The Echo Size Distribution of Precipitating Shallow Cumuli

PANU TRIVEJ

Department of Atmospheric and Oceanic Sciences, University of California, Los Angeles, Los Angeles, California

BJORN STEVENS

*Max Plank Institute of Meteorology, Hamburg, Germany, and Department of Atmospheric
and Oceanic Sciences, University of California, Los Angeles, Los Angeles, California*

(Manuscript received 13 April 2009, in final form 2 September 2009)

ABSTRACT

S-band radar surveillance scans of precipitating shallow convection are analyzed. The scans are complied from 52 days of near-continuous measurements in the winter trades of the North Atlantic during the Rain in Cumulus over the Ocean (RICO) field campaign. After being analyzed and filtered to exclude spurious returns, the scans are segmented to identify contiguous returns, or echoes. The echo size and reflectivity statistics are then analyzed. A new normalization method is developed to account for biases associated with the nonuniformity in the native radar grid. The echo area distribution robustly exhibits power-law scaling up until sizes of about 10 km^2, with a scaling exponent of about -1.1. At larger sizes the scaling behavior breaks down and varies more markedly across subsamples of the data. Conditional sampling suggests that the scaling behavior of the larger echoes does, however, approach that of the smaller echoes as echo coverage increases, which supports the idea of a limiting distribution. Departures from this limiting distribution are argued to reflect finite size effects, modulated by the presence of a capping inversion whose height and strength varies across the samples.

1. Introduction

Relaxation events of systems driven by the slow input of energy can be described by the concept of self-organized criticality. These events, or bursts, are assumed to be scale free (i.e., not characterized by any favored scale). Examples include continental plates releasing their energy in an earthquake and piles of rice or sand in an avalanche. The scale-free character of such bursts results in a power-law relationship between the number density n of events at some size or intensity a and the size or intensity:

$$n(a) \propto a^{-B}. \tag{1}$$

The Gutenberg–Richter law, for instance, states that the probability of an earthquake follows Eq. (1) with a being the seismic moment, the measure of the released energy (Gutenberg and Richter 1944).

The scale-free character is expressed by the fact that the probability of an event of intensity a_1 versus an event of an intensity a_2 depends only on the ratio a_1/a_2, not on the intensities themselves. The physics of an object or event that belongs to an ensemble that behaves in a power-law manner may therefore be expected to hold true for another object or event of different size, as long as the two sizes are within the power-law regime. This can be useful to modeling, as it suggests that every object or event within the asymptotic regime may be described similarly.

Recently Peters and collaborators used this framework to inform their measurements of rain events (Peters et al. 2002; Peters and Christensen 2002, 2005; Peters and Neelin 2006). Radiative fluxes destabilize the atmospheric column, which drives a flux of latent and sensible energy into the atmosphere. The latent energy is stored in the atmosphere as water vapor up to a point or threshold, beyond which it has to relax to saturation, and in so doing the energy is released in the bursts of rain events. Using a vertically pointing Doppler radar Metek Micro Rain

Corresponding author address: Panu Trivej, Department of Atmospheric and Oceanic Sciences, University of California Los Angeles, 405 Hilgard Ave., Box 951565, Los Angeles, CA 90095–1565.
E-mail: p2trivej@atmos.ucla.edu

Radar Disdrometer (MRR-2), Peters et al. (2002) measured the size spectrum and fall velocities of hydrometeors from which they derive a rain rate. They define the event intensity as the released water column in millimeters—that is, the time integral of the rain rate over a sequence of successive nonzero rain rates. They found Eq. (1) capable of describing rain events over the Baltic coast town of Zingst with B equal to 1.36. Measurements in the Basilicata region (southern Italy) from February 1998 to February 2002 yield a similar result (Telesca et al. 2004). Similar conclusions can be drawn to some extent from Tropical Rainfall Measuring Mission (TRMM) observations over the Amazon (Petersen et al. 2002) and radar and sounding analysis over Kwajalein Atoll (Holder et al. 2008).

The rain events studied by Peters et al. (2002) and subsequent collaborators include deep convection and frontal systems. Our study focuses on rain from shallow cumuli. The contribution of shallow convection to the net rainfall is often overlooked. While it is true that, individually, shallow clouds produce relatively little rain, they are climatologically common. Hence, their collective effects can be substantial. Petty (1999) found that 20% to 40% of the surface nondrizzle precipitation over the ocean east of Australia was associated with warm cloud tops (infrared temperature above 273 K); Short and Nakamura (2000) estimated shallow convection is responsible for about 20% of the total rainfall over the tropical oceans.

Does the echo area distribution of shallow convection fit into the paradigm explored by Peters and collaborators? We explore the relationship between the number of events and the size of their instantaneous radar echoes. The radar echoes are measured by a scanning S-band radar. At these wavelengths larger reflectivities are proportional to the sixth moment of the droplet spectrum and hence are sensitive to the presence of large drops. Ideally echoes are equivalent to Peters et al.'s scaleless bursts of the stored energy. In particular, to the extent that echo area is a good proxy for rain rate (Nuijens et al. 2009; Doneaud et al. 1984), one expects Eq. (1) to be applicable to an echo field with a denoting area and n number density.

The advantage of a radar study is that in the S band large drops that contribute to rainfall are especially effective scatterers of radiation; echoes have a stronger tie to precipitation than clouds. Echoes underline many satellite retrievals of surface rain. However, the remote sensors making these measurements do not resolve the smallest echoes or those with insufficient reflectivities. The footprints of satellite imagery (pixel size of measurement) tend to be larger than the typical horizontal extent of a great fraction of shallow cumuli. TRMM, for

example, has a sensitivity of about 18 dBZ and a horizontal resolution 4.3 km (Short and Nakamura 2000) and may not be able to measure a great variety of smaller convection entities. CloudSat is much more sensitive, measuring echoes as weak as -30 dBZ, but its footprint of 2.3 km is still larger than many cumulus clouds (Haynes and Stephens 2007). Furthermore, CloudSat suffers contamination from surface returns in its near-surface pixels, which complicates interpretations of its measurement. Hence, in addition to providing physical insight, characterizations of the echo area distribution from ground-based sensors provide a basis for filling gaps in the global record.

In this study we are mostly concerned with establishing the extent to which general or regular features emerge in the echo statistics. By "regularity," we refer to the reoccurrence of recognizable patterns over a geographical area regardless of the changing environment over the area. In addition to satisfying our curiosity, the establishment of regularity may also provide observational targets for theory and constraints for models and simulations. Luo and Liu (2007) and Neggers et al. (2003) showed that the distributions of clouds produced by large-eddy simulation (LES) correspond well with satellite measurements, thus adding credence to the simulations. By extending our investigation to radar echoes, we also hope to provide a basis for constraining the representation of microphysical processes in LES, a useful step in an attempt to better understand age-old questions such as the production of rain by warm processes.

We address the question of the character and regularity of the echo area distribution by analyzing the radar echoes collected during the Rain in Cumulus over the Ocean (RICO) field campaign. Section 2 discusses the methodology. Section 3 presents the results of our analysis, emphasizing characteristic features of the size distribution. Section 4 endeavors to explain why we think large echoes evince different behavior than small echoes. Section 5 provides a discussion of our results along with a summary.

2. Methodology

The RICO field study took place during the months of November 2004–January 2005 in the vicinity of Antigua and Barbuda in the northeast trades of the western Atlantic. The project is more fully summarized in Rauber et al. (2007) and briefly recapitulated here. The trades blew from the east-northeast at about 10 m s^{-1}, average sea surface temperatures were near 27°C, and a weak trade wind inversion was often evident near 800 hPa. RICO operations were centered around the National

Center for Atmospheric Research (NCAR) SPolKa dual polarization, dual wavelength (S and K band) radar system on Barbuda. The routine scan sequence consisted of 180° plan position indicator (PPI) volume scans and full (360°) surveillance scans. In this paper we study the data based on the surveillance scans. The surveillance scan are conducted every 20 min, at a fixed elevation of 0.5°. The beamwidth of the radar is 0.9°.

Ten consecutive surveillance scans are lumped together to form independent samples. This is equal to roughly 200-min sampling interval. Although the autocorrelation time scale for averaged properties within each scan is somewhat longer than our sampling interval, this sampling interval is much larger than the expected lifetime of individual events. Out of fifty-six days from 1 December to 25 January, two periods are characterized by heavier rain associated with higher than usual echo fraction. This quantity is the ratio between total echo area and the domain area. On average the echo fraction throughout the campaign is 0.02, but on those mentioned days, it is as high as 0.1. Higher echo area is the result of deeper convection. Such periods are poorly sampled. This fact combined with our underlying interest in shallow convection motivates us to exclude these days (13 and 15 December and 9 and 10 January) from our subsequent analysis. (See Nuijens et al. 2009 for a further discussion of this point.)

a. Preliminary data analysis

The raw data are interpolated onto a regular polar grid with a diameter of 300 km discretized radially into 984 and azimuthally into 540 pixels. This polar grid is the native grid of the radar instrument. Although it presents some difficulties, which will be later discussed, we choose not to interpolate the data onto a Cartesian grid because doing so would only introduce methodological issues. We also note that the 0.5° elevation angle of the surveillance scans that we analyze means that radar beam probed deeper into the cloud layer at larger distances from the radar. Possible biases resulting from this property of the scan are not investigated here, in part because such effects are mitigated by the beamwidth being nearly twice the elevation angle. The interested reader is referred to Snodgrass et al. (2009) for a further discussion of these issues.

The azimuthal pixel size $d\theta$ is $2\pi/540$ radians, and the radial dr is 0.15 km. The data are filtered to exclude the returns from and over islands, backscattering from the surface, insignificant echoes, birds, and regions of predominantly Bragg scattering as described by Nuijens et al. (2009).

A 7-dBZ threshold is used to distinguish Bragg from Rayleigh scattering. There is nothing magical about 7 dBZ;

rather, it reflects a conservative attempt to diminish possible artifacts from Bragg scattering, which, in the S band, increasingly dominates the signal at smaller reflectivities. Because small echoes can potentially be dimmer than this threshold, marking anything below 7 dBZ as background robs us some of the less reflective signals. Using one day of data, we also explored the sensitivity of our results to the specific choice of threshold. Although the distributions changed slightly—for instance, there were fewer large echoes for higher thresholds (as one would expect)—the scaling behavior we explore was not sensitive to the choice of threshold.

Filtered and thresholded images are then segmented. The process is to distinguish between different connected pixels and label each as a unique entity. Two cloudy pixels belong to the same echo if they are immediate neighbors either vertically or horizontally, but not diagonally. This is called the four-connected method, detailed in appendix A. An ''echo'' hereafter is defined as a four-connected set of reflective pixels.

Our analysis focuses on two macrophysical parameters: echo area and reflectivity. Echo area is measured by simply adding up the size of every pixel; mathematically,

$$a = \sum_i a_i, \qquad (2)$$

with a being the echo area and a_i the pixel area, which varies according to

$$a_i = r_i\, dr\, d\theta, \qquad (3)$$

where r_i is the radial distance of the ith pixel from the radar. Reflectivity Z can be summed in a similar manner except that ζ_i or reflectivity in mm^6 m^{-3} is an addable quantity unlike Z_i, reflectivity in dBZ. Mathematically,

$$Z = 10\log\sum_i \zeta_i = 10\log\sum_i 10^{0.1Z_i}. \qquad (4)$$

Three other parameters that occasionally crop up are the echo fraction ϕ, the number of pixels, and the distance from the radar. Here ϕ is computed by the sum of the area of all echoes divided by $\pi \times (150^2 - 30^2)$ km^2 in an annular sample volume whose outer radius is the extent of the radar domain and whose inner radius is 30 km. The reason for excluding echoes within 30 km of the radar will be discussed in section 2b. The number is counted in a straightforward manner, and the distance is the radial position of the center of mass with the effect of pixels being of different sizes taken into account.

Out of 51 RICO days, about 800 000 echoes are evident on the surveillance scans. Each day contains roughly six to seven samples—series of surveillance scans are

regularly interrupted by PPI scans—so each sample hosts on average around 3000 signals. A size range is divided into 100 bins (or size categories), with a width of one bin being the multiplicity of the previous bin. The bins depend on the maximum and minimum size of interest. Echoes are sorted according to areas. We count and report a number of echoes within each bin and then divide the number by the bin width to get the density. Echoes' reflectivities are averaged within each bin.

The number densities and reflectivities are plotted versus area on a log–log scale. Information from a sample or even a single day is insufficient to provide reliable statistics. Data among all samples—312 in total—are thus combined. We assume the relationships among parameters are unchanged through the period. This is not necessarily true, and its failure may in fact be responsible for the anomalous scaling of large echoes discussed later in section 3b and appendix C. Sampling errors are estimated from the standard deviations, but our data are so numerous that the 98% confidence interval for individual data points tends to imply a level of uncertainty on the order of the size of the symbols used to plot the mean of the data, and hence is not shown.

b. Further filtering

The averaged echo reflectivities are shown as a function of echo area in Fig. 1.

Below 0.25 km^2 reflectivities increase with decreasing areas. This runs counter to our intuition that large echoes contain more liquid water content, develop a higher number of large drops, and are more reflective, as is evident for $a > 1 \text{ km}^2$. To better understand this apparently anomalous behavior, Fig. 2 shows the reflectivity distribution of echoes in two area categories, with the first, near 0.02 km^2, being in the region of anomalous reflectivities and the second, near 2 km^2, being in the range where reflectivities increase with echo area.

The vertical dashed lines denote the average reflectivity. The average intensity of the 0.02 km^2 echoes is higher than that of the 2 km^2 echoes. Both distributions are skewed toward more reflective echoes, but the tail is more prominent for the 0.02 km^2 echoes. This tail consists of small but very reflective echoes; some are even more reflective than echoes 100 times their area, which leads to the curious feature (Fig. 1) wherein the average reflectivity exhibits a local maximum at 10^{-2} km^2.

The reflectivity distribution of the smallest echoes is not well sampled because most echoes are dimmer than 7 dBZ. Even though this fact can be used to explain why the average reflectivity of the 0.02 km^2 echoes is higher
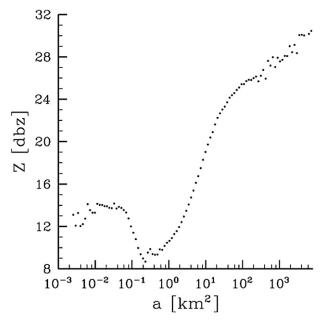

FIG. 1. Reflectivity–area relationship plotted on a lognormal scale. Data include echoes throughout the whole circular domain.

than the 2 km^2 echoes, it cannot explain why the breadth of the intensity distribution of small echoes is greater than that for large echoes (not shown), as a truncated distribution should have a smaller standard deviation. Because the intensity distribution of small echoes appears to be unduly influenced by our filtering procedure for Bragg echoes, they are excluded from this analysis.

Finally, we note that there appears to be an anomalous clustering of echoes near the radar as demonstrated by the radial distribution of echoes (Fig. 3).

However, in the absence of island influences the locations should be unbiased. That the region within 30 km evinces an unusually large number of returns suggests a possible near-field influence, either in the character of the measurement or in the natural environment. One possible cause of these returns is birds; Barbuda is a natural refuge for frigate birds, and returns from birds were often seen on the radar. While our initial filtering procedure attempted to remove such artifacts through the use of polarization or velocity anomalies (see Nuijens et al. 2009 for a full discussion), it may well be that such procedures were not entirely successful. Because we do not fully understand the anomalous returns, we simply consider only echoes located at ranges farther than 30 km from the radar. So doing yields a minimum echo area of about 0.1 km^2. The resulting data have the desirable property that the average echo intensity monotonically increases with echo area.
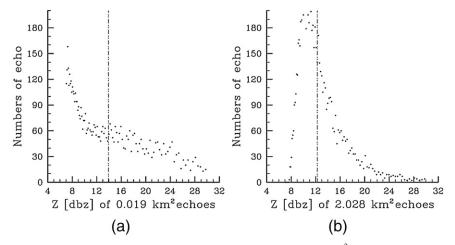
FIG. 2. Distribution of reflectivities for two areas (a) 0.02 and (b) 2 km^2. The vertical dashed line denotes the average reflectivities. Data include echoes throughout the whole circular domain.

c. The beam-filling effect

The procedure of recording and thresholding data on a polar grid introduces further biases that this section explores and for which we propose a remedy. The issue has to do with the fact that small echoes may only appear on a limited region or rings of the radar scan, hence biasing the statistics of these echoes. In this section we explore this issue and derive a normalization procedure that attempts to account for such biases in the measured data.

We use the phase "beam-filling effect" as a reference to two related effects regarding the resolution limit of the radar. One is the discreteness of the data field; pixilation changes the original echo areas. The smaller an echo is, the more pixilated it appears on a radar screen, and the beam-filling effect becomes increasingly important. Another is that the natural (polar) coordinate of the measurement results in pixels whose size depends on the distance from the origin. A small echo situated far away from the radar will be washed out in a pixel average. For example, a 7-dBZ echo smaller than 0.25 km^2 can never appear on the outermost annulus since the tiniest unit of measurement there is 0.25 km^2. Similarly a 7-dBZ 0.125 km^2 echo has only one-fourth of the whole domain as its visible field. By equating the area of a pixel at some distance r from the radar with the area of an echo, it follows that the maximum perceivable range r_p at which an echo can be measured is simply

$$r_p(a) = \frac{a}{d\theta\, dr}. \qquad (5)$$

An echo of a certain area, a, and the maximum perceivable range r_p can be single pixels whose size is equal to a, or double $a/2$ pixels, or triple $a/3$ pixels, and so on. Those a, $a/2$, $a/3$, ... pixels locate at the distance r_p, $r_p/2$, $r_p/3$, ... Small echoes will preferentially be confined to concentric rings whose outmost radius is r_p, as schematically illustrated in Fig. 4. At intermediate ranges echoes will be aliased onto different sizes.

Such effects are apparent in the data. Figure 5 shows the radial distribution of echoes belonging to three sizes: 0.13, 0.25, and 9.74 km^2. The 0.13 km^2 echoes have the outermost perceivable radius of 74 km, in accord with Eq. (5), the second perceivable annulus at 74/2 = 37 km and the third at 74/3 = 25, which is beyond the designated
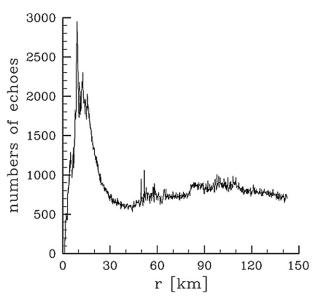

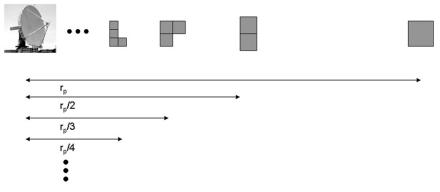FIG. 3. Radial distribution of signals.

FIG. 4. Schematic illustrating signals of the same size located at different radii and containing different numbers of pixels.

inner radius of the domain, 30 km. Thus, all the 0.13 km^2 echoes are clumped within either the narrow annular of 74 or 37 km as in Fig. 5a. The right annulus consists of single-pixel echoes and the left annulus of double-pixel echoes. The 0.25 km^2 echoes have the outermost perceivable radius of 147 km according to Eq. (5) (Fig. 5b). The next consecutive rings are at $147/2 = 74$ km, $147/3 = 49$ km, and $147/4 = 37$ km. The perceivable area is then the four concentric rings. Figure 5c shows the distribution of a size big enough for echoes to be perceived anywhere within the domain; the discreteness becomes irrelevant.

This tendency of the domain to bias the sampling of small echoes (by limiting the ranges at which they can be measured) depends on the minimum and maximum radius of the domain, which we denote by R_{min} (30 km) and R_{max} (150 km), respectively. The outermost perceivable radius might be bigger than 150 km, in which case the actual perceivable radius *within the domain* is r_p/k when $k = (r_p/R_{max})$. Similarly, the minimum annulus is r_p/j when $j = (r_p/R_{min})$. The perceivable area of an echo A_p is proportional to the sum of the perimeters of the rings:

$$A_p(a) \propto 2\pi \left[\frac{r_p(a)}{k} + \frac{r_p(a)}{k+1} + \frac{r_p(a)}{k+2} + \ldots + \frac{r_p(a)}{j} \right]. \quad (6)$$

If $n_*(a)$ is the total density of the measured echoes, the new density, $n(a)$, can be estimated through a normalization procedure that accounts for these biases, namely

$$n(a) = n_*(a) \frac{S_p(a)}{A_p(a)}. \quad (7)$$

Note that S_p is called an isolation factor, which accounts for the fact that even within the same perceivable area large entities are less abundant since they require a larger distance from the center of a neighboring echo to become distinguishable as individual entities. If a portion of one entity crosses over the other, they will be registered as one larger entity. For a sake of simplicity,
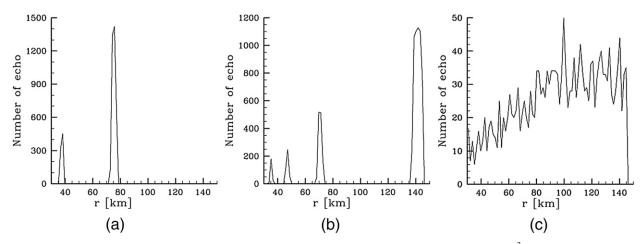

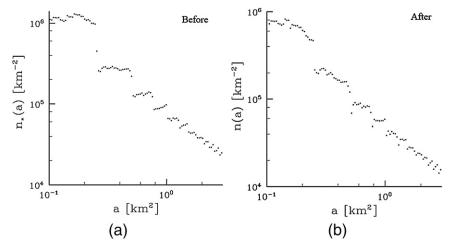FIG. 5. Radial distribution of echoes for three sizes: (a) 0.13, (b) 0.25, and (c) 9.74 km^2.

FIG. 6. The distribution (a) before the calibration, n_*, and (b) after Eq. (8) is applied, n.

we let $S_p \propto a$, but since r_p is proportional to a, Eq. (6) can be rewritten as

$$n(a) \propto \frac{n_*(a)}{\dfrac{1}{k} + \dfrac{1}{k+1} + \dfrac{1}{k+2} + \cdots + \dfrac{1}{j}}. \qquad (8)$$

Equation (8) can be thought of as the result of a deconvolution procedure designed to remove sampling artifacts that can be estimated for an ideal radar. It results in an estimate of the number density for every echo area so that our subsequent analysis is based on n rather than the raw densities n_*.

Figure 6 compares the distribution before and after (8) is applied. The discrete behavior of n_* in a small area in Fig. 6a is tempered by this calibration and turns into a more continuous distribution of n in Fig. 6b. That said, some artifacts remain, perhaps related to the fact that our radar is not ideal, and the signals that we sample are not fixed-reflectivity targets, which was assumed in the derivation of Eq. (8).

3. The echo area distribution and the scale break

The average echo area distribution for the entirety of our dataset is plotted on a log–log scale and presented in Fig. 7. Two regimes are evident. First is the tendency of echoes smaller than ~10 km^2 to follow a power-law scaling with an exponent $B \sim 1.06$. This is evident by the correspondence between the data and the dashed line. Arguably one finds a deviation from this scaling for the very smallest echoes, but we believe this deviation reflects the uncorrected sampling biases discussed at the end of section 2.

The average density of echoes larger than 10 km^2 shows a sharper falloff with area. Relative to the small

echoes, large echoes are disproportionately rare. The size demarcating the change in the scaling of small (<10 km^2) versus large (>10 km^2) echoes is called the scale break. In what follows we explore more closely the nature of the break and the nature of the distribution of large echoes.

a. The scale break

To more clearly identify the scale break and the nature of the scaling above this break, we introduce the compensated number density. Compensated densities are the differences between the actual data and the predictions, which in this case are the linear least squares
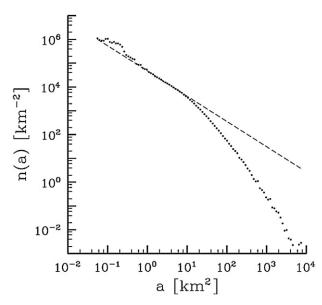

FIG. 7. Area distribution according to RICO's SPolK data. The dots are the average number densities among 312 samples. The dashed line is the least squares fitting from 1 to 10 km^2.
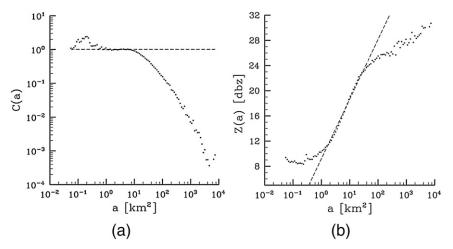
FIG. 8. Log–log plot between areas vs (a) compensated densities and (b) average reflectivities. The dashed line in (a) is the assumed power law. The dashed line in (b) is the least squares linear fit.

fit of the area distribution from 1 to 10 km^2 projected beyond 10 km^2. Mathematically the compensated density is $c(a)$, in which $n(a) = c(a) \times ka^{-B}$, where k, the proportional constant, and B are determined by a least squares fit in log space between n and a for $1 < a < 10$ km^2. Plotting our data in this manner more clearly illustrates the disparity between the black dots and the dashed line of Fig. 7.

If $n(a)$ is already a power law of the form a^{-B}, then $c(a)$ is unity and the plot will appear horizontal. This is the case as shown in Fig. 8a, the plot between sizes and compensated number density. The slope of the distribution (the dashed line in Fig. 7) is leveled down into a horizontal line in Fig. 8a. From 1 to 10 km^2, the densities align closely to the horizontal dashed line. This reinforces our idea that the power law is a good approximation for echoes smaller than 10 km^2.

The scale-free characteristic breaks down where the area distribution starts to deviate from the original power law: according to the figure, around 6 to 12 km^2.

This scale break does not appear to be an artifact caused by the finite domain size. Our numerical simulations (not shown) demonstrate that artifacts emerging from the representation of a power-law distribution on a domain of our size are entirely negligible at the size near the scale break. This is not surprising given the three orders of magnitude difference between the size of the domain, $\sim 10^4$ km^2, and the size at which the break appears, ~ 10 km^2.

Further evidence in support of the break is presented in Fig. 8b. The average echo intensity scales with the area of the echo. Bigger echoes are generally brighter, although the 7-dBZ threshold causes the intensity of the smallest echoes to asymptote to a value near this bound.

The dashed line is a linear fit. It is only to guide the eyes; we make no claim about the power-law relationship between areas and reflectivities. The reason for drawing the line is to emphasize the departure of large echoes from the general trend of small ones. The scale break is at the location of this departure, which is around 20 km^2, suggesting that the break demarcates not only a change in the behavior of the echo area distribution but also the intensity distribution. This break in the reflectivity plot is about twice as big as the break in the compensated density plot. It is possible that at the very threshold, around 10 to 20 km^2, reflectivities are less sensitive to the changing regime than number densities. Given that the break is both well resolved by the data and emerges in different fields, we believe it is justified to identify it as a real feature in the data rather than an artifact of our analysis.

b. The area distribution of large echoes

For echoes larger than the scale break, the area distribution lacks the robustness of the distribution apparent among echoes smaller than the break. The density–area relationship for the large echoes changes from one sample to the next, and a single function with fixed parameters is not suitable to describe distributions across samples. One example of such variations is in the relationship between the maximum echo area a_m within a sample and the echo fraction ϕ, which is the total area of echoes divided by the domain area. It is mathematically defined as $\pi \times (150^2 - 30^2)$ km^2. An environment with high echo fraction is prone to produce a big a_m, as evinced in Fig. 9.

The scatterplot compares the maximum area versus the echo fractions; ϕ was shown by Nuijens et al. (2009)
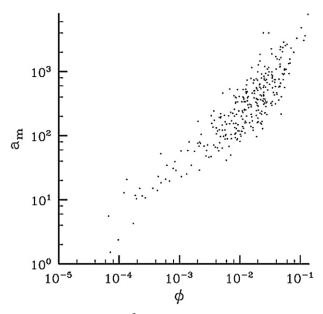
FIG. 9. Maximum area (km^2) vs the echo fraction ϕ for each sample on a log–log scale.

to scale with the environmental humidity. Moist environments, which are well known to be more conducive to convection, are also more conducive to producing large echoes. This motivates us to look at the distribution conditioned on the echo fraction as a proxy for the meteorological control on convection. Note that a_m is not the only parameter that changes as we look from sample to sample; even the distribution itself assumes different shapes.

Figure 10 plots conditional compensated densities of the large echoes; Fig. 10a shows samples with $0.001 < \phi < 0.003$ and Fig. 10b samples with $0.06 < \phi < 0.1$. Since $n(a) \propto c(a)a^{-B}$, $c(a)$ in Fig. 10 encapsulates the behavior of $n(a)$. Not only do the distributions have different upper bounds, but the shapes themselves are also dissimilar as the number of echoes falls off more rapidly with size for samples with smaller echo fractions (i.e., samples that are statistically less humid). While 100 km^2 echoes are rare among samples with relatively low echo fractions (Fig. 10a), the number of 100 km^2 echoes within samples with larger echo fractions is almost of the same order of magnitude as for the 10 km^2 echoes (Fig. 10b). These results lead us to hypothesize that the behavior of the large echoes is conditioned on some environmental (hidden) factor and thus is influenced by variations in this factor across our samples.

The actual shape of the distribution is a subject of much debate. The c in Fig. 8a, and by implication n, seems to relate to a in an exponential manner. Equally possible is the existence of yet another scale break that bookends another power-law regime. The shape that

best describes the distribution will be the result of the composition among data with different distributions of their own; thus, finding the best fit does not shed any light on the physical phenomenon. A more practical issue is how large echoes behave in the limiting cases, such as when ϕ approaches zero or one, and whether that can teach us anything about the nature of the break.

4. Factors controlling the different statistics of small and large echoes

We define the scale break as the point or size demarcating two regimes in the echo area distribution: one, corresponding to small echoes, with relatively uniform scaling and the other, corresponding to large echoes, with scaling that varies as a function of some unknown parameter. This section looks more deeply into how the echo distributions differ across this break and as a function of the echo fraction, ϕ, of a sample.

The differing statistics of large echoes may be the result of some intrinsic properties of convection, either microphysical or dynamical, that change once echoes exceed a certain size. For instance, microphysically one might argue that the echo area distribution is somehow sensitive to the development of an equilibrium drop size distribution associated with the development of a mature rain shaft, and that these large echoes then influence the overall echo fraction. Dynamically, perhaps the development of cold pools with deep convection qualitatively changes the underlying convective dynamics. On the other hand, given that our methodology already answered the shortcomings in sampling and observational strategy as best as possible, if the differences between the large and small echoes are not robust, perhaps this simply tells us that the differing statistics of the large echoes are a byproduct of the meteorological environment (e.g., the signature of the trade inversion suppressing large echoes).

To address this question, we formulated the null hypothesis such that the area distribution beyond the break forms a second power law, but one with a cutoff at some maximum size. This hypothesis is principally one of convenience, as it provides a basis for gaining insight into some of the questions we have posed for ourselves.

Contrary to the evidence presented in Figs. 7 and 8a, which argues that large echoes may be distributed exponentially, it can also be argued that the apparent exponential falloff in the power-law distribution for large echoes is due to variations in a_m across the samples, as illustrated with the help of Fig. 11.

Gray slant lines are the assumed power-law distributions of individual samples, and the dashed solid line is $n(a)$, the average distribution. (The actual height of the
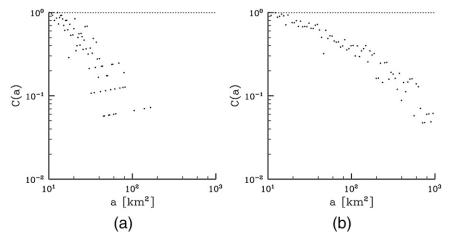
FIG. 10. Compensated density of large echoes averaged among samples with (a) $0.001 < \phi < 0.003$ and (b) $0.06 < \phi < 0.1$.

slant gray lines is irrelevant in this schematic diagram.) Only within regime A is the average consistent with the underlying distribution of the samples, as each sample spans the range of areas associated with this regime. Inside regime B, because of fewer samples, the average increasingly curves down as shown by the dashed line. At the extreme, in regime C, even though individual lines are straight there are so few samples that the average distribution falls well below the underlying distribution of the samples. While regime A and B are distinguishable from one another at the point where the average distribution starts to fall off from the power law, the difference between regime B and C is merely a matter of degree. The shape of the distribution inside regime B and C is further discussed in appendix C.

Thus, we argue that the underlying average echo area distribution, $n(a)$, is

$$n(a) \propto \begin{cases} a^{-B_1}, & \text{for} \quad a_0 < a < a_b, \\ a^{-B_2}, & \text{for} \quad a_b < a < a_u \end{cases}, \qquad (9)$$

where a_0 is the smallest echo size, which after the filtering in section 2 is 0.1 km^2; a_b is the location of the scale break; and a_u is the smallest size among many a_m of the averaged samples. Only under this size, a_u, is the distribution not biased by the sample-to-sample variability in a_m (regime A in Fig. 11). Appendix C investigates the distribution beyond a_u. The quantities B_1 and B_2 are the exponents of the average distribution as measured by the slope of the least squares fit between n and a on the log–log scale.

From Fig. 10, the compensated density plot, and our previous discussion, we learn that large echoes are rarer in proportion among samples with small echo fraction, ϕ, as compared to those with large ϕ. In other words, B_2

is a function of ϕ; $B_2 = B_2(\phi)$. In what follows we look more quantitatively at how B_1 and B_2 vary with ϕ and argue that B_1 is the limiting slope irrespective of echo size. The suggestion is that small echoes have reached their limiting distribution, and while large echoes do not behave intrinsically differently, they are limited or modulated by external factors, leading to apparently different statistics.

To evaluate how B_1 and B_2 depend on ϕ, we estimate these exponents by finding the power law that minimizes the square error over finite area ranges. The least squares fit is an intuitive (perhaps naive) way to parameterize the distribution and is explored first. In that scenario, B_1 is computed by fitting a line to the data between 1 and 8 km^2, and B_2 is obtained by fitting a line to the data between 20 and 100 km^2. The lower bound of 1 km^2 is chosen so that the distribution is unaffected by possible beam-filling effects. The choice of 100 km^2 corresponds
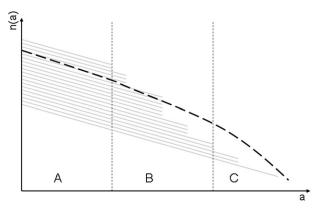

FIG. 11. Schematic diagram illustrating how data composition affects the average of area distributions. The dashed line is the average.

to the fact that we only include samples whose largest echoes a_m are greater than 100 km^2. In other words, 100 km^2 is a_u in Eq. (9). This ensures the distribution remains a power law without the departures due to compositing data whose maximum areas vary within the range over which we construct the power-law fit. Since the location of the scale break has not been determined except that it is in the vicinity of 10 to 20 km^2, the intermediate areas of 8 and 20 km^2 are conservative choices to avoid the influence of large echoes on the scaling of small ones, and vice versa.

The result of such a procedure is shown in Fig. 12. Crosses and squares trace B_1 and B_2, respectively.

In this case B_1 remains relatively constant as B_2 decreases and approaches B_1. The uncertainty in the estimate of B_1 is smaller than the size of the plotting symbol and thus not shown. Estimates of B_2 have substantial uncertainty except at echo fraction 0.01 to 0.05, where the samples are most numerous. There appears to be a tendency for B_2 to approach B_1 as ϕ increases, although we note that as $\phi \to 1$ the concept of an echo area distribution becomes meaningless.

The apparent convergence of the exponent of the distribution for large echoes B_2 to the exponent of the distribution for small echoes B_1 can be further illustrated by comparing $n(a)$ conditioned on ϕ for two regimes, one with a small echo fraction ($0.001 < \phi < 0.003$) and the other with a big echo fraction ($0.06 < \phi < 0.09$). Figure 13 compares these conditional distributions. The solid and dashed lines are the least squares fitting of the two power laws. The fits for small echoes (solid lines) are similar for both conditional distributions, in contrast to the fit for large echoes (dashed lines). As the echo fraction increases, the exponent becomes less negative and the break is less emphasized.
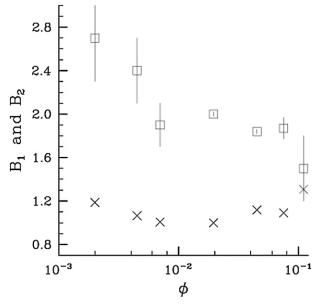
Similar behavior appears in reflectivity versus area (Fig. 14). The intensity of large echoes in the composition over samples with large ϕ departs less from the increasing tendency for small echoes than it does for the composition over samples with small ϕ. Yet the sampling error in the latter case is large given the paucity of large echoes at weak echo fraction. That said, the general trend is clear.

The main advantage of the least squares fit is its intuitiveness. However, the method is sensitive to a number of issues: a binning strategy, the choice of bin width, the locations of the first and last bins, and the bin centers. These systematic errors can be equal to or more than statistical ones (the error bars in Fig. 12). To circumvent these issues, we employ an alternative method to explore this question of the limiting distribution. This is called the direct power-law method, and it yields a similar conclusion to the least squares fit. A discussion of the direct power-law method is included in appendix B.

FIG. 12. Number of data and the exponents as measured by the least squares fit plotted against the echo fraction. Crosses represent B_1 and squares B_2.

Based on this we hypothesize a limiting echo area distribution of the form $n(a) \propto a^{-B_1}$ of which the range of applicability depends on the meteorological environment. This result is not trivial as there is a long list of reasons, some of which we have already discussed, as to why one might expect large echoes to evince different statistics than small echoes.

5. Summary and discussion

The average area distribution of radar echoes smaller than 10 km^2 is well described by a power law of the form $n(a) \propto a^{-B_1}$, with a denoting area and B_1 slightly greater than one; B_1 remains relatively invariant across subsamples conditioned on echo fraction, which earlier work has shown to correlate well with meteorological variability. The regularity evident in the distribution of small echoes should provide a basis for constraining models and simulations. It can further supplement measurements by devices incapable of resolving small echoes, such as are characteristic of modern satellite remote sensing. Similar exponents are apparent in analyses of cloud size distributions taken from visible imagery (Lovejoy and Mandelbrot 1985; Cahalan and Joseph 1989; Sengupta et al. 1990; Benner and Curry 1998; Zhao and Di Girolamo 2007). In their analysis of echo area distribution, Peters et al. (2008) found a much sharper falloff, with $B \simeq 2$. Given the differences in the analysis (planar cuts in our case, projected areas in theirs) and the differences in regimes (shallow versus deep convection), it remains unclear how substantive these differences are.
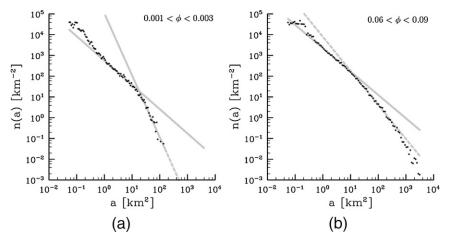
FIG. 13. Area distribution for (a) a very dry regime when echo fractions are between 0.001 and 0.003 and (b) a very wet one when echo fractions are between 0.06 and 0.09. Unlike in Fig. 7, here we assume the double power law. The solid and dashed gray lines are linear fitting for small and big echoes, respectively. The exponents for the small and big echoes are respectively (a) −1.18 and −3.04 and (b) −1.12 and −1.64.

On scales larger than 10 km^2 we find evidence of a pronounced departure from the scaling at smaller areas. The extent of the deviation and its character depend, however, on the meteorological environment. In convectively favorable situations, which we associate with periods of high echo fraction, departures are less pronounced. There is even some evidence that the scaling of large echoes approaches that of small echoes in the limit of high echo fractions; the two power laws merge, and the break disappears altogether. In their analysis of the rank of clusters as a function of size for various values of the water vapor, Peters et al. (2008) report similar behavior (albeit with a different power law) where for more humid conditions the distribution approaches a limiting power law. The similarity of these findings is compounded by the observed correlation between humidity and echo fraction in the RICO data (Nuijens et al. 2009).

The break thus appears to reflect meteorological controls on convection. An obvious cause of the break and the one we propose here is the separation between the moist, cool air near the surface and the dry, warm air in the free troposphere as manifested across the trade inversion. Around Barbuda at the time of the RICO campaign, an inversion, or hydrolapse, was evident between 2 and 3 km. The lifting condensation level or the
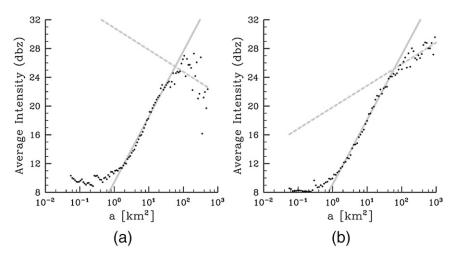

FIG. 14. Reflectivity–area relationship for (a) a very dry regime when echo fractions are between 0.001 and 0.009—note that this is wider than what we used for Fig. 13 since we need more data—and (b) a very wet regime when echo fractions are between 0.06 and 0.09. Also shown are the two least squares fits.
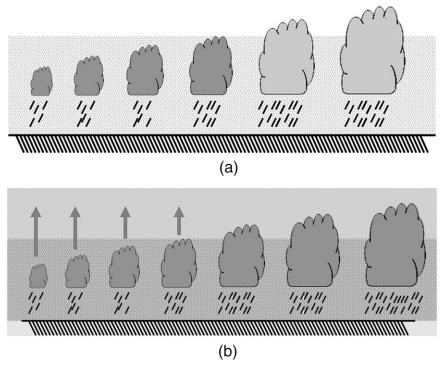
FIG. 15. Schematic showing the behavior of echoes in (a) a dry case and (b) a moist case. In (a), the shaded environment denotes moist air under the inversion; clear air above has no shade. The difference in cloud shading denotes different behaviors and area distributions. Darker shaded clouds are more abundant than lighter shaded ones. While darker shaded clouds are small and can fit under the inversion, lighter shaded ones are big and can penetrate the inversion. In (b), arrows indicate how moisture is deposited in the upper layer. The upper environment is now lightly shaded to indicate some amount of moisture. Big clouds become darker since they now are more abundant and similar to small clouds as the result of deposited moisture in the upper environment.

base of cloud layer is around 700 m (Rauber et al. 2007). The cloud layer height, h, is thus 1–2 km. Kuo et al. (1993) showed that the aspect ratio of cloud height and radius are about unity for almost all area ranges. (Although we have yet to explore if this finding also holds for the RICO data, we suspect it does.) If their result holds true for any convective entities, this layer height produces echoes whose area, $\pi h^2 \sim 10$ km^2, is near the size where we observe a break in the scaling. Hence, on days with weak convection the presence of a trade inversion effectively limits or inhibits the growth of clouds beyond a certain size.

Figure 15a helps illustrate our proposal. Small echoes contained within the cloud layer are scale free, following the self-organized criticality framework developed by Peters and colleagues. Unusually strong echoes that penetrate this layer can be expected to diminish rapidly in the presence of more stable and drier air. The physics that govern big and small echoes are not the same, nor are their area distributions. This is illustrated in the schematic diagram as big echoes having a lighter shade than small ones. The consistent physics of small echoes led to their distribution being approximately power law. The same might be true for large echoes, but that conclusion is less convincing.

As convection carries moisture away from the surface and evaporates it to the upper level, the environment moistens and the trade-wind layer deepens as illustrated by arrows in Fig. 15b (Stevens 2007). The behavior of echoes that penetrate the inversion start to depart less from small echoes as illustrated by their having the same shade. This is because the weaker inversion is less relevant or, equivalently, the properties of the two air masses are more similar; B_2 takes on values approaching B_1 in relatively convective environments (i.e., during periods when echo fractions are large).

Another way to consider this is that smaller B_2 implies the environment producing more large echoes. In suppressed conditions, a shallow, moist layer and a pronounced inversion impede the vertical development of convection; these echoes are mostly small and shallow (Fig. 15a). The moist cloud layer of Fig. 15b, on the other
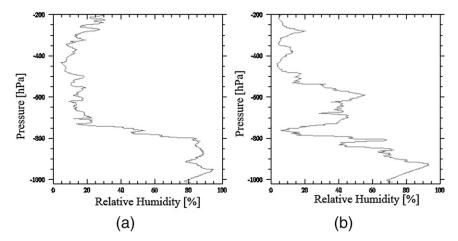
FIG. 16. Vertical structure of relative humidity as taken from the soundings on (a) 21 and (b) 23 December.

hand, is conducive to deeper convection and bigger echoes.

This conceptual model is supported by the vertical structure of relative humidity as taken from the sounding on 21 and 23 December. On 21 December (Fig. 16a), the environment above 700 hPa is drier and the maximum echo on that day is about 295 km^2. As more and more moisture is deposited into the upper layer, eventually the sounding evolves to the state shown in Fig. 16b where the humid layer around 700 hPa may support deeper convection, irrespective of the thin dry layer near 750 hPa. The maximum echo on 23 December reached 1104 km^2. Hence the humid environment is associated with larger echoes.

APPENDIX A

The Four-Connected Segmentation Method

Echoes can be segmented according to a four-connected or an eight-connected method. Two cloudy pixels belong to the same echo in the four-connected method if they are immediate neighbors either vertically or horizontally, but not diagonally. The eight-connected method allows diagonal neighbors to be of the same echo. Kuo et al. (1993) showed that the two methods do not yield

substantially different area distributions. This study engages the four-connected method.

The method is exemplified in Fig. A1. An initial cloudy pixel marked "0" is selected. Its four neighbors are examined in the clockwise direction from above, right, below, and left. The first that has not been visited is the next stop. The only major difference between our algorithm and that of Kuo et al. (1993) is that theirs employed a first-in-first-out waiting queue while ours uses a last-in-first-out. The pixels are examined and visited until the code finds a pixel with no unvisited neighboring pixel. We trace back our route to the first echo with an unvisited neighboring pixel. The process continues over again until no pixel has an unvisited neighboring pixel. That is when the entire echo has all been visited. The numbers in Fig. A1, from 1 to 5, mark five sequential stop spots where we retrace our steps. The end point, 6, is the same as the beginning. Since we end at the beginning, this process should be self-checked, and the last-in-first-out algorithm should not
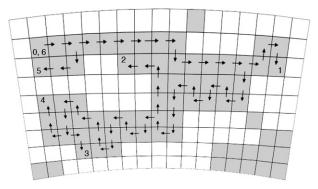

FIG. A1. Example of how segmentation is done. Gray pixels are cloudy; white is the background.

yield different result from the first-in-first-out. Kuo et al. (1993) and Wielicki and Welch (1986) discuss the segmentation algorithm in more detail.

APPENDIX B

The Direct Power-Law Method

Following Zhao and Di Girolamo (2007), we explore the behavior of scaling exponents as estimated by what they call the direct power-law method. It is motivated by the question "What can we learn from a single sample?" The name refers to how this method already assumes a power law, and the exponent is directly extracted without actually fitting or drawing any line.

The number density of echo within a single sample is assumed to be the double power law:

$$m(a) \propto \begin{cases} a^{-b_1}, & \text{for} \quad a_0 < a < a_b \\ a^{-b_2}, & \text{for} \quad a_b < a < a_m \end{cases}, \quad \text{(B1)}$$

where b_1 and b_2 are the hidden exponents, hidden because the statistic of one sample is too poor. The distribution of a sample, $m(a)$, is different from the average distribution, $n(a)$. Specifically, n is the average of m over 312 samples. We want to stress that Eq. (B1) is only an assumption. If we plot the distribution of a single sample, it will hardly look like any power law because of the poor statistics and not having enough data in one sample. It is not necessary for b to be the same as B. Only when the average does not suffer from the effect of data composition do the two confirm each another.

For a distribution following a power law, the average area, $\overline{A}$, and the number of echoes, N, between two limits, a_1 and a_2, can be written as

$$\overline{A} = \frac{1}{N}\frac{k}{2-b}(a_2^{2-b} - a_1^{2-b}), \quad N = \frac{k}{1-b}(a_2^{1-b} - a_1^{1-b}),$$
$$\text{(B2)}$$

so long as a_1 and a_2 lie within the same power-law regime. Unlike the least squares fit method, the direct power law is based on the distribution of individual sample, $m(a)$ not $n(a)$. Substituting N into $\overline{A}$ yields an implicit equation for b,

$$\overline{A} - \frac{1-b}{2-b}\frac{a_2^{2-b} - a_1^{2-b}}{a_2^{1-b} - a_1^{1-b}} = 0. \quad \text{(B3)}$$

Choosing appropriate $a_{1,2}$ and solving for b_1 (crosses) and b_2 (squares) numerically for every sample yields the results shown in Fig. B1 along with the regression lines representing the least squares estimation of these two exponents.
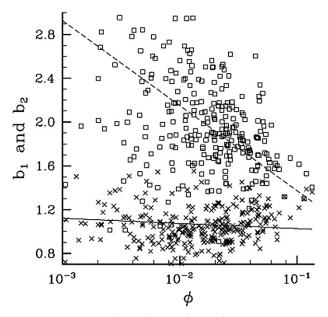

FIG. B1. Values of b_1 (crosses) and b_2 (squares) plotted vs the echo fraction. Regression lines fitted to the data are shown.

The small-scale exponent, b_1, takes on values between 0.8 and 1.6 and exhibits no clear dependence on the echo fraction. The large-scale exponent, b_2, however, appears to decrease with ϕ. Moreover, the regression line of b_1 is effectively constant and similar to the tendency line of B_1 in Fig. 12, while b_2 and B_2 appear to approach b_1 and B_1 respectively as ϕ increases. Hence, the two methods, the least squares fit and the direct power law, yield results that can be consistently interpreted.

APPENDIX C

The Departure from the Power Law

One possible cause for the exponential area distribution in Fig. 8a is our averaging across samples with different maximum areas, as demonstrated by Fig. 11. To the extent that our idea is correct, we should be able to mitigate the effects of meteorological variability by picking only samples whose maximum echo areas, a_m, are greater than some threshold.

Figure C1a shows $d(a)$, the deviations of $n(a)$ from the assumed second power law. Here $n(a) \propto d(a)a^{-B_2}$, in which B_2 is determined by fitting the line from 20 to 100 km^2. The deviated or the compensated spectra of the entire samples already begin to fall off from the power-law approximation at an area as small as 100 km^2. This suggests that the majority of samples have a maximum area around 100 km^2. On the other hand, Fig. C1b shows the compensated spectra averaged among samples with a_m greater than 10^3 km^2. These compensated spectra
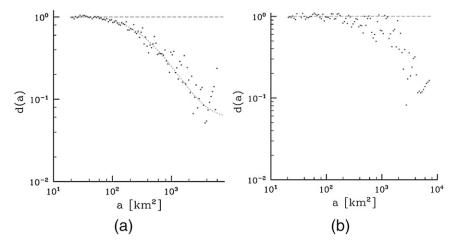
FIG. C1. Compensated densities of large echoes (a) for all samples and (b) on a condition that the maximum areas are only above 1000 km^2. The dotted line is the prediction based on Eq. (C5).

do not significantly deviate from the power law until after 1000 km^2, the area at which we condition the samples.

If one assumes the distribution of all individual samples, $m(a)$, to be a power law of the same exponent, b_2, and the distribution of the maximum areas, a_m, to be lognormal, then

$$H(a_m) = \frac{1}{\mathcal{X}\sqrt{2\pi}} \exp\left[-\frac{\log^2(a_m/\overline{a}_m)}{2\mathcal{X}^2}\right], \quad (C1)$$

with $\overline{a}_m$ being the average maximum size and $\mathcal{X}$ the standard deviation. Both $\overline{a}_m$ and $\mathcal{X}$ can be directly measured from all 312 samples; their values are 316 km^2 and 0.51, respectively. For any size a, the number of samples whose $a_m > a$ or samples that contain an echo of size a is

$$l(a) = L \times \text{erfc}\left[\frac{\log(a/\overline{a}_m)}{\mathcal{X}\sqrt{2}}\right], \quad (C2)$$

where L is 312, the total number of samples, and erfc is the complementary error function. Within region A of Fig. 11, $a \ll \overline{a}_m$, so erfc = 1 and $l(a) = L$; that is, every sample contains the echoes of size a. As we move closer to $\overline{a}_m$, we enter region B. There are fewer and fewer days that contribute to echoes of that size, and the value of erfc decreases. When $a = \overline{a}_m$, only half the samples are included. In region C, as a becomes larger, erfc approaches zero and eventually there is no echo—not in any sample—big enough to fit into that size.

The area distribution of echoes bigger than the scale break from a single sample is

$$m(a) \propto \begin{cases} a^{-b_2} & a_b < a < a_m, \\ 0 & a > a_m \end{cases} \quad (C3)$$

with the proportional constant varying from sample to sample. Note that this equation is the second half of Eq. (B1) in appendix B with an upper constraint. The average size density of the composited data is therefore

$$n(a) = \frac{l(a)m(a)}{L} \propto a^{-b_2} \text{erfc}\left[\frac{\log(a/\overline{a}_m)}{\mathcal{X}\sqrt{2}}\right]. \quad (C4)$$

We are interested in the compensated density, $d(a)$, according to $n(a) = d(a)a^{-B_2}$. Note that B_2 is not the same as b_2 but depends on the range over which we fit the distribution because the actual number densities curve down. This is as illustrated in Fig. C2.

The gray, solid, slanted line is the hidden size distribution of each sample. Because a_m varies according to Eq. (C1),
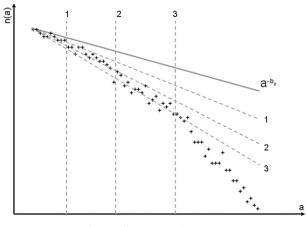

FIG. C2. Schematic showing how B_2 depends on the range over which we perform the least squares fit.

the composited distribution appears curved down from the actual distribution as shown by the plus signs. If we perform a least squares fit up to the vertical marking line 1, the fit will appear as the slant line 1. If the fit is performed up to 2 and 3, the result will be slant lines 2 and 3. With these fitting exponents, $-B_2$ is always more negative than the actual exponent of the distribution, $-b_2$, and the further we fit the distribution, the smaller is the exponent.

By fitting the distribution up to 100 km^2, we unintentionally assume there is no missing sample under 100 km^2 so the compensated density is the unity; the distribution under 100 km^2 is of the form $n(a) \propto a^{-B_2}$, and $d(a)$ is one. The distribution above 100 km^2 depends on how many data are missing from those that contain the echoes of the size, 100 km^2, which is $m(a) - m(100)$. Thus, the compensated density is

$$d(a; \mathcal{X}, \overline{a}_m) = \begin{cases} 1, & \text{for} \quad a < 100\,\text{km}^2 \\ \text{erfc}\left[\dfrac{\log(a/\overline{a}_m)}{\mathcal{X}\sqrt{2}}\right] - \text{erfc}\left[\dfrac{\log(100/\overline{a}_m)}{\mathcal{X}\sqrt{2}}\right], & \text{for} \quad a > 100\,\text{km}^2 \end{cases} \tag{C5}$$

where $d(a)$ is plotted on a log–log scale. The dotted line in Fig. C1a is in accord with Eq. (C5), while the dots are the true compensated densities. The line is supposed to be a theoretical counterpart of the dots, and their shapes are in rough agreement. This agreement supports our analysis here of the departure from the power law of big echoes.

REFERENCES

Benner, T. C., and J. A. Curry, 1998: Characteristics of small tropical cumulus clouds and their impact on the environment. *J. Geophys. Res.,* **103** (D22), 28 753–28 767.

Cahalan, R. F., and H. J. Joseph, 1989: Fractal statistics of cloud fields. *Mon. Wea. Rev.,* **117,** 261–272.

Doneaud, A. A., S. Ionescu-Niscu, D. L. Priegnitz, and P. L. Smith, 1984: The area-time integral as an indicator for convective rain volumes. *J. Climate Appl. Meteor.,* **23,** 555–561.

Gutenberg, B., and C. F. Richter, 1944: Frequency of earthquakes in California. *Bull. Seismol. Soc. Amer.,* **34,** 185–188.

Haynes, J. M., and G. L. Stephens, 2007: Tropical oceanic cloudiness and the incidence of precipitation: Early results from CloudSat. *Geophys. Res. Lett.,* **34,** L09811, doi:10.1029/2007GL029335.

Holder, C. T., S. E. Yuter, A. H. Sobel, and A. R. Aiyyer, 2008: The mesoscale characteristics of tropical oceanic precipitation during Kelvin and mixed Rossby–gravity wave events. *Mon. Wea. Rev.,* **136,** 3446–3464.

Kuo, K., R. M. Welch, R. C. Weger, M. A. Engelstad, and S. K. Sengupta, 1993: The three-dimensional structure of cumulus clouds over the ocean. 1. Structural analysis. *J. Geophys. Res.,* **98** (D11), 20 685–20 711.

Lovejoy, S., and B. B. Mandelbrot, 1985: Fractal properties of rain, and a fractal model. *Tellus,* **37,** 209–232.

Luo, Z., and C. Liu, 2007: A validation of the fractal dimension of cloud boundaries. *Geophys. Res. Lett.,* **34,** L03808, doi:10.1029/2006GL028472.

Neggers, R. A. J., H. J. J. Jonker, and A. P. Siebesma, 2003: Size statistics of cumulus cloud populations in large-eddy simulation. *J. Atmos. Sci.,* **60,** 1060–1074.

Nuijens, L., B. Stevens, and A. P. Siebesma, 2009: The environment of precipitating shallow cumulus. *J. Atmos. Sci.,* **66,** 1962–1979.

Peters, O., and K. Christensen, 2002: Rain: Relaxations in the sky. *Phys. Rev. E,* **66,** 036120, doi:10.1103/PhysRevE.66.036120.

——, and ——, 2005: Rain viewed as relaxational events. *J. Hydrol.,* **328,** 46–55.

——, and J. D. Neelin, 2006: Critical phenomena in atmospheric precipitation. *Nat. Phys.,* **2,** 393–396.

——, C. Hertlein, and K. Christensen, 2002: A complexity view of rainfall. *Phys. Rev. Lett.,* **88,** 018701, doi:10.1103/PhysRevLett.88.018701.

——, J. D. Neelin, and S. W. Nesbitt, 2008: Mesoscale convective systems and critical clusters. *J. Atmos. Sci.,* **66,** 2913–2924.

Petersen, W. A., S. W. Nesbitt, R. J. Blakeslee, R. Cifelli, P. Hein, and S. A. Rutledge, 2002: TRMM observations of intraseasonal variability in convective regimes over the Amazon. *J. Climate,* **15,** 1278–1294.

Petty, G. W., 1999: Prevalence of precipitation from warm-topped clouds over eastern Asia and the western Pacific. *J. Climate,* **12,** 220–229.

Rauber, M. R., and Coauthors, 2007: Rain in shallow cumulus over the ocean: The RICO campaign. *Bull. Amer. Meteor. Soc.,* **88,** 1912–1928.

Sengupta, S. K., M. S. Navar, D. W. Chen, R. M. Welch, and T. A. Berendes, 1990: Cumulus cloud field morphology and spatial patterns derived from high spatial resolution Landsat imagery. *J. Appl. Meteor.,* **29,** 1245–1267.

Short, D. A., and K. Nakamura, 2000: TRMM radar observations of shallow precipitation over the tropical oceans. *J. Climate,* **13,** 4107–4124.

Snodgrass, E. R., L. Di Girolamo, and R. M. Rauber, 2009: Precipitation characteristic of trade wind clouds during RICO derived from radar, satellite, and aircraft measurements. *J. Appl. Meteor.,* **48,** 464–483.

Stevens, B., 2007: On the growth of layers of nonprecipitating cumulus convection. *J. Atmos. Sci.,* **64,** 2916–2931.

Telesca, L., G. Colangelo, V. Lapenna, and M. Macchiato, 2004: On the scaling behavior of rain event sequence recorded in Basilicata region (southern Italy). *J. Hydrol.,* **296,** 234–240.

Wielicki, B. A., and R. M. Welch, 1986: Cumulus cloud field properties derived using Landsat satellite data. *J. Appl. Meteor.,* **25,** 261–276.

Zhao, G., and L. Di Girolamo, 2007: Statistics on the macrophysical properties of trade wind cumuli over the tropical western Atlantic. *J. Geophys. Res.,* **112,** D10204, doi:10.1029/2006JD007371.