



Precipitable water and the lognormal distribution

James Foster,¹ Michael Bevis,² and William Raymond^{3,4}

Received 3 October 2005; revised 1 March 2006; accepted 7 April 2006; published 4 August 2006.

[1] Histograms of precipitable water from radiosondes and zenith neutral delay estimated by GPS are found at many locations to follow a lognormal distribution. This observation is consistent with a theoretical expression for precipitable water based on moisture flux. Two unimodal cases are identified: the traditional lognormal distribution is commonly found in subtropical and temperate climates while tropical oceanic environments tend to exhibit a reversed lognormal form. Bimodal cases, formed by combinations of either or both of these unimodal distributions, are found where seasonal (e.g., monsoons) or climatic (e.g., El Niño) variations generate distinct precipitable water modes with rapid transitions between them. This connection with the lognormal distribution suggests a basis for the parameterization both of precipitable water in climate models as well as of the delay due to water vapor in atmospheric models used for space geodesy.

Citation: Foster, J., M. Bevis, and W. Raymond (2006), Precipitable water and the lognormal distribution, *J. Geophys. Res.*, *111*, D15102, doi:10.1029/2005JD006731.

1. Introduction

[2] The total column quantity of water vapor overlying a specific point on or above the Earth's surface, expressed as the height of an equivalent column of liquid water, is known as precipitable water (PW). In the last decade the Global Positioning System (GPS) has been used to measure PW with an accuracy comparable to that achieved using radiosondes [e.g., Duan *et al.*, 1996; Tregoning *et al.*, 1998; Gutman and Benjamin, 2001]. The recent appearance of multiyear time series of PW measurements derived from continuous GPS networks has prompted a resurgence of interest in the statistics of PW variability. One of the most basic properties of an environmental variable is its statistical distribution. Foster and Bevis [2003] studied the statistical distribution of PW in Hawaii and found it to be lognormal, or very nearly so, when variability is considered over a period of several years. They analyzed PW time series derived from a network of Global Positioning System (GPS) receivers based on the Big Island of Hawaii and from a suite of radiosonde profiles associated with the Lihue radiosonde station in Kauai.

[3] The geodetic measurement of PW is achieved indirectly. What is estimated directly from the observations collected at each GPS station is the zenith neutral delay (ZND) which is a measure of the propagation delay imposed by the neutral atmosphere on GPS (radio) signals reaching that station. This delay is expressed or parameterized as the

equivalent excess path length associated with a vertical path through the atmosphere. A ZND time series is transformed into a PW time series using measurements of surface temperature and pressure at each GPS station or by inferring these quantities using a numerical weather model or via objective analysis [Bevis *et al.*, 1992]. The ZND can be decomposed into the zenith hydrostatic delay (ZHD), proportional to the surface pressure [Davis *et al.*, 1985], and the zenith wet delay (ZWD) which is very nearly proportional to PW [Bevis *et al.*, 1992, 1994]. Since the temporal variability of ZND is known to be dominated by the variability of the ZWD, it is reasonable that Foster and Bevis [2003] found ZND in Hawaii is also lognormally distributed, or very nearly so.

[4] The lognormal nature of PW and ZND in Hawaii was a surprise to many geodesists (including us) who have typically assumed that PW and ZND have Gaussian statistics (largely for the purposes of computational convenience) if they have considered the statistical distribution at all. Also, even though many other meteorological and hydrological quantities (such as rainfall and cloud droplet size) have been characterized as lognormal, this association rarely occurs in the literature discussing the climatology of PW. Most likely this is because in many parts of the world the empirical probability density function for PW (and for ZND) is bimodal (or multimodal) or, for some other reason, has an appearance which is not strongly suggestive of lognormality. The lognormal character of PW and ZND in Hawaii is a local but not a global property of these quantities, but how common is lognormality, and is it possible that bimodal distributions of PW or ZND represent combinations of two lognormal populations?

[5] In this paper we use both GPS measurements and radiosonde measurements to survey the statistical distribution of PW and ZND in a variety of oceanic and continental settings. By including GPS measurements in addition to the direct PW measurements of radiosondes, we are able to

¹School of Ocean and Earth Science and Technology, University of Hawaii, Honolulu, Hawaii, USA.

²Department of Civil and Environmental Engineering and Geodetic Science, Ohio State University, Columbus, Ohio, USA.

³Cooperative Institute for Meteorological Satellite Studies, University of Wisconsin-Madison, Madison, Wisconsin, USA.

⁴Deceased 5 February 2003.

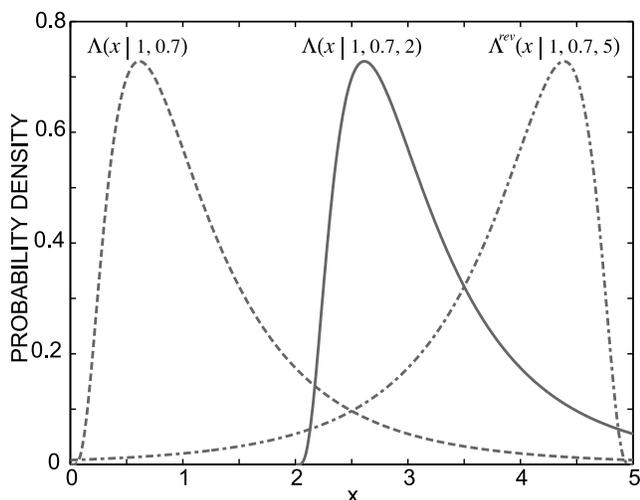


Figure 1. Family of theoretical lognormal probability distribution functions illustrating the three-parameter form of the lognormal and the reverse-lognormal.

confirm whether the ZND, the parameter of direct interest to space geodesists, reflects the same distributions as the PW. The ZND also provides an independent platform to corroborate the radiosonde results. In addition, the rapid global expansion of GPS networks means GPS data is available from many areas that have little or no coverage by traditional meteorological instruments.

2. Statistical Background

[6] The two continuous probability distributions most commonly used to describe atmospheric variables are the normal (N) and lognormal (Λ) distributions. Rain rate [Biondini, 1976; Sauvageot, 1994], cumulus cloud populations [López, 1977] and aerosol optical depth [O'Neill et al., 2000] are all found to have lognormal distributions. Examining relative humidity in the upper troposphere, Soden and Bretherton [1993] and Yang and Pierrehumbert [1994] find several instances of the lognormal distribution. A variate x is Λ -distributed if the variate $z = \log(x)$ is N -distributed. For a detailed discussion of Λ , see Aitchison and Brown [1957]; here we will simply give the equation for the PDF for variable X distributed according to the two-parameter $\Lambda(x|M,s)$:

$$\frac{1}{xs\sqrt{2\pi}} \times \exp\left\{-\frac{(\log \frac{x}{M})^2}{2s^2}\right\}, \quad 0 < x < \infty, \quad (1)$$

where M is the median and s is the geometric standard deviation (GSD). The geometric standard deviation (GSD) should not be confused with the more widely used arithmetic standard deviation (ASD), usually referred to simply as the standard deviation, which is most commonly associated with the N -distribution. The population mean is

$M \exp\left(\frac{s^2}{2}\right)$ and the population variance is $M^2 \exp(s^2)(\exp(s^2) - 1)$ (Figure 1). The maximum likelihood estimators of M and s are calculated using the logarithm of the data set: M is estimated by the exponential of its mean and s by its (arithmetic) standard deviation. More

sophisticated estimation methods may be necessary for data that are sparse and/or noisy.

[7] A slightly more general version of Λ is the three-parameter distribution $\Lambda(x|M, s, t)$. Here an extra term t is included as a “threshold” parameter. The threshold parameter allows the distribution to describe the situation where the variable has a nonzero lower bound. The threshold simply acts to translate the PDF along the x axis (Figure 1). The PDF for the three-parameter distribution is given by equation (1) above substituting the variate with $x' = x - t$. The PDF is now defined for $0 < x - t < \infty$ and the locations of the median and mean are shifted by t .

[8] By adopting a similar approach, it is possible to generate a “reverse” lognormal distribution. In the reverse case, the shape of the distribution is a mirror image of the normal form (Figure 1). Now the distribution is bounded by an upper value t with the long tail tending toward zero (or $-\infty$). The reverse lognormal case is described by defining a new variate $x' = t - x$ which is distributed according to the two-parameter distribution.

3. Data

[9] We use radiosonde profiles retrieved from the National Climatic Data Center (NCDC) CARDS (now IGRA) database and the Forecast Systems Laboratory (FSL). These profiles cover the period from 1973 through 2002 (see Table 1 for site-specific radiosonde model details) and report surface, standard, and significant levels. Not all sites have complete data coverage for this time period, but we select for this paper only sites for which there exist enough data (several thousand profiles) to allow us to examine the gross statistics of the precipitable water distribution. The temperature and dew point profiles were transformed to precipitable water using the equations from Buck [1981]. In order to examine the height variations the cumulative precipitable water profiles were all interpolated to a common set of elevations. Examining the amount of precipitable water in the last height interval of each profile allowed us to identify and exclude profiles that had a significant percentage of their total observed precipitable water in the final layer and so were considered to have terminated prematurely.

[10] The various corrections to the profiles due to model/instrument type that are necessary for detailed intercomparison [e.g., Gaffen and Elliott, 1993; Wang et al., 2002] were not made for this study. It is well known that there are biases in the relative humidity, and to a lesser degree temperature measurements between the various radiosonde models. These biases are generally most pronounced in cold, dry conditions [Elliott et al., 2002; Miloshevich et al., 2006], however, and as the bulk of the water vapor is in the lower troposphere, the impact of relative humidity biases are expected to be relatively small for the PW estimates for our sites, except perhaps the highest-latitude sites in winter months. Any such biases will introduce small errors into our estimates of the lognormal parameters. As we are most concerned with general patterns and gross long-term statistical analyses, however, the conclusions we draw will not significantly be affected by these errors.

[11] The International GPS Service (IGS) provides data from a global network of GPS stations. One of the primary parameters estimated during geodetic processing of GPS

Table 1. Radiosonde Models

| Site | WMO | Profile Dates | Effective Date | Radiosonde Model ^a |
|----------------|-------|---------------|----------------|-------------------------------|
| Mount Pleasant | 88889 | 1989–2002 | 1989 | VAISALA RS80 |
| | | | 1973 | Unknown |
| Koror | 91376 | 1973–2002 | 1986 | VIZ (Generic) |
| | | | 1988 | VIZ B |
| | | | 1995 | VAISALA RS80 |
| | | | 1973 | Unknown |
| Niamey | 61052 | 1973–2002 | 1982 | MESURAL FMO 1950A |
| | | | 1995 | VAISALA RS80 |
| | | | 1973 | U.S.W.B. ELECTRONIC(?) |
| Whitehorse | 71964 | 1973–2002 | 1988 | VIZ -SANGAMO |
| | | | 1988 | VIZ B |
| | | | 1988 | VAISALA (Generic) |
| | | | 1973 | VIZ (Generic)? |
| Funafuti | 91643 | 1973–2002 | 1993 | VAISALA RS80 |

^aFrom *Gaffen* [1996].

data is the Zenith Neutral Delay (ZND). Although, if surface meteorological data is available, the ZND can be further reduced to a direct estimate of precipitable water [Bevis *et al.*, 1992], most of the time the variation in ZND is almost entirely due to variation in PW and so the ZND can serve as a good proxy for PW [Bevis *et al.*, 1992]. During very dry conditions at high latitudes, the hydrostatic delay component of the ZND can experience variations that approach those due the PW. Caution should therefore be used in interpreting details in histograms of ZND from sites and times where these conditions apply. In the work of Foster and Bevis [2003] the ZND for GPS stations in Hawaii is found to closely mirror the lognormal distribution observed for the PW at these sites. As ZND is the parameter directly estimated in GPS positioning, and as the availability of surface meteorological observations at GPS sites is relatively sparse, we choose in this paper to examine the ZND rather than restricting ourselves to PW and those GPS sites for which reliable auxiliary meteorological data exist. The IGS tropospheric combination product provides ZND estimates every 2 hours for 1997 through 2002. Obvious outliers were removed from the final collated ZND data sets.

4. Results

[12] We examine first the gross statistical distributions of PW and ZND, taking no account of seasonality. Although it is expected that seasonal variations will strongly influence PW and ZND at most sites outside the tropics it is of some interest to investigate the histograms of the entire time series. If the time series is long enough that any annual cycles are well averaged, and pressure and temperature can each be reasonably approximated by a mean and variance then it can be inferred that any strong departure from a lognormal form indicates that the source regions for the moisture are strongly segregated and the regions are described by different means and variances. It is also of some interest to examine whether a single parameterization of PW can be adopted in models for some areas, without the need to adopt a seasonally varying set of parameters.

[13] We find that the statistical distributions of PW or ZND from many sites worldwide can be divided into three basic categories; lognormal, reverse-lognormal, and bimodal. The bimodal category can be further divided into three

subcategories; a bimodal-I type is constituted by two lognormal distributions, bimodal-II has both a lognormal and a reverse-lognormal distribution, while bimodal-III comprises two reverse-lognormal distributions. These categories are most easily defined through example, so we present here data from sites that best illustrate each category.

4.1. Lognormal Category

[14] Foster and Bevis [2003] have already shown that sites in Hawaii fall into the first, lognormal, category. Investigating contrasting climatic regimes, we find other sites that exhibit an approximate lognormal form. The GPS station BAHN, located near sea level in Bahrain (Figure 2), is in an arid subtropical coastal setting. The histogram of the 5 years of ZND estimates (Figure 3) clearly closely approximates the typical lognormal form with a rapid rise from the lower values and an elongated upper tail.

[15] Mount Pleasant (WMO 88889) in the Falkland Islands is a high-latitude site in the southern Atlantic. The PW distributions from the radiosonde profiles (Figure 4) show that the median falls quickly for the first kilometer before changing to decline more slowly above this level. Like all the radiosonde sites examined in this study, including those with reverse lognormal distributions, the profile of median values can be closely matched by an exponential function, in agreement with the common assumption that PW falls off roughly exponentially with height. The GSD starts at ~ 0.4 and increases almost linearly. The increasing GSD reflects the growing skewness of the distribution (as the skewness is also dependent on M this is not necessarily strictly true in general), even though the overall width of the histograms is decreasing.

4.2. Reverse-Lognormal Category

[16] The second category we identify is the “reverse-lognormal.” The distributions for sites in this category have a typical lognormal shape except that it is reversed, with the long tail on the lower end of the distribution and the rapid drop to zero probability at the upper end. Occurrences for this type of distribution appear to be more geographically/climatologically limited. Two of the clearest examples are shown in Figures 5 and 6 with radiosonde profiles from Koror in the Republic of Palau (WMO 91408) and GPS ZND estimates from NUSA, Solomon Islands. The surface distribution for the radiosonde station shows a gentle rise in

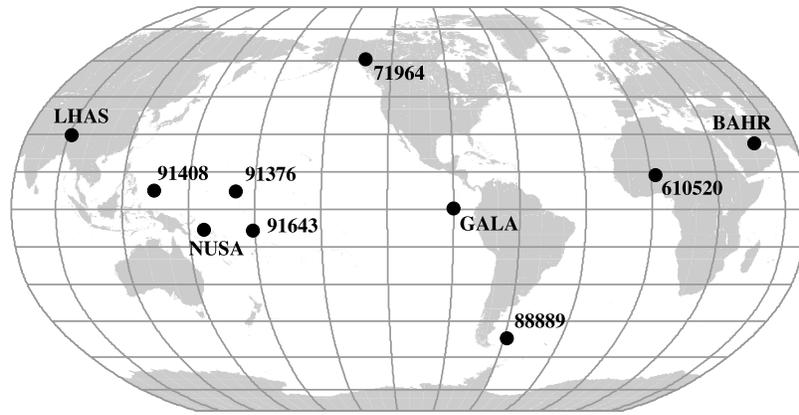


Figure 2. Location map of radiosonde and GPS sites used in this study.

occurrence for PW values above ~ 20 mm rising to a peak occurrence at ~ 50 – 55 mm before dropping rapidly to zero at ~ 70 mm. The ZND at NUSA shows a similar pattern but shifted to between ~ 2450 and ~ 2725 mm of delay. The radiosonde distributions retain the reverse-lognormal form with increasing elevation but shift to lower values. The GSD increases almost linearly from ~ 0.15 to >0.3 while the median declines more like a weak exponential function. At higher elevations the profiles become noisier and fit less well to the simple reverse-lognormal curves.

4.3. Bimodal Category

[17] In the final category the distributions are bimodal. Zhang *et al.* [2003] find that bimodality of water vapor is common in the tropical upper troposphere and examine its implications for the drying and mixing of air parcels. They use a conservative test for bimodality as they are unable to make any prior assumption for the expected distribution of water vapor in their analysis. Perhaps as a consequence of this, they do not find bimodality in the lower troposphere. As we have reason to expect that, over the long term, water vapor should in fact approximate a lognormal distribution (see Appendix A), we can examine distributions that are clearly not approaching a simple lognormal form from this

perspective. We believe that a large fraction of locations exhibit a bimodal type of distribution; however, in many cases the two (or perhaps more) modes have such similar parameters that it is difficult to distinguish them.

[18] Niamey, Niger, is an example of a well-defined bimodal distribution (Figure 7). The lower component is most easily modeled with a standard lognormal distribution. The upper mode is more ambiguous. Either a lognormal (indicating a type I bimodal) or a reverse-lognormal (i.e., a type II bimodal) distribution might reasonably fit the profile

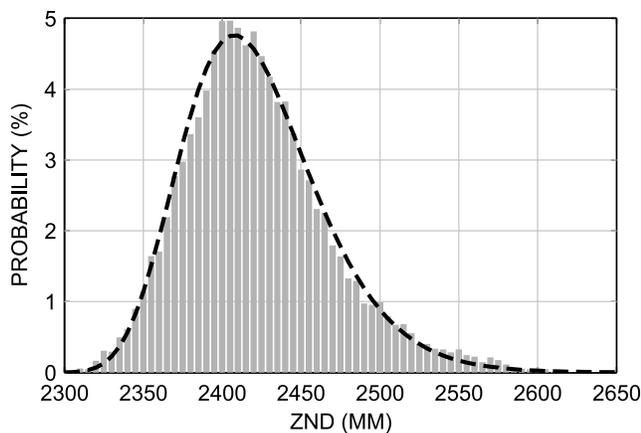


Figure 3. Histogram of zenith neutral delay estimates from GPS site BAHR, Bahrain. M , s , and t for the curve fitted to the histogram are 2417.0, 0.22, and 2222.0, respectively.

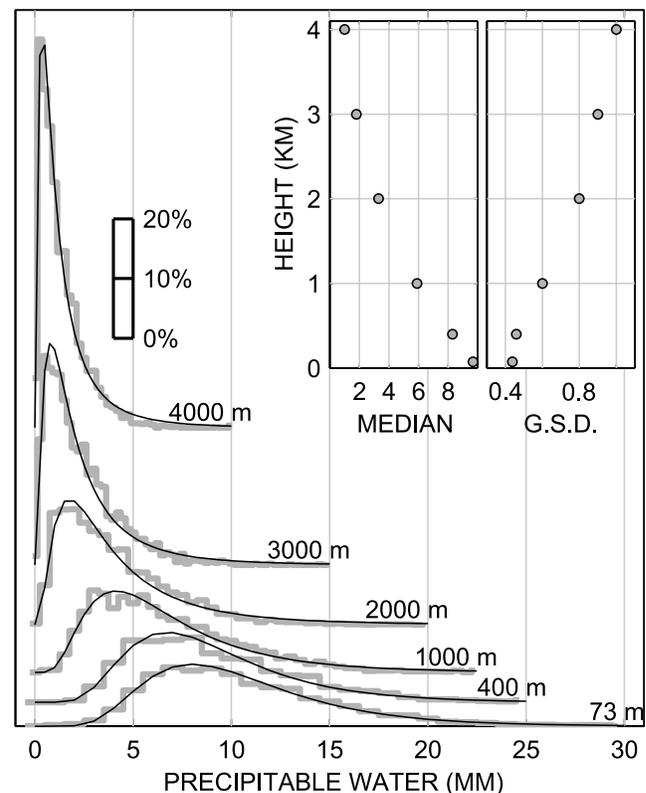


Figure 4. Histograms of precipitable water at different elevations from RAOBS launches at WMO 88889: Mount Pleasant, Falkland Islands. M and s for the curves fitted to each of the six histograms are 9.7, 0.438; 8.3, 0.459; 5.9, 0.60; 3.3, 0.80; 1.8, 0.90; 1.0, 1.0.

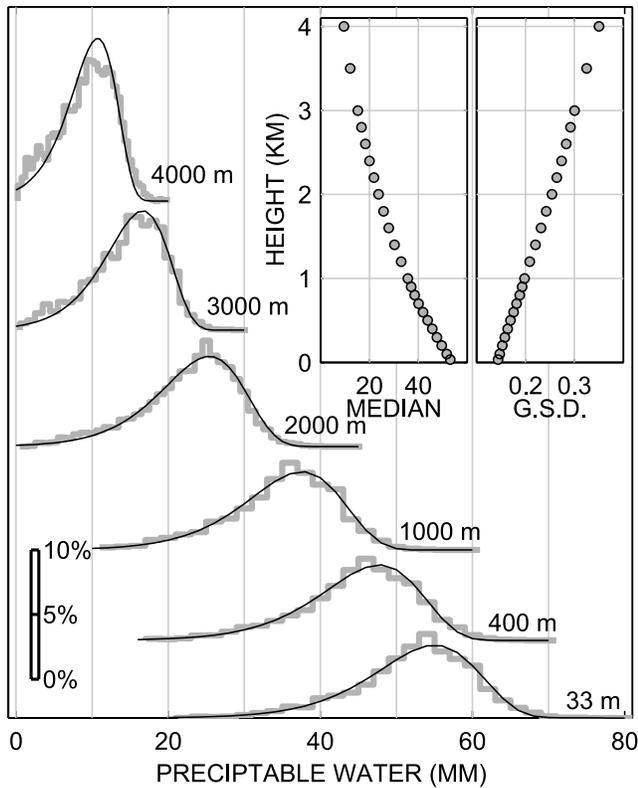


Figure 5. Histograms of precipitable water at different elevations from RAOBS launches at WMO 91376: Koror, Republic of Palau. M , s , and t for the fitted reverse-lognormal curves are 53.63, 0.290, 80.0; 46.09, 0.309, 70.0; 35.59, 0.291, 60.0; 23.08, 0.280, 45.0; 14.61, 0.299, 30.0; 8.90, 0.299, 20.0.

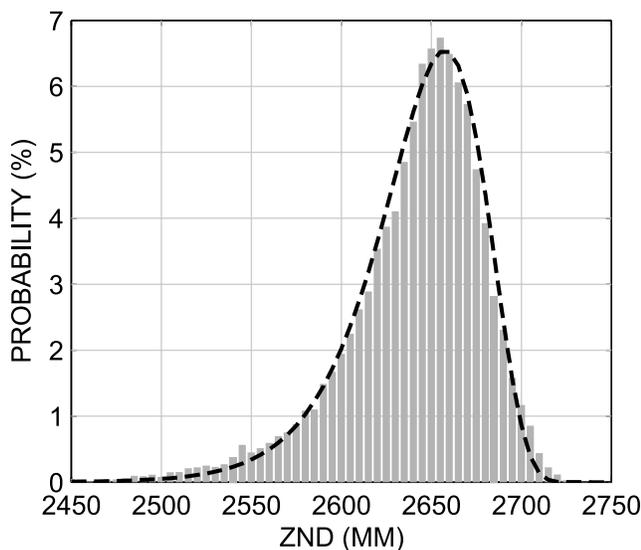


Figure 6. Histogram of zenith neutral delay estimates from GPS site NUSA, Solomon Islands. M , s , and t for the fitted reverse-lognormal curve are 2647.5, 0.33, and 2745.0, respectively.

as the overlap between the modes obscures the telltale asymmetry that would distinguish between them. The two modes remain distinct up to ~ 2000 m beyond which they become increasingly merged and the combination simply resembles a single lognormal distribution.

[19] The GPS station LHAS in Lhasa, Tibet also shows a clear bimodal distribution of ZND (Figure 8). In this case the rapid drop to zero probability on the upper end of the histogram suggests that this is most likely a type II bimodal, comprising a lower lognormal and an upper reverse-lognormal mode. The modes themselves are clearly (Figure 8b) related to the monsoonal seasonality at this site.

[20] An example of a type III bimodal is the GPS site GALA in the Galapagos Islands (Figure 9). Here the bimodality is not immediately obvious; indeed, at first glance it might even be thought to be normal/Gaussian. Examining the time series for this site, however, highlights the dramatic influence of El Niño with PW strongly elevated as the dropping of the westerly winds leads to a warming of the surrounding ocean and much increased convection and rainfall. The histogram of ZND estimates for all times excluding the 1997–1998 El Niño (Figure 9b) suggests a simple reverse lognormal distribution. The histogram during the El Niño episode (Figure 9c) is less well defined as it is a relatively short time period and a limited

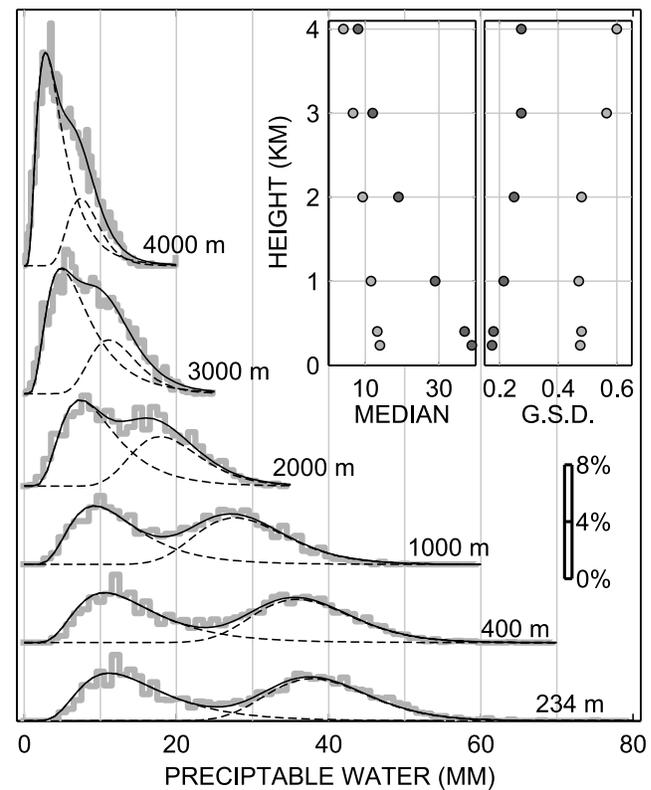
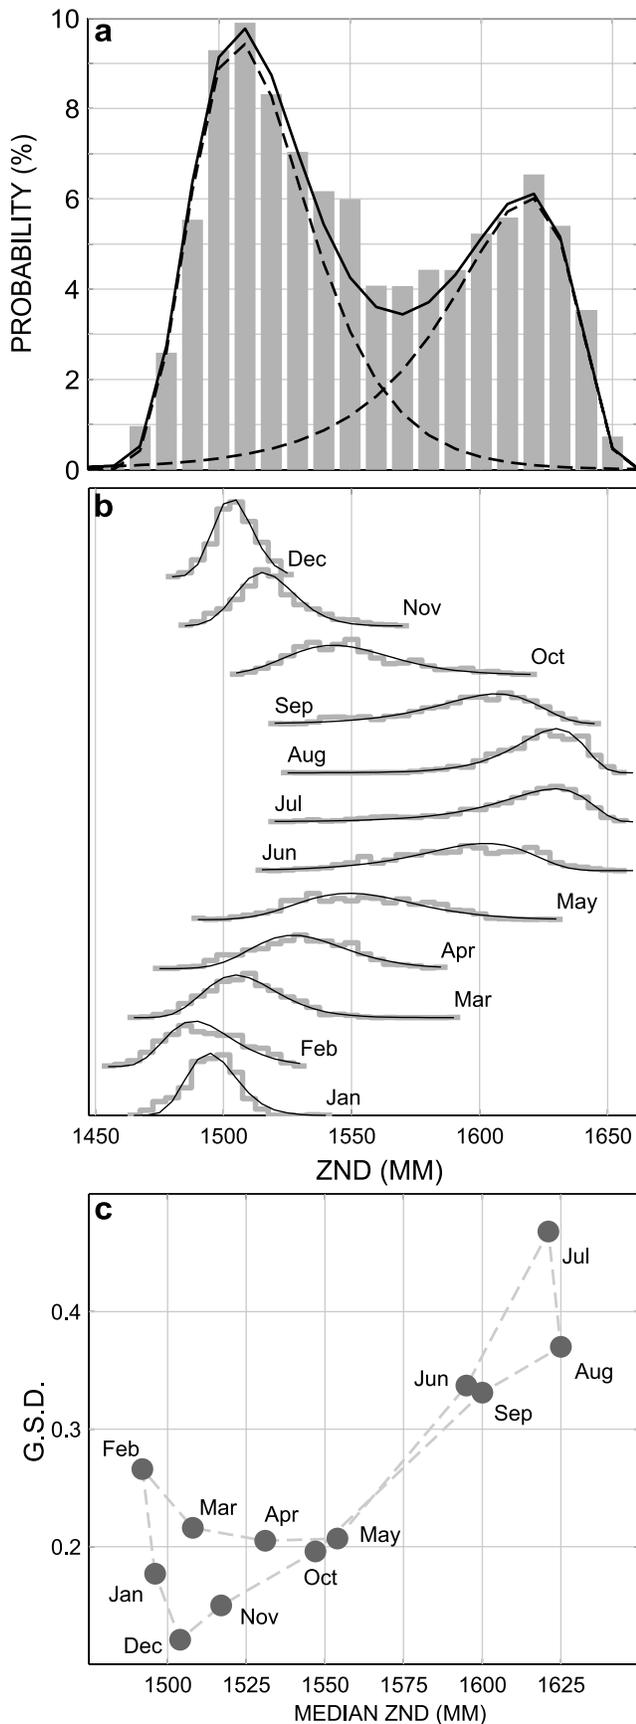


Figure 7. Histograms of precipitable water at different elevations from RAOBS launches at WMO 61052: Niamey, Niger. M and s for the fitted lognormal (lower mode) curves are 14.0, 0.475; 13.25, 0.480; 11.5, 0.470; 9.25, 0.480; 6.65, 0.565; 4.0, 0.600. M and s for the fitted lognormal (upper mode) curves are 39.0, 0.175; 37.0, 0.180; 29.0, 0.215; 19.0, 0.250; 12.0, 0.275; 8.0, 0.275.



number of samples, but it appears once again to be a reverse-lognormal.

4.4. Seasonal Variability

[21] The effect of seasonal periodicity is studied by examining histograms generated by stacking observations by month of year. For the sites in this study, the effect appears to vary with climatic zone. Bahrain and the west Pacific sites, for example, with their limited seasonality, have similar lognormal parameters for all months. Whitehorse, by contrast, appears lognormal at two different timescales. The histogram of its full time series resembles that of Mount Pleasant and in addition the months also appear individually lognormal (Figure 10), with large variations in the lognormal parameters. This makes it particularly notable that the aggregate histogram also appears lognormal despite the variability of the monthly distributions. It seems that in the aggregate case the lognormal form derives from the seasonal cycle itself. Plotting the two lognormal parameters M and s for each of Whitehorse's monthly distributions against each other reveals an interesting seasonal pattern. The winter months cluster with low median and high GSD, and as the summer approaches the median increases while the GSD drops. After the peak of summer the trend returns to the winter values; however, though it follows a similar trend, it has a slightly higher GSD for a given median value. This higher GSD indicates that although PW for these months has the same expected value as earlier in the year, it has higher variance. The fit of the lognormal curves to the histograms is noticeably weaker during the dry winter months. It is possible that this is simply a case where the assumption of approximate lognormality is less valid or that the monthly timescale is inappropriate. Examining ZND estimates from the GPS site at Whitehorse, however, reveals a much closer fit between the histograms and the derived lognormal curves. This suggests an alternate explanation that the RAOBS data is poor, with a dry bias, for some, or all, of the winter months in the data record. Whitehorse was using VIZ model radiosondes during most of the period for which we have data (Table 1). VIZ radiosondes are known to perform very poorly during dry and cold conditions [Elliott et al., 2002; Miloshevich et al., 2006], suggesting that the details of the winter results for Whitehorse be treated with caution. These misfits suggest that a divergence from a lognormal form might be used as an indicator of potential problems in a data set.

[22] Monthly histograms for TVLU illustrate the reverse-lognormal case (Figure 11). Here in the tropics there is very little seasonal variation: the austral winter months are

Figure 8. (a) Histogram of zenith neutral delay estimates from GPS site LHAS, Lhasa, Tibet. M , s , and t for the fitted lognormal and reverse-lognormal curves are 1515.0, 0.325, 1440.0, and 1605, 0.50, 1665.0 respectively. (b) Histograms by month of year. June, July, August, and September are all fitted with reverse-lognormal curves with $t = 1665$. All other months are fitted with lognormal curves with $t = 1440$. (c) Scatterplot of the parameters for the fitted lognormal curves for each month.

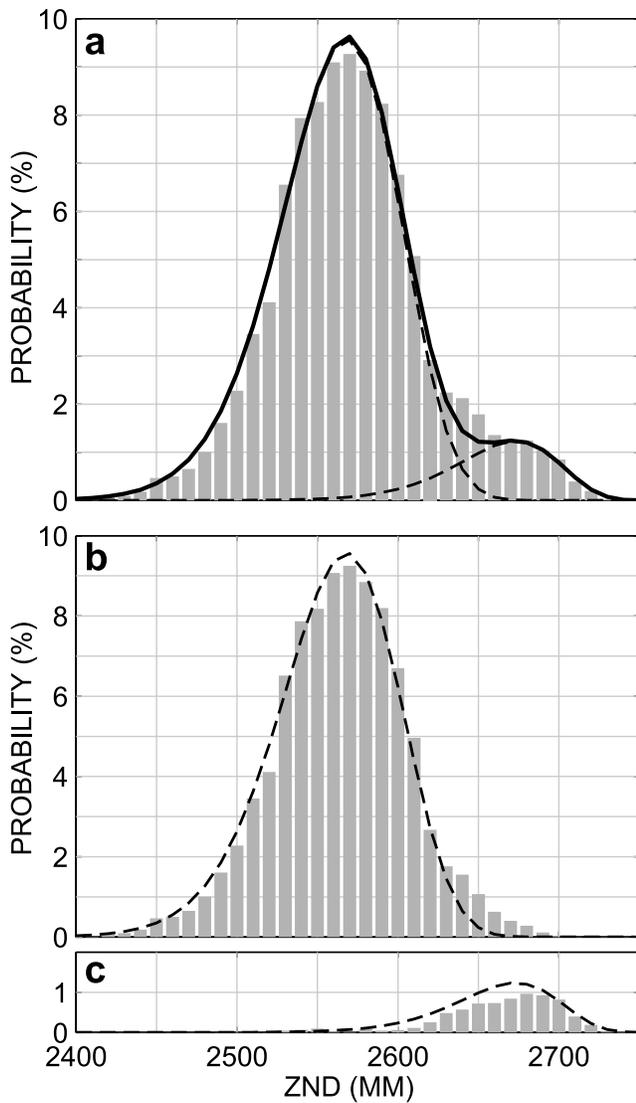


Figure 9. Histogram of zenith neutral delay estimates from GPS site GALA, Galapagos Islands. (a) Composite histogram. (b) Histogram for ZND estimates for all periods excluding the 1997–1998 El Niño. M , s , and t for the fitted reverse-lognormal curve are 2562.0, 0.16, 2800.0. (c) Histogram for ZND estimates during the 1997–1998 El Niño. M , s , and t for the fitted reverse-lognormal curve are 2665.0, 0.25, 2800.0.

slightly drier but the variation of the lognormal parameters is small and any trend is poorly defined.

[23] Although GALA, as noted above, derives its bimodality from a climatic event, the source of the bimodality for the other sites is typically the seasonality. This is illustrated by the monthly histograms from LHAS (Figure 8) which show the winter and summer months occupying limited ranges with a rapid transition between them. This is confirmed by the plot of GSD and median values showing the clustering into two regions. As the onset and termination of the monsoons is quite variable from year to year, the histograms for those months when the onsets and terminations typically occur may be less well defined as they

constitute a mixture of the modes as well as the transition periods themselves.

5. Discussion

[24] The time period over which observations need to be collected in order for the distribution to be approximately lognormal will depend on the timescales of perturbations in the data and especially on any intrinsic periodicities that must be averaged. The timescales for the principle non-periodic perturbations are those of synoptic and climatic events. Similarly, the two main periodicities are the diurnal and annual cycles. Depending on the time window over

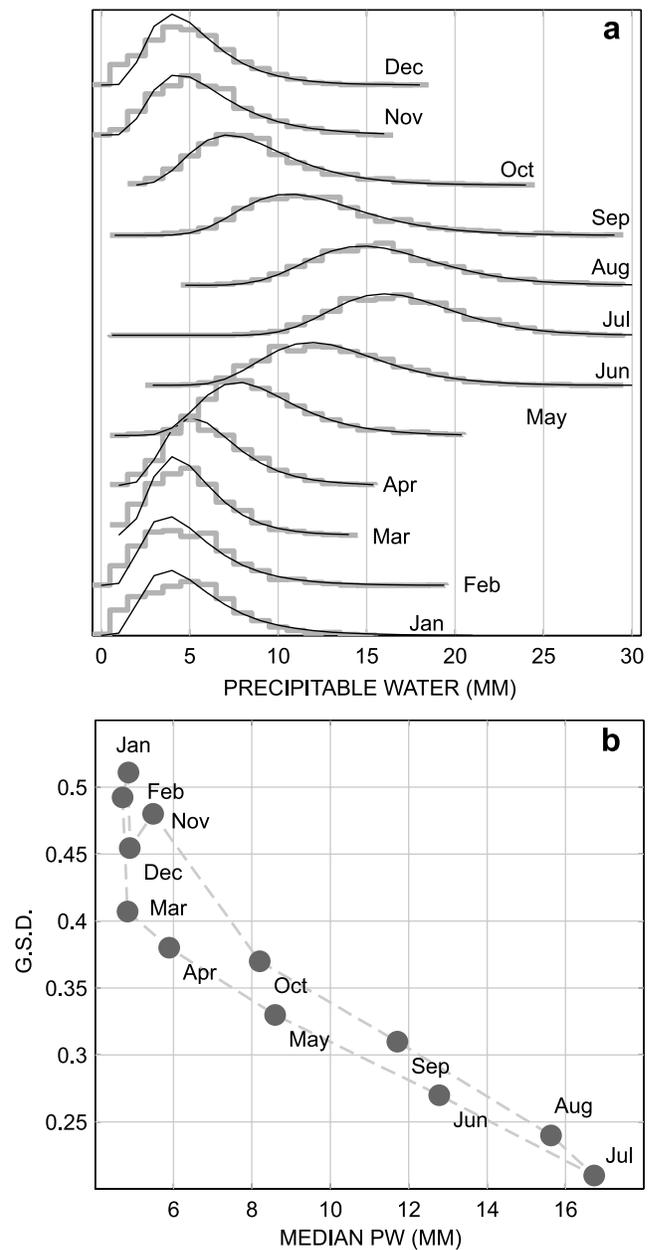


Figure 10. (a) Histograms of precipitable water for Whitehorse, Canada (WMO 71964), for each month of the year. (b) Scatter plot of the parameters for the fitted lognormal curves for each month.

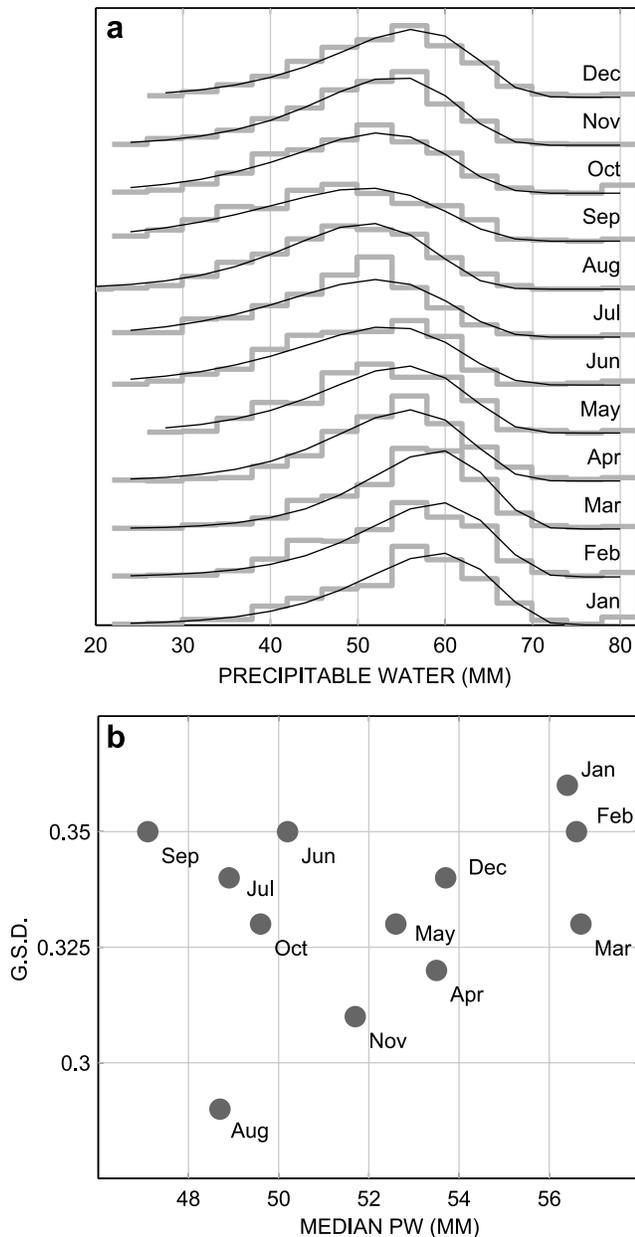


Figure 11. (a) Histograms of precipitable water for Funafuti, Tuvalu (WMO 91643), for each month of the year with fitted reverse-lognormal curves. All curves have $t = 80$. (b) Scatterplot of the parameters for the fitted lognormal curves for each month.

which observations are collected, the effects of these cycles and perturbations on the PDF may be very different. We have shown that GALA is well modeled as a bimodal PDF with modes for the El Niño and non-El Niño periods. In our time series we have only the 1997–1998 El Niño event. Were we to have sampled a period between El Niños, we might have concluded that it was a simple reverse-lognormal. Equally, it is possible that if we had an entire century’s worth of observations, covering multiple El Niño events of varying degrees of strength, the PDF might again appear to be a simple reverse lognormal. There is a potential conflict therefore between the period of time needed to accumulate

enough observations in order to satisfy the asymptotic tendency toward the lognormal distribution and the need for these observations to sample a statistically uniform source population. It is possible that both conditions are satisfied only over certain periods or time windows for a given data set. The characteristic periods for various meteorologic and climatological phenomena may place bounds on the intervals over which (simple) lognormality is observed. Similarly, the sampling interval(s) of the data being examined may also introduce intrinsic limits on the periods over which lognormality can be reliably identified or deduced.

[25] One possible interpretation of a simple lognormal form for a PW time series is that the source region(s) for the observed moisture is(are) effectively mixed over that time period (for further discussion, see Appendix A), as this ensures that one mean and variance can be used to describe the region(s). If, alternatively, the seasonality, for example, involves a distinct change in source region for moisture, and these source regions are effectively unmixed over the time periods of the observations, although each would be expected generate a lognormal distribution, if they have distinct means and variances the distributions will have different characteristic parameters. This will naturally lead to the bimodal (type I) distribution. This should be most clearly illustrated at locations that experience strong monsoon seasons with abrupt onsets and terminations, and Niamey and Lhasa (Figures 9 and 10) are indeed examples of this.

[26] Much more surprising than our observation that the many sites have a lognormal distribution of PW is the occurrence of the reverse-lognormal distribution. We are unable to find any previous mention of this form of distribution occurring and, from existing theory (see equation (A1)), it is difficult to understand its connection to moisture flux. One possible qualitative explanation for the occurrence of the reverse-lognormal derives from consideration of the conditions that lead to the genesis of the lognormal. A common conceptual description for the formation of a lognormal distribution is a system where the magnitude of perturbations of a parameter are proportional to the instantaneous magnitudes of the parameter being perturbed. That is, the greater the current value of PW, say, the greater the incremental change it would be expected to undergo. For the reverse-lognormal the converse must hold: the lower the PW, the greater the expected incremental change.

[27] An interesting related question concerns the physical interpretation of the threshold value needed to describe the reverse-lognormal. Whereas the two-parameter form of the lognormal distribution appears sufficient to describe the simple lognormal form of the PW in every case we have examined, the reverse-lognormal requires a nonzero value for t . We have simultaneously solved for t [Iwase and Kanefuji, 1994] in our fits for the single reverse-lognormal cases, but when modeling multiple curves we have attempted to restrict t to a common value (e.g., GALA, Figure 9). This reflects our belief that t probably represents a real physical limit for the atmosphere at each site. For the two parameter case where, by definition $t = 0$, this limit is simply a completely dehydrated atmosphere. The most obvious interpretation for t in the reverse-lognormal situa-

tion is that it corresponds to some maximum carrying capacity of the atmosphere: complete saturation of the atmospheric column up to some characteristic level for example. We attempted to find a general pattern for t by investigating the relationship between the t values estimated for the Funafuti radiosonde profiles and the column height of saturated air needed to produce those values of PW. The estimated value of t for the surface histogram corresponded to a column height of 4 km, given a physically reasonable surface temperature of 30°C; however, the t estimates for the higher elevation histograms required the top of the saturated column to be progressively higher, and so a physical interpretation for t remains elusive.

6. Conclusions

[28] Sites from various different climatic regimes from around the world illustrate that histograms of long time series of precipitable water, and, by extension, atmospheric delay, often closely approximate a lognormal distribution. This observation is supported by a theoretical study of moisture flux which suggests that, particularly for climatologically averaged data, one would expect observations of precipitable water to tend toward a lognormal distribution.

[29] The reverse-lognormal observed for PW (radiosondes) and ZND (GPS) in tropical, mostly oceanic, locations was not anticipated by any theory known to us. Whereas Λ is bounded below for PW ($PW = 0$), Λ^{rev} is bounded above. The formation of this unusual form of the lognormal distribution, and the nature of its upper bound is not well understood.

[30] The probability distribution functions for precipitable water can be categorized into three common forms: lognormal, reverse lognormal, and bimodal. The bimodal category can be further subdivided into three possible combinations: type I has lognormal + lognormal, type II has lognormal + reverse-lognormal, and type III has reverse-lognormal components. Reverse-lognormal distributions appear to be restricted to areas that experience almost total saturation, typically oceanic equatorial zones. Bimodality is most obvious where seasonality is both pronounced and distinct, such as some monsoonal zones. For most of the rest of the world the simple lognormal seems to be common, even where seasonality is quite pronounced, although it is possible that a multimodal description might be more appropriate for many sites.

[31] The distribution of water vapor and clouds is one of the key issues for climate models, and the refraction of radio waves due to water vapor is one of the primary sources of noise in space geodetic measurements. Recognizing that precipitable water and the delay it induces has a lognormal distribution has implications for how it might best be parameterized. As the lognormal distribution is asymmetric, constraints applied to lognormal parameters need to reflect this. One simple approach would be to parameterize precipitable water (or its delay) using the logarithm of precipitable water as this is normally distributed and so can be easily accommodated by the common least-squares approach. The expectation of lognormality might also provide a useful quality check both for measured and modeled precipitable water. Lognormality also affects the expectations for extreme events, with the skewness of the distri-

bution predicting extremely moist conditions as much more likely than would be expected if precipitable water were normally distributed.

Appendix A: Theoretical Derivation of Precipitable Water Lognormality

[32] Working from first principles, *Raymond* [2000a] derives an equation for moisture transport in terms of relative humidity and finds that the solution to that (differential) equation is of the form:

$$RH(t) = RH(0) \times \exp \left\{ - \int_{t_0}^t \left\{ \frac{\left[\frac{(6889.0952}{T} - 5.31) \frac{d(\ln T)}{dt} + \frac{d(\ln p)}{dt} \right]}{\left(1 - \frac{e_s}{p} \right)} \right\} dt \right\}, \quad (A1)$$

where the solution has been rewritten from *Raymond* to emphasize that it refers to the time domain and to clarify the terms. T and p are temperature and pressure and e_s is the saturation vapor pressure. The equation describes the changes in relative humidity a parcel of air experiences as it is advected through changing pressure and temperature fields. This solution for moisture transport should asymptotically approach a lognormal distribution under both adiabatic and diabatic conditions provided the integrand arguments are sampled from statistical populations with the same expected value and variance and provided the exponent asymptotically approaches a normal distribution [*Raymond*, 2000a].

[33] The exponent describes the time integrated perturbations about a mean or equilibrium condition for moisture as it is transported through the temperature and pressure fields. Natural perturbations about equilibrium conditions are known generally to tend toward a normal (Gaussian) distribution. It is possible that external forcing involving strong latent and sensible heating effects during diabatic activity might disrupt any limited time trend toward a Gaussian; however, long time series that contain many of these cycles may still converge to a near Gaussian distribution.

[34] The integrand is a complex function of pressure and temperature as the saturation vapor pressure depends only on T . In order for each of the components to be approximately described by a single expected value and variance, the timescale of any periodicities and excursions in the pressure and temperature fields must be either significantly longer or shorter than the time period being studied or the amplitude of these signals should be small relative to the variance so that they can be approximated as part of the variance. Although (A1) was derived in a Lagrangian framework, it should apply equally well to Eulerian measurements as they are sampling spatial differences from the same population over the time period. This should hold true provided that the source of the air parcels being measured and their transport history satisfy the same constraints above.

[35] Equation (A1) provides a theoretical grounding for the observations of lognormal relative humidity and is no doubt also connected to the observations of lognormality in

the other moisture related parameters. This result is qualitatively supported by the form of the Clausius-Clapeyron equation describing saturation water vapor pressure based on thermodynamic principals. The exponential of temperature in this equation is also suggestive that moisture is likely to take a lognormal form.

[36] In order to investigate the expected theoretical distribution for precipitable water (PW), we need to extend the derivation from this Lagrangian time evolution of relative humidity and consider the column integrated moisture (PW). We start with the definition of PW [Haltiner, 1957]:

$$PW(p) = \frac{1}{g} \int_{p_{top}}^p q dp. \quad (A2)$$

Here PW is proportional the integral of the mixing ratio q over the pressure interval p_{top} (top of the atmosphere or interval) to p and g is the acceleration of gravity. To assess (A2) and determine the expected vertical variation of q , we can use the modified power law for q [Raymond, 2000b], which in simplified form becomes

$$q(p) \approx q_0 \left(\frac{p}{p_0} \right)^{\lambda + \alpha}, \quad (A3)$$

where $q_0 = q(p_0, T_0)$ and $T_0 = T(p_0)$ where p_0 is the reference pressure level. The parameter $\lambda = \beta \left(\frac{6889.0952}{T_0} - 5.31 \right) - 1.0$, and β is defined by the relationship between pressure and temperature which approximately satisfies $\left(\frac{p}{p_0} \right) \beta = \left(\frac{T}{T_0} \right)$. The exponent α is introduced to compensate for unknowns is the relative humidity. The parameter β varies both spatially and temporally and, as it is weakly dependent on temperature, some prior knowledge of the vertical temperature profile is needed in order to numerically evaluate (A3). Here α also needs to be estimated in order to numerically evaluate the function. This requires knowing the relative humidity at p_{top} , the top of the column; otherwise one can use Smith's [1966] power law formulation for climatological PW estimates given the surface dew point.

[37] This power law relationship describes vertical profiles well in cases where inversions and isolated dry or moist layers are not pervasive [Raymond, 2000b]. Although these features are relatively common, they are quickly suppressed in time averaged climatological profiles.

[38] Substituting (A3) into (A2) and integrating over the pressure interval, we get

$$PW(p) \approx \frac{q_0 \left(p^{\lambda + \alpha + 1} - p_{top}^{\lambda + \alpha + 1} \right)}{g p_0^{\lambda + \alpha} (\lambda + \alpha + 1)}. \quad (A4)$$

Expanding q_0 [Haltiner, 1957] in terms of relative humidity RH_0 and saturation vapor pressure e_{s0} , where $RH_0 = RH(p_0, T_0)$ and $e_{s0} = e_s(T_0)$, namely

$$q_0 \approx \frac{0.622 RH_0 e_{s0}}{p_0}, \quad (A5)$$

allows us to relate PW to RH_0 , e_{s0} , and $\frac{p}{p_0}$ directly:

$$PW(p) \approx \frac{0.622 RH_0 e_{s0}}{g(\lambda + \alpha + 1)} \left[\left(\frac{p}{p_0} \right)^{\lambda + \alpha + 1} - \left(\frac{p_{top}}{p_0} \right)^{\lambda + \alpha + 1} \right]. \quad (A6)$$

Because RH_0 and e_{s0} are not independent, PW is strongly dependent upon RH_0 . Consequently, this then implies that if RH is distributed lognormally then PW should also be distributed lognormally.

[39] Looking at the detail of (A6) indicates that the vertical variation of PW is nearly proportional to $\left(\frac{p}{p_0} \right)^{\lambda + \alpha + 1}$.

Typical values for β and α are ~ 3 and $\sim \pm 1$ respectively, indicating that PW should generally decline rapidly with height as a power of the pressure ratio. For small pressure intervals

$$\left(\frac{p}{p_0} \right)^{\lambda + \alpha + 1} \approx 1 + \frac{(\lambda + \alpha + 1)(p - p_0)}{p_0} + \dots \quad (A7)$$

and is linear in $(p - p_0)$ for $p \approx p_0$. The hydrostatic approximation then implies that variation of PW with height is also approximately linear provided changes in pressure are small.

[40] Scatter about a lognormal distribution may be enhanced with height since spatial and temporal variations for β are more pronounced in the upper half of the troposphere [Raymond, 2000b]. Additionally, α adds to this variability since it ranges between positive and negative values with a magnitude slightly greater than unity.

[41] **Acknowledgments.** Our coauthor, William Raymond, died before the completion of this manuscript. We hope that the final version would have met with his approval and would like to dedicate this paper to him. We thank Brian Mapes, Siegfried Schubert, Stephen Businger, Jim Davis, and two other reviewers for their comments and suggestions. This manuscript is SOEST contribution 6759.

References

- Aitchison, J., and J. A. C. Brown (1957), *The Lognormal Distribution*, Cambridge Univ. Press, New York.
- Bevis, M., S. Businger, T. A. Herring, C. Rocken, R. A. Anthes, and R. H. Ware (1992), GPS meteorology: Remote sensing of atmospheric water vapor using the Global Positioning System, *J. Geophys. Res.*, *97*, 15,787–15,801.
- Bevis, M., S. Businger, S. Chiswell, T. A. Herring, R. A. Anthes, C. Rocken, and R. H. Ware (1994), GPS meteorology: Mapping zenith wet delays onto precipitable water, *J. Appl. Meteorol.*, *33*, 379–386.
- Biondini, R. (1976), Cloud motion and rainfall statistics, *J. Appl. Meteorol.*, *15*, 205–224.
- Buck, A. L. (1981), New equations for computing vapor pressure and enhancement factor, *J. Appl. Meteorol.*, *20*, 1527–1532.
- Davis, J. L., T. A. Herring, I. I. Shapiro, A. E. E. Rogers, and G. Elgered (1985), Geodesy by radio interferometry: Effects of atmospheric modeling on estimates of baseline length, *Radio Sci.*, *20*, 1593–1607.
- Duan, J., et al. (1996), GPS meteorology: Direct estimation of the absolute value of precipitable water, *J. Appl. Meteorol.*, *35*, 830–838.
- Elliott, W. P., R. J. Ross, and W. H. Blackmore (2002), Recent changes in NWS upper-air observations with emphasis on changes from VIZ to Vaisala radiosondes, *Bull. Am. Meteorol. Soc.*, *83*, 1003–1017.
- Foster, J., and M. Bevis (2003), Lognormal distribution of precipitable water in Hawaii, *Geochem. Geophys. Geosyst.*, *4*(7), 1065, doi:10.1029/2002GC000478.
- Gaffen, D. J. (1996), A digitized metadata set of global upper-air station histories, *Tech. Memo. ERL ARL-211*, NOAA, Silver Spring, Md.
- Gaffen, D. J., and W. P. Elliott (1993), Column water vapor content in clear and cloudy skies, *J. Clim.*, *6*, 2278–2287.

- Gutman, S. I., and S. G. Benjamin (2001), The role of ground-based GPS meteorological observations in numerical weather prediction, *GPS Solut.*, *4*, 16–24.
- Haltiner, G. J. (1957), *Dynamical and Physical Meteorology*, 470 pp., McGraw-Hill, New York.
- Iwase, K., and K. Kanefuji (1994), Estimation for 3-parameter lognormal distribution with unknown shifted origin, *Stat. Pap.*, *35*, 81–90.
- López, R. E. (1977), The lognormal distribution and cumulus cloud populations, *Mon. Weather Rev.*, *105*, 865–872.
- Miloshevich, L. M., H. Vömel, D. N. Whiteman, B. M. Lesht, F. J. Schmidlin, and F. Russo (2006), Absolute accuracy of water vapor measurements from six operational radiosonde types launched during AWEX-G, and implications for AIRS validation, *J. Geophys. Res.*, *111*, D09S10, doi:10.1029/2005JD006083.
- O'Neill, N. T., A. Ignatov, B. N. Holben, and T. F. Eck (2000), The lognormal distribution as a reference for reporting aerosol optical depth statistics: Empirical tests using multi-year, multi-site AERONET sun-photometer data, *Geophys. Res. Lett.*, *27*, 3333–3336.
- Raymond, W. H. (2000a), Moisture advection using relative humidity, *J. Appl. Meteorol.*, *39*, 2397–2408.
- Raymond, W. H. (2000b), Estimating water vapor profiles using a modified power law, *J. Appl. Meteorol.*, *39*, 1059–1070.
- Saastamoinen, J. (1972), Atmospheric correction for the troposphere and stratosphere in radio ranging of satellites, in *The Use of Artificial Satellites for Geodesy*, *Geophys. Monogr. Ser.*, vol. 15, edited by S. W. Henriksen et al., pp. 247–251, AGU, Washington, D. C.
- Sauvageot, H. (1994), The probability density function of rain rate and the estimation of rainfall by area integrals, *J. Appl. Meteorol.*, *33*, 1255–1262.
- Smith, W. L. (1966), Note on the relationship between total precipitable water and surface dew point, *J. Appl. Meteorol.*, *5*, 726–727.
- Soden, B. J., and F. P. Bretherton (1993), Upper tropospheric humidity from the GOES 6.7- μm channel: Method and climatology for July 1987, *J. Geophys. Res.*, *98*, 16,669–16,688.
- Tregoning, P., R. Boers, D. O'Brien, and M. Hendy (1998), Accuracy of absolute precipitable water vapor estimates from GPS observations, *J. Geophys. Res.*, *103*, 28,701–28,710.
- Wang, J., H. L. Cole, D. J. Carlson, E. R. Miller, K. Beierle, A. Paukkunen, and T. K. Laine (2002), Corrections of humidity measurement errors from the Vaisala RS80 radiosonde: Application to TOGA COARE data, *J. Atmos. Oceanic Technol.*, *19*, 981–1002.
- Yang, H., and R. T. Pierrehumbert (1994), Production of dry air by isentropic mixing, *J. Atmos. Sci.*, *51*, 3437–3454.
- Zhang, C., B. E. Mapes, and B. J. Soden (2003), Bimodality in tropical water vapor, *Q. J. R. Meteorol. Soc.*, *129*, 2847–2866.

M. Bevis, Department of Civil and Environmental Engineering and Geodetic Science, Ohio State University, Columbus, OH, USA.

J. Foster, School of Ocean and Earth Science and Technology, University of Hawaii, Honolulu, HI 96822, USA. (jfooster@soest.hawaii.edu)