RESEARCH ARTICLE

Fitting Cumulus Cloud Size Distributions From Idealized Cloud Resolving Model Simulations

Julien Savre[1] and **George Craig[1]**

[1]Physics Department, Meteorological Institute, Ludwig-Maximilians-Universität, Munich, Germany

Abstract Whereas it is now widely accepted that cumulus cloud sizes are power-law distributed, characteristic exponents reported in the literature vary greatly, generally taking values between 1 and >3. Although these differences might be explained by variations in environmental conditions or physical processes organizing the cloud ensembles, the use of improper fitting methods may also introduce large biases. To address this issue, we propose to use a combination of maximum likelihood estimation and goodness-of-fit tests to provide more robust power-law fits while systematically identifying the size range over which these fits are valid. The procedure is applied to cloud size distributions extracted from two idealized high-resolution simulations displaying different organization characteristics. Overall, power-laws are found to be outperformed by alternative distributions in almost all situations. When clouds are identified based on a condensed water path threshold, using power-laws with an exponential cutoff yields the best results as it provides superior fits in the tail of the cloud size distributions. For clouds identified using a combination of water content and updraft velocity thresholds in the free troposphere, no substantial improvement over pure power-laws can be found when considering more complex two-parameter distributions. In this context however, exponential distributions provide results that are as good as, if not better than power-laws. Finally, it is demonstrated that the emergence of scale free behaviors in cloud size distributions is related to exponentially distributed cloud cores merging as they are brought closer to each other by underlying organizing mechanisms.

Plain Language Summary Clouds constitute an important element of the climate system reflecting incoming solar radiation and emitting infra-red radiation that heats the atmosphere. The net radiative impact of clouds however depends on many factors including their size. It is thus of prime importance to characterize the size of clouds, in particular convective clouds, and understand the underlying processes controlling them. In this study, a numerical model is used to simulate two convective situations at horizontal resolutions providing a fine description of cloud processes. After identifying individual clouds and calculating their size, statistical methods are employed to characterize the cloud size distributions. Depending on the situation, cloud size distributions are found to be best represented by either power-laws with an exponential cutoff or exponential functions. Pure power-laws, which constitute the most popular model used to represent cloud size distributions, are generally found to yield poorer fits. Finally, it is demonstrated that power-laws in cloud size distributions emerge when individual cloud cores, that are exponentially distributed in size, are brought closer to each other and merge as the cloud ensemble organizes.

1. Introduction

The first Landsat satellite launched in 1972 provided high-resolution data that allowed the first large-scale analysis of cloud properties and their radiative impact on our climate. It was later recognized that spatial inhomogeneities in individual cloud scenes are an important factor determining cloud radiative properties. It quickly appeared that an essential step toward better predicting the radiative impact of broken cloud fields (cumulus and stratocumulus) consisted in understanding the large-scale characteristics, the organization, as well as the small-scale structure of the cloud ensembles including their fractal shapes and size distributions (Kuo et al., 1988; Welch et al., 1988; Wielicki & Welch, 1986).

Early observational studies have suggested that cloud size distributions may be approximated by various functional forms including exponential (Plank, 1969; Wielicki & Welch, 1986) and lognormal (Houze Jr. & Cheng, 1977; Lopez, 1977) distributions. Later, Parker et al. (1986), Welch et al. (1988), and Kuo et al. (1988) introduced the power-law as the functional form that best represents cumulus and stratocumulus cloud size distributions. Accordingly, the number of clouds n having a certain size l can be described by:

$$n(l) \propto l^{-b}, \tag{1}$$

where $\propto$ should be read "scales as," and b is a characteristic power-law exponent. Since then, power-laws (or derived forms thereof like power-laws with exponential cutoff or broken power-laws) have been universally recognized as the best functions to model cloud size distributions obtained from either satellite imagery (Benner & Curry, 1998; Bley et al., 2017; Koren et al., 2008; Kuo et al., 1993; Mieslinger et al., 2019; Senf et al., 2018; Sengupta et al., 1990; Welch et al., 1988; Wood & Field, 2011; Zhao & Di Girolamo, 2007), aircraft measurements (Benner & Curry, 1998; Jiang et al., 2008; Wood & Field, 2011), or high-resolution simulations (Barron et al., 2020; Dawe & Austin, 2012; Garrett et al., 2018; Heus & Seifert, 2013; Jiang et al., 2008; Neggers, Jonker, & Siebesma, 2003; Rieck et al., 2014; Senf et al., 2018; van Laar et al., 2019; Xue & Feingold, 2006).

One problem that stands out from the large body of literature existing on cumulus cloud size distributions is the large spread in fitted power-law exponents. Values varying between ~1 (Benner & Curry, 1998) to more than 3 (Bley et al., 2017; Senf et al., 2018) can indeed be found, although most estimates reported lie between 1.7 and 2.4. These values are not inconsistent with other natural phenomena exhibiting power-law behaviors and for which typical exponents vary between 2 and 3 (Clauset et al., 2009; Newman, 2005). Note that an exponent value of 2 constitutes an important threshold since power-laws with exponents $b \leq 2$ have undefined mean, while power-laws with $2 < b \leq 3$ have finite mean but undefined variance. An important consequence of this is that any given observable following a power-law distribution with exponent in the 2–3 range has a directly measurable mean, but the undefined variance makes the occurrence of extreme outliers plausible. Nonetheless, we may ask ourselves whether this rather large variability in reported exponents can be explained by underlying physical processes, environmental conditions, large-scale forcing, or perhaps factors depending on retrieval techniques (from detection thresholds used in satellite imagery to model configurations in numerical simulations).

So far, few studies have focused on trying to explain what controls the shape of cumulus cloud size distributions. Kuo et al. (1993) suggested that α values tend to increase as clouds become more mature, indicating that α may have a connection with the clouds life cycle. Similarly, Neggers et al. (2019) noticed that cumulus cloud size distributions in a continental environment and the corresponding fitted power-law exponents follow a certain diurnal cycle. This result again points to the fact that cloud size distributions and their exponents depend on the maturity of the cloud population, possibly through the physical processes contributing to its organization. Besides, it is also known that boundary layer properties, and in particular their depth, may influence the size of the biggest clouds identified (Rieck et al., 2014; van Laar et al., 2019), although no clear relationship with power-law exponents were reported. Finally, Sakradzija and Hohenegger (2017) suggested that the shape of cloud mass flux distributions (which were fitted using double Weibull distributions) is to a large extent controlled by the Bowen ratio, that is the ratio between sensible and latent heat fluxes at the surface.

While the aforementioned studies suggest that environmental and forcing conditions are indeed likely to affect cumulus cloud size distributions, we may still wonder if the broad range of power-law exponents reported in the literature solely reflects changes in underlying physical processes. Before examining the impact of these processes on power-law exponents, one should first make sure that the methods used to identify power-law scalings in cloud size distributions are statistically robust to avoid possible misinterpretations. In other words, the fitting procedures employed should be considered as an additional and potentially large source of uncertainties that should also be called into question. It is indeed now generally admitted that fitting power-law distributions is a particularly difficult problem (Clauset et al., 2009; Goldstein et al., 2004; White et al., 2008), especially considering the lack of significant data in the tail of the considered empirical distributions, as well as the narrow and indistinct power-law ranges often seen in the data [ideally, power-laws should be valid over at least two decades to be considered reliable (Stumpf & Porter, 2012)]. Besides, the improper use of certain statistical fitting methods may also yield important biases. A notable common mistake consists in relying solely on the apparent constant slope of empirical distributions in log-log coordinates (Goldstein et al., 2004; White et al., 2008), and fit power-laws using simple linear regression over subjectively determined data ranges. It can be shown in this context that linear regression in log-log coordinates can induce exponent errors as large as 40% when fitting a true power-law distribution having an exact exponent of 2.5 (Goldstein et al., 2004).

In this work, we seek to reassess the determination of power-law scalings in cloud size distributions by focusing on issues related to improper fitting. Although we chose to focus our work on this particular problem that is often overlooked in the atmospheric science literature, we do not ignore the fact that the shape and slope of cloud size

distributions may also be strongly impacted by physical processes and external conditions. It is claimed however that the analysis of the impacts of the latter on cumulus cloud sizes can only be adequately addressed with fitting methods able to reliably identify power-law scalings and robustly estimate the associated exponents. A fitting approach is thus introduced to alleviate this issue, inspired by a methodology proposed by Clauset et al. (2009) who described a complete procedure dedicated to fitting power-law distributions in empirical data. Although the procedure has previously been applied to fit rain cluster size distributions (Peters et al., 2010; Traxl et al., 2016), we wish here to emphasize following Clauset et al. (2009) that the use of such methods should always be preferred over simple linear regression techniques when estimating power-law fits.

The power of the fitting algorithm is demonstrated using cloud size distributions obtained from Cloud Resolving Model (CRM) simulations of two idealized convective situations: a diurnal transition from shallow to deep convection over land (Grabowski et al., 2006), and a case of maritime shallow convection in the trade-winds region (vanZanten et al., 2011). These cases were selected as they are representative of cloud populations evolving at different spatial and temporal scales, with organization patterns driven by different physical mechanisms. Besides, the degree of idealization adopted facilitates our analysis by suppressing variability introduced by large-scale transport or heterogeneous surface conditions. Our main objective is then to adapt the methodology proposed by Clauset et al. (2009) in order to evaluate the hypothesis that cloud sizes in these two case studies are indeed power-law distributed and give precise estimates of power-law exponents. Thanks to the robustness of the statistical methods on which it is based, it is further demonstrated that the algorithm may be utilized to identify and characterize the impact of physical processes (organization) on cumulus cloud sizes.

The paper is organized as follows. The complete fitting procedure and the statistical methods forming the basis of the fitting algorithm are first introduced in Section 2. In Section 3, we give details on the numerical experiments used to extract cloud size distributions and demonstrate the capabilities of the algorithm. Power-law best fits obtained in the two case studies are then presented and assessed in Section 4. Various alternative distributions are then evaluated in Section 5, and the role of convective organization in the emergence of power-law behaviors in cloud size distributions is discussed in Section 6. Finally, we give our conclusions in Section 7.

2. Fitting Cloud Size Distributions

The fitting algorithm introduced hereafter was implemented in python and is available upon request as described in the Open research section of this manuscript.

2.1. The *Clauset et al.* Method

Clauset et al. (2009) proposed a detailed procedure to identify power-law distributions in empirical data and estimate their exponents with great precision. The whole procedure was thoroughly tested by Clauset et al. (2009) for many classical empirical distributions exhibiting power-law behaviors.

The Clauset et al. (2009) method relies first and foremost on the idea that, when dealing with heavy-tailed data, best fits should be sought for complementary cumulative distribution functions (denoted CCDF) instead of probability density functions (PDF). The advantage of using CCDFs is indeed twofold. First, calculating CCDFs does not require binning the data, whereas PDFs can be very sensitive to the binning procedure. Second, heavy-tailed PDFs may be very noisy, especially in their tail, while the corresponding CCDFs appear to be smoother. Once CCDFs have been obtained from empirical data, the Clauset et al. (2009) method then makes use of a combination of Maximum Likelihood Estimation (MLE) and goodness-of-fit tests to determine power-law best fits and estimate the corresponding characteristic exponents.

In the following, a modified version of the method proposed by Peters et al. (2010) and Deluca and Corral (2013) is described and used. The modifications introduced allow the application of the procedure to truncated distributions for which both the lower and upper bounds of validity (denoted l_{min} and l_{max}) must be simultaneously estimated. The complete procedure can be summarized as follows (statistical methods are described in Section 2.2):

1. Best fit power-law exponents $\hat{\alpha}$ are estimated using MLE for all possible values of l_{min} and l_{max}.
2. The goodness of these fits is evaluated for each triplet $\{\hat{\alpha}, l_{min}, l_{max}\}$ using the Kolmogorov-Smirnov (KS) goodness-of-fit test.
3. The best fit triplet is the one that minimizes the KS statistics D (provided it is smaller than an arbitrary threshold of 0.1), while the range covered by the fit remains larger than half of the total cloud size range [this constitutes the main difference with respect to Deluca and Corral (2013)].
4. Once a best fit triplet has been obtained, the p-value associated with the KS statistics is computed using Monte-Carlo sampling following Clauset et al. (2009). The fit is accepted if the p-value is larger than an arbitrary threshold of 0.05.
5. The power-law fit can finally be compared to alternative distributions over the range l_{min}–l_{max} using dedicated metrics and an alternative statistical test for model selection.

It must be noted that all clouds having a size smaller than $5\Delta x$ were systematically ignored when fitting the simulated distributions. This choice was made to reflect the inherent limitation of any high-resolution model set by the spatial discretization. It is thereby acknowledged that only motions with scales larger than $\sim 5\Delta x$ can be faithfully resolved by the model.

2.2. Statistical Methods

The procedure described above makes use of several statistical methods summarized below. Considering first the MLE method, let's assume a set of empirical data $\mathbf{x} = (x_{min}, \ldots, x_{max})$ that we wish to approximate by a known distribution $p_{X|\Theta}$ described by N parameters $\Theta = (\theta_1, \ldots, \theta_N)$. Defining the log-likelihood function as:

$$\ell(\Theta, \mathbf{x}) = \ln \prod_{x=x_{min}}^{x_{max}} p_{X|\Theta}(x), \tag{2}$$

the set of parameters $\hat{\Theta}$ yielding the best fit is the one that maximizes $\ell(\Theta, \mathbf{x})$. For numerous standard distributions such as non-truncated power-law and exponential distributions, the optimal parameters maximizing $\ell(\Theta, \mathbf{x})$ can be found analytically by letting $\partial \ell(\Theta, \mathbf{x})/\partial \Theta = 0$. In the more general case however, $\hat{\Theta}$ cannot be reduced to a simple analytical formula, and it must be obtained from the numerical maximization of the function ℓ.

In our algorithm, $\hat{\Theta}$ values maximizing the log-likelihood ℓ are obtained for any couple $\{x_{min}, x_{max}\}$. The best overall fit, corresponding to an optimal range and parameter $\hat{\Theta}$, is the one that minimizes the KS statistics D, provided that D remains larger than 0.1. D is simply defined as the maximum absolute distance between the empirical and theoretical CCDFs over the considered range:

$$D = \sup_{x_{min} < x < x_{max}} \left| P_e(x) - P_{X|\hat{\Theta}}(x) \right|, \tag{3}$$

where P_e and $P_{X|\hat{\Theta}}$ denote the empirical and theoretical cumulative distributions.

Whereas lower D values intuitively indicate better fits, the statistical significance of this quantity depends strongly on the sample size. For this reason, the statistics must be complemented by a $p-$value computed from the PDF of D, and allowing one to either accept or reject the null hypothesis that the empirical data was indeed drawn from the considered theoretical distribution. In the situation where the distribution's parameters are estimated, Monte-Carlo sampling must be used (Clauset et al., 2009). First, n points (n being the number of points in the empirical data set) are drawn randomly from the estimated theoretical distribution $p_{X|\hat{\Theta}}$. Next, new values of the best fit parameters are obtained for the simulated data, and the associated KS statistics are computed. Repeating the procedure a sufficiently large number of times (2,500 in our case), the probability of occurrence of the original D value can be estimated: the null hypothesis is then rejected if this probability is lower than a predefined threshold set to 0.05.

Finally, Vuong's test (Vuong, 1989) is employed in order to compare the power-law fits to several plausible alternatives and select the best overall model. Considering a best fit distribution $p_{X|\hat{\Theta}}$ determined over the range

$x_{\min}$–$x_{\max}$, and any alternative distribution $p_{X|\widehat{\Phi}}$ described by a set of parameters $\widehat{\Phi}$ fitted over the same x range, we first define the likelihood ratio LR as:

$$LR = -2\ln\frac{\mathcal{L}\left(\widehat{\Theta}, \mathbf{x}\right)}{\mathcal{L}\left(\widehat{\Phi}, \mathbf{x}\right)} = -2\left[\ell\left(\widehat{\Theta}, \mathbf{x}\right) - \ell\left(\widehat{\Phi}, \mathbf{x}\right)\right],\tag{4}$$

$\mathcal{L}$ being the standard likelihood function. We can then define a normally distributed Z score as:

$$Z = -\frac{LR + (k_{\Theta} - k_{\Phi})\ln N}{2\sigma_{\ell}\sqrt{N}},\tag{5}$$

where k represents the number of degrees of freedom of each distribution tested, N is the sample size and σ_{ℓ} is the standard deviation associated with the log-likelihoods of each individual sample. In our situation, Vuong's test possesses two particularly interesting properties: (a) if we consider two distributions providing equivalently good fits, Vuong's test will always prefer the simplest model over the more complex one thanks to the explicit dependence of Z on k; (b) the Z score is relatively straightforward to interpret since it corresponds to the number of standard deviations by which a given statistics deviates from its mean. In the present case, the null (alternative) model, $p_{X|\widehat{\Theta}}\left(p_{X|\widehat{\Phi}}\right)$, will be preferred if Z is significantly positive (negative), with a Z threshold arbitrarily set to ± 1.28 (corresponding to a 80% confidence interval).

2.3. Theoretical Distributions

We introduce here six theoretical distributions that are thought to constitute adequate models representing empirical cloud size distributions. As discussed previously, our starting hypothesis is that cloud size distributions follow power-laws. The five other distributions presented constitute our alternative hypotheses. We limited our selection to single functions fitting the entire or part of the empirical cloud size distributions, thus excluding distributions like broken power-laws (Benner & Curry, 1998; Kuo et al., 1993) or double Weibull distributions (Sakradzija & Hohenegger, 2017). Similarly, only theoretical functions having at most two degrees of freedom have been considered.

The candidate distributions, $p_X(l)$ (l being the cloud size and X denoting one of the following distributions), are introduced hereafter and their choice briefly justified. The corresponding CCDFs are defined as:

$$P_X(l) = \int_l^{l_{\max}} p_X(l^*)dl^*.\tag{6}$$

In the above, $P_X(l)$ is the probability that clouds have size larger than or equal to l, l being bounded by $l_{\min}$ and $l_{\max}$. The use of truncated functions is here fundamental given that the cloud size distributions analyzed generally extend over relatively narrow size ranges (less than two decades), cloud sizes being naturally limited by the model's horizontal grid spacing Δx on the one hand, and the extent of the numerical domain on the other hand.

2.3.1. Power-Law Distribution

The truncated power-law distribution and corresponding CCDF are defined as:

$$p_{PL}(l) = C_{PL}(\alpha - 1)l^{-\alpha}\tag{7}$$

$$P_{PL}(l) = \frac{l^{1-\alpha} - l_{\max}^{1-\alpha}}{l_{\min}^{1-\alpha} - l_{\max}^{1-\alpha}}.\tag{8}$$

C_{PL} is a normalizing coefficient depending on $l_{\min}$ and $l_{\max}$, and α is the power-law exponent. As explained in introduction, power-laws still constitute the most common choice to fit cloud size distributions and, more generally, empirical data exhibiting heavy tails. One particularly appealing aspect of power-law distributions is the fact that among all generative mechanisms proposed over the years (Newman, 2005; Sornette, 2006), many rely on very intuitive arguments that may explain the shape of empirical cloud size distributions. Note that power-law distributions can be considered a special case of both the Weibull and cutoff power-law distributions presented hereafter.

2.3.2. Exponential Distribution

The exponential distribution and its corresponding CCDF defined between l_{min} and l_{max} are respectively given by:

$$p_E(l) = C_E \lambda \exp(-\lambda l) \tag{9}$$

$$P_E(l) = \frac{\exp(-\lambda l) - \exp(-\lambda l_{max})}{\exp(-\lambda l_{min}) - \exp(-\lambda l_{max})}, \tag{10}$$

with, again, C_E a normalizing coefficient and λ the exponential rate parameter. Exponential distributions occur commonly in physics and are generally associated with waiting times. In addition, and according to the principle of maximum entropy, the exponential distribution is also the most likely distribution (i.e., the one maximizing the likelihood) describing data with known mean, but no other constraint. In this context, the exponential distribution has been frequently used in atmospheric sciences as one of the simplest models capable of representing cloud size and mass flux distributions (Craig & Cohen, 2006; Plank, 1969; Wielicki & Welch, 1986).

2.3.3. Rayleigh Distribution

The Rayleigh distribution and its CCDF defined between l_{min} and l_{max} are given by:

$$p_R(l) = C_R \frac{l}{\sigma^2} \exp\left(-\frac{l^2}{2\sigma^2}\right) \tag{11}$$

$$P_R(l) = \frac{\exp\left(-\frac{l^2}{2\sigma^2}\right) - \exp\left(-\frac{l_{max}^2}{2\sigma^2}\right)}{\exp\left(-\frac{l_{min}^2}{2\sigma^2}\right) - \exp\left(-\frac{l_{max}^2}{2\sigma^2}\right)}. \tag{12}$$

C_R is the normalizing coefficient and σ the distribution's scale parameter. A Rayleigh cloud size distribution is obtained by first assuming that the cloud mass flux distribution follows an exponential distribution, as proposed by Craig and Cohen (2006). Further noting that the cloud mass flux is proportional to the cloud surface area l^2, it can be shown using a simple change of variable that l must follow a Rayleigh distribution. The Rayleigh distribution therefore arises by applying the principle of maximum entropy to the cloud mass fluxes (i.e., it is the most likely cloud size distribution given a known, imposed total mass flux). Note that the Rayleigh distribution constitutes a special case of the Weibull distribution introduced later with $\beta = 2$.

2.3.4. The Lognormal Distribution

The truncated lognormal distribution and its CCDF are given by:

$$p_{LN}(l) = C_{LN} \frac{1}{l v \sqrt{2\pi}} \exp\left[-\frac{(\log l - \mu)^2}{2v^2}\right] \tag{13}$$

$$P_{LN}(l) = \frac{\Phi\left(\frac{\log l - \mu}{v}\right) - \Phi\left(\frac{\log l_{max} - \mu}{v}\right)}{\Phi\left(\frac{\log l_{min} - \mu}{v}\right) - \Phi\left(\frac{\log l_{max} - \mu}{v}\right)} \tag{14}$$

with C_{LN} a normalizing factor, and μ and v the two parameters of the distribution characterizing the mean of the variable's logarithm and its standard deviation, respectively. Φ is here the complementary error function. Lognormal distributions are frequently met in various branches of natural and social sciences. The lognormal distribution is notably the distribution of the product of n independent and identically distributed variables, as n becomes large (multiplicative central limit theorem). These distributions are thus often appropriate to model systems in which multiplicative growth plays an important role [as assumed by Lopez (1977) for cloud sizes]. Besides, lognormal and power-law distributions may be easily confused when plotted side by side in log-log coordinates where both appear linear, in particular if the lognormal's variance v^2 is large (Mitzenmacher, 2003). For these reasons, lognormal distributions constitute one of the most obvious alternatives to which power-laws should be compared.

2.3.5. Power-Law Distribution With Exponential Cutoff

The truncated power-law distribution with exponential cutoff and its CCDF are defined as:

$$p_{PE}(l) = C_{PE}l^{-\kappa}\exp(-\theta l) \tag{15}$$

$$P_{PE}(l) = \frac{\gamma(1-\kappa, \theta l) - \gamma(1-\kappa, \theta l_{\max})}{\gamma(1-\kappa, \theta l_{\min}) - \gamma(1-\kappa, \theta l_{\max})} \tag{16}$$

with γ the upper incomplete gamma function defined by:

$$\gamma(a, b) = \int_{b}^{+\infty} x^{a-1}e^{-x}dx. \tag{17}$$

C_{PE} is the corresponding normalizing coefficient, and θ and κ the two parameters of the distribution (κ being the power-law exponent and θ the exponential cutoff scale). The power-law with exponential cutoff is introduced here as a natural extension of the power-law distribution, noting that power-law distributed data in physics often exhibit an exponential decay beyond a certain threshold (Clauset et al., 2009; Newman, 2005). This behavior may be explained in particular by finite-size effects in systems otherwise exhibiting scale free properties (Stauffer & Ahaorny, 2003). Power-law distributions with exponential cutoffs have notably been employed by Peters et al. (2009), Windmiller (2017), and van Laar et al. (2019) to fit cloud size distributions.

2.3.6. The Weibull Distribution

Finally, the truncated Weibull distribution and its CCDF are defined as:

$$p_W(l) = C_W \beta \eta l^{\beta-1}\exp\left(-\eta l^{\beta}\right) \tag{18}$$

$$P_W(l) = \frac{\exp\left(-\eta l^{\beta}\right) - \exp\left(-\eta l_{\max}^{\beta}\right)}{\exp\left(-\eta l_{\min}^{\beta}\right) - \exp\left(-\eta l_{\max}^{\beta}\right)} \tag{19}$$

with C_W the normalizing coefficient, and β and η the two parameters of the Weibull distribution. Although mechanisms generating Weibull distributions are not as simple and well understood as those generating power-laws, these distributions nonetheless often provide excellent fits to empirical distributions exhibiting heavy-tails with cutoffs. As mentioned by Laherrere and Sornette (1998), the Weibull distribution may indeed stem from the superposition of finite-size scalings and other deviations from pure power-laws (including e.g., subsampling). As such, the Weibull distribution may be viewed as a more general form of the power-law distribution with exponential cutoff. Besides, Weibull distributions were recently used to model cloud mass flux distributions (Sakradzija & Hohenegger, 2017; Sakradzija et al., 2015), noting that these functions appropriately represent physical systems exhibiting memory and aging effects: large clouds being generally longer lived than small clouds, the probability to find large clouds at any given time is increased compared to what a memoryless exponential distribution would predict.

3. Numerical Experiments

3.1. Model Configuration

The first experiment is based on the Large-scale Biosphere-Atmosphere (LBA) intercomparison study (Grabowski et al., 2006), with modifications similar to Böing et al. (2012) and Savre (2021). The initial potential temperature profile is taken from Grabowski et al. (2006), while the initial moisture content is modified with a relative humidity held constant and equal to 80% between the surface and 2.5 km, and decreasing linearly to 15% up to 18 km. Horizontal winds are initially set to 0 m s^{-1} everywhere. Although a diurnal cycle is imposed through time dependent surface fluxes in the original case description, we chose here to adopt constant fluxes equal to 161 and 343 W m^{-2} (sensible and latent fluxes respectively) following Böing et al. (2012). The numerical model is doubly periodic and uses a constant horizontal grid spacing of 100 m with $1,024 \times 1,024$ grid points. The domain extends vertically to 14,250 m with 180 grid points. The vertical grid spacing is held constant and equal to 25 m below 2 km, and increases geometrically above. A Rayleigh damping layer with characteristic time scale of 2 hr
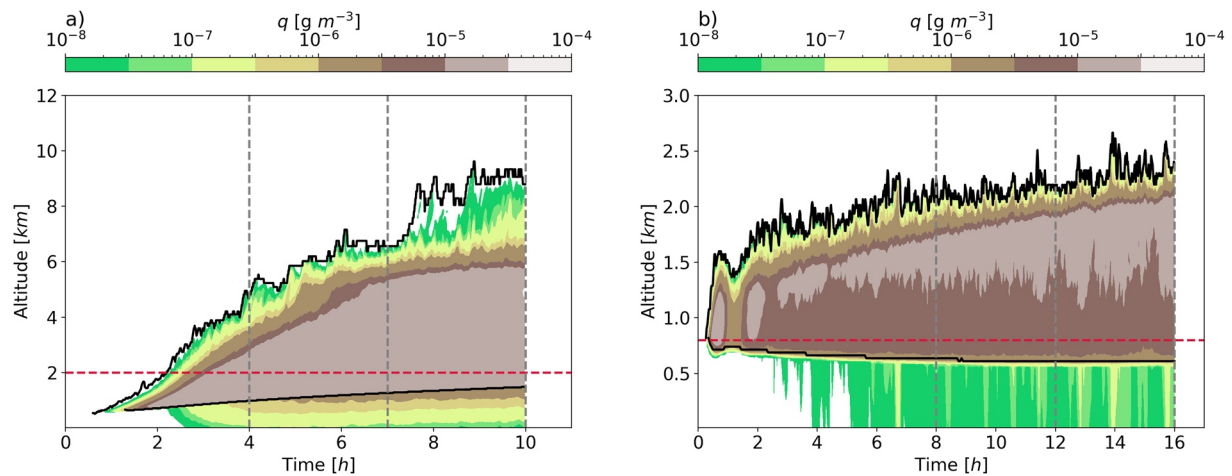
Figure 1. Time series of condensed water content q in the LBA (a) and RICO (b) simulations. Solid black lines denote the positions of the mean cloud base and maximum cloud top altitudes attained in each case. The vertical dashed lines indicate the locations in time when the cloud scenes and cloud size distributions are analyzed, whereas the horizontal red dashed lines denote the altitudes at which clouds were identified.

is prescribed above 12 km, and horizontal winds are nudged everywhere to their initial value with a time scale of 6 hr. The simulation was run for a total of 10 hr, but the first 4 hr were discarded.

The second case follows the Rain In Cumulus over the Ocean (RICO) model intercomparison (vanZanten et al., 2011). The overall setup follows closely the description given by vanZanten et al. (2011), with similar initial potential temperature, total water mixing ratio, horizontal wind, and large scale forcing profiles. No interactive radiation is employed, but net radiative cooling is accounted for in the prescribed forcing. Surface fluxes are computed using bulk formula with a prescribed surface temperature of 299.8 K, and surface water vapor assumed to be at saturation. The numerical domain extends over 20×20 km^2 in the horizontal plane, with a homogeneous grid spacing of 26 m (768×768 grid points). This is similar to the setup employed by Heus and Seifert (2013) despite a slightly smaller domain (25 km^2 in Heus and Seifert (2013)). The domain extends vertically up to 3.9 km with a constant spacing of 26 m. A Rayleigh damping layer with characteristic time scale of 2 hr is used above 3.2 km. The simulation was run for 16 hr, and the results analyzed after 8 hr.

Figure 1 displays the temporal evolution of the simulated cloud water mixing ratio q in the LBA and RICO cases. The LBA simulation exhibits a fast transition from shallow to deep, precipitating convection, with clouds reaching an altitude of about 4 km after only 3 hr. The cloud population then continues to evolve under the influence of subcloud layer organization (Böing et al., 2012), with clouds continuously deepening and eventually reaching an altitude of 9 km after 8 hr. After this, the cloud layer remains quasi-stationary despite the fact that the boundary layer slowly deepens. In contrast, the RICO case is mostly driven by the large-scale forcing held constant during the course of the simulation. After an initial spin-up time of about 4 hr, the cloud population evolves only slowly and clouds remain relatively shallow with cloud tops reaching at most 2.8 km in altitude, consistent with previous studies (Heus & Seifert, 2013; Radtke et al., 2021). The mean cloud base remains at a constant altitude of about 600 m and only light drizzle is observed in the subcloud layer and at the surface. Whereas cloud organization in the LBA case is primarily controlled by cold pool dynamics, vertical wind shear plays an important role in the RICO simulation with clouds aligning along parallel rolls (see Text S1 in Supporting Information S1 for an impression of the cloud population and organization in both cases).

Both experiments were carried out using the MISU-MIT Cloud and Aerosol (MIMICA) model solving anelastic governing equations for potential temperature, total water mass mixing ratio and momentum (Savre, 2021; Savre et al., 2014). The numerical methods employed in both cases are similar: scalar advection uses a flux-limited version of the Lax-Wendroff scheme, momentum advection uses fourth order central finite differences, and time integration is performed using a second order Runge-Kutta method. Turbulent mixing is parameterized using the Smagorinsky-Lilly closure, whereas no interactive radiation is necessary in either case. Different microphysical schemes are employed for the LBA and RICO simulations. In the first case, we use the simple one-moment microphysics from Grabowski (1998) which only distinguishes between precipitating and non-precipitating cloud

particles, and partitions liquid and ice based on a linear function of temperature. In the second case, the warm part of the more advanced two-moment microphysics scheme from Seifert and Beheng (2006) is used with a fixed cloud droplet number concentration of 70 cm^{-3}.

3.2. Cloud Identification

The first step toward estimating cloud size distributions is to define cloud objects and calculate their properties. Two different criteria are used to identify clouds: a cloud water mixing ratio (q) exceeding 10^{-3} g kg^{-1} combined with a vertical velocity (w) threshold of 0.1 m s^{-1} in the lower troposphere (hereafter referred to as the "updraft criterion"), or a Condensed Water Path (CWP, including both liquid and ice particles) exceeding 50 g m^{-2} in the LBA case, or 10 g m^{-2} in the RICO case (hereafter referred to as the "CWP criterion"). The updraft criterion is applied within horizontal slices extracted at 2,000 m and 800 m in the LBA and RICO cases, respectively. These two levels were selected as they consistently remain above cloud base, as seen on Figure 1, and allow the identification of most of the clouds rooted in the boundary layer while leaving cloud remnants at higher altitudes out of the analysis. For the CWP criterion, two different thresholds were chosen in the two cases to account for the different cloud depths.

Clouds are then identified as clusters of connected grid boxes respecting either criterion given above. Four point connectivity is considered, that is grid boxes must at least share a face to belong to the same cloud object. Once cloud objects are identified, their equivalent sizes are calculated as $l = \sqrt{\Delta x \Delta y N_{pts}}$, with N_{pts} the number of grid points belonging to a cloud.

The distinction between clouds identified from updraft and CWP criteria is here justified by the fact that both definitions correspond to distinct situations commonly met in the literature. The CWP criterion allows the identification of clouds as seen from above, by analogy with clouds retrieved from satellite imagery. In contrast, the updraft criterion is relevant in the context of convection parameterizations, these latter being concerned with the representation of upward transport in the lower troposphere.

4. Evaluating the Power-Law Hypothesis

We are initially only interested in fitting simulated cloud size distributions to power-law functions. The objective is to evaluate the null hypothesis that cloud size distributions are suitably modeled by power-laws using the *Clauset et al.* method introduced in Section 2.1.

4.1. Visual Inspection

Cloud size CCDFs extracted at three different times into the LBA and RICO simulations are shown in Figures 2 and 3, respectively. Each distribution is multiplied by the total number of clouds identified in each scene and therefore represents the number of clouds with size larger than or equal to a given value l. At this point, we are only concerned with the empirical distributions depicted as thick gray lines. Cloud size density distributions are also available as Text S2 in Supporting Information S1 and essentially exhibit the same features as the CCDFs.

In the LBA simulation, starting at 4 hr, clouds identified using a CWP threshold are generally bigger and twice as numerous as those determined from the updraft criterion (~6,000 with CWP as opposed to ~2,200 with the updraft criterion). At later times, the largest clouds have widths of 1.5 and 4 km when determined based on the updraft and CWP criteria, respectively. The number of clouds identified using CWP is also largely reduced to ~3,000. The emergence of much bigger clouds indicates the development of convective outflows and the possible merging of cloud objects as convection becomes deeper. After 10 hr, a clear linear scaling in log-log coordinates between 200 m and over 2,000 m is visible, and the largest clouds reach up to 9 km in size. Still no clear linear scaling is visible on the cloud size CCDF obtained from the updraft criterion.

As expected, cloud size CCDFs extracted from the RICO simulation (Figure 3) do not show the same evident temporal evolution as the ones obtained from the LBA case. The total number of clouds identified based on the updraft criterion is relatively constant, between 300 and 400, which constitutes a relatively small sample to guarantee the estimation of reliable fits. This is especially true if one is interested in fitting power-laws given the narrow size range over which the empirical distributions extend (up to ~700 m). Again, more clouds can be identified using a CWP threshold (between 600 and 700), with cloud objects reaching up to 3.3 km in size after 16 hr. No clear power-law behavior is evident, regardless of the cloud identification criterion considered.
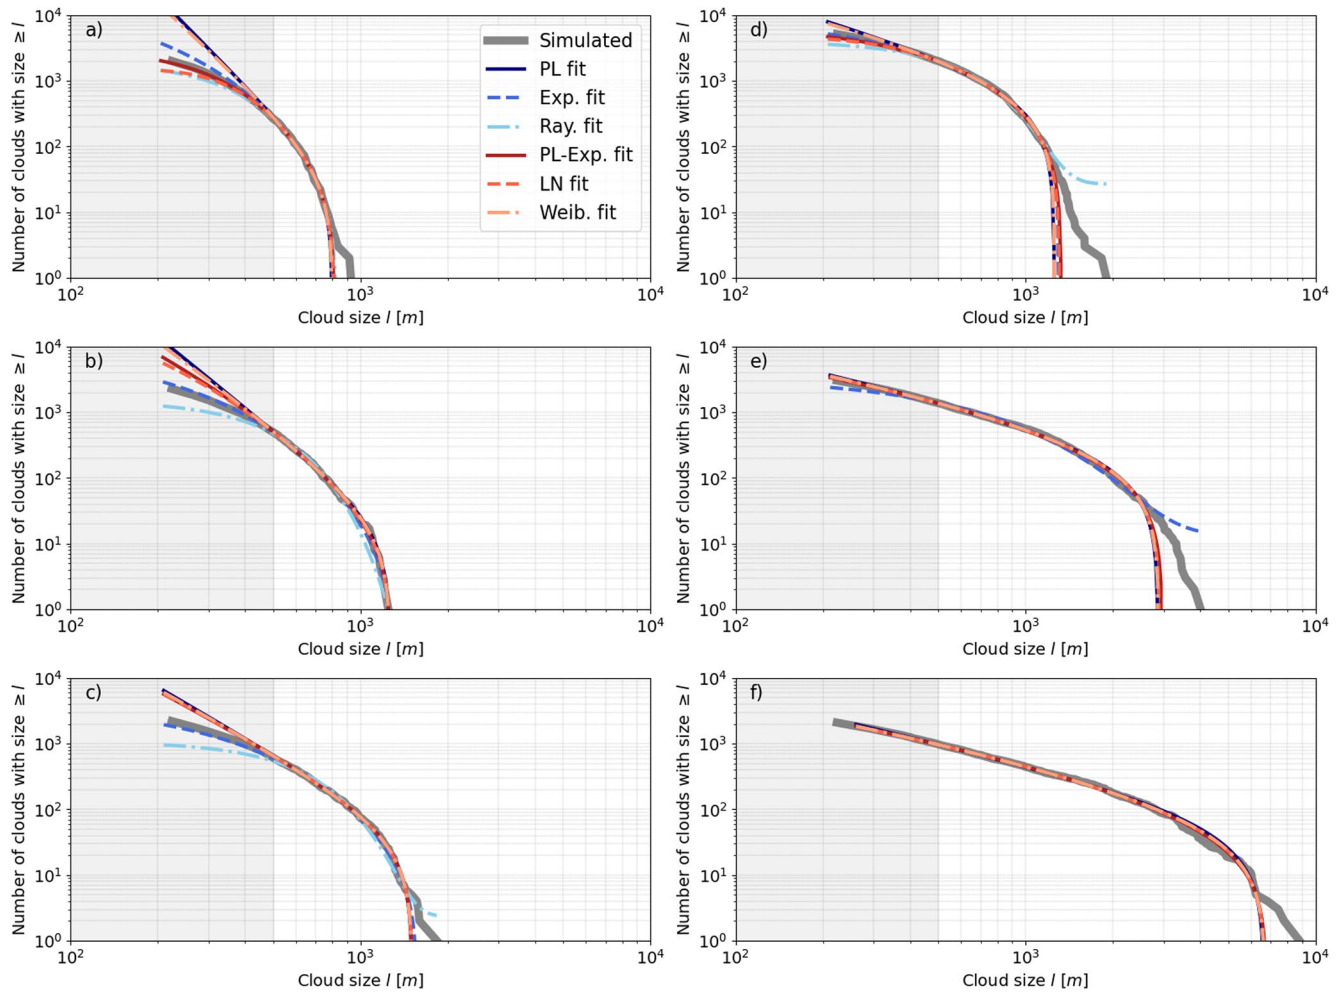
Figure 2. Cloud size CCDFs calculated for the LBA case, at three different times into the simulation. From top to bottom: at 4, 7, and 10 hr. Left column: clouds identified based on the updraft criterion at 2,000 m. Right column: clouds identified based on the CWP criterion. Best fits obtained with all theoretical distributions are also plotted. The corresponding best fit parameters are collected in Table 1. The range of sizes not included when fitting the cloud size CCDFs (clouds smaller than $5\Delta x$) is indicated by the gray shading.

Table 1
Best Fit Parameters Estimated From the Cloud Size CCDFs Obtained in the LBA Simulation (Fits Are Shown on Figure 2)

Time	$\hat{\alpha}$	$\hat{\lambda}$	$\hat{\sigma}$	$\hat{\mu}$-$\hat{\nu}$	$\hat{\eta}$-$\hat{\beta}$	$\hat{\kappa}$-$\hat{\theta}$
10 hr	3.57	4.04×10^{-3}	424.7	-5.32–2.17	0.045–43.1	3.25–0.37×10^{-3}
7 hr	4.58	6.06×10^{-3}	327.7	4.46–0.77	0.09–21.2	2.94–2.18×10^{-3}
4 hr	5.46	8.67×10^{-3}	248.6	5.90–0.35	0.10–22.1	-2.97–0.013
10 hr	2.02	$-$	$-$	1.31–2.42	0.13–3.63	1.81–0.12×10^{-5}
7 hr	2.02	1.88×10^{-3}	$-$	1.38–2.32	0.08–7.57	1.68–0.33×10^{-3}
4 hr	2.29	3.12×10^{-3}	415.7	6.12–0.60	0.16–3.09	-0.68–4.03×10^{-3}

Note. Cloud objects are identified based either on the updraft criterion (see text) at $z = 2,000$ m (top rows), or the CWP criterion (bottom rows). The parameters reported correspond to (from left to right): power-law, exponential, Rayleigh, lognormal, Weibull and cutoff power-law distributions.
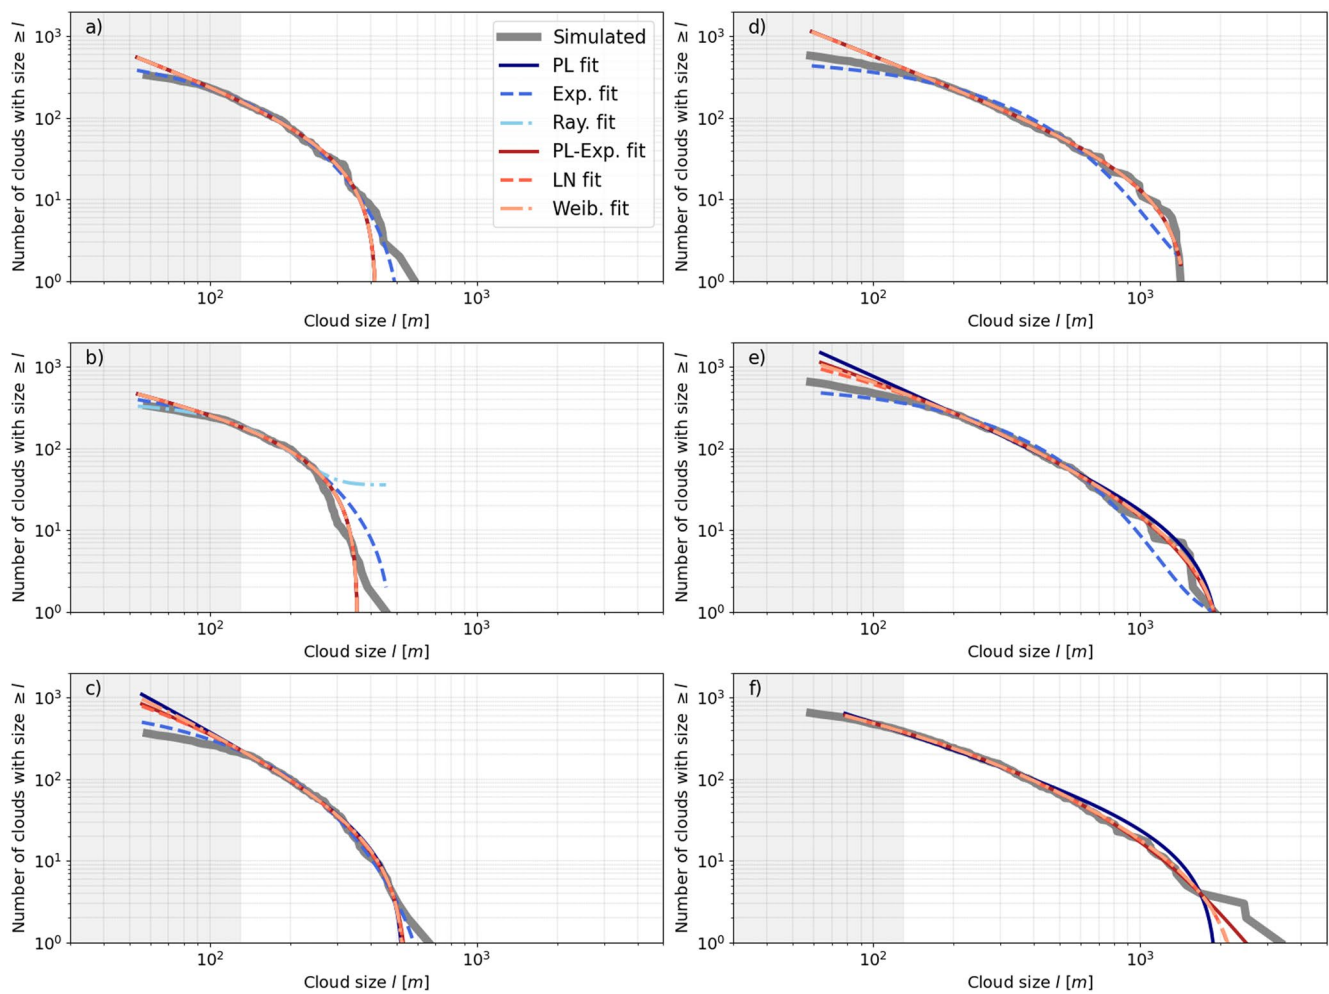
Figure 3. Same as Figure 2 for the RICO case. From top to bottom: at 8, 12, and 16 hr. The corresponding best fit parameters are collected in Table 2.

As highlighted above, the visual determination of power-law behaviors from empirical cloud size CCDFs is not always evident. Besides the fact that the considered distributions may eventually not obey power-law scalings at all, it must be stressed that truncated power-law distributions, generally reflecting the emergence of finite-size scalings, do not necessarily follow straight lines in log-log coordinates. This is all the more so true as the data ranges considered here (and in most similar studies) remain relatively narrow. It is indeed generally recommended that the fitted data extends over at least two decades to obtain reliable power-law fits (Stumpf & Porter, 2012).

Table 2
Same as Table 1 but for the RICO Simulation (Fits Are Shown on Figure 3)

Time	$\hat{\alpha}$	$\hat{\lambda}$	$\hat{\sigma}$	$\hat{\mu}$-$\hat{\nu}$	$\hat{\eta}$-$\hat{\beta}$	$\hat{\kappa}$-$\hat{\theta}$
16 hr	2.76	0.011	–	3.08–1.17	0.15–5.34	1.93–3.37×10^{-3}
12 hr	1.59	9.67×10^{-3}	104.8	-595–32	3.02×10^{-3}–192.5	1.59–0
8 hr	2.17	0.011	-	-300–16.1	5.17×10^{-3}–221.6	2.17–0
16 hr	2.04	–	–	3.74–1.40	0.34–0.57	1.56–1.25×10^{-3}
12 hr	2.46	4.40×10^{-3}	–	3.37–1.39	0.24–1.57	1.99–0.85×10^{-3}
8 hr	2.27	4.49×10^{-3}	–	-312–15.8	4.80×10^{-3}–258	2.27–0

The non-linearity of truncated power-laws in log-log coordinates can be illustrated by taking the logarithm of the left- and right-hand sides of Equation 8, yielding:

$$\log P_{PL}(l) = (1-\alpha)\log\left(\frac{l}{l_{\min}}\right) + \log\left[\frac{1 - \left(\frac{l_{\max}}{l}\right)^{1-\alpha}}{1 - \left(\frac{l_{\max}}{l_{\min}}\right)^{1-\alpha}}\right]. \tag{20}$$

The first term on the right-hand side is linear with respect to l in log-log coordinates, but the second term is clearly non-linear and arises because of the introduction of the upper size bound $l_{\max}$. Note that if empirical distributions indeed follow truncated power-laws defined by Equation 8, a pure linear behavior in log-log coordinates will still be observed as l approaches $l_{\min}$, and the extent of the linear part will be larger as α is increased.

4.2. Power-Law Best Fits

Figure 4 shows the temporal evolution of the best fit power-law exponents α obtained for the LBA and RICO cloud size CCDFs. The power-law hypothesis appears to be plausible (circles in Figure 4) at most times analyzed in both cases based on a p-value threshold arbitrarily set to 0.05. This result however merely suggests that the amount of samples used in each case to obtain the fits is sufficient given the constraints imposed by the algorithm (Section 2.1). This test alone is not designed to evaluate the quality of each fit, and dedicated statistics must be used to this end (see Section 5.2).

Considering only clouds defined based on CWP in the LBA case, fitted α values remain relatively constant in time over the whole period, in particular after 6 hr, with an average α computed from 6.5 hr onwards equal to 2.03. Larger overall exponents and more temporal variability are found over the same time period for clouds identified based on the updraft criterion, with a mean α of 4.12. To further evaluate the robustness of our estimates, Root Mean Square Errors (RMSE) can be computed for each fit using a standard bootstrapping procedure. The method involves resampling the empirical data with replacement, computing the corresponding best fit parameters for the resampled data, and repeating the fitting operation a sufficient number of times (here 2,500 times) to calculate reliable statistics. The average RMSE (represented by error bars on Figure 4) computed between 6.5 and 10 hr for clouds defined based on CWP remains low at about 0.06, indicating that the power-law scaling is a robust feature of these distributions. Again, more variability is found when cloudy updrafts are considered with a RMSE of 0.34 over the same period, making the power-law fits less reliable. In fact, the large variability observed on Figure 4 as well as the large RMSE could be interpreted as indications that cloud sizes are, in this situation, not power-law distributed.
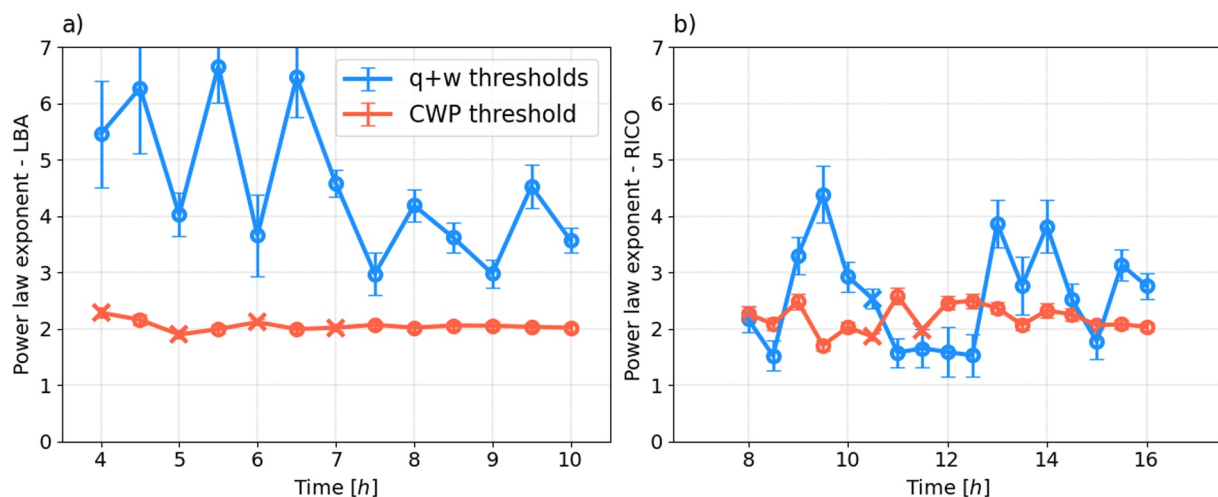

Figure 4. Time series of power-law exponents α estimated in the LBA (left) and RICO (right) cases. Uncertainty bars correspond to the root mean square errors computed from standard bootstrapping (see text). Estimates that satisfy the $p-$value condition (also obtained from standard bootstrapping, with a p threshold arbitrarily set to 0.05) are represented as circles. Estimates that do not satisfy the condition are represented as crosses.

In the RICO simulation, average best fit α values amount to 2.19 for clouds determined based on CWP and 2.58 for cloud identified based on the updraft criterion. The corresponding RMSE values are equal to 0.1 and 0.34, respectively. Again, both Figure 4 and the RMSE values reported previously indicate that the power-law fits determined from CWP clouds are more robust than those obtained with the other cloud criterion.

Power-law exponents estimated by our algorithm for CWP clouds during RICO are associated with relatively large temporal fluctuations (at least compared to the LBA case) despite the fact that the cloud population does not evolve significantly after 8 hr. This variability likely stems from fluctuations in the tail of the distributions, α being found to correlate particularly well with l_{max}. For example, for clouds identified based on CWP, α steadily increases from ~1.9 for $l_{max} \approx 1,000$ m to ~2.6 for $l_{max} \approx 2,500$ m. It can be further verified that the α RMSE can be significantly reduced when the fits are performed over fixed size ranges (a mean α of 2.07 with a RMSE of 0.07 were founds for fits performed over the range 150–800 m). This result is most likely indicative of a subsampling issue, the number of large clouds available within each instantaneous snapshot being insufficient to provide robust α estimates. A similar sampling issue within instantaneous shallow cumulus cloud scenes was discussed by Neggers et al. (2019) who noted increased cloud number variability when considering smaller spatial domains in which the largest clouds can be severely undersampled. Note that no such relationship between α and l_{max} was found in the LBA case which further supports our assumption as well as the reliability of the LBA fits.

Considering only situations in which the algorithm returns robust α estimates (i.e., for clouds determined based on CWP), most fitted exponents take values between 1.5 and 2.5 regardless of the case considered. For comparison, mean exponents reported in the literature generally range between ~1 and ~3.3, depending on the case and source of the data. For example, for deep convective cases (like LBA), exponents between ~1.9 (Rieck et al., 2014) and ~2.6–3.3 (Bley et al., 2017; Senf et al., 2018) were reported from high-resolution simulations. Exponents derived from high-resolution satellite imagery in similar situations vary between 1.7 (Kuo et al., 1993) and ~3.05 (Bley et al., 2017). Considering maritime shallow convection, our average α estimates are significantly smaller than the reference values (without cloud tracking) of 2.42–2.61 reported by Heus and Seifert (2013) for similar RICO simulations. Heus and Seifert (2013) however showed that α could vary greatly depending on model configuration (domain size, numerical advection, cloud detection thresholds…) although most of their estimates were still larger than 2.19. Exponents reported for other shallow convection simulations are generally smaller, from 1.7 (Neggers, Duynkerke, & Rodts, 2003; Neggers, Jonker, & Siebesma, 2003; Xue & Feingold, 2006) to 1.9 (Dawe & Austin, 2012), which was interpreted by Heus and Seifert (2013) as being a consequence of clouds remaining too shallow. Finally, satellite retrievals generally yield exponents ranging between 1.6 and 2.2 (Benner & Curry, 1998; Koren et al., 2008; Rodts et al., 2003; Wood & Field, 2011; Zhao & Di Girolamo, 2007), consistent with our data. The latter studies are in general particularly relevant as cloud size distributions extending over several orders of magnitudes [up to 5 decades in Wood and Field (2011)] can be retrieved from satellite imagery, therefore providing more reliable fits.

4.3. Linear Regression in Log-Log Coordinates

In general, trying to fit power-laws over size ranges simply guessed from visual inspection, and using linear regression in log-log coordinates are two common sources of fitting inaccuracies. To evaluate these two propositions, power-law exponents were recomputed using linear regression applied over varying size ranges for all empirical distributions analyzed previously.

To begin with, linear regression is applied over fixed size ranges determined visually from the distributions plotted in Figures 2 and 3. Logarithmic binning, which already reduces fitting errors a lot (Goldstein et al., 2004; White et al., 2008), is employed to estimate the power-law exponents. Exponents obtained from linear regression are compared to our MLE estimates on Figure 5. In the LBA case, linear regression is able to reproduce the power-law exponent predicted by MLE with relatively good accuracy, giving α values for clouds identified based on CWP and updraft criteria averaged between 6.5 and 10 hr of 2.11 and 4.11, respectively (comparable to the MLE estimates of 2.03 and 4.12). The corresponding RMSE computed over the same intervals are however twice as large at 0.09 and 0.75, making linear regression less reliable than MLE. The largest discrepancies between linear regression and MLE are found at earlier times, when no statistically significant power-law fit could be identified.

Mean estimates from linear regression computed for the RICO simulation are also found to remain close to the MLE values despite showing more variability and being less reliable. Mean exponents obtained using linear
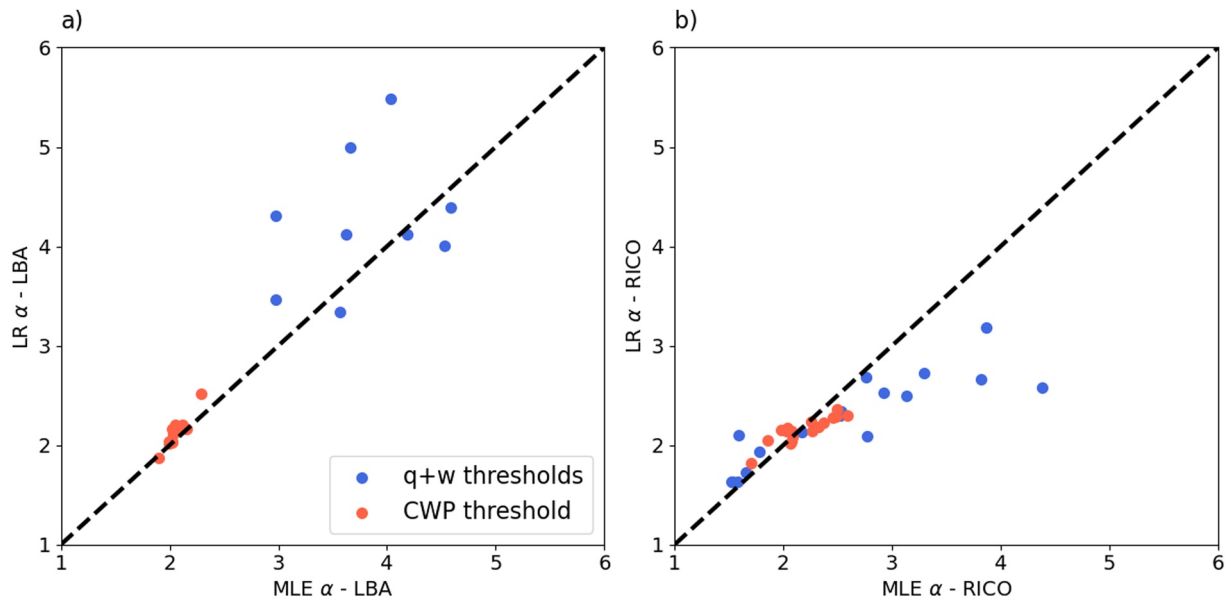
Figure 5. Power-law exponents estimated using the MLE based fitting algorithm for the LBA (left) and RICO (right) cases against equivalent estimates obtained from linear regression.

regression for clouds identified based on the CWP and updraft criteria indeed amount to 2.15 and 2.26 respectively, with mean RMSEs of 0.24 and 0.63.

To complement the analysis, best fit exponents obtained for CWP clouds using linear regression at time 10 hr in the LBA case and 16 hr in the RICO case were recomputed using varying lower and upper size bounds. Results are shown on Figure 6. For comparison, the upper and lower size bounds, l_{min}–l_{max}, returned by the fitting algorithm are equal to 500–6,829 m and 130–2,498 m in the LBA and RICO simulations, respectively. In both cases, α takes values between ∼1.7 and 2.4, again consistent with the range of exponents found in the literature and mentioned in Section 4.2. α generally increases with increasing l_{max} as linear regression tends to be overly
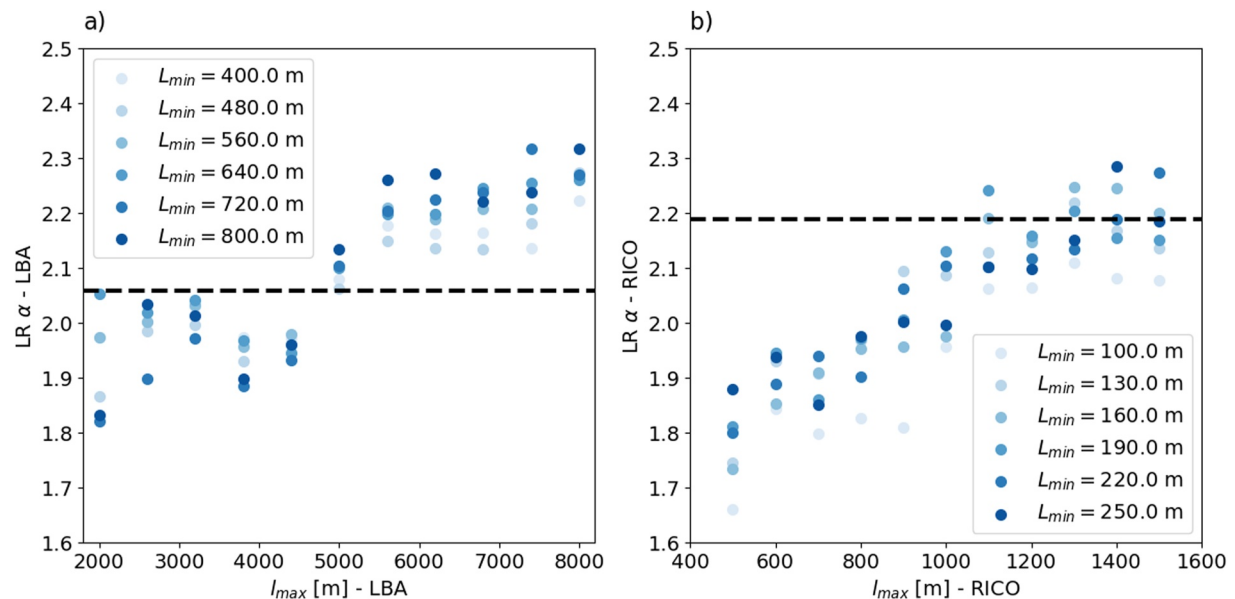

Figure 6. Sensitivity of linearly regressed exponents to the lower and upper size bounds for clouds identified based on CWP in the LBA (left) and RICO (right) cases. Best fits were obtained at the end of each simulation, that is at 10 and 16 hr respectively. The horizontal dashed lines correspond to the mean power-law exponents returned by the fitting algorithm based on MLE.
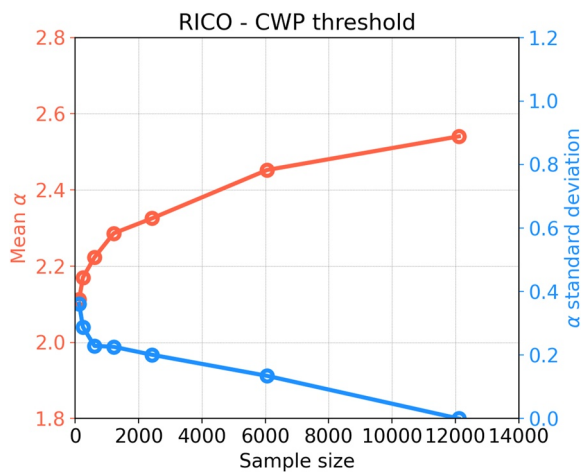
Figure 7. Sensitivity of best fit α estimates and their standard deviation to sample size in the RICO simulation for clouds identified based on CWP. A large cloud ensemble was obtained by collecting all clouds identified in the simulation within a 16 hr period. Bernoulli sampling was used to create each subsample with acceptance probabilities of {0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1}, and 1,000 samples were created in each case to produce significant statistics.

sensitive to the few, noisy points located in the tail of the distributions where the slope is steeper. Applying linear regression over too narrow size ranges is therefore likely to underestimate the fitted power-law exponents. In contrast, α does not show any systematic dependence on l_{min} although it is clear that only small changes in the arbitrarily selected minimum size bound can result in relatively large errors of the fitted exponents.

Overall, these results suggest that fitting power-laws using linear regression rather than MLE may yield very similar best fit exponents provided that: (a) the distributions indeed exhibit a very clear power-law behavior, and (b) the size range over which the fit is computed has been properly selected. This last point is particularly cirtical since the power-law range over which cloud size distributions are fitted is often determined by eye. This may explain, at least in part, the wide range of power-law exponents reported in the literature (see Section 4.2).

4.4. Subsampling and Finite Size Effects

As seen from the results discussed previously, finite size effects and, in particular, subsampling are two important factors likely to introduce large biases in the identification and evaluation of power-law scalings in empirical cloud size distributions. Note that sampling issues were specifically addressed, although in a slightly different manner, by Neggers et al. (2019).

In the present study, subsampling errors are expected to be particularly large for the RICO data where the amount of clouds available in each scene can be as low as ∼200. To estimate the impact of sample size on power-law best fit estimates, we first collect all clouds identified in the RICO simulation over a 16 hr period (17 cloud scenes), and apply the fitting algorithm to subsamples of different sizes. A series of 7 sampling experiments are conducted using Bernoulli sampling with acceptance probabilities of {0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1}, and 1,000 independent samples were created in each case to obtain reliable statistics. Mean α values and standard deviations obtained for each sample size using the fitting algorithm are plotted on Figure 7 for clouds identified based on CWP. Note that cloud size CCDFs never exhibit power-law scalings when using the updraft criterion, even when considering the entire cloud population obtained from the 17 cloud scenes analyzed.

Clearly, subsampling tends to underestimate the power-law exponent α. When considering a sample size of ∼600 clouds, which corresponds approximately to the amount of clouds available in each individual cloud scene, a mean α value of 2.22 is found, consistent with the mean estimate of 2.19 (Section 4.2), but with a comparatively large standard deviation of 0.22. The estimated exponent then increases (and the associated standard deviation decreases) when increasing the sample size, up to $\alpha = 2.54$ when the whole cloud population available is considered. This estimate is now in better agreement with the range of exponents reported by Heus and Seifert (2013) in their RICO simulations (2.42–2.61). This sensitivity of α to sample size is relatively large compared to what we might expect from MLE (see e.g., White et al. (2008)), but this may be explained by the narrow size range over which power-law fits generally hold that exaggerates small errors resulting from fitting biases (including subsampling). Note that the identification of power-law behaviors in empirical distributions obtained from subsampled data, and therefore the reliability of the estimated exponents, is all the more questionable that data sampled from pure power-law distributions may not even be power-law distributed themselves (Stumpf et al., 2005).

Finite size effects may act on top of this issue to further limit the relevance and precision of estimated power-law exponents. By restricting the size of the biggest clouds allowed in the simulation, the use of small numerical domains may indeed produce too narrow cloud size distributions which then enhances potential errors made when estimating power-law fits. A similar issue was discussed by Heus and Seifert (2013). The authors indeed found in their RICO simulations that the largest clouds sampled were smaller when the domain size was reduced, and that the estimated power-law exponents were impacted (although slightly). Besides, if not by directly affecting the power-law size range, the domain size may also impact the spatial organization of the cloud ensemble and thus

indirectly influence cloud size distributions and their fits. The impact of domain size on convective organization is indeed now a well documented issue, in particular in the context of mesoscale aggregation (Wing et al., 2018).

5. Alternative Distributions and Model Selection

5.1. Visual Comparison

Figures 2 and 3 allow us to visually compare best fits obtained for all distributions proposed in Section 2.3. All alternative fits were obtained using MLE applied over the same optimal size ranges as those found for power-laws. Corresponding best fit parameters are reported in Tables 1 and 2.

The first thing that stands out from Figures 2 and 3 is that the alternative fits are often undistinguishable from the power-law fits. This is particularly true when considering the fits obtained with the more complex, multi-parameter distributions, that is with the lognormal, Weibull and cutoff power-law distributions. Only in rare instances do the more complex alternatives provide visibly better fits than power-laws (see e.g., Figures 3e and 3f). Clearly, unambiguously identifying true power-laws from empirical data cannot be done based simply on visual inspection.

More obvious discrepancies can be seen when comparing the power-law fits to the ones obtained with exponential and Rayleigh distributions. The latter rarely constitutes an acceptable alternative, whereas the former only provides reasonably good fits when the size of the biggest clouds remains limited and when disregarding the tails of the empirical distributions.

5.2. Objective Comparison and Model Selection

In order to compare the individual best fits shown on Figures 2 and 3, we first introduce an RMSE-like score calculated between empirical and fitted distributions, defined as:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1,N} \left[P_e(l_i) - P_{X|\hat{\Theta}}(l_i) \right]^2}, \tag{21}$$

with P_e and $P_{X|\hat{\Theta}}$ the empirical and best fit CCFDs, and N the number of samples over which the fits were obtained. In the LBA case, RMSE values reported on Figure 8 confirm our previous observations, namely that the two-parameter distributions (lognormal, Weibull and cutoff power-law) generally provide fits as good as, but rarely substantially better than, power-laws. In the RICO case, the two-parameter distributions generally perform better than power-laws when considering CWP clouds, but RMSEs obtained with power-laws and the two-parameter alternatives are again relatively close for clouds identified based on the updraft criterion. Overall, RMSE does not give enough details to identify the model that provides the best approximation to the empirical cloud size distributions.

As an alternative to the RMSE score above, we now introduce a new statistics, denoted S_{cc}, representing the relative cloud cover error made when modeling cloud size distributions using one of the theoretical functions defined in Section 2.3. S_{cc} is computed as:

$$S_{cc} = \frac{\sum_{k=1}^{N} \tilde{l}_k^2 - \sum_{k=1}^{N} l_k^2}{\sum_{k=1}^{N} l_k^2}, \tag{22}$$

where l indicates the size of the N empirical clouds available in each scene, whereas $\tilde{l}$ is used to refer to the size of a cloud randomly drawn from a given best fit distribution $p_{X|\hat{\Theta}}$. S_{cc} constitutes a particularly relevant criterion to evaluate the various cloud size models provided that the total cloud cover and, by extension, the convective mass flux are two important parameters used to characterize convective cloud ensembles.

As shown on Figure 9, the power-law and exponential distributions generally yield the largest cloud cover errors of all the distributions tested. Power-law fits tend to overestimate the total cloud cover while exponential fits underestimate it. This result may be explained by the fact that, in all situations depicted, half of the total cloud cover is accomplished by the 10%–20% largest clouds in the population (see Table 3). This means that small fitting errors made on the large cloud side of the spectrum (i.e., clouds larger than the size thresholds indicated in Table 3) lead to large errors
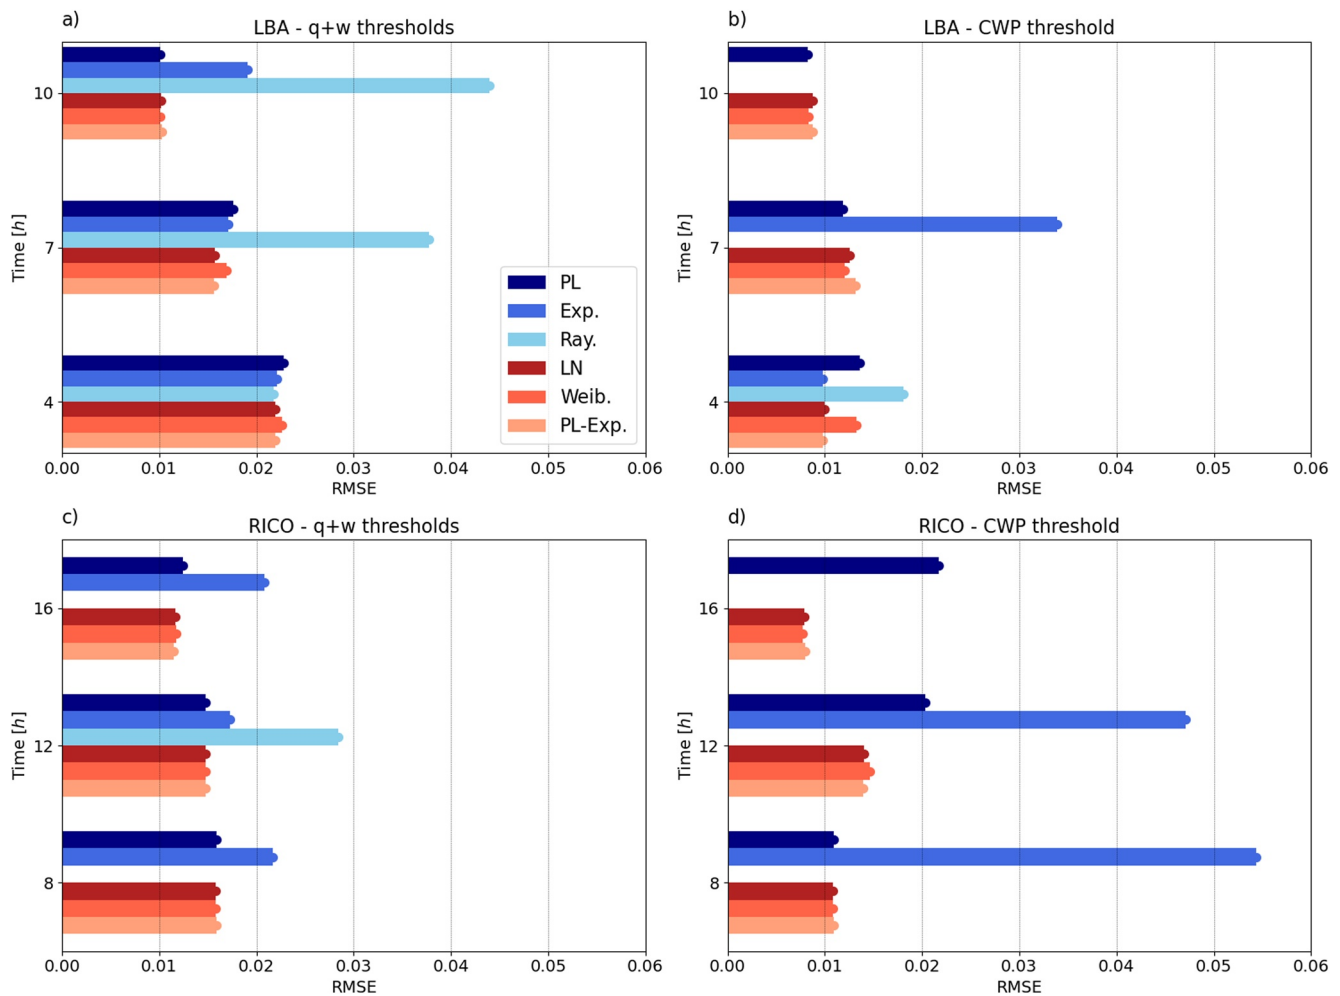
Figure 8. RMSE errors computed between the empirical and fitted distributions for the LBA (top row) and RICO (bottom row) simulations, and clouds identified based on the updraft (left) and CWP (right) criteria. All alternative fits were obtained on the same optimal size range as the power-law fits.

in predicted cloud cover. Since these large clouds constitute the tail of the distributions, it can be concluded that the best model according to S_{cc} should be the one providing the best fit near the tail of the distributions, that is where the exponential and power-law distributions are the least reliable. None of the other alternative distributions tested stands out as a better model, as they all outperform the others under specific conditions. Note that the Rayleigh distribution yields particularly good results in all situations (at least when a fit can be found) since, by definition, the Rayleigh distribution is the most likely distribution (in a maximum entropy sense) describing the known total cloud cover.

In general, Vuong's test constitutes a more robust tool to decide which model should be selected among available alternatives. Z scores are computed as described in Section 2.2 and shown on Figure 10 choosing the power-law as the reference to which all alternative models are compared. Elements depicted in red (blue) indicate that the Z score is larger (smaller) than 1.28 (-1.28). When Z is significantly large (larger than 1.28), the power-law model can be regarded as the best of the two models. In contrast, when Z is significantly small (smaller than -1.28), the alternative constitutes the best model. White squares on Figure 10 indicate Z values lying within the ± 1.28 confidence interval and none of the two models tested can be said to be better than the other.

According to Vuong's test, exponential distributions may constitute viable alternatives to power-laws when clouds are identified based on the updraft criterion. Even though they do not provide significantly better results than power-laws in terms of RMSE or S_{cc} scores, they do not yield considerably worse results either such that the two models can be regarded as being equally good. In contrast, for CWP clouds, exponential fits are now generally
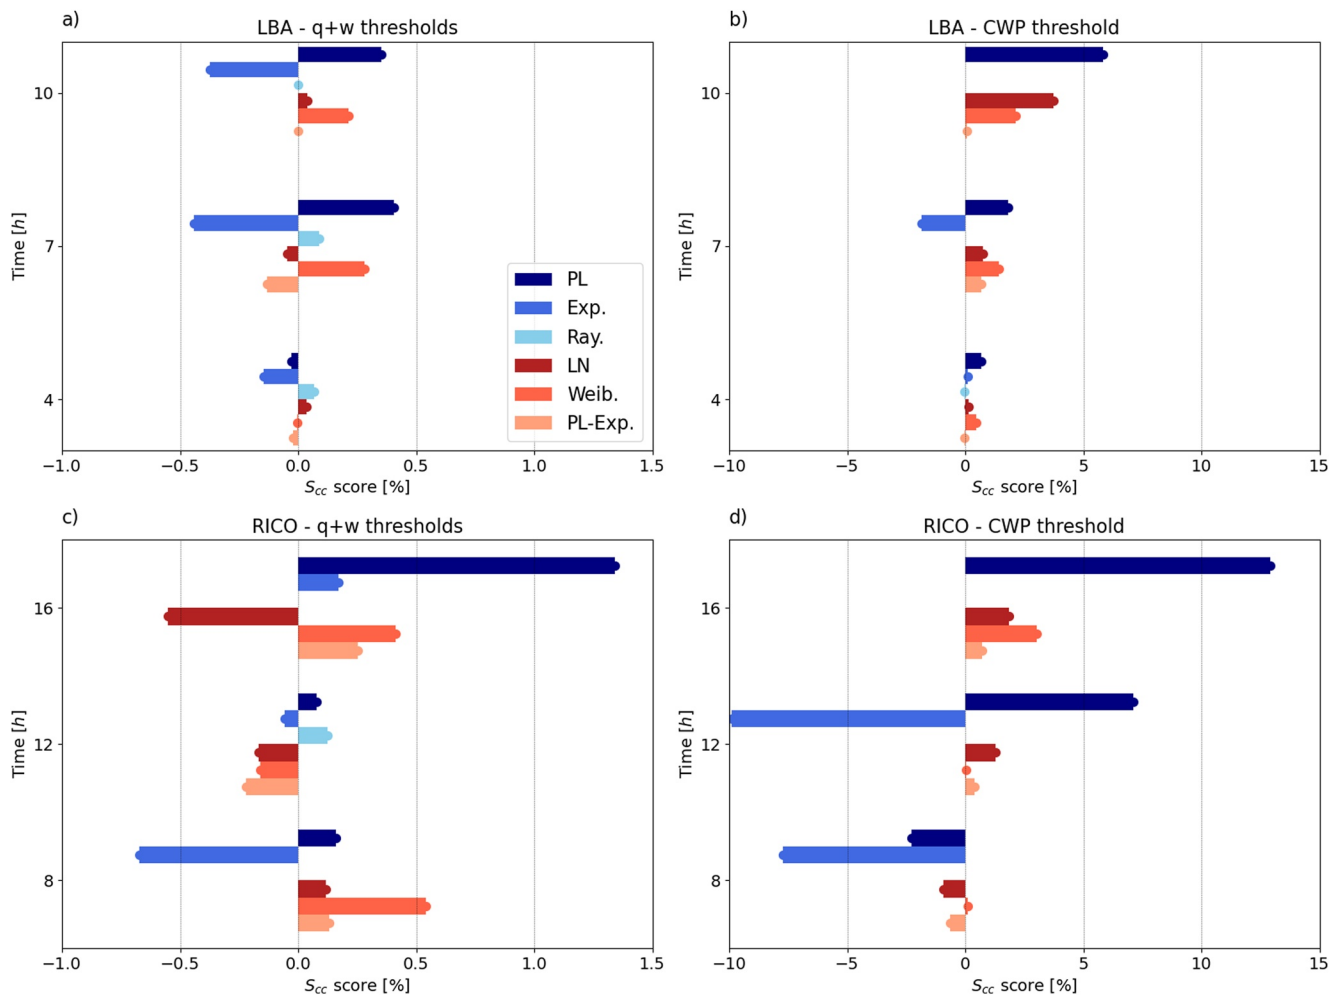
Figure 9. Relative total cloud cover error S_{cc} computed between the empirical and fitted distributions obtained in the LBA (top row) and RICO (bottom row) simulations, and for clouds identified based on the updraft (left) and CWP (right) criteria. All alternative fits were obtained on the same optimal size range as the power-law fits.

less reliable than the power-law ones and power-law distributions clearly constitute a better choice. In the LBA case, the more complex distributions tested (lognormal, Weibull and cutoff power-law) rarely constitute superior alternatives to power-laws, except for clouds identified based on a CWP threshold at 4 hr where the empirical data aren't adequately modeled by a power-law. This result is consistent with the RMSE values shown on Figure 8. In contrast, in the RICO case, power-law fits are often outperformed by the two-parameter distributions, in particular when clouds are identified based on CWP.

5.3. Fitting Without Size Restriction

While the results presented previously are instructive, it may in fact be more useful to seek best fits for the entire cloud size ranges available instead of limiting the analysis to a small subset of the entire cloud population. In the following, we therefore estimate best fits including all clouds larger than $5\Delta x$ available in each scene.

Relative cloud cover errors (S_{cc}) computed in this context are shown on Figure 11. S_{cc} is now in general much larger than when the fits are obtained on subsets of the whole cloud populations (see Figure 9). Power-law fits significantly overestimate the total cloud cover in all cases and consist-

Table 3
In Each Case Reported, 50% of the Total Cloud Cover Is Accomplished by Clouds With Sizes Smaller Than the Indicated Size

Case	Time	$q + w$ thresholds	CWP threshold
LBA	10 hr	84%–714 m	88%–1,934 m
	7 hr	83%–557 m	83%–1,371 m
	4 hr	79%–412 m	81%–762 m
RICO	16 hr	82%–208 m	89%–385 m
	12 hr	80%–201 m	87%–372 m
	8 hr	82%–193 m	88%–360 m

Note. The percentages indicate the fraction of the total cloud population having sizes smaller than the indicated size.
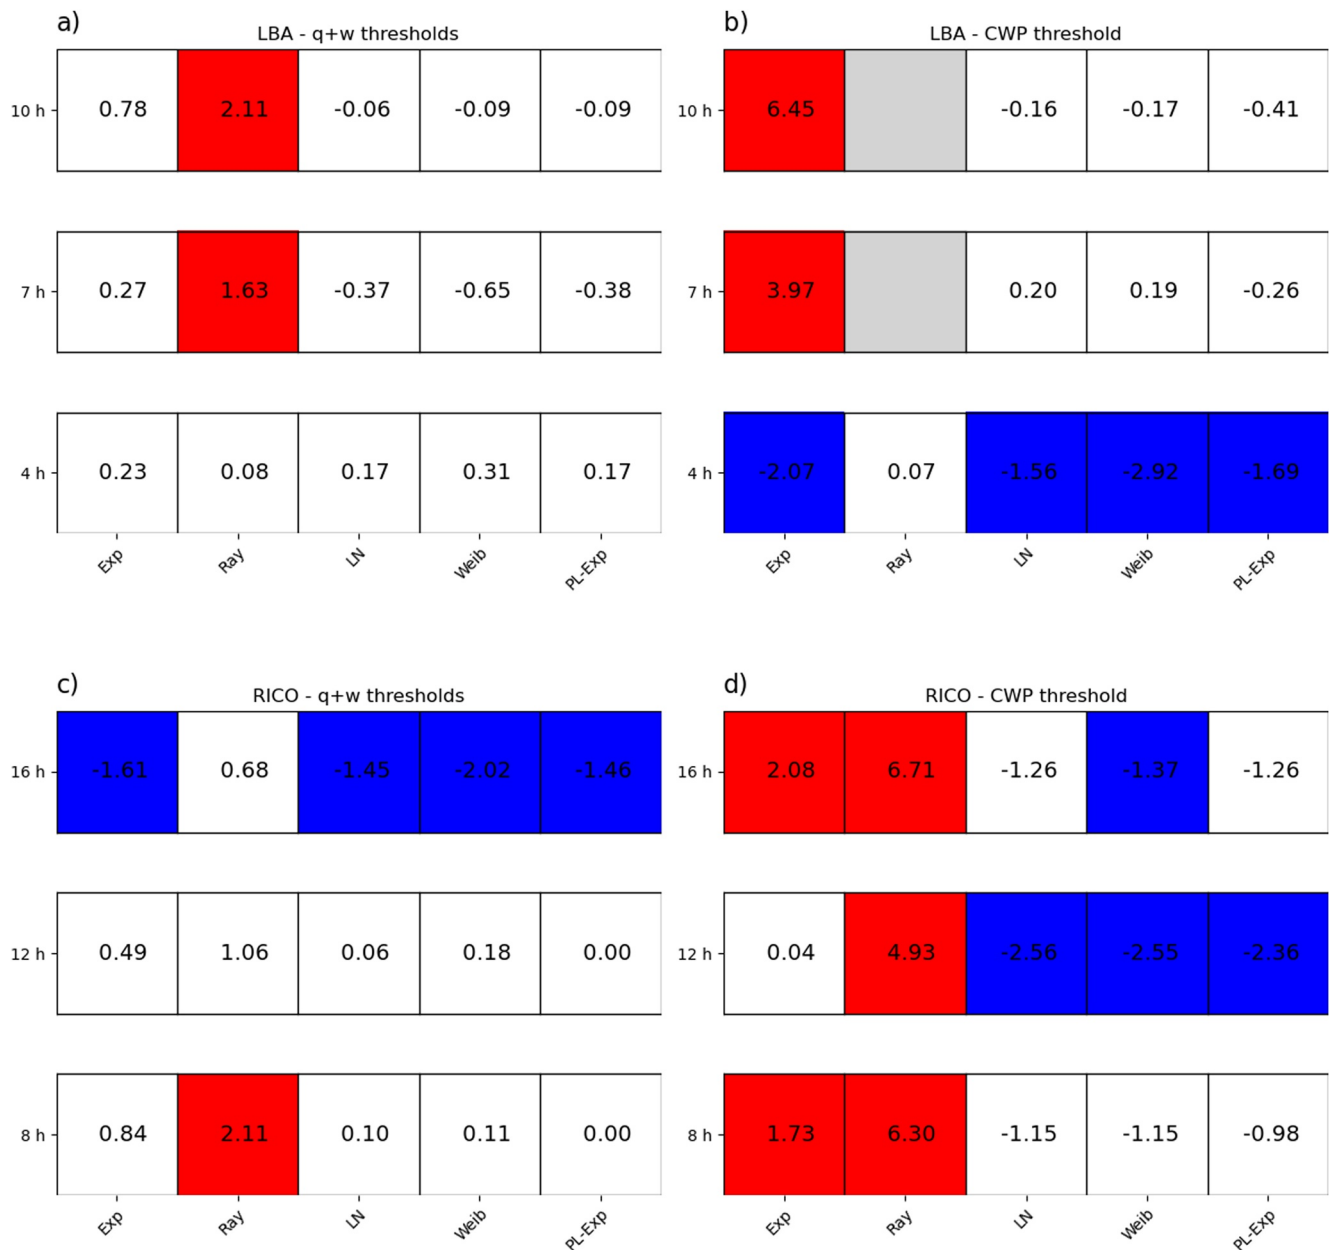
Figure 10. Results from Vuong's test (Vuong, 1989) applied in the LBA (top) and RICO (bottom) simulations for clouds identified based on the updraft (left) and CWP (right) criteria. Vuong's test is applied considering the power-law fits as the reference to which all other alternative models are compared. Colors should be understood as follows: Z values smaller than the -1.28 confidence level are in blue and indicate that the alternative distribution constitutes a better model than the power-law, Z values larger than the 1.28 confidence level are in red and indicate that the power-law distribution is the best model in a given situation. Values between -1.28 and 1.28 indicate that neither the power-law nor the alternative can be judged as better than the other. Gray cells indicate that no acceptable fit could be found with the alternative distribution. Assuming the Z statistics follows a standard normal distribution, 80% of the data is contained within the ± 1.28 confidence interval.

ently yield the worst results of all the distributions tested. For clouds identified based on the updraft criterion, Figure 11 indicates that the empirical cloud size distributions are best modeled by exponential functions than by power-laws. The lognormal and cutoff power-law distributions however tend to minimize S_{cc}, regardless of the case and time considered. For clouds identified based on CWP, power-laws again overestimate the total cloud cover by up to almost 50% despite the fact that clear power-law scalings could be identified on the empirical distributions (in particular in the LBA case). Again, either the lognormal or cutoff power-law distributions yield the smallest cloud cover errors.
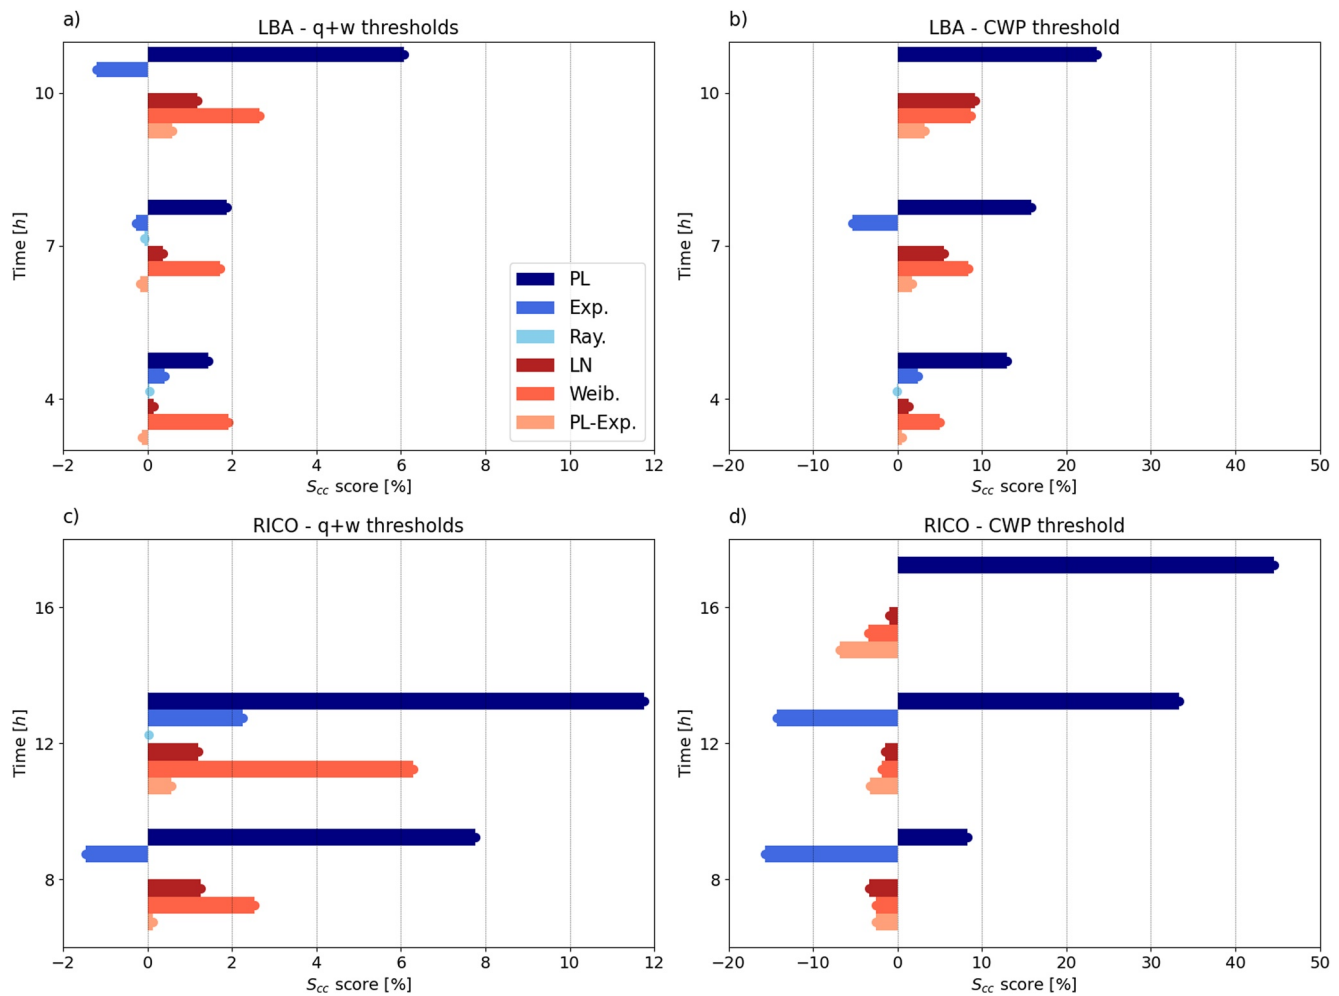
Figure 11. Same as Figure 9 except that best fits are now obtained with each model by considering the entire cloud size range available in each situation (ignoring all clouds smaller than $5\Delta x$).

To understand these results, it must be noted that power-law distributions, even when truncated, consistently overestimate the number of large clouds in the tail of the distributions (recall that the simulated total cloud cover is dominated by the few large clouds). On the opposite, the fact that cutoff power-laws generally yield the best fits overall (they tend to minimize S_{cc} in most cases) suggests that the fast decay in the tail of the distributions can be reasonably well approximated using exponential functions. For this reason, it is recommended not to use uncorrected power-law distributed cloud sizes when trying to model the convective cloud cover or mass flux. Cutoff power-laws and, to some extent, exponential distributions constitute in this context superior options.

These results are generally confirmed by the RMSE score and Vuong's test (not shown) indicating that both the power-law and exponential fits are systematically outperformed by the two-parameter distributions (Weibull, lognormal and cutoff power-law), although none of these more complex distributions stands out as a better model overall. According to these tests, exponential distributions can again be considered as a better model than power-laws for clouds identified based on the updraft criterion, but they do not constitute adequate alternatives for CWP clouds.

Summarizing the results presented in this section, it appears that power-laws rarely constitute the best model to represent the empirical cloud size CCDFs extracted from the LBA and RICO simulations. Considering clouds determined by the updraft criterion, exponential distributions are generally as good as, if not better than power-laws, and should be preferred when trying to represent the total cloud cover (or total mass flux) in the lower troposphere. In this situation, the two-parameter distributions also yield very good results, but

the improvements made over the simpler power-law and exponential distributions are rarely substantial. In contrast, for CWP clouds, alternative two-parameter distributions provide significantly better fits than power-laws in many situations, and should especially be preferred when one is interested in modeling the total cloud cover. In this context, power-laws with an exponential cutoff seem to provide the best alternative to pure power-laws.

6. The Emergence of Power-Law Scalings in Cloud Size Distributions

We wish here to give tentative arguments to explain the emergence of power-law scalings (generally exhibiting an exponential decay in the tail) in cloud size distributions by drawing an analogy between convective cloud ensembles and percolating systems approaching criticality (Stauffer & Ahaorny, 2003). Note that percolation theory bears of lot of similarities with the theory of self-organized criticality (SOC) already invoked to explain behaviors observed in convective cloud ensembles (Peters & Neelin, 2006; Peters et al., 2009, 2010; Teo et al., 2017; Windmiller, 2017; Yano et al., 2012).

Approximating each cloud as a circle of random radius appearing quasi-randomly in space, the cloud ensemble can be mathematically described by the Boolean model which considers the set of overlapping discs whose centers are determined by a spatial Poisson point process (Meester & Roy, 1996). When discs overlap, they form clusters, and exact as well as asymptotic results can be obtained for cluster size distributions. In particular, it can be demonstrated that the bulk of the cluster size distribution exhibits a scale free behavior over a range that increases when the system is brought closer to criticality, while the distributions tail in the subcritical phase decays exponentially (Meester & Roy, 1996), two properties shared with some of the cloud size distributions analyzed in this work. This analogy is further supported by the close similarity between our best fit α exponent found in the LBA case for CWP clouds (2.03) and the theoretical cluster size distribution exponent for two-dimensional continuum percolation of $181/97 \approx 2.05$ (Gawlinski & Stanley, 1981). Based on this remark, we argue here that the emergence of power-law cloud size distributions results from the progressive merging and overlapping of cloud cores, which may themselves exhibit exponential size distributions, as the system is drawn toward a critical state.

To further support this discussion, we first take a closer look at the temporal evolution of fitted cloud size distribution statistics obtained in the LBA case (cloud size CCDFs from the RICO simulation did not reveal significant changes in time). Figure 12 displays the RMSE score defined by Equation 21 for both exponential and power-law fits as a function of time. All fits were obtained considering all clouds larger than $5\Delta x$. For CWP clouds, the exponential fit produces lower RMSE scores than power-law fits early in the simulation, until 5 hr. A clear transition is then seen starting at 5.5 hr with power-law RMSEs continuously decreasing to 0, while exponential RMSEs increase dramatically. This transition from nearly exponential to mostly power-law cloud size distributions
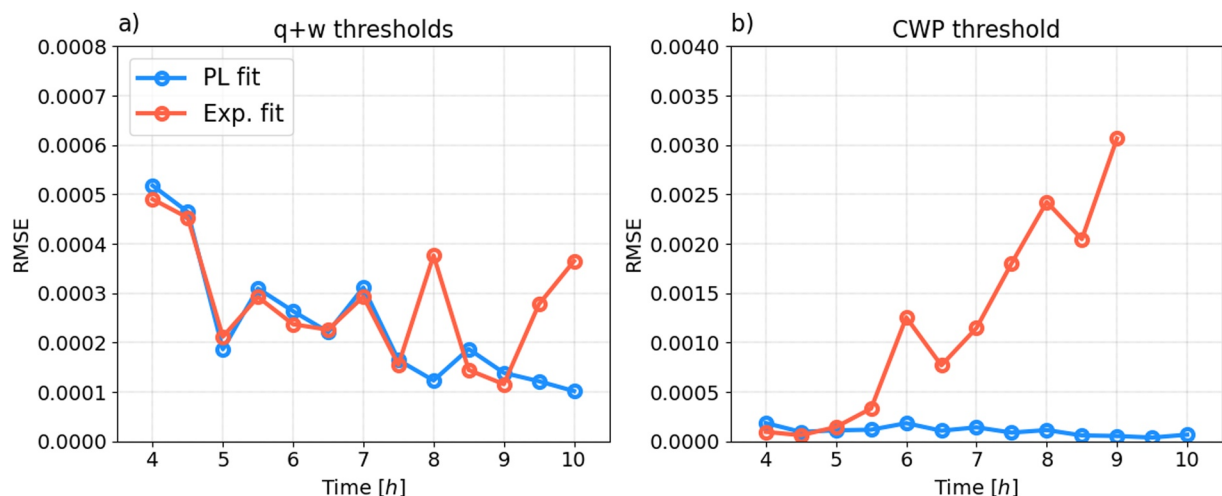

Figure 12. Time evolution of RMSE values computed according to Equation 21 for power-law (blue) and exponential (red) best fits obtained in the LBA case using either the updraft (left) or the CWP (right) criterion. The fits were performed on the entire cloud size range available in each case.
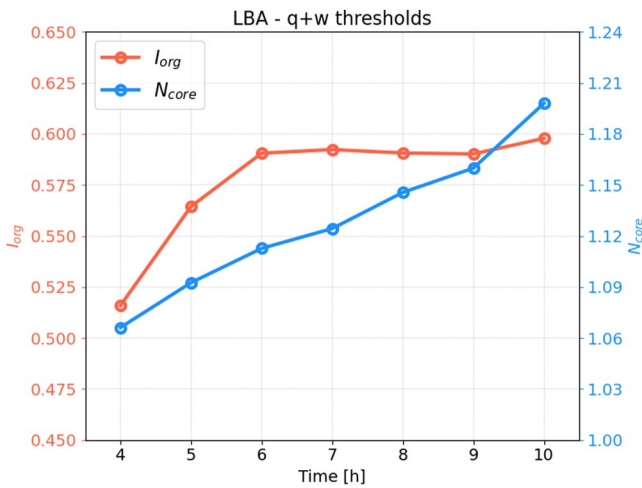
Figure 13. I_{org} index and mean core number per cloud plotted as functions of time in the LBA case. Clouds are here only identified based on the updraft criterion.

is accompanied by a constant increase in the size of the largest clouds found in the domain. For clouds identified based on the updraft criterion, no such transition between exponential to power-law cloud size distributions is found.

Cloud merging can be further analyzed by splitting each individual cloud into distinct convective cores. This is done by first identifying core centers as local vertical velocity maxima with $w > 3$ m s^{-1}, provided that each maximum is separated by a distance of at least $5\Delta x$ from its nearest neighbor. Each point belonging to a given cloud is then associated to the nearest core center belonging to the same cloud (Savre & Craig, 2023). Figure 13 displays the temporal evolution of the mean number of cores per cloud thus identified, N_{core}. Clearly, N_{core} increases steadily with time which confirms our intuition that as time progresses, clouds become more likely to merge and form larger clouds. Interestingly, the increased cloud merging frequency doesn't necessarily occur concomitantly with an increase in cloud clustering, quantified for example, by the I_{org} index proposed by Tompkins and Semie (2017). This index relies on the comparison of the nearest neighbor distances calculated for the empirical cloud population and a spatially random Poisson process with equivalent intensity. As seen from Figure 13, I_{org} in the LBA case increases from a value close to 0.5 at 4 hr (indicating spatial randomness), to ~0.59 at 6 hr (indicating spatial clustering). I_{org} then remains approximately constant until the end of the simulation while cloud merging continues to intensity. Note that, in the RICO simulation, I_{org} remains close to 0.6 throughout the analysis period, but the mean number of cores per cloud remains always close to 1 (not shown). Cloud merging therefore rarely happens in this case despite a high degree of organization, which again supports our hypothesis.

Besides, core size CCDFs can also be computed and corresponding exponential and power-law best fits can be obtained, as done in Section 4 (the corresponding best fit parameters are reported in Table 4). In general, core sizes preferentially follow exponential distributions in any situation, even when cloud sizes are power-law distributed, as for example, when clouds are identified based on a CWP threshold, after 10 hr. This visual impression is confirmed by all goodness-of-fit metrics introduced in the previous sections (not shown).

Overall, these results tend to support the idea that cloud core merging constitutes the main mechanism leading to the emergence of cloud size distributions exhibiting scale-free behaviors. While clustering clearly plays a role in increasing the probability that clouds will merge by bringing them closer to each other, our results suggest that it is not the sole mechanism at play. It may indeed be possible that the progressive moistening of the troposphere throughout the convective layer contributes to the increased merging frequency and to the formation of wider clouds exhibiting power-law size distributions via a moisture-entrainment feedback as described for example, in Derbyshire et al. (2004); Del Genio and Wu (2010).

7. Conclusion

Cumulus cloud sizes have been shown to be power-law distributed for more than 30 years. These power-laws can be characterized by their exponent whose values reported in the literature generally vary over a relatively broad range, from 1 to over 3. Although we do not ignore the fact that variations in the shape of cloud size distributions may stem from underlying physical processes (organization), external forcings or environmental conditions, it is here surmised that at least part of this variability can be explained by the use of improper fitting procedures.

In order to unambiguously identify and characterize power-law behaviors in cloud size distributions, a fitting algorithm was introduced, inspired by the works of Clauset et al. (2009), and applied to two distinct simulated convective situations. The method relies on robust statistical methods including Maximum Likelihood Estimation and goodness-of-fit tests, and allows the simultaneous determination of the characteristic power-law exponents, as

Table 4

Same as Table 2 but for Best Fits Obtained Considering Cloud Core Size CCDFs Computed for the LBA Case (Fits Are Shown on Figure 14)

Time	$\hat{\alpha}$	$\hat{\lambda}$
10 hr	2.70	3.24×10^{-3}
7 hr	3.12	4.23×10^{-3}
4 hr	3.99	6.28×10^{-3}
10 hr	2.07	1.49×10^{-3}
7 hr	2.27	2.18×10^{-3}
4 hr	2.79	3.84×10^{-3}

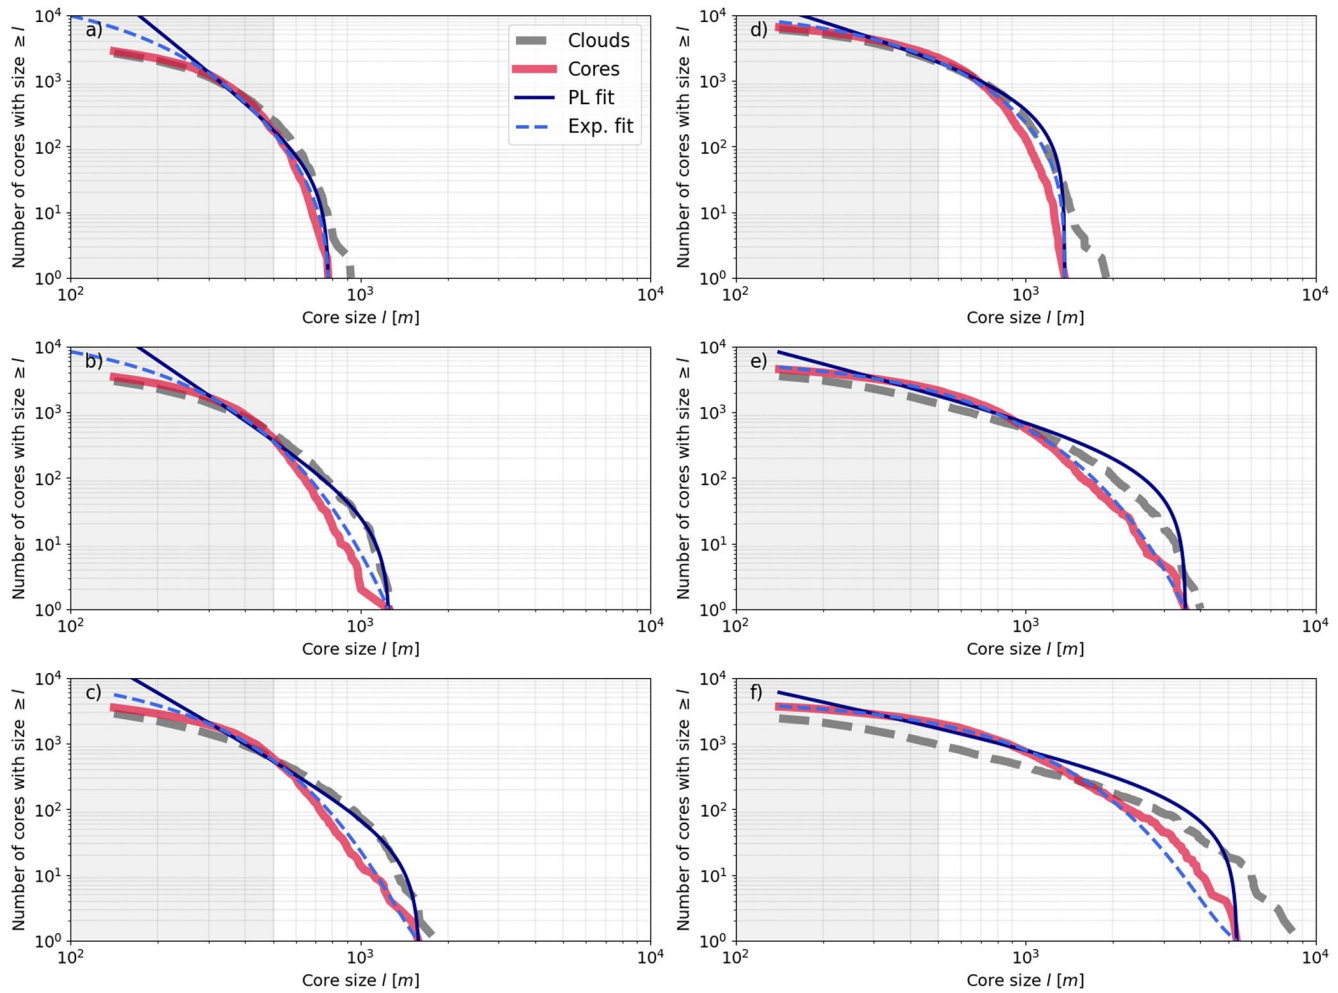*Note.* Best fits were here obtained considering power-law and exponential distributions only.

Figure 14. Same as Figure 2 but showing empirical CCDFs and corresponding best fits obtained for cloud core sizes (see text for details regarding the core identification procedure). Results are shown for the LBA case, and best fits were obtained considering power-law and exponential distributions applied to the entire core size ranges.

well as the optimal size ranges over which the fits are valid. The objective of this work was then three-fold: (a) reexamine the general hypothesis that cloud sizes are power-law distributed; (b) identify potential alternatives to power-law distributions; (c) demonstrate that important information regarding the processes shaping the simulated cloud ensembles may be learned from the systematic use of the fitting algorithm.

Cloud size distributions extracted from both continental deep convection and maritime shallow convection simulations were shown to follow power-laws as long as clouds were identified based on a Condensed Water Path (CWP) threshold (corresponding to clouds being observed and identified from above). The two cases analyzed in this study yielded power-law exponents of 2.03 and 2.19, in agreement with most estimates reported in the literature. Although the use of linear regression in log-log coordinates yielded comparable values, it should be noted that these latter were associated with uncertainties at least twice as large as those obtained with our algorithm. Cloud size distributions extracted near the base of the cloud layer using thresholds in condensed water content and vertical velocity were not found to exhibit clear power-law behaviors: although the algorithm was able to find power-law fits in these situations, their quality was generally too poor to affirm the existence of power-law scalings. It must be stressed that errors related to, in particular, the small cloud size samples available are also thought to hamper the reliable identification of power-law fits in this case.

Considering updraft clouds identified in the lower troposphere, exponential distributions were found to yield fits at least as good as power-laws. Because doubts remain concerning the validity of the power-law scalings identified in these situations by the algorithm, the exponential distribution can here be considered as the better model.

In addition, although better fits could be obtained with two-parameter alternative distributions like lognormal or Weibull distributions, the improvements made over power-law and exponential fits were not substantial enough to justify the increased complexity associated with these models. Considering CWP clouds for which power-law scalings are more evident, if one is interested in parameterizing the total cloud cover characterizing the cloud ensemble, correctly representing the fast decay often seen in the tail of the empirical size distributions was found to be crucial. In this context, power-laws with an exponential cutoff were found to yield superior fits.

Whereas the results summarized above apply generally to most cloud scenes extracted from the two simulations, the shape of the cloud size distributions computed in the continental convection case was shown to evolve with time despite the absence of a clear diurnal cycle. While cloud sizes were indeed found to be exponentially distributed at earlier times, a transition to power-law scalings (with an exponential decay) was observed as deeper clouds started to develop. It was argued that this transition is likely driven by (or happens simultaneously with) changes in the organization of the cloud and subcloud layers. Indeed, as deeper clouds grow, cold pools also become broader thus constraining the future generation of convective clouds to develop in more restricted areas. The result is that clouds are generally brought closer to each other, they are also more likely to merge and power-law scalings emerge as a consequence of the increased fraction of merging clouds (in analogy with percolating systems in a continuum). Exponential tails then generally persist as a manifestation of finite-size effects. This hypothesis could be verified by demonstrating that the size distributions of individual cloud cores (obtained by splitting cloud objects) generally follow exponentials, even when cloud sizes are clearly power-law distributed.

Finally, we would like to again emphasize that not enough is known about the processes (organization of the cloud population, environmental conditions) controlling the shape and characteristic exponents of observed or simulated cloud size distributions. Whereas more studies are encouraged in this direction, such a knowledge can only be gained by minimizing the errors made otherwise through improper fitting. To correctly interpret changes in power-law cloud size distributions it is indeed absolutely necessary to be able to unambiguously identify power-law behaviors (or reject the power-law hypothesis when it's not appropriate), and reduce as much as possible errors and inaccuracies made when estimating best fit power-law exponents. In this context, we strongly recommend the use of reliable fitting algorithms like the one introduced in this study.

Data Availability Statement

The fitting algorithm introduced in this work has been implemented as a python module (*pyCloudFit*) and was made available, upon request, at the following address: https://bitbucket.org/julien_savre/pycloudfit/src/master/. All data (cloud size distributions) produced and analyzed during this work is openly available at https://doi.org/10.5281/zenodo.7005140 (Sacre, 2022). Any updates will also be posted on Zenodo.

References

Barron, N. R., Ryan, S. D., & Heus, T. (2020). Reconciling chord length distributions and area distributions for fields of fractal cumulus clouds. *Atmosphere*, *11*(8), 364–374. https://doi.org/10.3390/atmos11080824

Benner, T., & Curry, J. (1998). Characteristics of small tropical cumulus clouds and their impact on the environment. *Journal of Geophysical Research*, *27*(D22), 28753–28767. https://doi.org/10.1029/98jd02579

Bley, S., Deneke, H., Senf, F., & Scheck, L. (2017). Metrics for the evaluation of warm convective cloud fields in a large-eddy simulation with meteosat images. *Quarterly Journal of the Royal Meteorological Society*, *143*(705), 2050–2060. https://doi.org/10.1002/qj.3067

Böing, S. J., Jonker, H. J. J., Siebesma, A. P., & Grabowski, W. W. (2012). Influence of the subcloud layer on the development of a deep convective ensemble. *Journal of the Atmospheric Sciences*, *69*(9), 2682–2698. https://doi.org/10.1175/JAS-D-11-0317.1

Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, *51*(4), 661–703. https://doi.org/10.1175/1520-0469(1998)055⟨3440:EADINS⟩2.0.CO;2

Craig, G. C., & Cohen, B. G. (2006). Fluctuations in an equilibrium convective ensemble. Part I: Theoretical formulation. *Journal of the Atmospheric Sciences*, *63*(8), 1996–2004. https://doi.org/10.1175/JAS3709.1

Dawe, J. T., & Austin, P. H. (2012). Statistical analysis of an LES shallow cumulus cloud ensemble using a cloud tracking algorithm. *Atmospheric Chemistry and Physics*, *12*(2), 1101–1119. https://doi.org/10.5194/acp-12-1101-2012

Del Genio, A. D., & Wu, J. (2010). The role of entrainment in the diurnal cycle of continental convection. *Journal of Climate*, *23*(10), 2722–2738. https://doi.org/10.1175/2009JCLI3340.1

Deluca, A., & Corral, A. (2013). Fitting and goodness-of-fit of non-truncated and truncated power-law distributions. *Acta Geophysica*, *61*(6), 1351–1394. https://doi.org/10.2478/s11600-013-0154-9

Derbyshire, S. H., Beau, I., Bechtold, P., Grandpeix, J.-Y., Piriou, J.-M., Redelsperger, J. L., & Soares, P. M. M. (2004). Sensitivity of moist convection to environmental humidity. *The Quarterly Journal of the Royal Meteorological Society*, *130*(604), 3055–3079. https://doi.org/10.1256/qj.03.130

Garrett, T., Glenn, I., & Krueger, S. (2018). Thermodynamic constraints on the size distributions of tropical clouds. *Journal of Geophysical Research: Atmospheres*, *123*(16), 8832–8849. https://doi.org/10.1029/2018JD028803

Gawlinski, E. T., & Stanley, H. E. (1981). Continuum percolation in two-dimensions: Monte Carlo tests of scaling and universality for non-interacting discs. *Journal of Physics A: Mathematical and General*, *14*(8), L291–L299. https://doi.org/10.1088/0305-4470/14/8/007

Goldstein, M., Morris, S., & Yen, G. (2004). Problems with fitting to the power-law distribution. *The European Physical Journal B*, *41*(2), 255–258. https://doi.org/10.1140/epjb/e2004-00316-5

Grabowski, W. W. (1998). Toward cloud resolving modeling of large-scale tropical circulations: A simple cloud microphysics parameterization. *Journal of the Atmospheric Sciences*, *55*(21), 3283–3298. https://doi.org/10.1175/1520-0469(1998)055⟨3283:TCRMOL⟩2.0.CO;2

Grabowski, W. W., Bechtold, P., Cheng, A., Forbes, R., Halliwell, C., Khairoutdinov, M., et al. (2006). Daytime convective development over land: A model intercomparison based on LBA observations. *The Quarterly Journal of the Royal Meteorological Society*, *132*(615), 317–344. https://doi.org/10.1256/qj.04.147

Heus, T., & Seifert, A. (2013). Automated tracking of shallow cumulus clouds in large domain, long duration large-eddy simulations. *Geoscientific Model Development*, *6*(4), 1261–1273. https://doi.org/10.5194/gmd-6-1261-2013

Houze, R., Jr., & Cheng, C.-P. (1977). Radar characteristics of tropical convection observed during gate: Mean properties and trends over the summer season. *Monthly Weather Review*, *105*(8), 964–980. https://doi.org/10.1175/1520-0493(1977)105<0964:rcotco>2.0.co;2

Jiang, H., Feingold, G., Jonsson, H., Lu, M.-L., Chuang, P., Flagan, R., & Seinfeld, J. (2008). Statistical comparison of properties of simulated and observed cumulus clouds in the vicinity of Houston during the Gulf of Mexico atmospheric composition and climate study (GoMACCS). *Journal of Geophysical Research*, *113*(D13), D13205. https://doi.org/10.1029/2007JD009304

Koren, I., Oreopoulos, L., Feingold, G., Remer, L., & Altaratz, O. (2008). How small is a small cloud? *Atmospheric Chemistry and Physics*, *8*(14), 3855–3864. https://doi.org/10.5194/acp-8-3855-2008

Kuo, K.-S., Welch, R., & Sengupta, S. (1988). Structural and textural characteristics of cirrus clouds observed using high spatial resolution landsat imagery. *Journal of Applied Meteorology*, *27*(11), 1242–1260. https://doi.org/10.1175/1520-0450(1988)027<1242:satcoc>2.0.co;2

Kuo, K.-S., Welch, R., Weger, R., Engelstad, M., & Sengupta, S. (1993). The three-dimensional structure of cumulus clouds over the Ocean: 1. Structural analysis. *Journal of Geophysical Research*, *11*(D11), 20685–20711. https://doi.org/10.1029/93jd02331

Laherrere, J., & Sornette, D. (1998). Stretched exponential distributions in nature and economy: "fat tails" with characteristic scales. *The European Physical Journal B*, *2*(4), 525–539. https://doi.org/10.1007/s100510050276

Lopez, R. (1977). The lognormal distribution and cumulus cloud populations. *Monthly Weather Review*, *105*(7), 865–872. https://doi.org/10.1175/1520-0493(1977)105<0865:tldacc>2.0.co;2

Meester, R., & Roy, R. (1996). *Continuum percolation* (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511895357

Mieslinger, T., Horvath, A., Buehler, S. A., & Sakradzija, M. (2019). The dependence of shallow cumulus macrophysical properties on large-scale meteorology as observed in aster imagery. *Journal of Geophysical Research: Atmospheres*, *124*(21), 11477–11505. https://doi.org/10.1029/2019JD030768

Mitzenmacher, M. (2003). A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, *1*(2), 226–251. https://doi.org/10.1080/15427951.2004.10129088

Neggers, R. A. J., Duynkerke, P., & Rodts, S. (2003). Shallow cumulus convection: A validation of large-eddy simulation against aircraft and Landsat observations. *Quarterly Journal of the Royal Meteorological Society*, *129*(593), 2671–2696. https://doi.org/10.1256/qj.02.93

Neggers, R. A. J., Griewank, P. J., & Heus, T. (2019). Power-law scaling in the internal variability of cumulus cloud size distributions due to subsampling and spatial organization. *Journal of the Atmospheric Sciences*, *76*(6), 1489–1503. https://doi.org/10.1175/JAS-D-18-0194.1

Neggers, R. A. J., Jonker, H. J. J., & Siebesma, A. (2003). Size statistics of cumulus cloud populations in large-eddy simulations. *Journal of the Atmospheric Sciences*, *60*(8), 1060–1074. https://doi.org/10.1175/1520-0469(2003)60

Newman, M. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, *46*(5), 323–351. https://doi.org/10.1080/00107510500052444

Parker, L., Welch, R., & Musil, D. (1986). Analysis of spatial inhomogeneities in cumulus clouds using high spatial resolution Landsat data. *Journal of Applied Meteorology and Climatology*, *25*(10), 1301–1314. https://doi.org/10.1175/1520-0450(1986)025<1301:aosiic>2.0.co;2

Peters, O., Deluca, A., Corral, A., Neelin, J. D., & Holloway, C. E. (2010). Universality of rain event size distributions. *Journal of Statistical Mechanics: Theory and Experiment*, *11030*(11), P11030. https://doi.org/10.1088/1742-5468/2010/11/P11030

Peters, O., & Neelin, J. (2006). Critical phenomena in atmospheric precipitation. *Nature Physics*, *2*(6), 393–396. https://doi.org/10.1038/nphys314

Peters, O., Neelin, J. D., & Nesbitt, S. (2009). Mesoscale convective systems and critical clusters. *Journal of the Atmospheric Sciences*, *66*(9), 2913–2924. https://doi.org/10.1175/2008JAS2761.1

Plank, V. G. (1969). The size distribution of cumulus clouds in representative Florida populations. *Journal of Applied Meteorology and Climatology*, *8*(1), 46–67. https://doi.org/10.1175/1520-0450(1969)008⟨0046:TSDOCC⟩2.0.CO;2

Radtke, J., Mauritsen, T., & Hohenegger, C. (2021). Shallow cumulus cloud feedback in large eddy simulations? Bridging the gap to storm-resolving models. *Atmospheric Chemistry and Physics*, *21*(5), 3275–3288. https://doi.org/10.5194/acp-21-3275-2021

Rieck, M., Hohenegger, C., & van Heerwaarden, C. (2014). The influence of land surface heterogeneities on cloud size development. *Monthly Weather Review*, *142*(10), 3830–3846. https://doi.org/10.1175/mwr-d-13-00354.1

Rodts, S. M. A., Duynkerke, P. G., & Jonker, H. J. J. (2003). Size distributions and dynamical properties of shallow cumulus clouds from aircraft observations and satellite data. *Journal of the Atmospheric Sciences*, *60*(16), 1895–1912. https://doi.org/10.1175/1520-0469(2003)060⟨1895:SDADPO⟩2.0.CO;2

Sakradzija, M., & Hohenegger, C. (2017). What determines the distribution of shallow convective mass flux through a cloud base? *Journal of Atmospheric and Solar-Terrestrial Physics*, *74*(8), 2615–2632. https://doi.org/10.1175/jas-d-16-0326.1

Sakradzija, M., Seifert, A., & Heus, T. (2015). Fluctuations in a quasi-stationary shallow cumulus cloud ensemble. *Nonlinear Processes in Geophysics*, *22*(1), 65–85. https://doi.org/10.5194/npg-22-65-2015

Savre, J. (2021). Formation and maintenance of subsiding shells around non-precipitating and precipitating cumulus clouds. *The Quarterly Journal of the Royal Meteorological Society*, *147*(735), 728–745. https://doi.org/10.1002/qj.3942

Savre, J. (2022). Data presented in "Fitting cumulus cloud size distributions from idealized cloud resolving model simulations" [Dataset]. Zenodo. https://doi.org/10.5281/zenodo.7005140

Savre, J., & Craig, G. (2023). The Sensitivity of Convective Cloud Ensemble Statistics to Horizontal Grid Spacing in Idealized RCE Simulations. *Journal of the Atmospheric Sciences*, *80*(5), 1267–1284. https://doi.org/10.1175/JAS-D-22-0164.1

Savre, J., Ekman, A. M. L., & Svensson, G. (2014). Technical note: Introduction to MIMICA, a large-eddy simulation solver for cloudy planetary boundary layers. *Journal of Advances in Modeling Earth Systems*, *6*(3), 630–649. https://doi.org/10.1002/2013MS000292

Seifert, A., & Beheng, K. D. (2006). A two-moment cloud microphysics parameterization for mixed-phase clouds. Part I: Model descroption. *Meteorology and Atmospheric Physics*, *92*(1–2), 45–66. https://doi.org/10.1007/s00703-005-0112-4

Senf, F., Klocke, D., & Brueck, M. (2018). Size-resolved evaluation of simulated deep tropical convection. *Monthly Weather Review*, *146*(7), 2161–2182. https://doi.org/10.1175/mwr-d-17-0378.1

Sengupta, S., Welch, R., Navar, M., Berendes, T., & Chen, D. (1990). Cumulus cloud field morphology and spatial patterns derived from high spatial resolution Landsat imagery. *Journal of Applied Meteorology*, *29*(12), 1245–1267. https://doi.org/10.1175/1520-0450(1990)029<1245:ccfmas>2.0.co;2

Sornette, D. (2006). *Critical phenomena in natural sciences* (2nd ed.). Springer.

Stauffer, D., & Ahaorny, A. (2003). *Introduction to percolation theory* (2nd ed.). Taylor & Francis.

Stumpf, M. P. H., & Porter, M. A. (2012). Critical truths about power laws. *Science*, *335*(6069), 665–666. https://doi.org/10.1126/science.1216142

Stumpf, M. P. H., Wiuf, C., & May, R. M. (2005). Subnets of scale-free networks are not scale-free: Sampling properties of networks. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(12), 4221–4224. https://doi.org/10.1073/pnas.0501179102

Teo, C.-K., Huynh, H.-N., Koh, T.-Y., Cheung, K., Legras, B., Chew, L., & Norford, L. (2017). The universal scaling characteristics of tropical oceanic rain clusters. *Journal of Geophysical Research*, *122*(11), 5582–5599. https://doi.org/10.1002/2016jd025921

Tompkins, A., & Semie, A. (2017). Organization of tropical convection in low vertical wind shears: Role of updraft entrainment. *Journal of Advances in Modeling Earth Systems*, *9*(2), 1046–1068. https://doi.org/10.1002/2016MS000802

Traxl, D., Boers, N., Rheinwalt, A., Goswami, B., & Kurths, J. (2016). The size distribution of spatiotemporal extreme rainfall clusters around the globe. *Geophysical Research Letters*, *43*(18), 9939–9947. https://doi.org/10.1002/2016gl070692

van Laar, T. W., Schemann, V., & Neggers, R. A. J. (2019). Investigating the diurnal evolution of the cloud size distribution of continental cumulus convection using multiday les. *Journal of the Atmospheric Sciences*, *767*(3), 729–747. https://doi.org/10.1175/JAS-D-18-0084.1

vanZanten, M. C., Stevens, B., Nuijens, L., Siebesma, A. P., Ackerman, A. S., Burnet, F., et al. (2011). Controls on precipitation and cloudiness in simulations of trade-wind cumulus as observed during rico. *Journal of Advances in Modeling Earth Systems*, *3*, 1–19. https://doi.org/10.1029/2011MS000056

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, *57*(2), 307–333. https://doi.org/10.2307/1912557

Welch, R., Kuo, K., Wielicki, B., Sengupta, S., & Parker, L. (1988). Marine stratocumulus cloud fields off the coast of southern California observed using Landsat imagery. Part I: Structural characteristics. *Journal of Applied Meteorology*, *27*(4), 341–362. https://doi.org/10.1175/1520-0450(1988)027<0341:mscfot>2.0.co;2

White, E., Enquist, B., & Green, J. (2008). On estimating the exponent of power-law frequency distributions. *Ecology*, *89*(4), 905–912. https://doi.org/10.1890/07-1288.1

Wielicki, B., & Welch, R. (1986). Cumulus cloud properties derived using Landsat satellite data. *Journal of Climate and Applied Meteorology*, *25*(3), 261–276. https://doi.org/10.1175/1520-0450(1986)025<0261:ccpdul>2.0.co;2

Windmiller, J. (2017). *Organization of tropical convection (Unpublished doctoral dissertation)*. Ludwig-Maximilians-Universität München.

Wing, A. A., Emanuel, K., Holloway, C. E., & Muller, C. (2018). Convective self-aggregation in numerical simulations: A review. In R. Pincus, D. Winker, S. Bony, & B. Stevens (Eds.), *Shallow clouds, water vapor, circulation, and climate sensitivity* (pp. 1–25). Springer International Publishing. https://doi.org/10.1007/978-3-319-77273-8_1

Wood, R., & Field, P. (2011). The distribution of cloud horizontal sizes. *Journal of Climate*, *24*(18), 4800–4816. https://doi.org/10.1175/2011jcli4056.1

Xue, H., & Feingold, G. (2006). Large-eddy simulations of trade wind cumuli: Investigation of aerosol indirect effects. *Journal of the Atmospheric Sciences*, *63*(6), 1605–1622. https://doi.org/10.1175/jas3706.1

Yano, J.-I., Liu, C., & Moncrieff, M. (2012). Self-organized criticality and homeostasis in atmospheric convective organization. *Journal of the Atmospheric Sciences*, *69*(12), 3449–3462. https://doi.org/10.1175/JAS-D-12-069.1

Zhao, G., & Di Girolamo, L. (2007). Statistics on the macrophysical properties of trade wind cumuli over the tropical wester Atlantic. *Journal of Geophysical Research*, *112*, D10204.