Progress in Probability, Vol. 60, 185–210
© 2008 Birkhäuser Verlag Basel/Switzerland

A Hierarchical Bayes Model for Combining Precipitation Measurements from Different Sources

Tamre P. Cardoso and Peter Guttorp

Abstract. Surface rain rate is an important climatic variable and many entities are interested in obtaining accurate rain rate estimates. Rain rate, however, cannot be measured directly by currently available instrumentation. A hierarchical Bayes model is used as the framework for estimating rain rate parameters through time, conditional on observations from multiple instruments such as rain gauges, ground radars, and distrometers. The hierarchical model incorporates relationships between physical rainfall processes and collected data. A key feature of this model is the evolution of drop-size distributions (DSD) as a hidden process. An unobserved DSD is modeled as two independent component processes: 1) an AR(1) time-varying mean with GARCH errors for the total number of drops evolving through time, and 2) a time-varying lognormal distribution for the size of drops. From the modeled DSDs, precipitation parameters of interest, including rain rate, are calculated along with associated uncertainty. This model formulation deviates from the common notion of rain gauges as "ground truth"; rather, information from the various precipitation measurements is incorporated into the parameter estimates and the estimate of the hidden process. The model is implemented using Markov chain Monte Carlo methods.

Mathematics Subject Classification (2000). 62P12, 62M09, 86A10.

Keywords. Hierarchical Bayes, MCMC, Precipitation, Rain rate, Drop-size distribution.

1. Introduction

Surface rainfall is an important environmental variable that is incorporated across many areas of study, including meteorology, climatology, agriculture, land use, and hydrology. These various fields of study require estimation of precipitation at a range of temporal and spatial scales. Hydrologists and land use planners may be interested in short-term rainfall in relatively small regional areas for studies

involving flood and flash flood forecasting. Researchers in climatology or agriculture may be interested in climatic studies that focus on weekly, monthly, or annual totals over large spatial extents.

Although of considerable interest, surface precipitation rates and amounts are difficult to estimate. Unlike other atmospheric variables, rainfall can display extreme heterogeneity in space and time. Both the occurrence of precipitation and the rate at which it falls may be highly variable, even within a single rain event. Due to the importance of and difficultly with estimation, there has been much research in the area of precipitation measurement and estimation.

Precipitation is measured using many different instruments. Some instrumentation measures rainfall directly, while others make indirect measurements of quantities that can be related to rainfall. All of the instrumentation, however, make indirect measurements of the usual quantity of interest, namely, rain rate. The most common instruments include rain gauges and ground-based scanning radar. Distrometer in situ measurements and satellite deployed instrumentation are also used, although less commonly. Each of the instruments has inherent strengths and weaknesses, which affect the observed measurements and subsequent estimation of surface rainfall.

There exist many algorithms, each with its own set of assumptions, for calculating rainfall related parameters from the instrumental observations. Many times the observations obtained from one instrument are used in the calculation of estimates based on the observations from another instrument. Rain gauge data are often considered to be "ground truth" and are used to adjust the estimates based on other instruments. This is done despite the known fact that observations from rain gauges are inherently prone to errors ([12] and [8]) and that they are no more "true" than the observations from other instruments.

Surface precipitation can be defined by a population of falling drops. The distribution of the size diameters of the drops characterizes the population behavior; thus, the drop-size distribution (DSD) forms a basic descriptor in the modeling of rain microphysics. The DSD is used to compute a variety of "derived" properties of rainfall through mathematical relationships; common computed properties include water content, rain rate, and reflectivity. DSD data, however, are not collected on a routine basis due to the expense of distrometer instruments and the limited spatial coverage of an individual instrument.

The literature is full of modeling approaches for the estimation of surface rainfall, including empirical statistical models ([24] and [21], for example) and a variety of stochastic models ([3], [16], [6], [15], and [22]). More recently, hierarchical Bayesian models have been used in a number of environmental applications ([17], [29], [4], [5], [30], and [9]).

In this paper, using a hierarchical Bayesian approach, we develop a model for estimating the underlying distribution of drop sizes at a site. This then can be used to calculate different quantities related to precipitation (such as rain rate, the instantaneous rate of precipitation across a unit area), and to calibrate various measurement devices (such as tipping rain gauges and radar reflectivity). The

advantage with this approach is that there is no need to think of a particular kind of measurement as "ground truth." Rather, the ground truth is the estimated distribution, and the methodology allows researchers to assess quantitatively the uncertainty associated with that distribution.

2. Hierarchical model

DSD spectra can be used to calculate many rainfall parameters. The availability of such spectra, however, is often limited and surface rainfall estimates usually rely on other data sources. Model development is focused on an unobserved process model for the evolution of drop-size distributions through time. Information contained in more widely accessible data is used to augment the parameter estimates for the hidden process. Specifically, these data include rain gauge and ground radar observations.

This section provides a summary of key characteristics of DSD spectra that are used in formulating the model for the unobserved process. This is followed by a conceptual description of the model and then model details. Model components are developed in the time domain only, assuming a single, fixed spatial location. The model is a flexible five-stage Bayesian hierarchical formulation that is based on a hidden model for generating rainfall drop-size distributions. Estimates of the unobserved process, denoted by $N(D)_t$, are the primary quantity of interest. The first stage of the model specifies a measurement error model for the observational data, radar data denoted by Z and rain gauge data denoted by G, both of which are observations of functions of $N(D)_t$ with error. The second stage of the model incorporates a time series formulation of the unobserved DSD process. Time series parameters and temporal dynamic terms of the DSD evolution process are modeled in the third stage. The Bayesian formulation is completed in stages four and five by specifying priors on model parameters. Figure 1 shows a schematic that summarizes the relationships among stages of the hierarchical model specification.

For readability, the model is presented in a different order than the component stages of the hierarchical framework. The complete model for the hidden process, including temporal dynamics and priors on related model parameters is presented first, followed by a description of the observational data components.

2.1. Empirical analysis of DSD data

Exploratory data analysis was conducted on a set of one-minute drop size spectra collected using a Josh-Waldvogel distrometer (JWD) from a site in Eureka, California from January through March 1999. Distrometers are direct measuring devices that collect observations of the ground rainfall drop-size distribution at a spatial point location with high temporal resolution. The JWD is a electro-mechanical spectrometer that separates raindrops entering a 50 cm^2 sampling area into 20 different drop-size classes using a 20-channel pulse-height analyzer. Drops with diameters between $0.3 - 5.0$ mm are partitioned into the first 19 bins. The last bin contains all drops with diameter >5 mm [28].

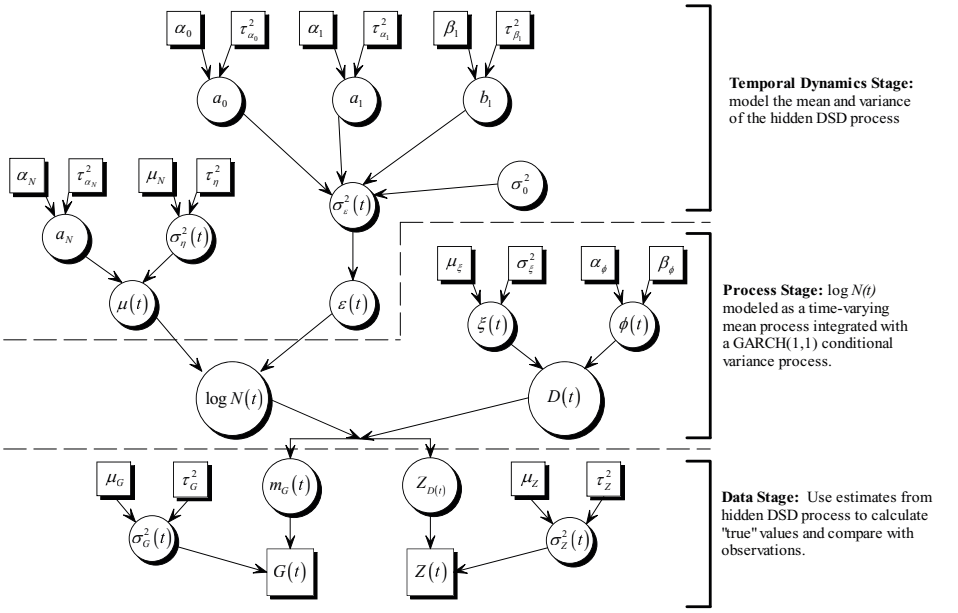

FIGURE 1. Summary of model stages showing the relationship among the data, process, and temporal dynamics stages. Rectangles represent observed data and fixed hyperparameters. Circles represent unobserved parameters.

Rain rates calculated for each of the one-minute spectra were used to identify rainfall events. Seven events spanning early February to mid-March, ranging from 12 to 24 hours in duration, were chosen for evaluation. Barplots of the binned drops and time series of the total number of drops/minute were evaluated for each event. Two general observations were made:

1. The distributions of drop-sizes were generally unimodal and right-skewed. Based on visual inspection and five number summaries, the shapes of the distributions of the one-minute DSD spectra were similar both within and across rainfall events. Figure 2 shows barplots for some representative one-minute DSDs.
2. The total number of drops from one-minute to the next was highly variable both within and across events. Evaluation of log-transformed time series data suggested an underlying auto-regressive process. Figure 3 shows some time series plots for total number of drops.

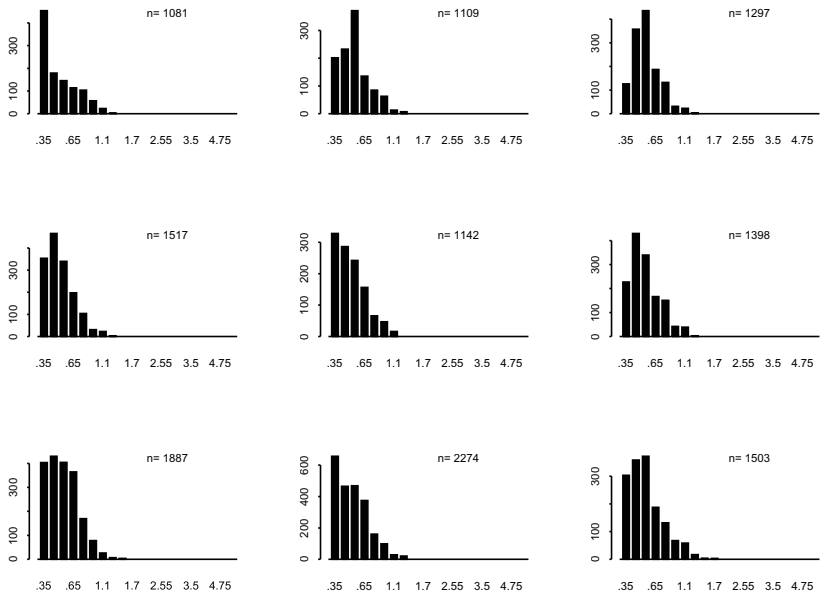

FIGURE 2. Barplots for nine consecutive minutes of distrometer DSD data from a Eureka rainfall event; n is the total number of drops in each one-minute interval.

2.2. Unobserved process

A DSD, denoting the number of drops per mm diameter bin interval and per m^3 of air, is well defined as

$$N(D_i)_t = \frac{n(D_i)_t}{AtV(D_i)\Delta D_i} \tag{1}$$

where $n(D_i)_t$ is the observed number of drops (raw counts) in each diameter bin interval D_i for sampling interval t, A is the sample surface area in m^2, t is the length of sample collection in seconds, $V(D_i)$ is the terminal fall speed in still air in m/s, and ΔD is the drop diameter interval in mm.

Now, consider defining the number of drops over all diameters D as

$$n(D)_t = n_t \times P(D)_t \tag{2}$$

where n_t is the total number of raw counts from a JWD instrument over all drop diameters for the time interval t and is assumed independent of $P(D)_t$, a probability density function defining the shape of the DSD over the interval t; $P(D)_t$ gives the probability of finding a drop within a diameter D to $D + \Delta D$.

Substituting (2) into (1) yields

$$N(D_i)_t = \frac{n_t \times P(D_i)_t}{AtV(D_i)\Delta D_i}. \tag{3}$$

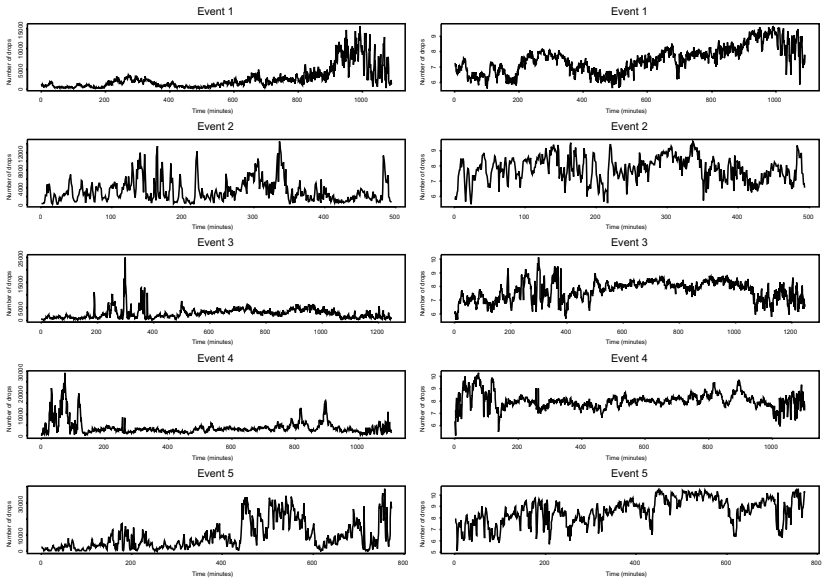

FIGURE 3. Time series of the one-minute total number of drops (left) and log transformed total number of drops (right) for each of five Eureka rainfall events.

Next, define N_t as the total number of drops per m^3 of air. Substitution into (3) gives the DSD formulated in terms of total counts as

$$N(D_i)_t = \frac{N_t \times P(D_i)_t}{\Delta D_i}. \tag{4}$$

For a rainfall event of duration T, the unobserved states over all drop diameters $N(D)_t$ will be estimated through models for N_t and $P(D)_t$; $t = 1, \ldots, T$. To be consistent with commonly collected distrometer data, each t will represent a one-minute interval.

The total number of drops is treated as independent of shape of the DSD spectra. Sauvageot and Lacaux [18] suggest the possibility that the total number of drops in a DSD varies with rain rate and this total can be thought of as independent of a generic shape function that reflects an equilibrium between coalescence and collisional breakup processes. Models for the two independent components of $N(D)_t$ are detailed below.

2.2.1. Model for total number of drops, $\mathbf{N_t}$. Let N_t denote the total number of drops at time t, $t = 0, 1, 2, \ldots, T$. To account for the serial correlation and burst variability observed in many of the time series of total number of drops in contiguous DSD spectra, we introduce a model composed of a time-varying mean process integrated with a generalized autoregressive conditional heteroscedasticity

(GARCH) conditional variance process. GARCH provides conditional models that allow for time-varying changes in the distributions that govern each data point, as opposed to assuming a constant variance model that applies equally anywhere in the series.

Let

$$TN_t = \ln N_t = \mu_{N_t} + \epsilon_t.$$

The ϵ_t's are an independent sequence of correlated errors from an unknown distribution. The ϵ_t values, conditional on all previous values are assumed to be normally distributed:

$$\epsilon_t | \epsilon_{t-1}, \epsilon_{t-2}, \ldots \sim N(0, \sigma^2_{\epsilon_t}).$$

Based on a analysis of data from several rainfall events, we selected a $GARCH(1,1)$ model, which provides a single lagged error square term and one autoregressive variance term to capture the "burst" variability observed in many of the total number of drop time series. Under a $GARCH(1,1)$ model

$$\sigma^2_{\epsilon_t} = a_0 + a_1 \epsilon^2_{t-1} + b_1 \sigma^2_{\epsilon_{t-1}}$$

where the parameters a_0, a_1, and b_1 are estimated, subject to the constraints

$$a_0 > 0,\ a_1, b_1 \geq 0, \quad \text{and} \quad a_1 + b_1 \leq 1 \quad [1].$$

By substitution,

$$\begin{aligned}
\epsilon_1 | \epsilon_0 &\sim N(0, a_0 + a_1 \epsilon^2_0 + b_1 \sigma^2_{\epsilon_0}) \\
\epsilon_2 | \epsilon_1, \epsilon_0 &\sim N(0, a_0 + a_1 \epsilon^2_1 + b_1 \sigma^2_{\epsilon_1}) \\
&\vdots \\
\epsilon_t | \epsilon_{t-1}, \epsilon_{t-2}, \ldots, \epsilon_1, \epsilon_0 &\sim N(0, a_0 + a_1 \epsilon^2_{t-1} + b_1 \sigma^2_{\epsilon_{t-1}}).
\end{aligned}$$

Given the conditional distribution of $\epsilon_t | \epsilon_{t-1}, \epsilon_{t-2}, \ldots, \epsilon_1, \epsilon_0$, the joint probability for ϵ is:

$$\begin{aligned}
P(\epsilon_t, \epsilon_{t-1}, \ldots, \epsilon_0) &= P(\epsilon_t | \epsilon_{t-1}, \epsilon_{t-2}, \ldots, \epsilon_0) P(\epsilon_{t-1} | \epsilon_{t-2}, \epsilon_{t-3}, \ldots, \epsilon_0) \\
&\quad \times \ldots \times P(\epsilon_0) \\
&= \left[\prod_{i=1}^{t} P(\epsilon_{t+1-i} | \epsilon_{t-i}, \epsilon_{t-i-1}, \epsilon_{t-i-2}, \ldots, \epsilon_0) \right] P(\epsilon_0) \\
&= \left[\prod_{i=1}^{t} \frac{1}{\sqrt{2\pi}\sigma_{\epsilon_{t+1-i}}} \exp\left\{ -\frac{1}{2\sigma^2_{\epsilon_{t+1-i}}} (\epsilon^2_{t+1-i}) \right\} \right] \\
&\quad \times \frac{1}{\sqrt{2\pi}} \exp\left\{ -\frac{1}{2} (\epsilon^2_0) \right\}.
\end{aligned}$$

Assume that conditional on μ_{N_t} and the parameters $\theta_N = \{a_0, a_1, b_1, \sigma^2_{\epsilon_0}\}$, the $\ln N_t$'s are independent and are distributed as $N(\mu_{N_t}, \sigma^2_{\epsilon_t})$. Then, the joint distribution of $\ln N$ is:

$$\begin{aligned} P(TN) &= \left[TN_1, \ldots, TN_T \,\middle|\, \mu_{N_1}, \ldots, \mu_{N_T}; \sigma^2_{\epsilon_1}, \ldots, \sigma^2_{\epsilon_T}\right] \\ &= \left[\prod_{t=1}^{T} \frac{1}{\sqrt{2\pi}\sigma_{\epsilon_t}} \exp\left\{ -\frac{1}{2\sigma^2_{\epsilon_t}}(TN_t - \mu_{N_t})^2 \right\}\right] \\ &\quad \times \frac{1}{\sqrt{2\pi}\sigma_{\epsilon_0}} \exp\left\{ -\frac{1}{2\sigma^2_{\epsilon_0}}(TN_0 - \mu_{N_0})^2 \right\}. \end{aligned}$$

Set $TN_0 = \mu_{N_0}$; this reduces $P(\epsilon_0)$ to the constant term $\frac{1}{\sqrt{2\pi}\sigma_{\epsilon_0}}$ and yields:

$$P(TN) = \left[\prod_{t=1}^{T} \frac{1}{\sqrt{2\pi}\sigma_{\epsilon_t}} \exp\left\{ -\frac{1}{2\sigma^2_{\epsilon_t}}(TN_t - \mu_{N_t})^2 \right\}\right] \times \frac{1}{\sqrt{2\pi}\sigma_{\epsilon_0}}.$$

Since the GARCH parameters are constrained to be greater than zero, each coefficient is assigned a Lognormal prior:

$$\begin{aligned} a_0 &\sim \text{Lognormal}(\alpha_0, \tau^2_{\alpha_0}) \\ a_1 &\sim \text{Lognormal}(\alpha_1, \tau^2_{\alpha_1}) \\ b_1 &\sim \text{Lognormal}(\beta_1, \tau^2_{\beta_1}). \end{aligned}$$

The hyperparameters α_0, α_1, β_1, $\tau^2_{\alpha_0}$, $\tau^2_{\alpha_1}$, and $\tau^2_{\beta_1}$ are taken as fixed, with values that may vary by application.

The mean for TN_t, μ_{N_t}, is modeled as a first-order autoregressive process. Let $\mu_{N_t} = a_N \mu_{N_{t-1}} + \eta_{N_t}$ where η_{N_t} is a random normal process with mean 0 and variance σ^2_η. The parameters are $\theta_\mu = \{a_N, \sigma^2_\eta\}$. Conditional on a_N and $\mu_{N_{t-1}}$, $\mu_{N_t} \sim N(a_N \mu_{N_{t-1}}, \sigma^2_\eta)$ with $\sigma^2_\eta \sim IG(\mu_\eta, \tau^2_\eta)$. The parameter a_N needs to be estimated. The prior distribution for a_N is defined as $a_N \sim N(\alpha_N, \tau^2_{\alpha_N})$. The hyperparameters α_N, $\tau^2_{\alpha_N}$, μ_η, and τ^2_η are taken as fixed, with values that may vary for different types of rainfall events (e.g., light rainfall versus heavy rainfall).

2.2.2. DSD shape, $\mathbf{P(D)_t}$. A number of distributions have been used to model DSDs including the gamma distribution proposed by Ulbrich [26], Ulbrich and Atlas [27], and Willis and Tattelman [32], in addition to the lognormal distribution ([18] and [22]). A three parameter normalized gamma distribution ([31], [25], and [10]), along with the Weibull distribution ([20], [11] and [19]) have also been used to model DSDs. We considered exponential, lognormal and gamma distributions for modeling the DSD shape. An exponential distribution tended to underestimate the number of drops in the 0.3 mm to about 1 mm diameters, while overestimating drops in the < 0.3 mm bin. The use of a gamma distribution was dismissed due to constraints on available computing resources and the prohibitive run time that was required due to excessive required calls to evaluate the Γ function. Thus, the DSD shape, is modeled as a lognormal distribution with time varying parameters

ξ_t and ϕ_t^2. Conditional on ξ_t and ϕ_t^2 the D_ts are assumed independent yielding the joint distribution for D as

$$[D|\xi_t, \phi_t^2] = \prod_{t=1}^{T} \frac{1}{\sqrt{2\pi} d_t \phi_t} \exp\left\{ - \frac{(\ln d_t - \xi_t)^2}{2\phi_t^2} \right\}$$

where $\xi_t \sim N(\mu_\xi, \sigma_\xi^2)$; $\phi_t^2 \sim IG(\alpha_\phi, \beta_\phi)$. The hyperparameters μ_ξ, σ_ξ^2, α_ϕ, and β_ϕ are taken as fixed, with values that may vary for different rainfall events.

2.3. Data components

The hierarchical model incorporates rain gauge and ground radar observations, which are two of the most common data sources for measuring surface rainfall. Since, given a DSD, the instantaneous (true) rain rate at a given time and the derived (true) radar reflectivity can be calculated, we develop a measurement error model for both the rain gauge and radar observations.

2.3.1. Gauge observations, G. Rain gauges measure the amount of rainfall accumulating over a fixed time at a point in space. If the amount of rainfall measured at a point is measured without error, the gauge measurement, G, can be equated to the appropriate time integral of rain rate. The predominant error in gauge measurements is a systematic bias induced by winds, which often results in an underestimate of the surface rainfall. Thus, define a measurement error model for gauges as a function of the true rain rate and wind speed (taking other systematic biases as negligible). Define a gauge observation as

$$G_t = \left\{ c(w_t) \int_{t-\Delta t}^{t} R_s ds \right\} + \epsilon_{G_t}$$

where R_t is the derived instantaneous rain rate at time t; $c(w_t)$ are gauge type-specific coefficients ($c \leq 1$) primarily based on wind speed, w; and ϵ_{G_t} are independent measurement errors modeled as $N(0, \sigma_G^2)$; $\sigma_G^2 \sim IG(\mu_G, \tau_G^2)$.

Given a DSD, the instantaneous rain rate for a given time t can be derived based on meteorological principles as a function of the third power of drop diameter [2]. The rain rate R, for a given t, is estimated as

$$R_t = c_R \frac{\pi}{6} \int_0^\infty D^3 V(D) N(D)_t dD$$

where D is the drop diameter in mm; $V(D)$ is a deterministic function for the terminal velocity for drops with diameter D; $N(D)_t$ is the DSD giving the number of drops per mm diameter D per m^3 at time t; and c_R is an additional constant that accounts for units.

Given the model for $N(D_i)$ (equation 4), summing over all ΔD, and substituting for $N(D)_t$ yields

$$R_t = c_R \frac{\pi}{6} \left[\int_0^\infty D^3 V(D) [N_t f(D)_t] dD \right] \tag{5}$$

where N_t is the total drop concentration at time t and $f(D)_t$ is the pdf for D_t defined through $P(D)_t$.

Assuming this independent Gaussian measurement error model, and conditional on rain rate as derived from a DSD and parameter $\theta_G = \{\sigma_G^2\}$, the elements of G are independent and in each case $G_t \sim N(m_{G_t}, \sigma_G^2)$, where m_{G_t} is the integral, over the time interval between gauge measurements, of the instantaneous rain rate defined above.

2.3.2. Ground radar observations, Z. Reflectivity values are related to the intensity of reflected power that is bounced off of raindrops in the atmosphere. Reflectivity will vary depending on the number of and the size of the drops. The equivalent radar reflectivity observations from ground radar, Z, are a function of the average returned power and range of the radar scan [2]. If there are no measurement errors, Z will conform to the reflectivity derived the DSD. To accommodate various factors that can lead to errors when estimating Z, we construct a simple measurement error model for equivalent radar reflectivity as a function of the derived reflectivity and random errors; $Z_t = Z_{D_t} + \epsilon_{Z_t}$, where Z_{D_t} is the meteorologically derived reflectivity at time t and ϵ_{Z_t} are independent measurement errors modeled as $N(0, \sigma_Z^2)$; $\sigma_Z^2 \sim IG(\mu_Z, \tau_Z^2)$.

The treatment of derived reflectivity, Z_D, is similar to that for rain rate. Given a DSD, Z_D can be derived as the sixth moment of the DSD [2]. Derived reflectivity at time t is estimated as:

$$Z_{D_t} = c_Z \int_0^\infty D^6 N(D)_t dD$$

where D is the drop diameter in mm; $N(D)_t$ is the current drop spectra giving the number of drops per mm per m^3, at time t, with diameter D; and c_Z is a constant related to the sample volume and equals one for a one-m^3 sample.

As for gauge observations, summing over all ΔD and substituting $N_t \times P(D)_t$ for $N(D)_t$ yields:

$$Z_{D_t} = c_Z \left[\int_0^\infty D^6 [N_t f(D)_t] dD \right]$$

where variables are as previously defined. In practice, Z_{D_t} is estimated as a sum over binned drop sizes. Assuming this independent Gaussian measurement error model and conditional on Z_D derived from the DSD and parameter $\theta_Z = \{\sigma_Z^2\}$, the elements of Z are independent and in each case $Z_t \sim N(Z_{D_t}, \sigma_Z^2)$.

2.4. Model summary

We use the components presented above to construct the hierarchical model formulation. At the highest level we assume, based on a set of conditional independence assumptions, that given a DSD ($N(D)$), the gauge (G) and radar (Z) observations are conditionally independent with the following factorization:

$$\begin{aligned} [N(D), G, Z] &= [N(D)][G, Z|N(D)] \\ &= [N(D)][G|N(D)][Z|N(D)]. \end{aligned}$$

TABLE 1. Hierarchical Model Summary

Variables	**Conditional Probabilities for Model Components**[a]
Observational Data	$[G_t \mid N(D)_t, \theta_G]$[b] $[Z_t \mid N(D)_t, \theta_Z]$
Hidden Process	$[TN_t \mid \mu_{N_t}, \theta_N]$[c] $[D_t \mid \theta_D]$
Temporal Dynamics for mean of TN	$[\mu_{N_t} \mid \theta_\mu]$
Model Parameters	$[\theta_G, \theta_Z, \theta_N, \theta_D, \theta_\mu \mid \theta_H] = [\theta_G \mid \theta_{H_G}][\theta_Z \mid \theta_{H_Z}][\theta_N \mid \theta_{H_N}][\theta_\mu \mid \theta_{H_\mu}][\theta_D \mid \theta_{H_D}]$ where θ_H is the collection of hyperparameters and $\theta_G = \{\sigma_G^2\} = [\sigma_G^2 \mid \mu_G; \tau_G^2]$ $\theta_Z = \{\sigma_Z^2\} = [\sigma_Z^2 \mid \mu_Z; \tau_Z^2]$ $\theta_N = \{a_0, a_1, b_1, \sigma_0^2\} = [a_0 \mid \alpha_0; \tau_{\alpha_0}^2][a_1 \mid \alpha_1; \tau_{\alpha_1}^2][b_1 \mid \beta_1; \tau_{\beta_1}^2][\sigma_0^2]$ $\theta_\mu = \{a_N, \sigma_{\eta_t}^2\} = [a_N \mid \alpha_N; \tau_{\alpha_N}^2][\sigma_{\eta_t}^2 \mid \mu_\eta; \tau_\eta^2]$ $\theta_D = \{\lambda_t\} = [\lambda_t \mid \mu_\lambda; \sigma_\lambda^2]$
Hyper-parameters	$\theta_H = \{\mu_G, \tau_G^2, \mu_Z, \tau_Z^2, \alpha_0, \alpha_1, \beta_1, \tau_{\alpha_0}^2, \tau_{\alpha_1}^2, \tau_{\beta_1}^2, \alpha_N, \tau_{\alpha_N}^2, \mu_\eta, \tau_\eta^2, \mu_\lambda, \sigma_\lambda^2\}$

[a] The bracket notation, $[\cdot|\cdot]$, is used as a shorthand for denoting conditional distributions.
[b] Recall that $N(D)_t$ is general notation for $N_t \times P(D)_t$.
[c] Recall that $TN = \ln(N_t)$.

The components $[G|N(D)]$ and $[Z|N(D)]$ represent likelihoods based on the gauge and radar data, respectively. $[N(D)]$, represents the prior probability for the DSD process. The conditional probabilities for the full model, including the parameters and hyperparameters, are summarized in Table 1.

Estimated model parameters are obtained by using a Markov Chain Monte Carlo (MCMC) implementation to sample from the posterior distribution of the hierarchical model. The appropriate Markov chain is constructed using a single-component Metropolis-Hastings algorithm based on the hierarchical structure of the model and on the conditional probabilities associated with the various model components.

Starting values, hyperparameter values, the length of burn-in, and the number of iterations must be determined for each application. Starting values are random, based on the prior distributions of the parameters; the distributions can be altered by the choice of the fixed hyperparameters. The required number of iterations for a given application depends on the length of the burn-in, autocorrelation of the parameters, and the number of samples needed for parameter estimation. Convergence is monitored using visual analysis of plotted parameters in two ways: 1) by performing multiple runs using different random starts; and 2) running multiple chains using the same starting values.

3. Application to Eureka rainfall

We applied the model developed in the previous sections to a site in Eureka, California. The Eureka site has three sources of available rainfall data: distrometer, rain gauge, and ground radar data. The distrometer data were collected using a JWD instrument. The distrometer was located near the Eureka radar site at $40°48'N/124°09'W$. Each data file contained one hours worth of raw bin counts collected at one-minute intervals, where raw counts are the actual number of drops passing through the instrument in each one-minute interval. The raw bin counts were converted to one-minute DSD spectra counts, N_{D_t} ($\mathrm{m}^{-3}\mathrm{mm}^{-1}$), by adjusting raw counts for the instrument collection area, time, and terminal fall speed, and drop diameter.

A set of processed 10-minute distrometer data was also available. The 10-minute data are drop spectra that had been averaged over 10 contiguous one-minute periods. The observed 10-minute drop spectra were used to calculate rain rates and reflectivity values that were used in place of gauge and radar observations in model verification runs.

Hourly accumulation data (TD-3240) were obtained from the National Climatic Data Center (NCDC) on-line service for a gauging station located close to the radar and distrometer instruments at $40°49'N/124°10'W$. Data were collected using Fischer-Porter precipitation gauges with automated readouts. Precipitation data from the NCDC database are quality checked and edited, as necessary, by an automated and a manual edit [7]. The gauge data were used in conjunction with the distrometer data to identify rainfall events.

Level II WSR-88D radar scans were ordered from NCDC for the Eureka radar site located at ($40°29'54''N/124°17'31''W$), for periods coinciding with rainfall events identified by gauge and distrometer data. There was a radar scan approximately every six minutes for several elevation increments. The spherical coordinate Level II data were interpolated to a 2 km $\times$ 2 km Cartesian coordinate system of equivalent reflectivity (dBZ) values using standard programs for NEXRAD radar. The interpolated data were only available for one of the identified events; the availability of radar data defined, and limited, the time period over which the model was implemented.

Appropriate reflectivity values needed to be extracted from the 2 km $\times$ 2 km gridded values to correspond with the fixed spatial location for the model runs. Reflectivity in the time domain was extracted from the spatial data by first identifying the closest pixel located vertically above the distrometer and nearby rain gauge. Given the uncertainty in the vertical flow of rainfall above the surface of the earth to a point on the ground, an equivalent reflectivity was calculated as the aerial average of first and second order neighbors using the identified pixel and the eight neighboring pixels. Issues of beam blockage along the path of the radar to the location of the distrometer were considered to be negligible due to the gradually increasing nature of the coastal terrain of Northern California. Therefore, the reflectivity values were extracted from the lowest radar beam (tilt elevation of 0.5°).

3.1. Hyperparameters, initial values, constraints

We used exploratory analysis of distrometer data from five February rainfall events to come up with informative means and variances for the fixed hyperparameters. The constant hyperparameter values are summarized in Table 2 with model results.

In addition to the hyperparameter values, several $t = 0$ starting values were specified as follows: $\epsilon_0 = 0$, $\sigma^2_{\epsilon_t} = 0.001$, and $\mu_{N_0} = 3.0$. We based initial conditions for the remaining parameters on random draws from the appropriate prior distributions.

We introduced cutoffs for constraints due to rainfall dynamics, instrumentation and statistical considerations. Constraints based on rainfall dynamics included upper bounds for three parameters:

1. We limited the total logged number of drops in a one-minute interval to 6.0, about 2.5% above the maximum value observed from the distrometer data. This corresponds to a total of $\exp(6.0) \approx 403$ raw count drops when summed across the 20 drop diameter bins. This cutoff was applied to the parameters TN_t and μ_{N_t}.
2. We limited the upper bound for drop diameter size such that $P(D \geq 5\text{ mm}) \approx 0.000001$. This limitation was imposed since the rain event of interest represented light rainfall. The distrometer data had no raw counts beyond a mean drop diameter of 3.15 mm.

Proposals for parameter updating were drawn from uniform distributions centered at the current parameter values. This implementation made it possible to draw negative-valued proposals, depending on the current value of the parameter. Constraints based on statistical considerations thus included lower bounds on the variance parameters (> 0) and $\lambda_t > 0$. Required model constraints were also placed on the three GARCH coefficients.

3.2. Estimating $\mathbf{m_G}$ and $\mathbf{Z_D}$

At each time period for which there were gauge or radar observations, the data components of the model required computation of the two quantities m_{G_t} and Z_{D_t}, the true equivalent gauge value and the true derived reflectivity, respectively. Both of these values are deterministic quantities based on the current state of the parameters that define the hidden DSD process.

We calculate the true equivalent gauge value as the sum over the previous 60 minutes of each of the instantaneous rain rates in mm/min derived from the state of the DSD for each point in time. The instantaneous rain rates defined by equation 5 are approximated by summing over the 20 drop diameter bins that would typically be observed using a JWD instrument as follows:

$$R_t = \frac{3.6\pi}{6000} \sum_{i=1}^{20} \left[\frac{\exp\{TN_t\} \times P(D_i)_t}{\Delta D_i} \right] D_i^3 V(D_i) \Delta D_i$$

where $P(D_i)_t$ is based on the cumulative distribution of a lognormal distribution with parameters ξ_t and ϕ_t^2. The terminal fall speed, $V(D_i)$, was calculated as

$3.78D_i^{0.67}$ [22], where D_i is the mean drop diameter for bin i. Then,

$$m_{G_t} = \sum_{i=t-60+1}^{t} \frac{R_i}{60}$$

True derived reflectivity is calculated as an average, over the previous 10 minutes, of the sixth moment of the DSD. As with the equivalent gauge calculation, we approximate derived reflectivity by summing over the 20 drop diameter bins using

$$Z_t = \sum_{i=1}^{20} \left[\frac{\exp\{TN_t\} \times P(D_i)_t}{\Delta D_i} \right] D_i^6 \Delta D_i.$$

Then,

$$Z_{D_t} = \frac{\sum_{i=t-10+1}^{t} Z_i}{60}$$

where Z_{D_t}, in units of $\mathrm{mm}^6\mathrm{m}^{-3}$, are converted to dBZ as $10\log_{10}(Z_{D_t})$.

3.3. Tuning, thinning, and convergence

Tuning the chain was optimized by making many small runs (5000 to 15000 iterations in length) and monitoring the acceptance rates for each of the parameters. Adjustments to the acceptance rates were made by either increasing or decreasing the upper and lower bounds on the uniformly distributed parameter proposals. After much experimentation, bounds on the uniform proposals were adjusted to yield acceptance rates ranging from about 0.42 on the low end to about 0.85 on the high end. Most parameters had acceptance rates in the range of 0.60 - 0.65. While these values exceed the commonly recommended values of 0.15 to 0.5 for updating multidimensional components [14], we found that mixing and convergence performed better at these higher levels.

We used thinning to obtain roughly independent samples from the joint posterior distribution. We used *gibbsit* [13] to get estimates of the ideal spacing between iterations and the number of burn-in iterations, as well as the number of iterations needed for a desired precision. In most cases, a thinning value of 150 was more than sufficient to obtain approximately independent samples although, on occasion, a parameter thinning value would hit 300 or so.

We monitored convergence by visual inspection of plotted output. The number of burn-in iterations recommended by *gibbsit* was generally low, and always < 5000. Convergence was also assessed by looking at the path of several runs using both the same and different starting values.

3.4. Results

Results for two types of model runs are summarized in the following sections. Since there is no way to specifically validate the model in terms of comparing outputs to "true" values, model verification was conducted by assessing the ability of the model to realistically reproduce targeted signals and integrated parameters associated with surface rainfall. For model verification we replaced NEXRAD radar

observations with derived reflectivity calculated from 10-minute distrometer data. Rain gauge data were omitted for this assessment. The number of inputs were maximized by using a 10-minute time step, thus insuring a data value at each time step. The full model runs use no distrometer data, but rather the gauge and radar data are combined with the model for the hidden DSD process to calibrate the parameter estimates with all available information.

3.4.1. Model verification. The verification run includes 49 derived reflectivity observations computed from the one set of distrometer data. The MCMC run was ended after 280,000 iterations following 100,000 burn-in iterations. Based on graphical output, all parameters converged within the 100,000 burn-in period. Results for a typical model verification run showing time independent parameter estimates and time varying parameter estimates (using the 800 MCMC samples generated after applying a thinning value of 350) are shown in Tables 2 and 3, respectively. Further, using the posterior parameter estimates for the hidden process we can compare calculated integrated rainfall parameters with values obtained from the 10-minute distrometer data. Figures 4 and 5 show the comparisons for rain rate and reflectivity for some representative samples from the model run. Note that although we compare the model output with quantities derived from distrometer

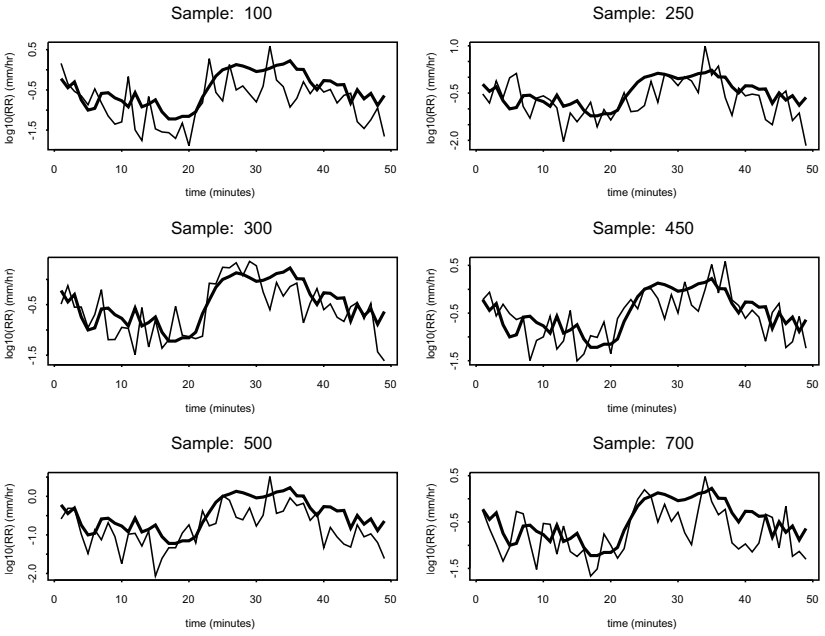

FIGURE 4. Time series posterior $\log_{10}$ rain rate estimates calculated for samples from the validation run (light lines) with rain rate estimates from the distrometer data (heavy lines).

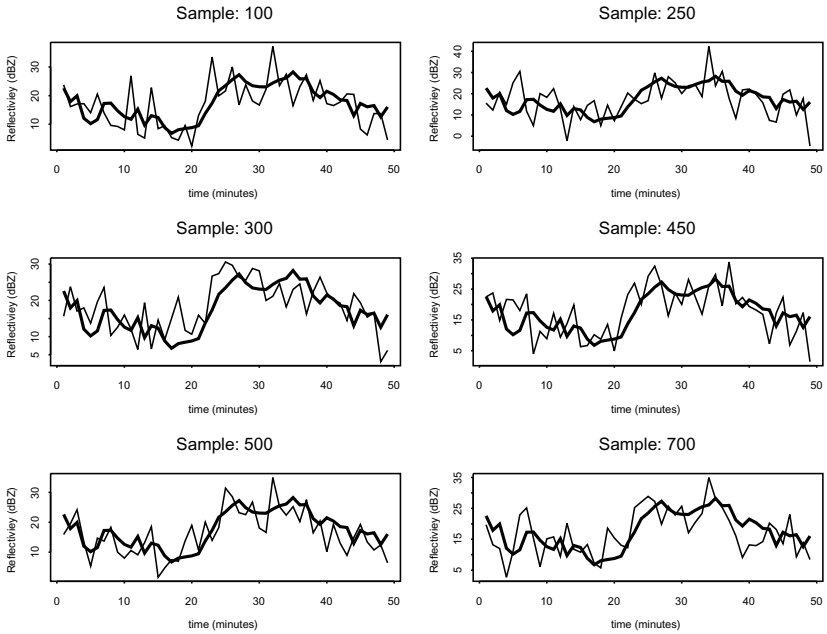

FIGURE 5. Time series posterior reflectivity estimates calculated for samples from the validation run (light lines) with reflectivity estimates from the distrometer data (heavy lines).

data, we do not consider the distrometer data as ground truth. Distrometer data are subject to measurement errors and our distrometer data set represents a single realization of the process.

The posterior means of the three GARCH coefficients are all lower than those specified for the prior distributions. The variability for the coefficients is slightly less than that for the priors. The mean of the AR(1) coefficient is slightly higher than that specified in the prior. The variance parameter for the AR(1) errors, σ_η^2, is considerably higher than that specified in the prior, as is the measurement error variance for radar (σ_Z^2). The mean values for TN and μ_N are about the same and posterior estimates of TN are close to expected counts although the model tends to underestimate the counts. In general, the DSD distribution parameters ξ_t and ϕ_t^2 are more variable than estimates based on the distrometer data. On average, the posterior estimates of ξ_t are similar to the values estimated from the distrometer data. The model tends to slightly underestimate ϕ_t^2.

The posterior estimates for both integrated parameters (rain rate and reflectivity) are more variable than values calculated from the distrometer data realization. This is expected since these estimates include the distrometer data realization with some uncertainty. Aside from the extra variability, the model estimates are

a good match with the integrated quantities calculated from distrometer observations. The model does a good job of capturing the overall means and the signals present in the time series.

3.4.2. Full model run. The full model run includes 9 hourly rain gauge observations and 92 effective reflectivity observations from ground radar. The MCMC run was ended after 245,000 iterations following a burn-in period of 135,000 iterations. After applying a thinning value of 350, there were 700 MCMC samples for posterior estimates. Convergence for the full model run is much slower than that for the verification run, taking up to about 60,000 iterations for a few parameters. With the exception of the AR(1) variance term, σ_η^2, all parameters appear to converge by 60,000 iterations. The σ_η^2 variance hardly moves until 60,000 iterations, at which point the value starts bouncing around and then settles down at about 0.5, followed by more intermittent bouncing. This pattern could be due to slow mixing or identifiability issues.

Results for a typical full model run showing time independent parameter estimates and time varying parameter estimates are shown in Tables 2 and 3, respectively, along side the verification run results. Time series results for integrated

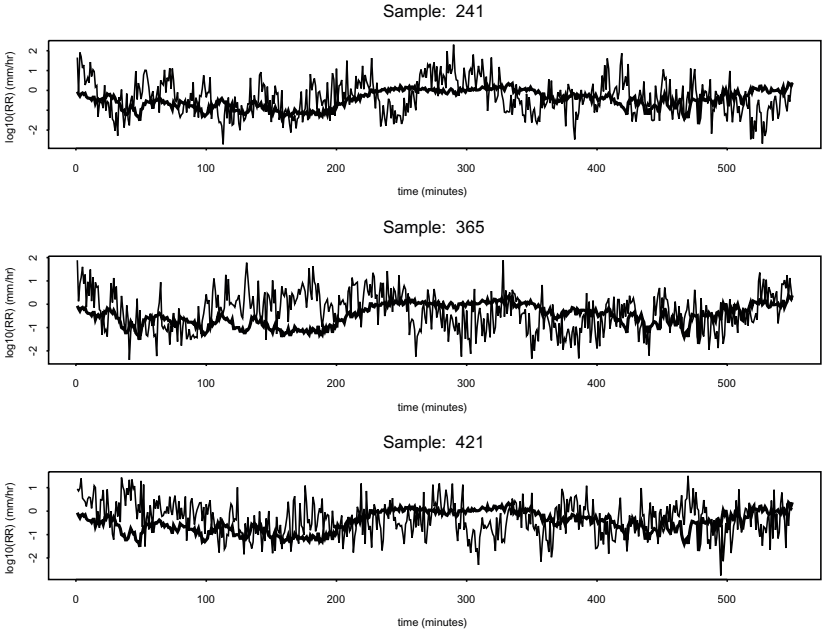

FIGURE 6. Time series plots of full model run posterior $\log_{10}$ rain rate estimates calculated for selected samples (light lines) with logged rain rate values based on distrometer data (heavy lines).

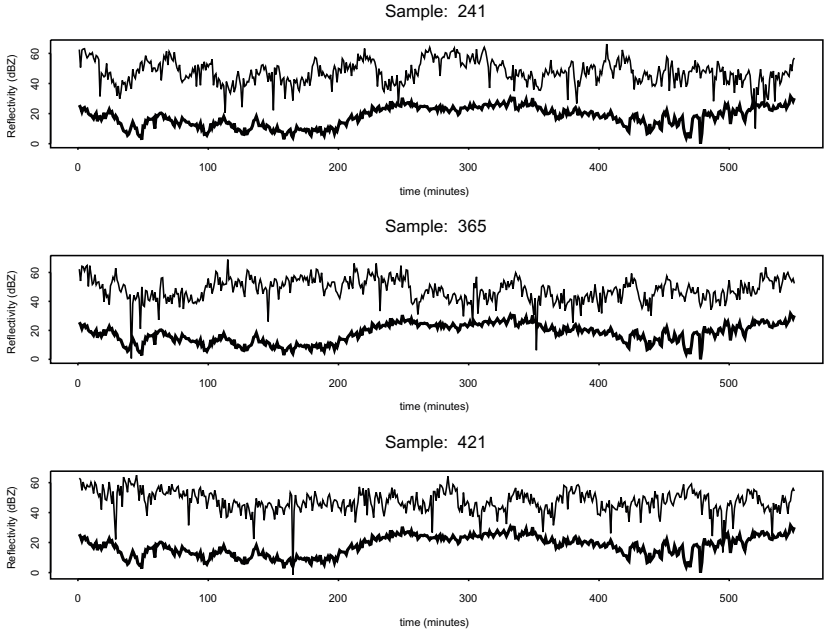

FIGURE 7. Time series full model run posterior reflectivity estimates calculated for selected samples (light lines) with reflectivity values based on distrometer data (heavy lines).

parameters calculated from the estimated parameters for the hidden DSD process are shown in Figures 6 through 9.

Except for a_0, the posterior estimates for the GARCH parameters are similar to those from the verification run. The generally larger standard errors reflect the greater uncertainty in the one-minute model, which does not have data at every time step. The mean error terms, ϵ_t, are not significantly different from zero based on 95% confidence intervals for the means. The posterior distributions are relatively symmetric and approximately normal, which is in line with the GARCH model for ϵ_t. The within time posterior means for TN_t and μ_{N_t} are essentially the same. The posterior distributions for TN_t are nearly symmetrical and appear to be approximately normally distributed. The representative ξ_t distributions all have means close to the mean for the prior distribution. The posterior distributions are fairly symmetric; values for ξ_t are variable, ranging from -1.82 to 1.01 and do contain the observed estimates based on the one realization of distrometer data. The representative ϕ_t^2 distributions all have means close to 0.18, the mean of the prior distribution. The posterior distributions are right-skewed; values for ϕ_t^2 are variable, ranging from 0.05 to 1.03 and do contain the observed estimates calculated from the distrometer data.

TABLE 2. MCMC Summary statistics for time independent parameters.

Model Parameter	Prior Mean	Prior SD	MCMC Posterior Mean (SE)	
			Verification Run	Full Run
GARCH Coefficients				
a_0	0.023	0.0062	0.0080 (0.00012)	0.049 (0.0012)
a_1	0.242	0.1007	0.200 (0.00082)	0.164 (0.0020)
b_1	0.742	0.1119	0.677 (0.00099)	0.684 (0.0031)
AR(1) Process				
a_N	0.95	0.04	0.981 (0.00018)	0.990 (0.0002)
σ_η^2	0.005	0.0007	0.418 (0.0023)	0.478 (0.0033)
Measurement Error				
σ_G^2	5.0	1.02	—	5.82 (0.044)
σ_Z^2	2.86	1.429	31.68 (0.132)	104.01 (0.865)

The in-series variability for TN_t, ξ_t, and ϕ_t^2 is greater than that for the observed distrometer data. Except for ϕ_t^2, the time series values estimated from the distrometer data are nearly centered in the posterior model results. As with the verification run, the full model underestimates the DSD shape parameter, ϕ_t^2.

From the posterior parameter estimates for the hidden DSD process, we can estimate the integrated values for rain rate and derived reflectivity. The time series of calculated rain rates show higher variability than that produced from the one realization of distrometer data. The large variability makes it difficult to compare the signals, though there appears to be some amount of agreement in the gross signal features for rain rate (Figure 6).

The time series plots of derived reflectivity also show higher variability than that produced from one realization of distrometer data, though not to the same extent as for rain rate (Figure 7). The model output consistently overestimates distrometer derived values. While the cause of the bias is unknown, it's possible that the observed bias indicates poor calibration of the radar. A mean 1 normalized plot for reflectivity (Figure 8) removes the bias, but basically shows that the fitted reflectivity estimates are stationary. This is contrary to what is observed for reflectivity values derived from the distrometer data, and may reflect a lack of

TABLE 3. MCMC summary statistics for time varying parameters for three time points. For the verification run, times are $t = 9, 24, 39$. For the full run, $t = 90, 235, 391$.

	Verification Run		**Full Model Run**	
Model Parameter	**Observed Value**	**Posterior Mean (SE)**	**Observed Value**	**Posterior Mean (SE)**
Total # Drops				
TN_{t1}	6.41	6.00 (0.026)	4.37	5.00 (0.056)
TN_{t2}	7.59	6.36 (0.029)	5.90	5.28 (0.055)
TN_{t3}	6.64	6.29 (0.028)	5.08	5.04 (0.051)
Mean # Drops				
$\mu_{N_{t1}}$	—	6.00 (0.026)	—	5.01 (0.052)
$\mu_{N_{t2}}$	—	6.34 (0.028)	—	5.25 (0.049)
$\mu_{N_{t3}}$	—	6.30 (0.026)	—	5.05 (0.048)
Errors for TN_t				
ϵ_{t1}	—	−0.0001 (0.008)	—	−0.017 (0.0214)
ϵ_{t2}	—	0.005 (0.008)	—	−0.016 (0.0221)
ϵ_{t3}	—	0.012 (0.010)	—	−0.023 (0.0217)
Error Variance				
$\sigma^2_{\epsilon_{t1}}$	—	0.050 (0.0025)	—	0.320 (0.0077)
$\sigma^2_{\epsilon_{t2}}$	—	0.071 (0.0036)	—	0.329 (0.0084)
$\sigma^2_{\epsilon_{t3}}$	—	0.080 (0.0040)	—	0.319 (0.0075)
Lognormal Scale Parameter Prior mean: −0.50; SD: 0.400				
ξ_{t1}	−0.54[a]	−0.68 (0.009)	−0.56[b]	−0.49 (0.015)
ξ_{t2}	−0.59	−0.46 (0.009)	−0.61	−0.54 (0.015)
ξ_{t3}	−0.63	−0.49 (0.009)	−0.63	−0.53 (0.015)
Lognormal Shape Parameter Prior mean: 0.182; SD: 0.0857				
ϕ_{t1}	0.316[c]	0.155 (0.0018)	0.26[d]	0.18 (0.003)
ϕ_{t2}	0.338	0.176 (0.0023)	0.34	0.19 (0.003)
ϕ_{t3}	0.371	0.166 (0.0023)	0.37	0.19 (0.004)

[a] Estimated from 10-minute distrometer data
[b] Estimated from 1-minute distrometer data
[c] Estimated from 10-minute distrometer data
[d] Estimated from 1-minute distrometer data

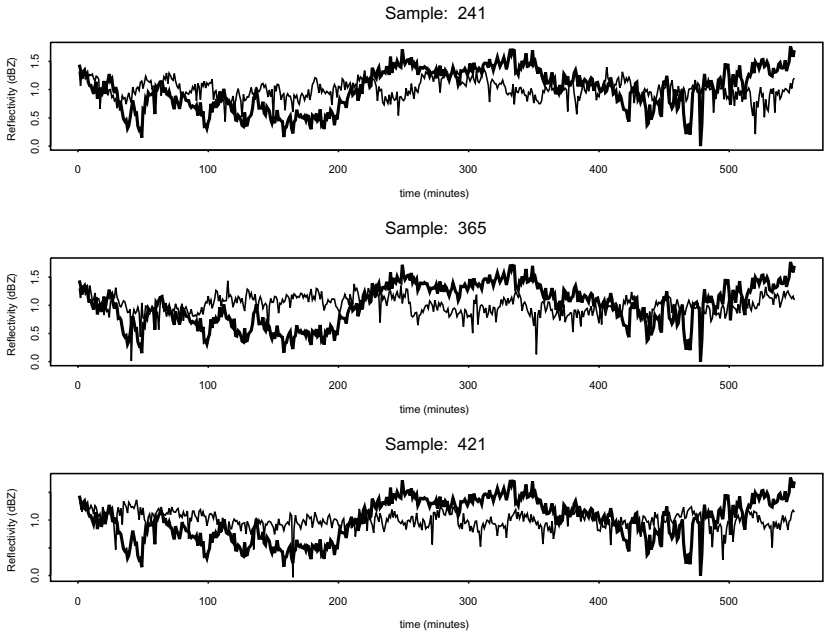

FIGURE 8. Time series full model run posterior reflectivity (normalized to mean 1) estimates calculated for selected samples (light lines) with reflectivity values based on distrometer data (heavy lines).

structure in parameters used to model the DSD shape. In general, a comparison of the signals between any two of the time series is difficult since we are trying to compare multiple realizations based on model output with one realization based on a set of distrometer data.

An alternate visualization of the calculated rain rate parameter based on model output displays histograms for rain rates from 12 points along the time series (Figure 9). Corresponding values calculated from the distrometer data are also shown. The individual histograms contain the distrometer-based rain rates, although the distrometer-based values are in the lower ends of the distributions for most time points.

4. Discussion

The model results for the Eureka rainfall are promising in that the use of hierarchical Bayes methods make it possible to integrate information from multiple data sources for purposes of estimating rain rate and related parameters. While we are unable to validate the posterior estimates of TN_t, ξ_t, and ϕ_t^2, the main parameters

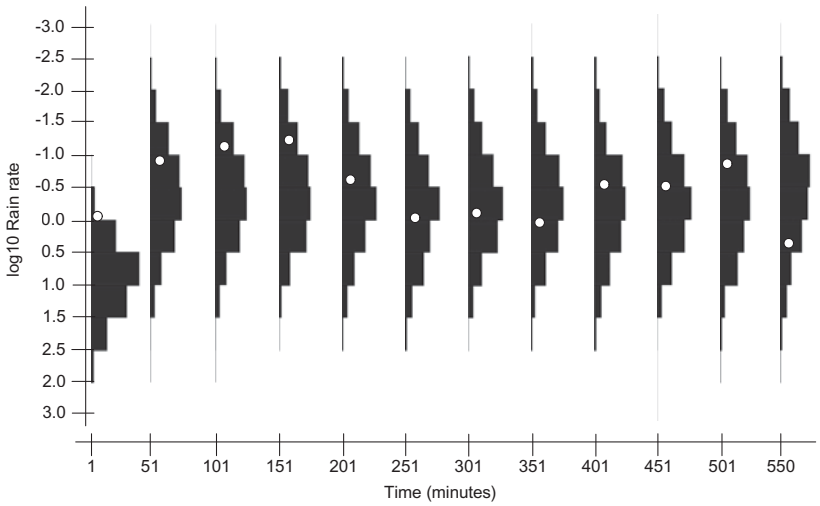

FIGURE 9. Frequency histograms of the full model run posterior distribution of $\log_{10}$ rain rate estimates for selected time points within the 550 minute modeled event. All histograms are scaled with equivalent y-axes. The white circles correspond to the values of rain rate calculated from one realization of distrometer data.

for the unobserved DSD process, we can compare their values with the one available distrometer realization. For TN_t, many of the MCMC realizations, though not all, produce time series signals similar to that observed for the distrometer realization with model output that is always more variable. This higher variability is expected, and is due in part to the higher uncertainty that is introduced when observations from two different data sources (gauge and radar) are considered.

The DSD scale and shape parameters capture the gross features of the drop spectra, but the time series for ξ_t and ϕ_t^2 are many times more variable than those produced from distrometer estimates. A comparison of autocorrelations for these parameters based on distrometer data versus model output suggests that there is more structure in the time evolution of the DSD shape than what is accounted for in the current model. Estimates of the parameters for the lognormal shape distribution based on distrometer data are highly autocorrelated out to about 25 minutes. The estimates based on the hierarchical model are essentially uncorrelated, corresponding with the implemented model parametrization for DSDs.

Ultimately the estimated DSD parameters are used to produce estimates of the integrated quantities rain rate and derived reflectivity. The verification runs do a good job of capturing the features of these quantities; the full model runs capture the gross features, but again produce higher variable time series output and overestimates of derived reflectivity. Aside from poor calibration of the radar, the bias seen in derived reflectivity could be due to a number of other factors;

most notably, differences may be related to sampling errors associated with distrometer data. Derived reflectivity is more adversely affected by small differences in drop-size distributions since Z_D is a function of D^6 while R is a function of D^3. Either overestimation of larger drops by the model, or under recording of larger drops by the distrometer could lead to these observed differences. Smith et al. [23] showed a resulting skewness in the sampling distributions for DSD spectra when the samples are small (size in raw counts). This skewness tends to result in an underestimate of the integrated quantities. Verification runs produced better consistency between model derived reflectivity and distrometer derived values; use of 10-minute distrometer data helps compensate for distrometer sampling errors that arise due to the small sample volume collected in a one-minute interval.

There is no measurable truth when it comes to determining rain rates at the earth's surface. For each instrument deployed in the field, there are claims and beliefs as to how well each performs their respective tasks, in light of uncertainty. Instead of using rain gauge data as "ground truth", the posterior model distributions capture multiple sources of uncertainty that reflect uncertainty associated with one or more sources such as: the inability to directly measure quantities of interest; measurement and sampling errors associated with the multiple instruments; sparse gauge data; and, incorrect or inadequate assumptions and specifications related to equations used to represent the underlying physical processes.

As a proof of concept, we were able to incorporate Eureka rainfall data into a hierarchical Bayes model to generate estimates for the modeled parameters. Most of the estimated parameter values seem to track to some extent the information contained in the single distrometer realization. Using existing model runs as baselines we can further characterize the observed uncertainties by making stepwise modifications to model stages or by incorporating additional data. In the short term, a temporal dynamic stage should be developed for the evolution of DSD through time. Also, we should run the model for other rainfall events. The particular Eureka event consisted of relatively light rainfall. Determining how well the model performs for heavier rainfall events will provide some idea as to the portability/usability of this model under a variety of rainfall scenarios. In the longer term, some ideas for expanding and using this modeling approach include:

1. If applicable, incorporate wind corrections for gauge observations in future model runs. Gauge catchments are inherently affected by winds; a wind effect would be one component of the overall uncertainty observed in the model output.
2. Incorporate a vertical correction model component for radar reflectivity observations, based on an underlying process model, to compensate for the fact that ground-based radar provides information about precipitation above the earth's surface. Include data from vertical pointing radar to augment the vertical correction model component.
3. Expand the time only model to a space-time model. Such a model could provide a tool that allows for the creation of spatial rainfall maps with un-

certainty estimates, while also providing the opportunity for more expanded analysis of the component uncertainties.

4. Use model output to estimate better Z-R relationships dynamically.
5. Use model output to assess the value added by using other types of precipitation measurements (or by removing some instruments).

Acknowledgment

We would like to thank Dr. Sandra Yuter, who provided the distrometer and NEXRAD radar data sets, for her help with data pre-processing and her invaluable insights related to instrumentation, meteorological properties of precipitation, and relationships between the two.

References

[1] C. Alexander. Volatility and correlation: measurement, models, and applications. In C. Alexander, editor, *Risk Management and Analysis. Vol.* 1. John Wiley Sons Ltd, 1998.

[2] Louis J. Battan. *Radar observation of the atmosphere*. University Chicago Press, Chicago, revised edition, 1973.

[3] Thomas L. Bell. A space-time stochastic model of rainfall for satellite remote-sensing studies. *Journal of Geophysical Research*, 92(D8):9631–9643, 1987.

[4] L. Mark Berliner. Hierarchical Bayesian modeling in the environmental sciences. *Allgemeines Statistisches Archiv*, 84:141–153, 2000.

[5] L. Mark Berliner, Richard A. Levine, and Dennis J. Shea. Bayesian climate change assessment. *Journal of Climate*, 13(21):3805–3820, 2000.

[6] D.R. Cox and Valerie Isham. A simple spatial-temporal model of rainfall. *Proceedings of the Royal Society of London*, A415:317–328, 1988.

[7] Greg Hammer and Pete Steurer. Data documentation for hourly precipitation data td-3240. Technical report, National Climatic Data Center, February 11, 1998.

[8] E. Hanna. How effective are tipping-bucket raingauges? A review. *Weather*, 50(10): 336–342, 1995.

[9] B. Hrafnkelsson. Hierarchical modeling of count data with application to nuclear fall-out. *Journal of Environmental and Ecological Statistics*, 10:179–200, 2003.

[10] A.J. Illingworth and M.T. Blackman. The need to represent raindrop size spectra as normalized gamma distributions for the interpretation of polarized radar observations. *Journal of Applied Meteorology*, 41:286–297, 2002.

[11] H. Jiang, M. Sano, and M. Sekine. Weibull raindrop-size distribution and its application to rain attenuation. *IEE Proceedings-Microwaves, Antennas and Propagation*, 144(3):197–200, 1997.

[12] E.L. Neff. How much rain does a rain gage gage? *Journal of Hydrology*, 35:213–220, 1977.

[13] A. Raftery and S. Lewis. Implementing mcmc. In Gilks W., S. Richardson, and D. Spiegelhalter, editors, *Markov Chain Monte Carlo in Practice*. Chapman Hall, London, 1996.

[14] G.O. Roberts. Markov chain concepts related to sampling algorithms. In Gilks W., S. Richardson, and D. Spiegelhalter, editors, *Markov Chain Monte Carlo in Practice*. Chapman Hall, London, 1996.

[15] Ignacio Rodriguez-Iturbe, D.R. Cox, and V. Isham. A point process model for rainfall: further developments. *Proceedings of the Royal Society London*, A417:283–298, 1988.

[16] Ignacio Rodriguez-Iturbe and Peter S. Eagleson. Mathematical models of rainstorm events in space and time. *Water Resources Research*, 23(1):181–190, 1987.

[17] J.A. Royle, L.M. Berliner, C.K. Wikle, and R. Milliff. A hierarchical spatial model for constructing wind fields from scatterometer data in the Labrador sea. In *Case Studies in Bayesian Statistics IV*. Springer-Verlag, 1998.

[18] Henri Sauvageot and Jean-Pierre Lacaux. The shape of averaged drop size distributions. *Journal of the Atmospheric Sciences*, 52(8):1070–1083, 1995.

[19] M. Sekine, S. Ishii, I. Hwang, and S. Sayama. Weibull raindrop-size distribution and its application to rain attenuation from 30 ghz to 1000 ghz. *International Journal of Infrared Millimeter Waves*, 28:383–392, 2007.

[20] Matsuo Sekine and Goran Lind. Rain attenuation of centimeter, millimeter and submillimeter radio waves. In *Proceedings of the 12th European Microwave conference*, pages 584–589, Helsinki, Finland, 1982.

[21] T. Skaugen, J.D. Creutin, and L. Gottschald. Reconstruction and frequency estimates of extreme daily areal precipitation. *Journal of Geophysical Research*, 101(D21):26287–25295, 1996.

[22] James A. Smith. Marked point process models of raindrop-size distributions. *Journal of Applied Meteorology*, 32:284–296, 1993.

[23] James A. Smith and Richard D. De Veaux. The temporal and spatial variability of rainfall power. *Environmetrics*, 3(1):29–53, 1992.

[24] R. Stern and R. Coe. A model fitting analysis of daily rainfall data. *Journal of the Royal Statistical Society A*, 147(Part 1); 1–34, 1984.

[25] J. Testud, E. Le Bouar, E. Obligis, and M. Ali-Mehenni. The rain profiling algorithm applied to polarimetric weather radar. *Journal of Atmospheric and Oceanic Technology*, 17:332–356, 2000.

[26] Carlton W. Ulbrich. Natural variations in the analytical form of the raindrop size distribution. *Journal of Climate and Applied Meteorology*, 22:1764–1775, 1983.

[27] Carlton W. Ulbrich and David Atlas. Rainfall microphysics and radar properties: analysis methods for drop size spectra. *Journal of Applied Meteorology*, 37:912–923, 1998.

[28] A. Waldvogel. The n0 jump of raindrop spectra. *Journal of the Atmospheric Sciences*, 31:1067–1078, 1974.

[29] Christopher K. Wikle. Hierarchical Bayesian space-time models. *Environmental and Ecological Statistics*, 5:117–154, 1998.

[30] Christopher K. Wikle, R.F. Milliff, D. Nychka, and M. Berliner. Spatiotemporal hierarchial Bayesian modeling: tropical ocean surface winds. *Journal of the American Statistical Association*, 96(454):382–397, 2001.

[31] Paul T. Willis. Functional fits to some observed drop size distributions and parameterization of rain. *Journal of Atmos. Sci.*, 41:1648–1661, 1984.

[32] Paul T. Willis and Paul Tattelman. Drop-size distributions associated with intense rainfall. *Journal of Applied Meteorology*, 28:3–15, 1989.

Tamre P. Cardoso and Peter Guttorp
Department of Statistics
Box 354322, University of Washington
Seattle, WA 98195-4322
e-mail: tamre@u.washington.edu

e-mail: peter@stat.washington.edu