Field (R_{djk})								
		CP	EF	JFK	LGA	NW	NB	NM	SI	WT
Anchor Stations (A_{dkt})	CP	3.29	1.70	2.00	2.73	0.79	1.68	0.67	0.98	0.87
	EF	0.56	2.56	NA	0.13	0.49	0.46	1.02	0.31	0.48
	JFK	3.25	1.70	2.03	2.71	0.79	1.68	0.65	0.89	0.88
	LGA	3.29	1.70	2.00	2.73	0.79	1.68	0.67	0.98	0.87
	NW	1.44	1.22	1.48	1.23	2.01	1.92	1.30	2.20	1.73
	NB	1.27	0.94	0.57	1.04	0.22	2.61	0.40	0.86	1.81
	NM	2.29	1.03	1.52	1.77	0.44	1.99	2.41	1.76	0.77
	SI	0.65	2.08	0.58	0.32	NA	1.46	0.23	3.66	1.62
	WT	1.84	1.32	1.64	1.56	1.72	2.07	1.62	2.03	1.92



400

Figure A1: Empirical cumulative distribution function of the 1-hr ten-year return period event from SF hierarchical models for the nine stations (Central Park, LGA, New Milford – top to bottom left panel; Essex Fells, New Brunswick, Staten Island - top to bottom, middle panel; JFK, Newark, Watchung - top to bottom, right panel). The Empirical cumulative distribution function derived from the observed 1-hr ten-year return period field is also shown. The band depicts the 99% credible interval of the posterior simulations.

405



415 **Figure A2:** Empirical cumulative distribution function of the 24-hr ten-year return period event from SF hierarchical models for the nine stations (Central Park, LGA, New Milford – top to bottom left panel; Essex Fells, New Brunswick, Staten Island - top to bottom, middle panel; JFK, Newark, Watchung - top to bottom, right panel). The Empirical cumulative distribution function derived from the observed 24-hr ten-year return period field is also shown. The band depicts the 99% credible interval of the posterior simulations.