Long Temporal Autocorrelations in Tropical Precipitation Data and Spike Train Prototypes

Tristan H. Abbott,[1] Samuel N. Stechmann,[2,3] J. David Neelin[4]

September 20, 2016

Key Points.

- Power-law ranges in precipitation autocorrelation reflect emergent behavior from exiting and reentering the precipitating regime
- Dry-spell processes play a key role in producing the inter-event correlations that control long precipitation autocorrelations
- Precipitation time series can be viewed as spike trains with short events and long dry spells
- **Index terms**: 3354 Precipitation, 3270 Time series analysis, 3371 Tropical convection, 3265 Stochastic Processes, 3235 Persistence, memory, correlations, clustering

Abstract:
Temporal precipitation autocorrelations drop slower than exponentially at long lags, and there is a range from tens to thousands of minutes where it is relevant to ask if a scale-free process might underlie the long autocorrelations. A simple stochastic model in which precipitation appears as variable-length spikes provides a reasonable prototype for this behavior. In both observations and the model, separating the component of the autocorrelation within wet events from the inter-event contribution suggests long autocorrelation behavior is primarily associated with the latter. When precipitation spikes are short compared to dry events, a true power law is obtained with analytical exponent -0.5 and precipitation autocorrelation is determined by dry-spell model parameters. In more realistic cases, wet-spell termination is also important. Although a variety of apparent power law exponents can be obtained for different parameters, the fundamental long-lag process appears to be that of the inter-event correlation.

[1]Department of Computer Sciences, University of Wisconsin-Madison, Madison, Wisconsin, USA.
[2]Department of Mathematics, University of Wisconsin-Madison, Madison, Wisconsin, USA.
[3]Department of Atmospheric and Oceanic Sciences, University of Wisconsin-Madison, Madison, Wisconsin, USA.
[4]Department of Atmospheric and Oceanic Sciences, University of California-Los Angeles, Los Angeles, California, USA.

0094-8276/16/$5.00

1. Introduction

Single-column temporal autocorrelations in tropical precipitation data decay much more quickly than column-integrated water vapor [*Holloway and Neelin*, 2010] yet simultaneously display slower than exponential decay at lags of tens to thousands of minutes. Ranges of scale-free, power-law behavior in spatial autocorrelation have been noted in studies of tropical convective onset that examine statistics motivated by the analogy to critical phase transitions [*Peters and Neelin*, 2006]. Temporal precipitation autocorrelations exhibit a power law-like form at subhourly to multi-day lags, so a natural question is whether their slow decay can be explained by a temporal analog to the scale-free range in spatial precipitation autocorrelations.

Because the transition into and out of strong convection occurs on time scales of minutes to days, time series and models with high temporal resolution are important for understanding the transition to strong convection and precipitation statistics that arise from it. Work on simple statistical models of precipitation (sometimes called weather generators) has historically focused primarily on daily precipitation [*Wilks and Wilby*, 1999]. Temporal autocorrelations in daily precipitation models are typically assumed negligible [*Katz*, 1977; *Buishand*, 1977; *Foufoula-Georgiou and Lettenmaier*, 1987], and the timescales are too long for our application. Adaptations of daily weather generators based on Markov chains for hourly precipitation data [*Katz and Parlange*, 1995] rely on tuning transition probabilities to fit observations and are somewhat difficult to interpret physically. Other sub-daily precipitation models are physically based but primarily address subgrid scale variations in more complex models (e.g. *Lin and Neelin* [2000]; *Majda and Khouider* [2002]), or model multiple air columns (e.g. *Khouider et al.* [2010]; *Hottovy and Stechmann* [2015]) or advection processes through dry spells with eventual condensation (e.g. *Pierrehumbert* [1998]; *Galewsky et al.* [2005]; *O'Gorman and Schneider* [2006]; *Pierrehumbert et al.* [2007]; *Sukhatme and Young* [2011]; *O'Gorman et al.* [2011]). In this study, we use a first-passage-time model developed by *Stechmann and Neelin* [2014] (hereafter the SN14 model) that produces single-column precipitation data with sub-minute granularity and is based on simple physical assumptions about the relationship between column water vapor and precipitation near the transition to strong convection.

In what follows, we first establish a set of connections between temporal autocorrelations in observational data and the SN14 model (Section 2). Next, we focus on long lags, where the slower-than-exponential decay appears (Section 3), and develop an analytically accessible model idealization for long-range autocorrelations (Sections 4 and 5.1). Finally, we use the idealization to analyze the SN14 model autocorrelation (Section 5.2). We conclude by discussing the relationship between autocorrelation behavior and various atmospheric processes.

2. Autocorrelations from Observations and Simple Models

First, we compare temporal autocorrelations in observational tropical precipitation records and in the SN14 model. For a time series P_t with overall mean μ and overall variance σ^2, we calculate its autocorrelation at lag τ as

$$R(\tau) = \frac{1}{N}\sum_{i=1}^{N}\frac{(P_{t_i}-\mu)\,(P_{t_i+\tau}-\mu)}{\sigma^2}. \quad (1)$$

In brief, precipitation autocorrelations in observations and the SN14 model contain several key similarities: they decay slower than exponentially and have apparent power law ranges with similar apparent exponents.

2.1. Observations

The observational precipitation data used in this paper are from 1 minute average optical gauge time series recorded on Manus and Nauru Islands. The Manus time series contains data from August 27, 2004 to September 14, 2012, and the Nauru time series contains data from September 15, 2003 to November 19, 2011. To reduce seasonal effects, autocorrelations were computed in 3-month nonoverlapping windows (corresponding to December-February, March-May, June-August, September-November seasons, excluding partial segments at the start and end of each time series) and averaged at each lag over all segments. Additional data analysis details are given in the SI.

The Manus and Nauru autocorrelation functions are shown in Figure 1. Holloway and Neelin [*Holloway and Neelin*, 2010] previously identified an apparent power law range with an estimated exponent of -0.85 in the autocorrelation of the Nauru optical gauge data. The Nauru autocorrelation is reproduced in Figure 1b along with a plot of the autocorrelation of the Manus optical gauge data (Figure 1a). Short-lag exponential fits to autocorrelation functions (shown as dot-dashed lines in the figures, with time scales 21 and 15 min. respectively) confirm that, at long lags, the Nauru and Manus autocorrelations decay slower than exponentially. Both autocorrelations contain ranges in which the decay is roughly power law with apparent exponents between -0.7 and -1.1. The autocorrelations for lags longer than about 1 day exhibit more complex behavior including small apparent peaks, but it is the overall background decay that is of interest here. Removing diurnal cycles from the observational time series had a negligible effect on their autocorrelations, as detailed in the S.I. At longer lags, seasonal signals affect the observational autocorrelations. In particular, the Nauru autocorrelation plateaus at long lags if the autocorrelation is calculated for the entire time series rather than seasonal segments. The autocorrelations for the full Nauru and Manus time series are given in the S.I. in Figure S1.

2.2. Model

The SN14 model is a first-passage-based two state model for the evolution of column water vapor (CWV) $q(t)$ and precipitation. The behavior of CWV is governed by a set of stochastic differential equations:

$$\frac{dq}{dt} = E_* + D_0\dot{\xi} \quad \text{if non-precipitating} \quad (2)$$

$$\frac{dq}{dt} = -P_* + D_1\dot{\xi} \quad \text{if precipitating.} \quad (3)$$

E_* and P_* are constant and represent small-scale moisture fluxes. Forcing by large-scale moisture convergence is represented by a white noise term with variance D_0^2 in the non-precipitating state and variance D_1^2 in the precipitating state. Precipitation begins in the SN14 model when CWV rises above a critical value q_c, and precipitation ceases when CWV falls below a second critical value $q_{np} < q_c$. We use b to denote $q_c - q_{np}$. The precipitation produced by the model is 0 in the non-precipitating state and P_* in the precipitating state.

Precipitation statistics in SN14 model share key similarities with observations. In particular, the SN14 model produces dry-spell duration, wet-spell duration, and event size probability density functions (PDFs) that are power laws with exponential cutoffs. Cutoff locations are set by time scales that depend on model parameters:

$$t_{Sw} = \frac{b^2}{2D_1^2} \quad \text{(wet-spell short-}t\text{ cutoff)} \quad (4)$$

$$t_{Lw} = \frac{2D_1^2}{P_*^2} \quad \text{(wet-spell long-}t\text{ cutoff)} \quad (5)$$

$$t_{Sd} = \frac{b^2}{2D_0^2} \quad \text{(dry-spell short-}t\text{ cutoff)} \quad (6)$$

$$t_{Ld} = \frac{2D_0^2}{E_*^2} \quad \text{(dry-spell long-}t\text{ cutoff).} \quad (7)$$

The dry-spell duration PDF is

$$p_{t0}(t) = \sqrt{\frac{t_{Sd}}{\pi}}\exp\left(\sqrt{\frac{4t_{Sd}}{t_{Ld}}}\right)\exp\left(-\frac{t_{Sd}}{t}\right)\exp\left(-\frac{t}{t_{Ld}}\right)t^{-3/2}. \quad (8)$$

The PDFs for wet-spell durations and event sizes likewise have a $t^{-3/2}$ power law and differ only in the location of the exponential cutoffs and in the normalization constant. The SN14 model also produces mean precipitation pickup and a precipitation variance spike for q near the critical value for onset. SN14 model time series were generated using the Euler-Maruyama method [e.g. *Higham*, 2001]. Details of the simulations and autocorrelation computations are given in the SI.

At suitable parameter settings, SN14 precipitation autocorrelations are similar in key ways to autocorrelations of observational precipitation data. Figure 1c displays the autocorrelation of SN14 precipitation data for one such parameter setting: $P_* = 10$ mm h^{-1}, $E_* = 0.3$ mm h^{-1}, $D_1^2 = 64$ mm^2 h^{-1}, $D_0^2 = 8$ mm^2 h^{-1}, $q_c = 65$ mm, and $q_{np} = 62$ mm. An exponential fit to short lags (with time scale 13 min.) confirms that, like the observational autocorrelations, the SN14 model autocorrelation decays slower than exponentially. Furthermore, the SN14 autocorrelation for these parameter settings has an approximately power law range with an apparent exponent similar to that seen in observational autocorrelations.

3. Intra- and Inter-Event Contributions to Autocorrelations

We now ask: to what extent do different subprocesses control the decay in precipitation autocorrelations? To distinguish between wet-spell subprocesses (represented by P_* and D_1^2 in the SN14 model) and dry-spell subprocesses (represented by E_* and D_0^2 in the SN14 model), we employ an

intra-/inter-event decomposition of the precipitation autocorrelation.

The intra-event component $R_a(\tau)$ of an autocorrelation is the contribution to the autocorrelation from pairs of points inside the same precipitation event. Let $X(t, t+\tau) = 1$ if no data points for times in the range $[t, t+\tau]$ have precipitation rates of less than 0.3 mm/h, and 0 otherwise. Then the intra-event autocorrelation is

$$R_a(\tau) = \frac{1}{N} \sum_{i=1}^{N} \frac{X(t_i, t_i+\tau)\,(P_{t_i} - \mu)\,(P_{t_i+\tau} - \mu)}{\sigma^2} \quad (9)$$

The inter-event component $R_r(\tau)$ is the contribution from all other pairs of points:

$$R_r(\tau) = \frac{1}{N} \sum_{i=1}^{N} \frac{(1 - X(t_i, t_i+\tau))\,(P_{t_i} - \mu)\,(P_{t_i+\tau} - \mu)}{\sigma^2} \quad (10)$$

For Manus and Nauru, we calculated an overall autocorrelation decomposition by splitting the time series into 3-month, nonoverlapping seasonal segments as in section 2.1, calculating $R_a(\tau)$ and $R_r(\tau)$ for each segment, and averaging at each τ over every segment's R_a and R_r. The SI contains additional details about the methods used to calculate autocorrelation decompositions.

The intra-event component of a precipitation autocorrelation decomposition is a clear measure of the extent to which correlations from within a single precipitation event control the overall autocorrelation. The inter-event component contains contributions from pairs of points in different wet-spells, in different dry-spells, in the same dry-spell, and split between a wet-spell and a dry-spell. For us, the key attribute of the inter-event autocorrelation is that it represents an autocorrelation that is controlled largely by dry-spell processes.

Plots of $R_a(t)$ and $R_r(t)$ are shown in Figure 1 for Manus (d), Nauru (e), and the SN14 model (f). Observed decompositions from Manus, Nauru, and the SN14 model share several key characteristics. First, the intra-event component dominates at short lags. Second, there is an overlap region at intermediate lags where both intra- and inter-event components make significant contributions to the overall autocorrelation. Finally, the inter-event component dominates at long lags. In the model, those ranges appear to be set by the cutoffs given in Equations 4-7. The intra-event component dominates approximately between the wet-spell short-time cutoff and the wet-spell long-time cutoff, and the inter-event component dominates approximately between the dry-spell short-time cutoff and the dry-spell long-time cutoff. Apparent power law ranges appear initially in the overlap region and persist into the region in which the inter-event component dominates. The most important features are that the intra-event component decays quickly at long lags and that long range power law correlations within both observed time series are controlled by inter-event interactions.

Because of the substantial similarities between the autocorrelations and autocorrelation decompositions in observations and the SN14 model, the SN14 model provides a prototype with which to study the processes that produce a slower than exponential decay in observed autocorrelations. The model is not without limitations; in particular, it does not have synoptic, intraseasonal or seasonal variations, and thus the autocorrelation at intraseasonal lags has a relatively simple cutoff that terminates the approximate power law range. However, focusing on the inter-event component of the model autocorrelation component provides insight into the degree to which scale free behavior produces the slower than exponential decay.

4. Parameter Dependence and Limiting Cases of the SN14 Model

In this section, we establish that approximate power law ranges appear in model output autocorrelations for a variety of model parameters and then introduce a limiting case that isolates the inter-event contribution to the overall autocorrelation and displays true scale free behavior at the expense of some realism.

The observation-like model autocorrelation shown in sections 2 and 3 was produced using model parameters that set the event duration cutoffs at $t_{Sw} = 4.2$ min, $t_{Lw} = 76.8$ min, $t_{Sd} = 33.8$ min, and $t_{Ld} = 10,670$ min. That autocorrelation is reproduced in Figure 2a. Fixing t_{Sw} and increasing (resp. decreasing) the other cutoffs stretches (resp. compresses) the autocorrelation without removing the power law range and without significantly altering the qualitative properties of the autocorrelation decompositions (see Figure 2a-c). However, changing cutoff locations does not preserve the apparent exponents of power law ranges. Generally, moving the cutoffs closer together increases the apparent exponent, and moving the cutoffs farther apart decreases the apparent exponent. (Details about exponent estimations and parameter settings for Figure 2 are given in the SI.)

In all of Figures 2a-c, the intra-event component of the autocorrelation decays quickly and power law ranges are produced in large part by the inter-event component. Although the inter-event autocorrelation appears to be the primary long-lag process, it is not immediately obvious to what extent it is a scale-free process for realistic parameters. The presence of a power law with an exponent that is invariant with cutoff locations is a defining characteristic of a scale-free process [e.g., *Peters and Neelin*, 2009], but in these figures, moving the event duration cutoffs affects the apparent exponent of the power law range.

Lastly, consider a limiting case of the SN14 model that isolates the inter-event component. Fixing t_{Sw} and t_{Sd} at the values for Figure 2a and bringing t_{Lw} close to t_{Sw} removes most of the overlap region and produces an apparent power law range entirely dominated by the inter-event component. Taking t_{Ld} out to long lags extends the power law; once t_{Ld} is sufficiently large, the exponent of the power law stabilizes near -0.5 (see Figure 2d). This well-separated case is the only case that we found that yields numerical results that are consistent with true scale freedom. An analytic theory for scale freedom in this limiting case is developed in the next section.

5. Point-process Idealization

In numerical experiments, precipitation autocorrelations in the SN14 model have power laws with stable exponent -0.5 when (i) precipitation events are short compared to dry spells and (ii) dry spell cutoffs are well separated. An analytic argument that supports the numerical experiments can

be can be developed by considering a point-process idealization of the SN14 model. This idealization treats precipitation events as single points in time with infinitely large precipitation rates. Additional details about its development are given in the SI. We will show that in this limit, the behavior of the model depends only on dry regime parameters E_* and D_0^2.

5.1. Spike Trains, Point Processes, and Precipitation

Durations of observed precipitation events are on average significantly shorter than the durations of intervening dry-spells [*Pierrehumbert et al.*, 2007; *Peters et al.*, 2010; *Deluca and Corral*, 2014]. Time series with short-duration events, or "spike trains", are common in studies of neuron dynamics, where high-potential periods in a neuron are short compared with intervening periods at the neuron's resting potential [see e.g. *Gerstner and Kistler*, 2002], and calculating autocorrelations of neuronal spike trains can be done by approximating the spike trains as temporal point processes, i.e. as sums of Dirac deltas.

A point-process idealization of the SN14 model, where precipitation time series become sums of Dirac deltas

$$P(t) = b \sum_{t_i} \delta(t - t_i), \tag{11}$$

is useful for studies of long-range autocorrelations in precipitation time series. Additional details about the idealization are given in the SI. Figure 3 shows sample precipitation time series from observations at Manus and Nauru, from a run of the full SN14 model, and from a run of the SN14 point-process idealization (with spikes scaled to 10 mm h^{-1}). Over 30 day periods (shown on the main axes), individual precipitation events in the observed time series appear reasonably spike-like, and the time series from the full SN14 model is virtually indistinguishable from the spike train idealization. Over 1000 minute periods (shown in inset plots), the limitations of the spike train idealization are obvious: both observed time series and full SN14 time series contain precipitation events of substantial duration. Beyond these time scales, the point-process idealization appears to be a valuable tool for studying medium-to-long-range autocorrelations in precipitation data.

5.2. SN14 Point Process Autocorrelation

An analytical technique for calculating the autocovariance of a temporal point process is presented in *Gerstner and Kistler* [2002]. Here, it is used along with asymptotics to derive a true power law with exponent -0.5 in the autocovariance of the SN14 point-process idealization. This section contains a summary of our results; a more detailed derivation is given in the SI.

For the SN14 point-process idealization (11), its autocovariance $R_*(\tau)$ has a Fourier transform of

$$\hat{R}_*(\omega) = \frac{b^2}{\sqrt{t_{Sd} t_{Ld}}} \Re \left\{ \frac{1 + \hat{p}(\omega)}{1 - \hat{p}(\omega)} \right\} + 2\pi \frac{b^2}{t_{Sd} t_{Ld}} \delta(\omega). \tag{12}$$

See, e.g., *Gerstner and Kistler* [2002] for a derivation. Here, $\hat{p}(\omega)$ denotes the Fourier transform of the interspike interval distribution and $\Re$ denotes the real part of a complex variable. For the SN14 model, the interspike interval is governed by the dry-spell duration PDF $p_{t0}(t)$ given in Equation 8, and its Fourier transform is

$$\hat{p}(\omega) = \exp\left[2\sqrt{\frac{t_{Sd}}{t_{Ld}}} \left(1 - \sqrt{1 - i\omega t_{Ld}} \right) \right] \tag{13}$$

[see e.g. *Johnson et al.*, 1994]. Note that because the wet-spell cutoffs are implicitly taken small relative to the dry-spell cutoffs, the spike train idealization depends only on dry-spell parameters, i.e., only the dry-spell cutoffs appear in Equations 12 and 13.

Our asymptotics proceed by assuming sufficient separation between dry-spell cutoffs t_{Sd} and t_{Ld} and examining the regime where t_{Ld} is large relative to ω^{-1} and t_{Sd} is small relative to ω^{-1}. This corresponds roughly to the properties of the cutoffs in the limiting case presented in Section 4. It is also a regime in which scale-free behavior is expected: if t_{Sw}, t_{Lw}, and t_{Sd} are all very small and t_{Ld} is very large, there is a large range of lags that should not be influenced by any time scale.

Under these assumptions, it can be shown that

$$\hat{p}(\omega) \approx 1 + (1 - i)\sqrt{2 t_{Sd} \omega}, \tag{14}$$

including the leading-order term 1 and the next-order term $(1 - i)\sqrt{2 t_{Sd} \omega}$. Inserting this form of $\hat{p}(\omega)$ into (12) leads to

$$\hat{R}_*(\omega) \approx \frac{b^2}{\sqrt{2 t_{Sd}^2 t_{Ld}}} \omega^{-1/2}, \;\; \omega > 0. \tag{15}$$

See the supplement for a derivation. Performing an inverse Fourier transform using Equation 15 shows that for sufficiently separated t_{Sd} and t_{Ld}, the autocovariance of the SN14 point-process idealization contains a range with a $\tau^{-1/2}$ power law:

$$R_*(\tau) \approx b^2 \sqrt{\frac{2}{t_{Sd}^2 t_{Ld}}} \tau^{-1/2} + C, \;\; t_{Sd} << \tau << t_{Ld}. \tag{16}$$

The exponent of this power law is consistent with the results of the numerical experiments described in Section 4, and the presence of an analytic power law for lags that are very different from any time scale indicates that power laws in the autocovariance of the SN14 point-process idealization are produced by scale free behavior.

6. Concluding Discussion

Examining precipitation temporal autocorrelations in the range of lags at which slower-than-exponential decay appears, the overall behavior in tropical observed time series is reasonably reproduced in the simple SN14 model. In this system, a stochastic representation of mesoscale to large-scale moisture convergence forces column water vapor variations across a threshold for precipitation onset. Separating the autocorrelation into a contribution from within precipitation events and an inter-event contribution, we find in both observed and model time series that longer lags in the range of interest are controlled primarily by inter-event correlations.

When viewed over the time scales corresponding to those lags, precipitation time series look like spikes separated by long dry spells, and this can be examined as a simplifying limit in the model. Although formally treating precipitation events as spikes with infinitesimal width is not a good approximation for autocorrelations at intermediate lags, it allows for an analytically tractable solution for long-range

autocorrelations with power-law decay. Under this idealization, with certain assumptions about dry spell processes, the SN14 model displays true scale-free behavior that leads to a power-law decay in temporal autocorrelations with exponent -0.5. While one might tend to expect precipitation statistics to be controlled primarily by processes within precipitation events, this result suggests that the slow decay in temporal precipitation autocorrelations is controlled primarily by dry-spell processes. More generally, it draws attention to the fact that dry spell processes can play important roles in precipitation statistics if those statistics describe relationships between distinct wet spells.

Physically, this scale-free limit with short precipitation events occurs when variability of large- to meso-scale moisture convergence is the dominant dry-spell process and moisture sink by precipitation yields very short wet spells. In realistic settings, evaporation, moisture convergence and small-scale precipitation processes each play a role in determining behavior in both precipitating and non-precipitation regimes. Accordingly, realistic autocorrelations are more complex than a simple scale-free process and reflect an emergent behavior from exiting and re-entering the precipitation regime. Nevertheless, these results indicate that dry-spell processes play a key role in the inter-event correlations that produce slower-than-exponential decay at long lags in realistic precipitation autocorrelations.

These conclusions have some practical implications for studies of convective parameterizations. Slower-than-exponential decays in precipitation autocorrelations are not products of convection alone; rather, they arise from interactions between large- to meso-scale processes and the small-scale processes, which remain poorly understood. In the context of a general circulation model (GCM), these correspond to interactions between the large-scale model and sub-grid-scale convective parameterizations. Evaluation of assumptions that convective parameterizations make about these interactions is an ongoing effort in systems from single-column models to cloud resolving models [e.g., *Kuang and Bretherton*, 2006; *Daleau et al.*, 2015]. Because precipitation autocorrelations represent emergent behavior from the interactions between convective and non-convective processes, examining temporal precipitation autocorrelations in GCMs under different circumstances could provide an additional means to constrain aspects of that interaction. Furthermore, it may be possible to extend the approach here to spatial correlations in models like that of *Hottovy and Stechmann* [2015], in existing convective parameterizations and GCMs, or to other natural phenomena that can be represented as point processes and produce slowly decaying temporal autocorrelations.

Acknowledgments. This research is supported in part by Office of Naval Research grant ONR MURI N00014-12-1-0912 (SNS), ONR Young Investigator Award ONR N00014-12-1-0744 (SNS and THA), National Science Foundation grants NSF DMS-1209409 (SNS) and AGS-1540518 (JDN), a Sloan Research Fellowship (SNS) and National Oceanic and Atmospheric Administration grant NA14OAR4310274 (JDN). Precipitation data are from the U.S. Department of Energy Atmospheric Radiation Measurement (ARM) Climate Research Facility Tropical West Pacific field campaign and are available at www.archive.arm.gov. We thank K. Schiro for assistance with this dataset. Computational resources and assistance were provided by the UW-Madison Center For High Throughput Computing (CHTC) in the Department of Computer Sciences. The CHTC is supported by UW-Madison, the Advanced Computing Initiative, the Wisconsin Alumni Research Foundation, the Wisconsin Institutes for Discovery, and the National Science Foundation, and is an active member of the Open Science Grid, which is supported by the National Science Foundation and the U.S. Department of Energy's Office of Science. The authors thank two anonymous reviewers for helpful comments.

References

Buishand, T. (1977), Stochastic Modeling of Daily Rainfall Sequences, *Mededlingen Land- bouwhogeschool Wageningen*, *77*(3).

Daleau, C. L., R. S. Plant, S. J. Woolnough, S. Sessions, M. J. Herman, A. Sobel, S. Wang, D. Kim, A. Cheng, G. Bellon, P. Peyrille, F. Ferry, A. P. Siebesma, and B. Van Ulft (2015), Intercomparison of methods of coupling between convection and large-scale circulation: 1. Comparison over uniform surface conditions, *Journal of Advances in Modeling Earth Systems*, *8*.

Deluca, A., and Á. Corral (2014), Scale invariant events and dry spells for medium-resolution local rain data, *Nonlin. Processes Geophys.*, *21*(2), 555–567.

Foufoula-Georgiou, E., and D. Lettenmaier (1987), A Markov renewal model for rainfall occurences, *Water Resources Research*, *23*, 875–884.

Galewsky, J., A. Sobel, and I. Held (2005), Diagnosis of Subtropical Humidity Dynamics Using Tracers of Last Saturation, *Journal of the Atmospheric Sciences*, *62*, 3353–2267.

Gerstner, W., and W. M. Kistler (2002), *Spiking Neuron Models: Single Neurons, Populations, Plasticity*, Cambridge University Press.

Higham, D. J. (2001), An Algorithmic Introduction to the Numerical Simulation of Stochastic Differential Equations, *SIAM Review*, *43*(3), 525–546.

Holloway, C. E., and J. D. Neelin (2010), Temporal Relations of Column Water Vapor and Tropical Precipitation, *Journal of the Atmospheric Sciences*, *67*(4), 1091–1105.

Hottovy, S., and S. N. Stechmann (2015), A Spatiotemporal Stochastic Model for Tropical Precipitation and Water Vapor Dynamics, *Journal of the Atmospheric Sciences*, *72*, 4721–4738.

Johnson, N. L., S. Kotz, and N. Balakrishnan (1994), *Continuous Univariate Distributions*, Wiley.

Katz, R. W. (1977), Precipitation as a Chain-Dependent Process, *Journal of Applied Meteorology*, *16*(7), 671–676.

Katz, R. W., and M. B. Parlange (1995), Generalizations of chain-dependent processes: Application to hourly precipitation, *Water Resources Research*, *31*(5), 1331–1341.

Khouider, B., J. Biello, and A. J. Majda (2010), A Stochastic Multicloud Model for Tropical Convection, *Communications in Mathematics Sciences*, *8*(1), 187–216.

Kuang, Z., and C. S. Bretherton (2006), A Mass-Flux Scheme View of a High-Resolution Simulation of a Transition from Shallow to Deep Convection, *Journal of the Atmospheric Sciences*, *63*, 1895–1909.

Lin, J. W.-B., and J. D. Neelin (2000), Influence of a Stochastic Moist Convective Paramaterization on Tropical Climate Variability, *Geophysical Research Letters*, *27*(22), 3691–3694.

Majda, A. J., and B. Khouider (2002), Stochastic and Mesoscopic Models for Tropical Convection, *Proceedings of the National Academy of Sciences*, *99*(3), 1123–1128.

O'Gorman, P., and T. Schneider (2006), Stochastic Models fo the Kinematics of Moisture Transport and Condensation in Homogenous Turbulent Flows, *Journal of the Atmospheric Sciences*, *63*, 2992–2005.

O'Gorman, P., N. Lamquin, T. Schneider, and M. S. Singh (2011), The Relative Humidity in an Isentropic Advection-Condensation Model: Limited Poleward Influence and Properties of Subtropical Minima, *Journal of the Atmospheric Sciences*, *68*, 3079–3093.

Olver, F. W., et al. (Eds.) (2010), *NIST Handbook of Mathematical Function*, Cambridge University Press.

Peters, O., and J. D. Neelin (2006), Critical Phenomena in Atmospheric Precipitation, *Nature Physics*, *23*, 393–396.

Peters, O., and J. D. Neelin (2009), Atmospheric Convection as a Continuous Phase Transition, *International Journal of Modern Physics B*, *23*(28-29), 5453–5465.

Peters, O., A. Deluca, A. Corral, J. D. Neelin, and C. E. Holloway (2010), Universality of Rain Event Size Distributions, *Journal of Statistical Mechanics: Theory and Experiment*, *2010*.

Pierrehumbert, R. T. (1998), Lateral mixing as a source of subtropical water vapor, *Geophysical Research Letters*, *25*, 151–154.

Pierrehumbert, R. T., H. Brogniez, and R. Roca (2007), On the Relative Humidity of the Atmosphere, in *The Global Circulation of the Atmosphere: Phemonema, Theory, Challenges*, edited by T. Schneider and A. Sobel, pp. 143–185, Princeton University Press.

Stechmann, S., and J. D. Neelin (2014), First-Passage-Time Prototypes for Precipitation Statistics, *Journal of the Atmospheric Sciences*, *71*, 3269–3291.

Sukhatme, J., and W. R. Young (2011), The advection-condensation model and water-vapor probability density functions, *Quarterly Journal of the Royal Meterological Society*, *137*, 1561–1572.

Wilks, D., and R. Wilby (1999), The weather generation game: a review of stochastic weather models, *Progress in Physical Geography*, *23*(3), 329–357.

Corresponding author: Samuel N. Stechmann, Department of Mathematics, University of Wisconsin-Madison, 480 Lincoln Dr, Madison, WI 53706, USA. (stechmann@wisc.edu)

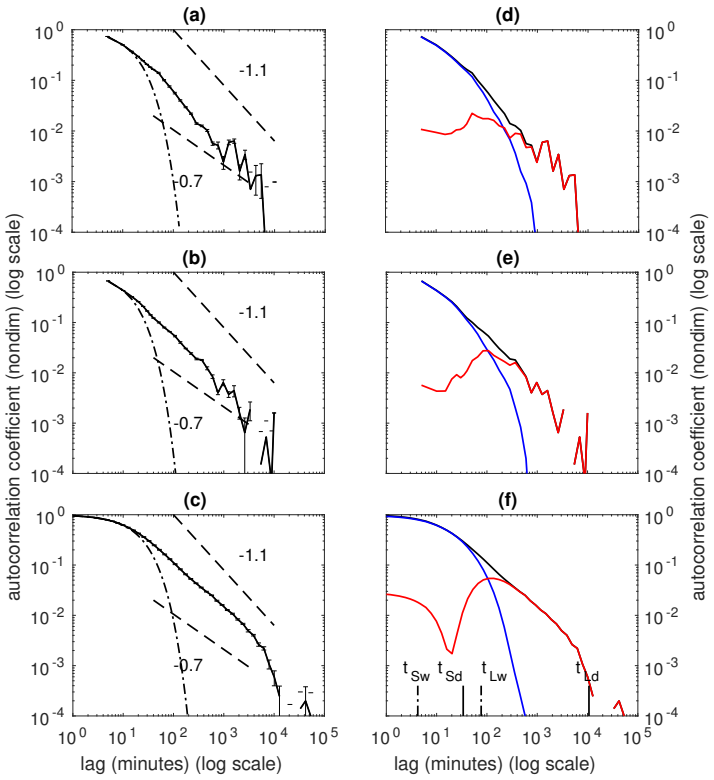

Figure 1. Temporal autocorrelations and autocorrelation decompositions for observed precipitation time series from Manus (a,d) and Nauru (b,e) and precipitation time series from an SN14 two state model simulation (c,f). Note that autocorrelations and autocorrelation decompositions for Manus and Nauru are averages of autocorrelations for 3 month seasonal segments. Because of the time series length limit imposed by the use of 3 month segments, Manus and Nauru autocorrelations are shown only out to lags of 10^4 minutes. On the left, solid black lines indicate autocorrelations after binning, error bars show the standard error of the mean for each bin, dashed black lines show the power laws $t^{-0.7}$ and $t^{-1.1}$, and dot-dashed black lines show an exponential fit from lag 0 to lag 10 minutes. On the right, black lines show the full autocorrelation of the time series, blue lines show the intra-event component, and red lines show the inter-event component. On the lower right, dashed (resp. solid) black ticks indicate the wet-spell (resp. dry-spell) duration cutoffs of the two state model.

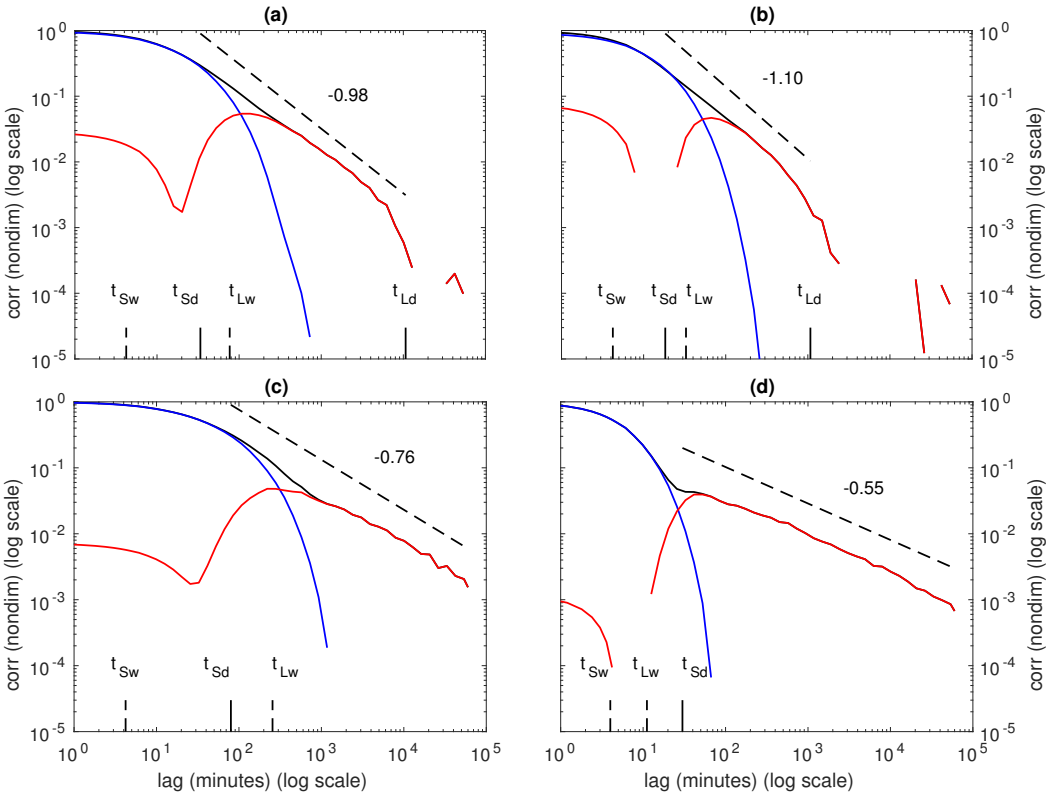

Figure 2. Parameter dependence of the two-state model autocorrelation. Solid black lines indicate full autocorrelations, blue lines indicate intra-event components, and red lines indicate intra-event components. Dashed (resp. solid) black ticks indicate the wet-spell (resp. dry-spell) duration cutoffs. Dashed black lines illustrate power laws with exponents estimated with a least squares linear regression on the full autocorrelation between t_{Sd} and t_{Ld}. Relative to (a), time scales are closer together in (b) and farther apart in (c). (d) shows a limiting case with short wet spell time scales. Exact model parameter settings for each figure are given in the SI.

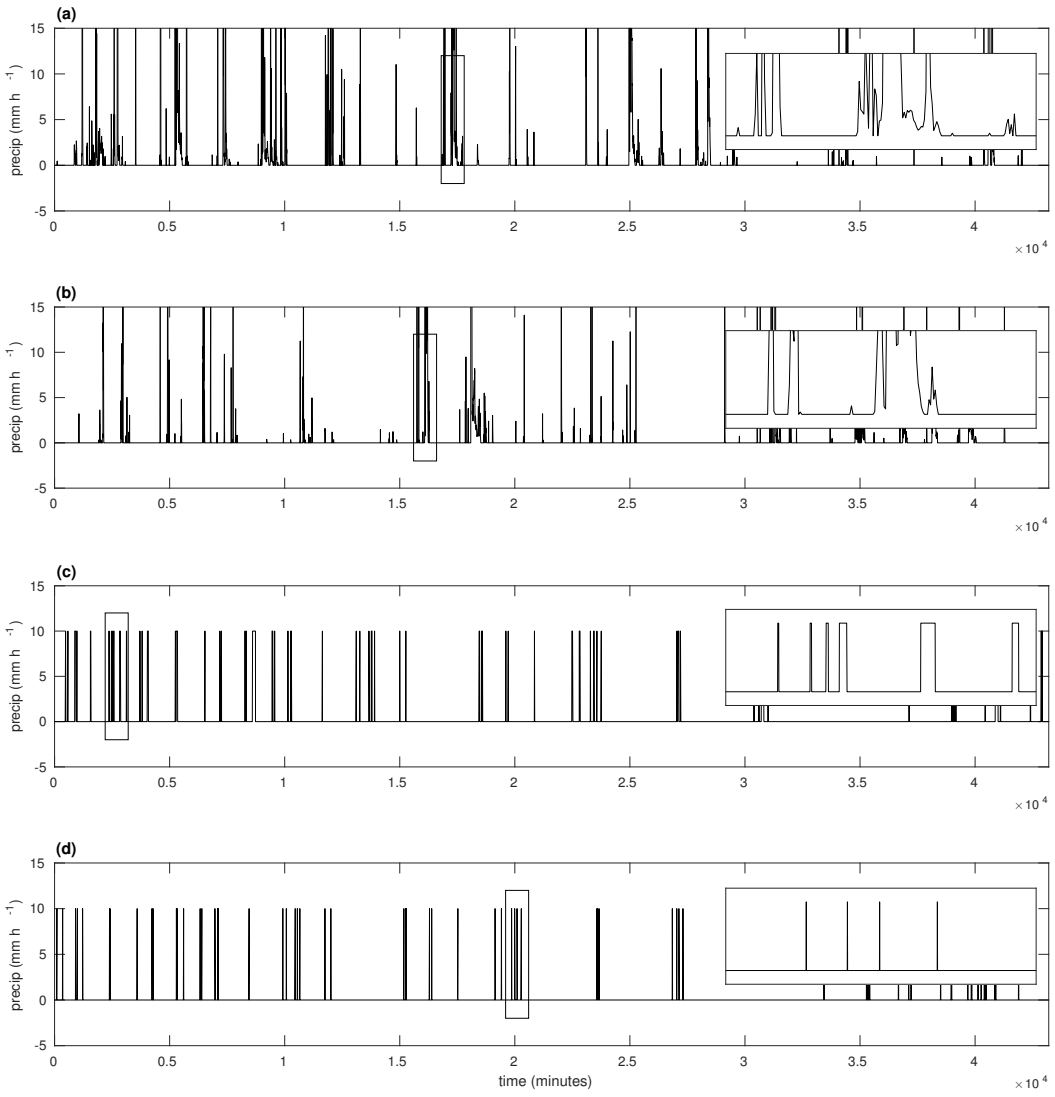

Figure 3. Sample precipitation time series from observations at (a) Manus Island and (b) Nauru Island, from (c) a run of the full SN14 model, and from (d) a run of the SN14 spike train idealization. In the spike train idealization, spikes are rescaled to 10 mm h^{-1}. Inset plots show details of the 1000 minute periods enclosed by boxes on the main plots. The full SN14 time series was produced using the same parameters as Figure 2a, and the simulation that produced the spike train time series used the same dry-spell parameters.