A Stochastic Model for Diagnosing Subtropical Humidity Dynamics with Stable Isotopologues of Water Vapor

JOSEPH GALEWSKY

Department of Earth and Planetary Sciences, University of New Mexico, Albuquerque, New Mexico

DAVID RABANUS

Atacama Large Millimeter Array Observatory, San Pedro de Atacama, Chile

(Manuscript received 15 June 2015, in final form 9 December 2015)

ABSTRACT

The humidity of the free troposphere can be modeled, to first order, in terms of cold-point dehydration, followed by moistening via mixing with boundary layer air. The relative balance between these processes is of prime interest for understanding interannual variability of humidity and for understanding the water vapor feedback. Measurements of water vapor isotopic composition can provide quantitative constraints on these processes. The authors developed a stochastic model that parameterizes water vapor isotopic composition in terms of these processes and fit the model parameters to data from the Chajnantor Plateau, Chile (23°S). For August–November 2012, the average mixing ratio was 1680 ppmv, with mean water vapor δD of −234‰ and mean deuterium excess of 21‰. The data were best fit by an asymmetric last-saturation distribution with mean last-saturation mixing ratio r_s of 391 (+45, −75) ppmv, a median r_s of 368 (+45, −75) ppmv, and a mean mixing fraction between the freeze-dried air and moist boundary layer air of 6.6% (+0.6%, −1.2%). Measurements from August to November 2014 had an average mixing ratio of 2210 ppmv, an average δD of −220‰, and an average deuterium excess of 14‰. The last-saturation distribution for this period was less skewed than for 2012, with an average r_s of 520 (+42, −75) ppmv and a median r_s of 507 (+25, −75) ppmv. The mean mixing fraction for 2014 was 8.8%(+0.8%, −1.8%). The results show that the moistening in 2014, relative to 2012, requires increases in both the last-saturation mixing ratio and the postcondensation moistening and illustrate the utility of isotopic measurements for constraining the processes governing subtropical humidity.

1. Introduction

Global tropospheric water vapor amounts are widely expected to increase in a warming climate, leading to a positive feedback arising from the enhanced greenhouse effect (Held and Soden 2000). The water vapor feedback is the strongest global feedback in general circulation models (GCMs) (Soden and Held 2006), and observations, model results, and theory show a consistent and strong link between increases in surface temperature and increases in water vapor mixing ratio (WVMR) (e.g., Dessler and Davis 2010).

Our understanding of the water vapor feedback and, more generally, the humidity of the free troposphere has been substantially clarified by the so-called last-saturation concept, which approximates the WVMR of air by its saturation value when it was last in a cloud (Pierrehumbert et al. 2006; Hurley and Galewsky 2010b). This concept has proved very useful, especially for understanding the controls on subtropical humidity, which is influenced by far-field processes and nonlocal mixing (Galewsky et al. 2005; Cau et al. 2007; Wright et al. 2010). To first order, the humidity of the subtropics can be understood in terms of drying via condensation to the coldest point encountered by air parcels traveling to the subtropics and the subsequent moistening that may occur. Changes in the balance between cold-point drying and subsequent moistening can arise from a variety of processes and are difficult to disentangle.

Denotes Open Access content.

Corresponding author address: Joseph Galewsky, Department of Earth and Planetary Sciences, University of New Mexico, 221 Yale Blvd. NE, Albuquerque, NM 87131.
E-mail: galewsky@unm.edu

DOI: 10.1175/JAS-D-15-0160.1

Many studies that make use of the last-saturation concept rely in part on reanalysis or GCM output (Pierrehumbert 1998; Dessler and Sherwood 2000; Galewsky et al. 2005; Cau et al. 2007). Given the importance of the water vapor feedback, it would be useful to identify an independent dataset that could be used to constrain the processes that influence free-tropospheric humidity. In principle, measurements of the stable isotopic composition of water vapor could provide such a dataset. The most readily measured isotopologues of water vapor are HDO and $H_2^{18}O$, and their relative abundances in free-tropospheric water vapor are largely set during condensation. While it is straightforward to conceive of the links between last-saturation thinking and water vapor isotopic composition, few measurements have been brought to bear on the problem. Galewsky and Hurley (2010) used an idealized GCM to place water vapor isotopic composition into a last-saturation framework, but they focused on the use of a single isotopologue and there were few reliable measurements available at that time that could have been used to constrain the model.

Since 2012, we have been making measurements of water vapor isotopic composition at the Atacama Large Millimeter Array (ALMA) on the Chajnantor Plateau (Fig. 1), an extremely arid site at an elevation of 5 km in the subtropical Chilean Andes (latitude 23°S). We now have nearly continuous data for two dry seasons, from the years 2012 and 2014, providing us with the first dataset suitable for more fully realizing the potential of water vapor isotopologues to constrain the dynamics of subtropical water vapor.

Here we present a new stochastic model of water vapor isotopic composition that can be applied to large datasets. The best-fitting model parameters for a given dataset are found using a genetic algorithm approach. We will show that the model can generate an excellent fit to the large dataset we have collected on the Chajnantor Plateau and that it can provide constraints on the changes in the relative balance between drying and moistening processes that govern the interannual variability of humidity.

2. Methods

For background material on the stable isotopic composition of water vapor and on the Chajnantor Plateau site, the reader is referred to our previous publications based on this dataset (Samuels-Crow et al. 2014; Galewsky and Samuels-Crow 2015, 2014) and the citations therein.

Mixing ratio and isotopic composition were measured using a Picarro L2130 cavity ringdown spectroscopy (CRDS) analyzer. Details on the instrument and the standardization techniques used are provided in Samuels-Crow et al. (2014) and Galewsky and Samuels-Crow (2015, 2014).

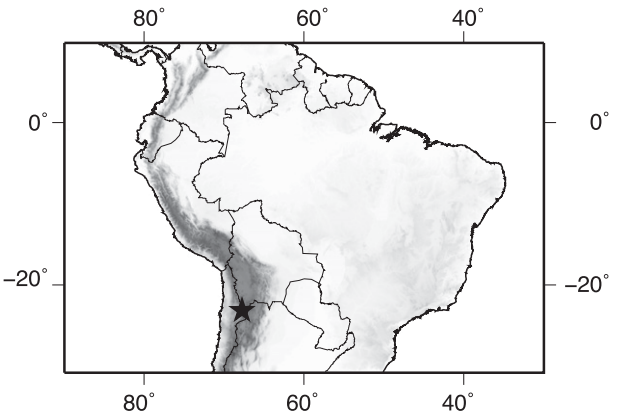

FIG. 1. Location map showing the Chajnantor Plateau (black star) and topography (gray shaded). Longitude in degrees west.

The 1σ uncertainty in 5-min averages of measurements increased at lower mixing ratio but was 5‰, 0.4‰, and 6‰ in δD, δ^{18}O, and deuterium excess (d-excess), respectively, at 110 ppmv and 0.4‰, 0.1‰, and 0.5‰ in δD, δ^{18}O, and d-excess, respectively, at 2750 ppmv.

The water vapor isotopic analyzer we originally deployed on Chajnantor suffered a hardware failure in February of 2013 and was replaced in July of 2014 with an identical model, but with an updated water vapor retrieval to correct for the bias in water vapor concentration measurements reported in Samuels-Crow et al. (2014). We compared the mixing ratio reported by this analyzer with the humidity measurements from the ALMA operations site and confirmed that the mixing ratio bias had been corrected. The Picarro mixing ratios were within about 2% of those measured at the array operations site, a difference that we determined to be insignificant for the analysis presented below. The replacement isotopic analyzer was subjected to the same concentration dependence correction and standardization as our original system, so we believe that the data obtained by the two instruments are comparable.

We used the ERA-Interim dataset (Dee et al. 2011) from the European Centre for Medium-Range Weather Forecasts (ECMWF) for the temperature profile used in the Rayleigh distillation calculation and for diagnosis of interannual variability in upper-tropospheric temperature, middle-tropospheric subsidence, and storm-track activity.

3. The stochastic model

In this section, we present the stochastic model and describe the techniques we use for fitting the model parameters to data. In section 4a, we then use the model in a forward modeling context to explore the model's sensitivities and then in section 4b and beyond, we

TABLE 1. Physical and model parameters.

Parameters	Description
Physical	
r_s	Mixing ratio at last-saturation point
r_m	Mixing ratio at moist source
LCL	Lifting condensation level
f	Mixing fraction between r_s and r_m
S_i	Saturation ratio with respect to ice
δD_i, dxs_i	Initial boundary layer water vapor δD and d-excess
δD_m, dxs_m	Water vapor δD and d-excess at the moist source
Model	
α_{r_s}, β_{r_s}	Beta distribution shape parameters for r_s
$r_s^{\max}$	Maximum possible value of r_s
α_f, β_f	Beta distribution shape parameters for f
μ_s, σ_s	Mean and standard deviation for saturation ratio
μ_{r_m}, σ_{r_m}	Mean and standard deviation for mixing ratio at moist source
σ_δ	Standard deviation for initial and moist source vapor δD
σ_{dxs}	Standard deviation for initial and moist source vapor d-excess
μ_{δ_i},	Mean δD for initial vapor
μ_{dxs_i}	Mean d-excess for initial vapor
μ_{δ_m},	Mean δD for vapor at moist source
μ_{dxs_m}	Mean d-excess for vapor at moist source

actually fit the model parameters to the data. Our guiding philosophy in developing this model was to find the simplest set of processes that fits the data. Most of the process representations we use are fairly simple, and one can easily envision more complex formulations. While such formulations may be fruitful for future study, here we focus on the most parsimonious model we could identify that fits the data.

a. Parameters

Our model generates a large number of synthetic measurements by subjecting each synthetic measurement to Rayleigh distillation to a random cold-point temperature and ice supersaturation (e.g., Jouzel and Merlivat 1984). Each point is then moistened by a random mixing fraction with a moisture source. Additional parameters are the boundary layer isotopic composition prior to ascent and the isotopic composition of the moisture source. The physical parameters and their associated model parameters are summarized in Table 1, and the model is illustrated in a cartoon in Fig. 2.

The initial water vapor isotopic composition prior to ascent is parameterized as an initial water vapor δD and d-excess ("A" in Fig. 2). Explicitly parameterizing the d-excess allows us to directly restrict the d-excess to plausible values. The $\delta^{18}O$ value is easily computed from the d-excess. Both parameters are modeled here as a normal distribution, with a specified mean and standard deviation.

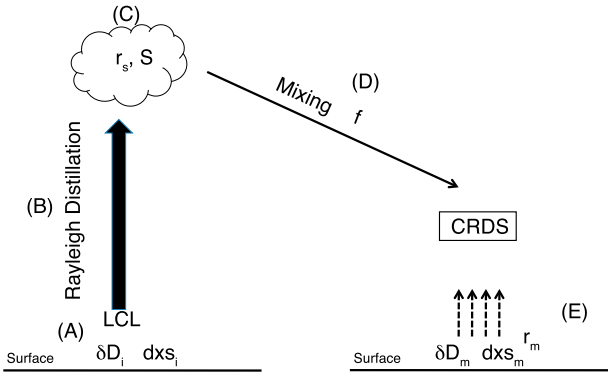

FIG. 2. Cartoon illustrating the physical parameters of the model and the sequence of events that generate a synthetic measurement at the CRDS site.

Rayleigh distillation ("B" in Fig. 2) is initiated at a random lifting condensation level (LCL) between 900 and 775 hPa. Any fixed LCL in this range can fit the data, and no variability in the LCL is required. Rayleigh distillation occurs along the sounding shown in Fig. 3, which is a tropical sounding retrieved from ERA-Interim averaged over August–November 2012 between longitudes −130° and −120° and latitudes 0° and −10°. The most important characteristics of the sounding are the initial condensation temperature (set by the LCL) and final condensation temperature (set by the final saturation mixing ratio r_s). We experimented with subtropical and extratropical soundings and found that they could only fit subsets of the data. The use of a tropical sounding does not imply that all of the condensation occurs in the deep tropics, but some degree of condensation at temperatures encountered in the upper tropical troposphere is required to fit the data. The use of a single tropical sounding was sufficient to fit the data, so we did not pursue more complex formulations involving multiple soundings. Along the same lines, the Rayleigh distillation framework with supersaturation (Jouzel and Merlivat 1984) is relatively simple, but because it fits the data we did not explore the use of more complex condensation models.

Two of the most important parameters in the entire model are related to the in-cloud processes at the last-saturation point ("C" in Fig. 2), and they are the final saturation mixing ratio r_s and the saturation ratio with respect to ice S_i. The results are very sensitive to the probability distributions used for both of these parameters. An asymmetric distribution for r_s was required to fit the data, and we found the best results were obtained with a beta distribution scaled by the maximum potential value of r_s. The beta distribution is given by

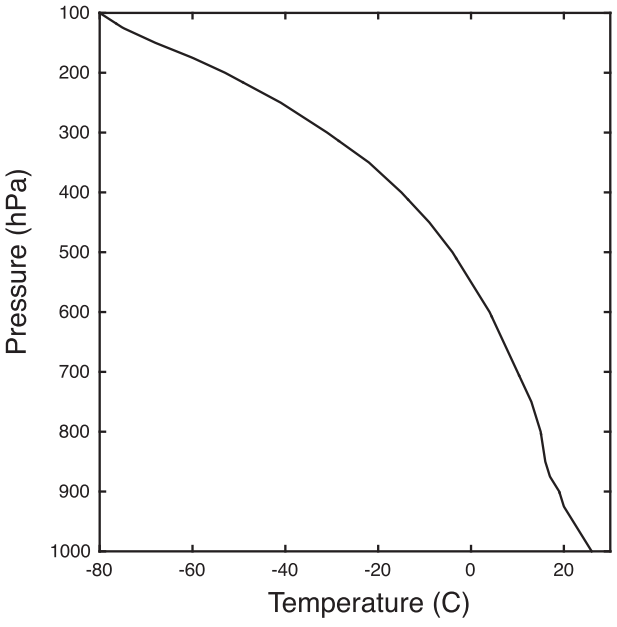

FIG. 3. Sounding used for the Rayleigh distillation component of the stochastic model. Average tropical Pacific sounding from August through November 2012 from ERA-Interim.

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}, \quad 0 \le x \le 1; \quad \alpha,\beta > 0, \qquad (1)$$

where α and β are the shape parameters and B is the beta function, which is given by

$$B(\alpha,\beta) = \int_0^1 t^{\alpha-1}(1-t)^{\beta-1}\,dt. \qquad (2)$$

While the last-saturation point is explicitly parameterized as a mixing ratio r_s, it can be cast in terms of a temperature of last-saturation T_s or as a pressure of last-saturation p_s, which we do below for continuity with the broader literature on last-saturation thinking.

Studies of the global distribution of supersaturation have shown that an exponential distribution fits satellite data (Gettelman et al. 2006). We found that no exponential distribution of supersaturation could fit our isotopic data. The d-excess values were especially sensitive to the supersaturation, and we found that a normal distribution for S_i was required to fit the data. This result is consistent with the central limit theorem, which states that the distribution of the average of a large number of independent, identically distributed variables will be approximately normal, regardless of the underlying distribution. The normal distribution of supersaturation in our model may arise from averaging the measurements as well as from the cumulative effect of many perturbations arising from in-cloud and transport processes (Karcher 2012).

The moistening after last saturation ("D" in Fig. 2d) is as important as the last-saturation process itself for fitting the observations. The moist source ("E" in Fig. 2) is parameterized by normal distributions of the humidity r_m and the isotopic composition. While the isotopic compositions of the initial vapor and the moist source may be different, we found that a single set of standard deviations was sufficient to capture the required variability. The mixing fraction f is the percentage of air that the moist source contributes to the final air mass. It is drawn from a beta distribution, and we found that an asymmetric distribution was required to fit the data. We fixed the average mixing ratio of the moist source r_m to 20 000 ppmv, with a standard deviation of 30 ppmv, in order to avoid a nonunique trade-off between f and r_m that can confound interpretation of the results. Defining r_m and f in this way provides an ultimate end member for the mixing line but does not necessarily imply that all moistening occurs via direct transport from the marine boundary layer. Any process generating air masses at intermediate points on the mixing curve would be consistent with this formulation, although the final mixing fraction would be correspondingly higher for such processes [see Fig. 2 and associated discussion in Galewsky and Samuels-Crow (2015) for more details on this point]. Other sources of moistening can include the outflow from shallow convection, evaporation, and sublimation of cloud condensate and evaporation from the land surface.

b. *Fitting the model to data*

We used genetic algorithm techniques to seek the set of model parameters that minimized the mismatch between the observations of mixing ratio and isotopic composition and the simulated dataset. Genetic algorithms are a class of optimization techniques based on a natural selection–like process inspired by biological evolution. The algorithm modifies an initial population of individual solutions and at each step the algorithm selects individuals from the current population as parents to produce the children for the next generation. Over successive generations, the population evolves toward an optimal solution as determined by a fitness function (Beasley et al. 1993a,b).

A candidate solution consists of a genome that spans the parameters of the problem. The initial population is generated by a random-number generator and can be constrained as required by the problem. The fitness of each individual is determined by a fitness function, and individuals with higher fitness are assigned a greater probability of reproducing. A mating pool is generated and offspring are generated from each mating pair by swapping randomly selected portions of each parent's

genome. The offspring are then subjected to random mutation. The average fitness of the next generation will be higher than previous generations, and the process is repeated for as many generations as are necessary for the population to converge to a steady-state population set.

Other techniques for global optimization include particle swarm (Kennedy 1995) and simulated annealing (Kirkpatrick et al. 1983), which we tested and found yield similar results to the genetic algorithm approach.

The selection of a fitness function is key to finding the optimal solution. We explored several different test statistics for comparing two-dimensional datasets, including simple metrics like the root-mean-squared (RMS) difference between the observed and simulated probability densities and the mean absolute error (MAE). We also tested more sophisticated techniques including the 2D kernel-based test of Duong et al. (2012), the nonparametric test of equality between two copulas described by Rémillard and Scaillet (2009), and the paired-sample Kolmogorov–Smirnoff test as implemented by Peacock (1983). The different techniques yielded consistent and comparable results for our application, but the results presented below are based on the Kolmogorov–Smirnoff test. This test also returns a p value for testing the null hypothesis that the simulated and observed datasets were generated from the same underlying distribution. The fitness function was the average Kolmogorov–Smirnoff statistic from the joint distribution of δD and mixing ratio and the joint distribution of d-excess and mixing ratio.

To keep the problem computationally tractable, we explored the number of simulated points required to obtain reliable results and found that 1000 simulated points yielded the same best-fitting model parameters as simulations with many more (10 000) points, and the results presented below are for 1000 simulated points.

There are several potential sources of uncertainty in the fitted parameter estimates. First, the measurements of mixing ratio and isotopic composition are subject to uncertainties that may influence the results. As described above, the reported mixing ratios have uncertainties of up to 4%, δD has uncertainties of up to 5‰, and d-excess has uncertainties of up to 6‰. These are the most extreme measurement uncertainties, which occur only at the lowest mixing ratios. Experiments in which the mixing ratios and isotopic measurements were randomly perturbed by up to these maximum uncertainties yielded model parameters that were nearly identical to those presented below (not shown), suggesting that measurement error is not biasing the results.

We also considered the influence of small perturbations to the best-fitting parameters. We systematically subjected each parameter to progressively larger perturbations until the results were statistically significantly different from the observations at the 95% confidence level. This technique was used to determine the upper and lower bounds for each parameter reported below.

4. Results

In the previous section, we presented an overview of the stochastic model and the techniques used for fitting the model parameters to the data. In this section we apply the model in two ways. First, we use it in a forward modeling context to illustrate some of its sensitivities and then we apply it to the data from Chajnantor.

a. Idealized cases

To explore the behavior of the model before trying to fit the model parameters to the data, we first present several forward modeling results. Figure 4 (red lines) shows some results of the model with parameters that yield results representative of isotopic data from subtropical settings (Noone et al. 2011; Galewsky et al. 2011; Hurley et al. 2012; Samuels-Crow et al. 2014) and illustrates some of the links between the modeled parameters and the resulting water vapor isotopic composition. The contours show the estimated probability densities for the water vapor δD and d-excess. In this example, the average r_s was 353 ppm, which corresponds to $T_s = -41$°C at $p_s = 235$ hPa. The air parcels were subjected to a skewed distribution of moistening, with an average mixing fraction between the moist end member and the last-saturation point of 7%. The average values for δD and d-excess are −233‰ and 20‰, respectively, and the average final mixing ratio is 1775 ppmv. This case is broadly representative of measurements made in subtropical high-altitude settings (Samuels-Crow et al. 2014).

With the blue lines in Fig. 4, we perturbed the reference case to lower the average r_s to 161 ppmv and the average T_s to −53°C but we held the mixing fraction (and all other parameters) constant. In this scenario, the average δD was reduced to −201‰, the average d-excess was increased to 22‰, and the average final mixing ratio was reduced to 1590 ppmv. These results are readily understood. The d-excess at low mixing ratios is expected to increase as a result of enhanced condensation, while the mixing between the very dry upper-tropospheric air and the moist source will act to raise the overall δD values.

This framework allows us to explore one of our goals for the use of water vapor isotopic measurements: can they be used to constrain the processes that govern moistening or drying of the subtropics? We used two idealized cases that moistened the reference case in Fig. 4 to identical average mixing ratios (2220 ppmv) but used different processes to reach that state. In the first case (Fig. 5, red lines), the average mixing fraction is

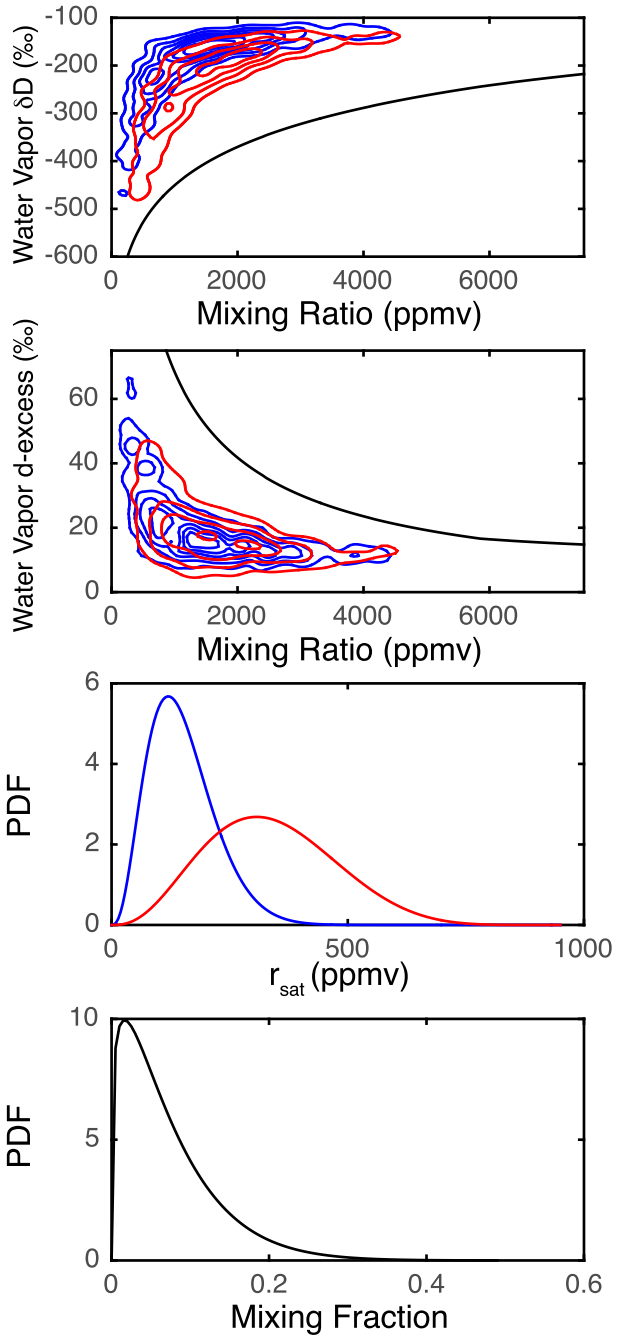

FIG. 4. Model output for reference simulation (red; $\overline{r_s} = 353$ ppmv, $\overline{T_s} = -41°C$, $\overline{P_s} = 235$ hPa) and reference simulation perturbed with a lower r_s (blue; $\overline{r_s} = 161$ ppmv, $\overline{T_s} = -53°C$, $\overline{P_s} = 200$ hPa). (a) Water vapor δD [contours show probability density; contour interval (c.i.) = 7.5×10^{-7}]; (b) as in (a), but for water vapor d-excess (c.i. = 5×10^{-6}); (c) PDF of cold-point (saturation) mixing ratios; and (d) PDF of mixing fraction between dry UT air and moistening source. The solid lines in (a) and (b) indicate a single Rayleigh distillation pathway.

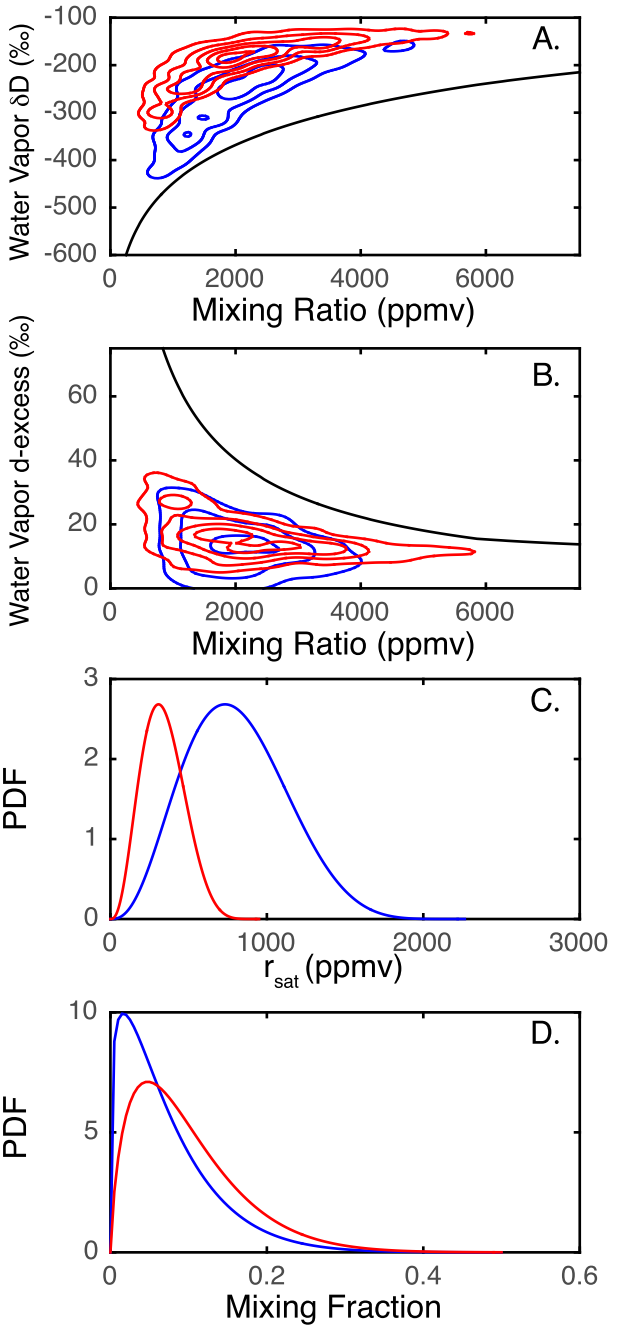

FIG. 5. As in Fig. 4, but moistened to 2200 ppmv by an increase in mixing fraction (red) and by an increase in saturation mixing ratio (blue).

increased from 7% to 9.5%, while the last-saturation parameters were held constant. In the second case (Fig. 5, blue lines), the mixing fraction distribution from the control case was held constant, but the average last-saturation mixing ratio was increased to 575 ppmv.

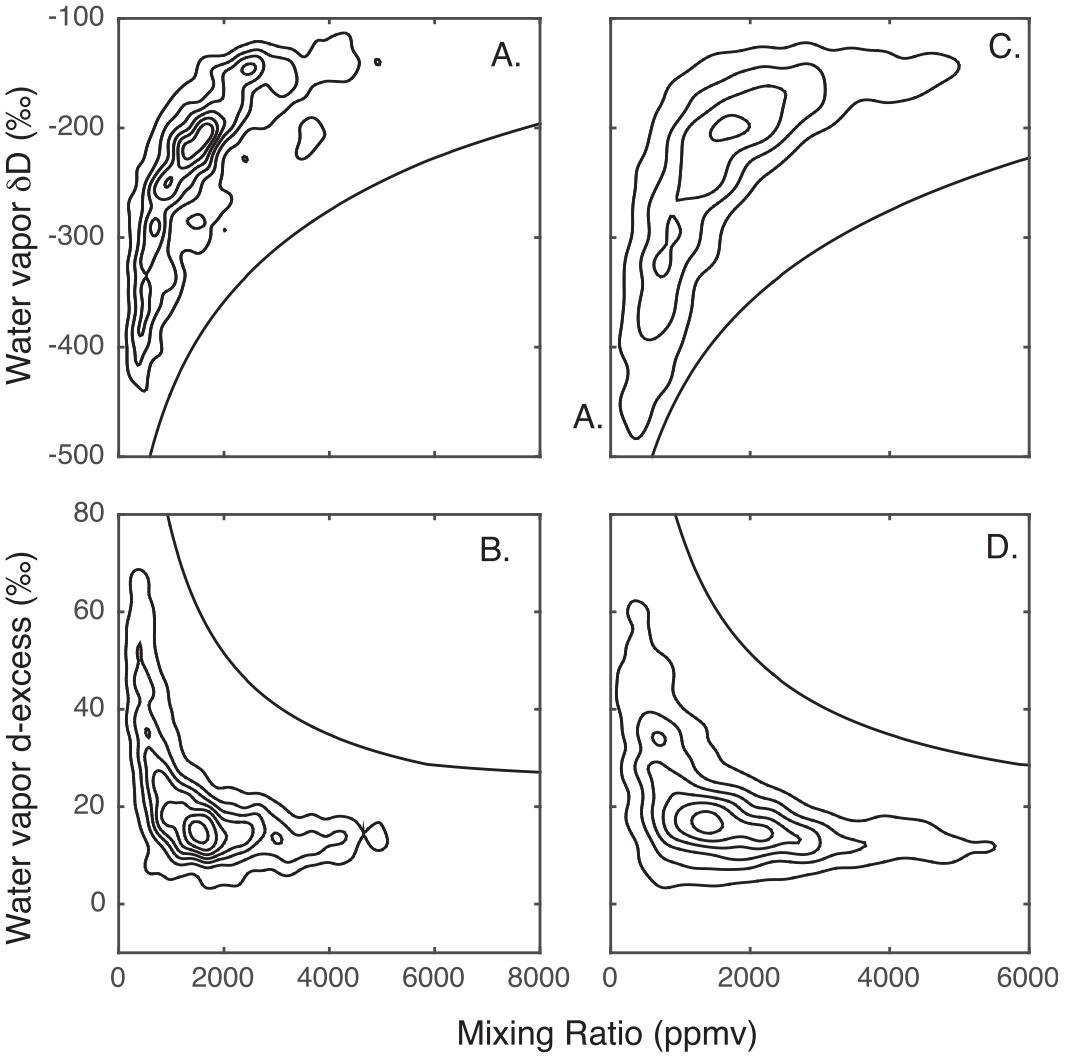

FIG. 6. (a),(b) Probability density estimates for the August–November 2012 measurements and (c),(d) the best-fitting model output for water vapor δD and d-excess. Contour interval is 7.5×10^{-7} for δD and 5×10^{-6} for d-excess. Representative Rayleigh distillation curves shown for reference.

When moistening is effected by enhanced mixing, the δD values are moved along a mixing line toward moister values and correspondingly higher δ values (Fig. 5a). Conversely, when the moistening is effected by increasing the last-saturation temperature (Fig. 5e), the δ values remain closer to the Rayleigh curve. If one only had measurements of mixing ratio, the processes governing the humidity would be difficult to disentangle. The addition of isotopic measurements, however, makes it clear that two different processes are at play. The joint distributions of mixing ratio and isotopic composition are clearly different between the two cases and are formally different at the 99% confidence level using the Kolmogorov–Smirnoff test. These idealized results show that there is some potential for using water vapor isotopologues for disentangling the mechanisms that govern changes in free-tropospheric humidity.

b. *The 2012 dry season*

In the previous section, we used forward modeling to illustrate some of the features of the stochastic model. In this section, we turn to the problem of fitting model parameters to the data and assessing the uncertainties in the fitted parameters. Figure 6 shows the probability density estimates for the water vapor δD (Fig. 6a) and d-excess measurements (Fig. 6b) from August through November 2012 plotted against the water vapor mixing ratio. The mean observed mixing ratio r for the period was 1680 ppmv. The mean δD was −234‰, and the mean d-excess was 21‰. Consistent with previous

TABLE 2. Modeled parameters for 2012 and 2014.

Parameters	2012	2014
Physical		
r_s	391 (+43, −76) ppmv [0.24 (+0.03, −0.05) g kg^{-1}]	520 (+42, −75) ppmv [0.32 + (0.02, −0.05) g kg^{-1}]
f	6.6% ± 1%	8.8%(+0.8%, −2%)
S_i	1.14 ± 0.005	1.16 ± 0.003
δD_i	−92‰ ± 9‰	−110‰ ± 5‰
dxs_i	28‰ ± 3‰	17‰ ± 2‰
δD_m	−118‰ ± 6‰	−104‰ ± 11‰
dxs_m	12‰ ± 1‰	13‰ ± 0.5‰
Model		
α_{r_s}	2.48 (+0.15, −0.26)	3.44 (+0.14, −0.36)
β_{r_s}	6.05 (+1.19, −0.52)	5.54 (+0.77, −0.48)
r_s^{max}	1346 (+140, −129) ppmv	1359 (+83, −67) ppmv
α_f	1.31 (+0.09, −0.14)	1.23 (+0.08, −0.14)
β_f	18.66 (+2.16, −1.27)	12.78 (+1.77, −0.50)

studies, the data are suggestive of condensation to a relatively low last-saturation temperature, followed by mixing with a moister source. The d-excess data were presented in Samuels-Crow et al. (2014) and are consistent with last saturation under ice-supersaturated conditions.

The genetic algorithm found a best-fitting set of parameters that yielded simulated data points consistent with the observations (Figs. 6c and 6d), with average simulated δD of −241‰, average simulated d-excess of 20‰, and an average simulated mixing ratio of 1677 ppmv. The last-saturation and moistening distributions (Fig. 8) are skewed, with a mean r_s of 391 (+45, −75) ppmv [0.24 (+0.03, −0.05) g kg^{-1}], median r_s of 368 (+45, −75) ppmv, and a mean mixing fraction f of 6.6%(+0.6%, −1.2%). The other modeled parameters and their uncertainties are given in Table 2.

About 1% of r_s occurs below 50 ppmv, which is consistent with the upper tropical troposphere, just below the base of the tropical tropopause layer (TTL) (Fueglistaler et al. 2009). Simulations conducted with an extratropical sounding (not shown) were unable to fit the observations, suggesting that some of the water vapor reaching Chajnantor has been processed through temperatures as cold as those encountered near the tropical tropopause layer.

c. *The 2014 dry season and comparison with 2012*

The 2014 dry season (Figs. 7a and 7b) was moister than 2012 (2210 ppmv in 2014 versus 1680 ppmv in 2012), with higher average δD (−220‰ in 2014 versus −234‰ in 2012) and lower average d-excess (14‰ in 2014 versus 21‰ in 2012).

The best-fitting model (Figs. 7c, 7d, and 8) captured the observed moistening and changes in isotopic composition well, with simulated values for the average mixing ratio, δD, and d-excess of 2234 ppmv, −224‰, and 13‰, respectively. The last-saturation distribution for 2014 was less skewed than for 2012, with an average r_s of 520 (+42, −75) ppmv [0.32 (+0.02, −0.05) g kg^{-1}] and a median r_s of 507 (+25, −75) ppmv. The average f was 8.8%(+0.8%, −1.8%). The parameter estimates and their uncertainties are shown in Table 2.

To further explore the processes that governed the moistening between 2012 and 2014, we performed idealized experiments in which we held r_s to 2012 values and used the genetic algorithm to see if other parameters could be found to fit the 2014 data and a similar experiment in which f was held to 2012 values. In neither case could the genetic algorithm find a set of parameters that fit the 2014 data, suggesting that the changes in both r_s and f were required to fit the 2014 data. We can thus conclude that the moistening in 2014 was effected by an increase in both r_s and f, compared to 2012.

The modeling results suggest that the moistening in 2014 required an increase in the average r_s of about 130 ppmv and an increase in the average f of about 2%. Lower-tropospheric subsidence in the southeast Pacific was weaker in 2014 than in 2012 (Fig. 9a), a result that is consistent with the increase in f in 2014. Assessing the change in T_s requires knowledge of the pressure of last saturation, which is not explicitly modeled. We can diagnose the average pressure of last saturation from the stochastic model, however, and find that the average pressure of last saturation was 235 hPa in 2012 and 250 hPa in 2014, and that the average change in last-saturation temperature between 2012 and 2014 was about 3°C.

The upper troposphere over the South Pacific was, on average, slightly warmer during 2014 than 2012 (Fig. 9b). While the average increase in T_s between

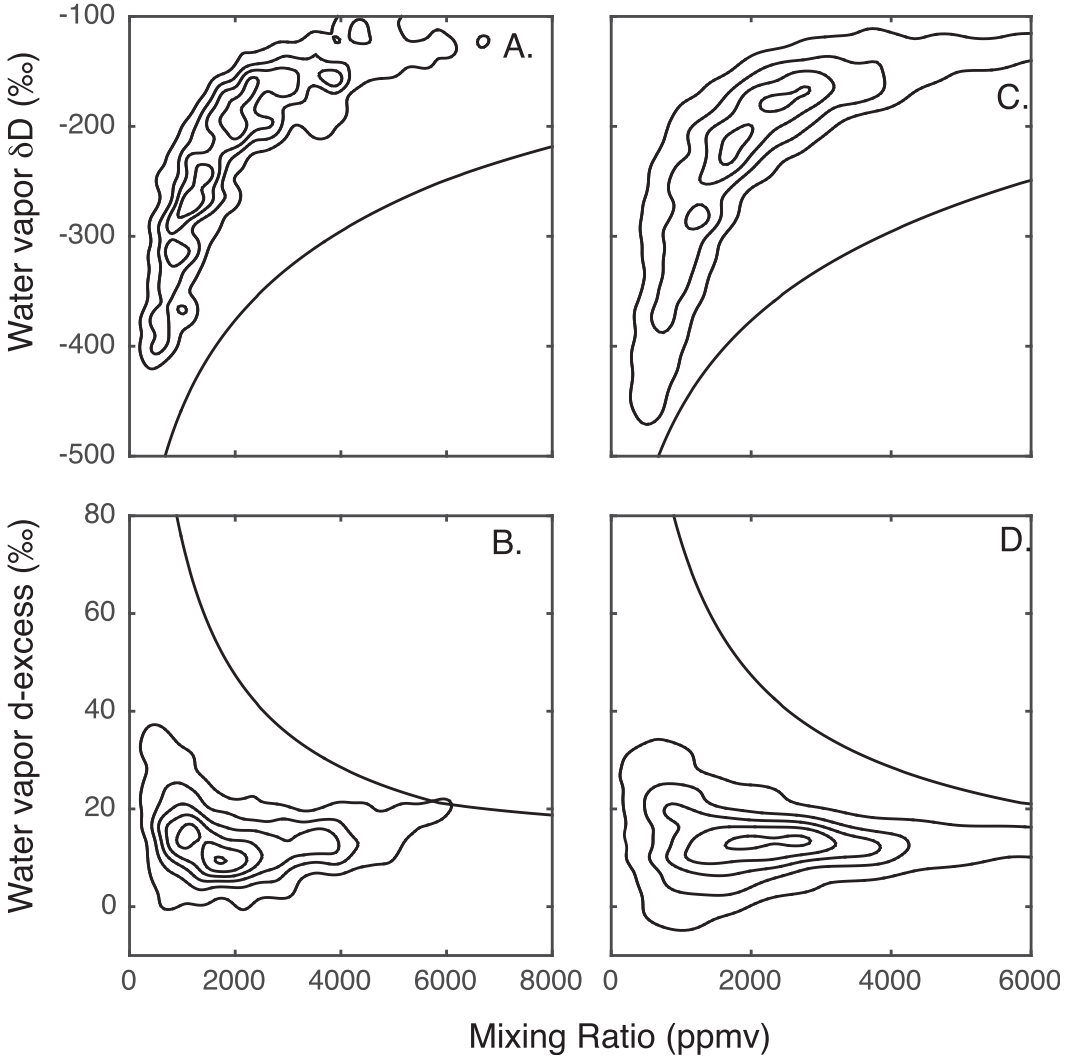

FIG. 7. As in Fig. 6, but for August–November 2014.

2012 and 2014 was about 3°C, the upper troposphere over the South Pacific warmed by less than 1.5°C, implying that a change in circulation was required to increase the average T_s. To assess the circulation changes that may have controlled the increase in T_s, we diagnosed changes in midlatitude storm-track activity in the ERA-Interim for August–November 2012 and 2014 using a 24-h difference filter applied to the 6-hourly pressure-level meridional winds in the southern Pacific Ocean, where $vv = [v(t + 24\,\mathrm{h}) - v(t)]^2$ (Wallace et al. 1988), a diagnostic often used for analysis of storm tracks. Figure 10 shows that South Pacific storm-track activity in 2014 was narrower than in 2012 and was enhanced between latitudes −60° and −50° and between 250 and 500 hPa. These changes are consistent with a decrease in T_s and an increase in the pressure of last saturation and with previous studies of the influence of storm-track variability on subtropical humidity (Hurley and Galewsky 2010b,a).

5. Discussion

Most previous studies that use the last-saturation framework rely, in one way or another, on some sort of air parcel trajectory analysis (Pierrehumbert 1998; Dessler and Sherwood 2000; Galewsky et al. 2005; Cau et al. 2007; Hurley et al. 2012). While there is no explicit representation of air parcel trajectories in our model, the changes in last-saturation PDFs described above do implicitly account for any circulation changes, in the sense that the PDFs are fitted to the data. Thus, the method does not ignore circulation changes despite the fact that circulation is not explicitly included. The potential for circulation changes to affect humidity is

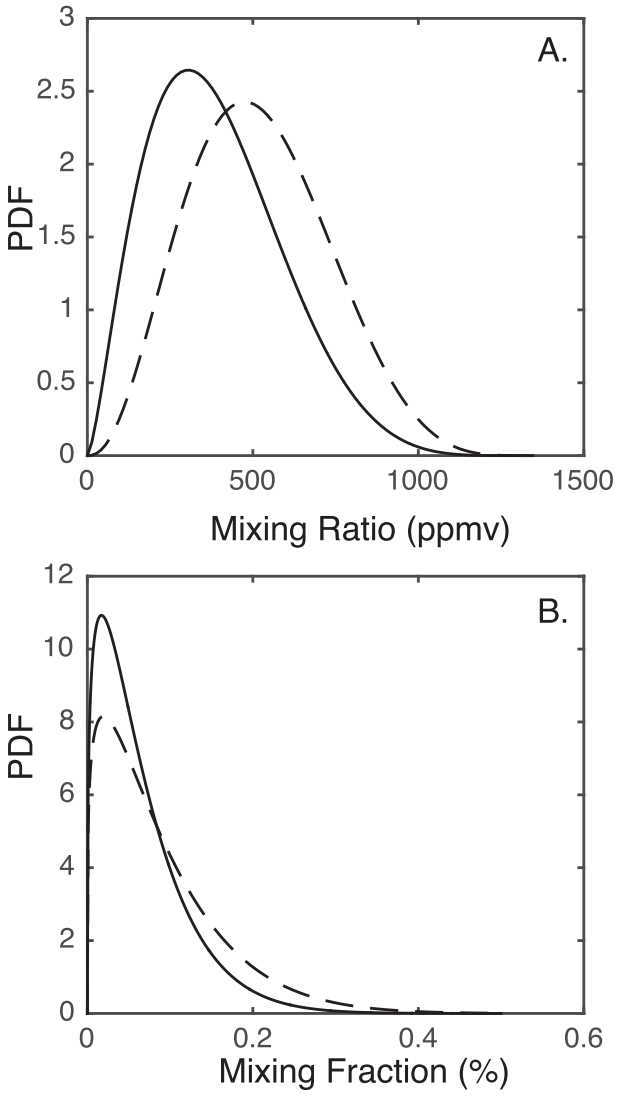

FIG. 8. Estimated probability densities for (a) the last-saturation mixing ratio r_s and (b) the mixing fraction f for 2012 (solid line) and 2014 (dashed line).

important and an understanding of such changes can provide a coherent framework for interpreting changes in humidity. For example, Hurley and Galewsky (2010a) showed how a shift in last-saturation positions to higher and colder regions during El Niño years, relative to La Niña years, could explain the increased aridity in the North Pacific subtropics during El Niño years. In their diagnosis of global warming GCM output, Hurley and Galewsky (2010b) showed that the moistening associated with the water vapor feedback may be partially offset by a shift of extratropical last-saturation sites upward and farther poleward. In the current modeling framework, condensation takes place entirely within a Rayleigh distillation framework along a single temperature profile, and such changes in last-saturation locale would all be mapped onto the single tropical sounding used here. Because we use a profile derived from ERA, we can extract the distribution of last-saturation pressures and diagnose their changes, but those are tied to the tropical sounding. In principle, one could extend this modeling framework to allow for some fraction of Rayleigh distillation to occur along different soundings, and that may allow for improved resolution of changes in last-saturation location; however, because Rayleigh distillation primarily depends on the initial and last-saturation temperatures and only weakly depends on pressure, the potential for fully diagnosing the influence of changes in circulation on humidity using this technique is probably limited.

Postcondensation moistening is modeled here as a mixing process between the freeze-dried air from the upper troposphere with a moist air source ($r_m = 20\,000$ ppmv or $12\,\mathrm{g\,kg^{-1}}$). The selection of this value of r_m is somewhat arbitrary, but it must be fixed in our model in order to avoid a nonunique trade-off between the mixing fraction f and the mixing ratio of the moist source r_m that would make it difficult to compare different parameter estimates. Simulations without any moistening at all (not shown) yield average final mixing ratios of only 410 ppmv and average final δD of $-514‰$, which are completely incompatible with the observations. Clearly, the data require some degree of moistening in order to fit the observations. Our approach is analogous to the approach used in Galewsky et al. (2005), who used a "source tracer" to represent water vapor from the boundary layer that was not processed through clouds. This process was understood in terms of a cross-isentropic turbulent water vapor flux in Schneider et al. (2006) and Couhert et al. (2010) and was modeled in terms of a remoistening process in Cau et al. (2007). Our results require mixing fractions of a few percent and are generally consistent with these previous studies. A potential source of moistening ignored in all of these studies (including the present study) is the evaporation or sublimation of condensate, a process postulated by Sun and Lindzen (1993) as being important for moistening the middle troposphere. The idealized modeling study of Wright et al. (2009) also showed the potential importance of condensate evaporation for setting the humidity of the troposphere. The theoretical study of Bolot et al. (2013) suggest that supercooled droplet evaporation in stratiform anvils should generate strongly negative water vapor d-excess ($-100‰$), an effect we do not see in the current dataset. While we cannot rule out some role for condensate evaporation as a moistening source, our study does not require its explicit inclusion in order to fit the observations.

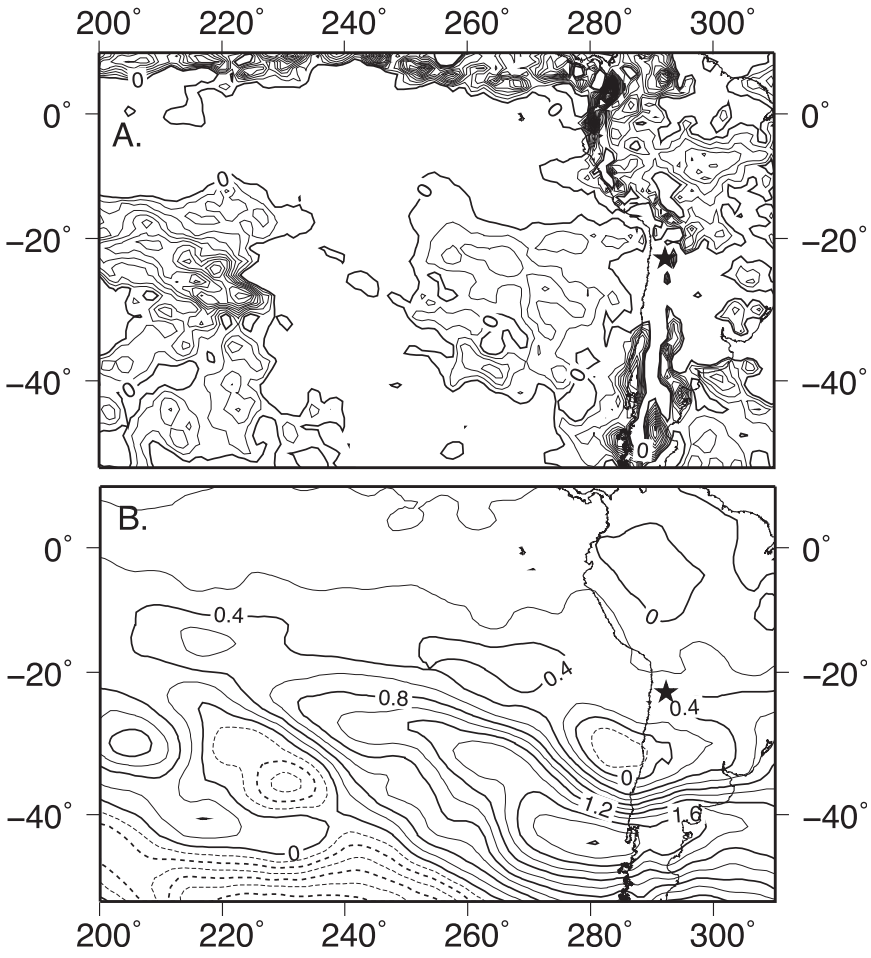

FIG. 9. (a) Difference in August–November average 650-hPa vertical velocity between 2014 and 2012, derived from ERA-Interim. Contour interval is 0.01 hPa s^{-1}; only negative contours (indicating weaker subsidence) are shown. (b) As in (a), but for temperature at 225 hPa; contour interval is 0.2 K and negative contours are dashed. Star indicates location of Chajnantor.

Another potential source of moistening could be sublimated snow or ice from the land surface on the Chajnantor Plateau. There was no snow cover on the Chajnantor Plateau during the periods presented here, so any moistening from the local land surface would have been derived from soil moisture or permafrost. Measurements of local surface fluxes would be a useful complement to the current dataset and would help resolve this particular question, but we do not at present have such data. In any case, none of the results presented here require a local land surface source of water vapor.

We have shown how measurements of water vapor isotopic composition can be used to constrain some key processes that govern free-tropospheric humidity. In particular, the method presented here can be used to constrain the distributions of last-saturation mixing ratio and postcondensation moistening. These are important parameters for diagnosing changes in humidity associated with ENSO and with the water vapor feedback, and we suggest that ongoing monitoring of water vapor isotopic composition may thus complement reanalysis and GCM-based diagnoses and provide a largely independent way to constrain the processes governing free-tropospheric humidity.

6. Conclusions

Our goal in this study was to use measurements of water vapor isotopic composition to constrain the last-saturation and mixing parameters that govern subtropical humidity. We applied a new stochastic model to two dry seasons of data and found that the model can reproduce the observations with reasonable fidelity and that the fitted parameters provide a coherent, physically

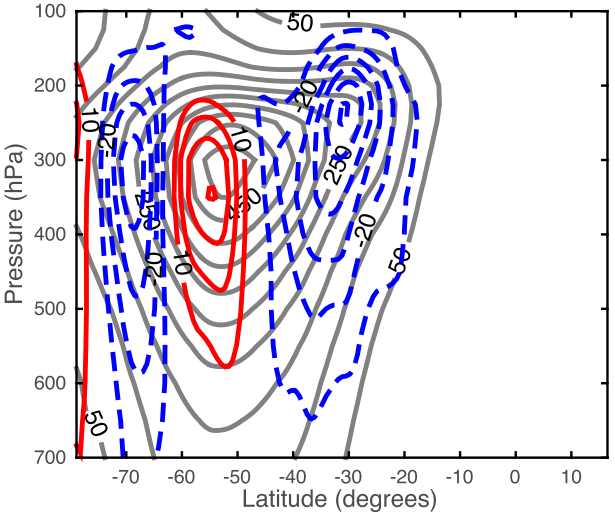

FIG. 10. Changes in August–November storm-track activity between 2014 and 2012 computed from ERA-Interim with a 24-h difference filter applied to meridional winds (vv, $m^2 s^{-2}$). Gray contours indicate an average of the two years; red and blue contours show 2014 increases and decreases, respectively, relative to 2012.

consistent picture of the differences between the two seasons. For the period August–November 2012, we found that the data were best fit by an asymmetric last-saturation distribution with mean r_s of 391 (+45, −75) ppmv [0.24 (+0.03, −0.05) g kg^{-1}], median r_s of 368 (+45, −75) ppmv, with a mean mixing fraction between the freeze-dried air and a moist source of 6.6%(+0.6%, −1.2%). The last-saturation distribution for the August–November 2014 period was less skewed than for 2012, with an average r_s of 520 (+42, −75) ppmv [0.32 (+0.02, −0.05) g kg^{-1}]and a median r_s of 507 (+25, −75) ppmv. The mean mixing fraction for 2014 was 8.8%(+0.8%, −1.8%). The results show that the moistening in 2014, relative to 2012, required increases in both the last-saturation mixing ratio (and hence the last-saturation temperature) and in the mixing fraction between the freeze-dried air and moist boundary layer air. These results show that measurements of water vapor isotopic composition can be used to provide quantitative constraints on some of the key processes governing subtropical humidity and can help disentangle some of the processes that are difficult to diagnose with measurements of mixing ratio alone.

Acknowledgments. This study is based on data collected at the Atacama Large Millimeter/Submillimeter Array (ALMA), an international astronomy facility which is a partnership of Europe, North America, and East Asia in cooperation with the Republic of Chile. We thank the staff of ALMA for their generous support of this project, especially Richard Hills, Joaquin Penroz, and Jim Murray. We also thank Kimberly Samuels-Crow, Dylan Ward, Alex Lechler, Alec Tunner, and Lauren Vargo for field assistance and Chris Rella, Danthu Vu, and Kate Dennis from Picarro, Inc. for their technical support. The data used in this study are available at https://repository.unm.edu/handle/1928/151. This project was supported by NSF-AGS Award 1158582 to JG.

REFERENCES

Beasley, D., R. R. Martin, and D. R. Bull, 1993a: An overview of genetic algorithms: Part 1, fundamentals. *Univ. Comput.*, **15**, 58–69.

——, D. R. Bull, and R. R. Martin, 1993b: An overview of genetic algorithms: Part 2, research topics. *Univ. Comput.*, **15**, 170–181.

Bolot, M., B. Legras, and E. J. Moyer, 2013: Modelling and interpreting the isotopic composition of water vapour in convective updrafts. *Atmos. Chem. Phys.*, **13**, 7903–7935, doi:10.5194/acp-13-7903-2013.

Cau, P., J. Methven, and B. Hoskins, 2007: Origins of dry air in the tropics and subtropics. *J. Climate*, **20**, 2745–2759, doi:10.1175/JCLI4176.1.

Couhert, A., T. Schneider, J. Li, D. E. Waliser, and A. M. Tompkins, 2010: The maintenance of the relative humidity of the subtropical free troposphere. *J. Climate*, **23**, 390–403, doi:10.1175/2009JCLI2952.1.

Dee, D. P., and Coauthors, 2011: The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quart. J. Roy. Meteor. Soc.*, **137**, 553–597, doi:10.1002/qj.828.

Dessler, A. E., and S. C. Sherwood, 2000: Simulations of tropical upper tropospheric humidity. *J. Geophys. Res.*, **105**, 20 155–20 163, doi:10.1029/2000JD900231.

——, and S. M. Davis, 2010: Trends in tropospheric humidity from reanalysis systems. *J. Geophys. Res.*, **115**, D19127, doi:10.1029/2010JD014192.

Duong, T., B. Goud, and K. Schauer, 2012: Closed-form density-based framework for automatic detection of cellular morphology changes. *Proc. Natl. Acad. Sci. USA*, **109**, 8382–8387, doi:10.1073/pnas.1117796109.

Fueglistaler, S., A. E. Dessler, T. J. Dunkerton, I. Folkins, Q. Fu, and P. W. Mote, 2009: Tropical tropopause layer. *Rev. Geophys.*, **47**, RG1004, doi:10.1029/2008RG000267.

Galewsky, J., and J. V. Hurley, 2010: An advection-condensation model for subtropical water vapor isotopic ratios. *J. Geophys. Res.*, **115**, D16116, doi:10.1029/2009JD013651.

——, and K. Samuels-Crow, 2014: Water vapor isotopic composition of a stratospheric air intrusion: Measurements from the Chajnantor Plateau, Chile. *J. Geophys. Res. Atmos.*, **119**, 9679–9691, doi:10.1002/2014JD022047.

——, and ——, 2015: Summertime moisture transport to the southern South American Altiplano: Constraints from in-situ measurements of water vapor isotopic composition. *J. Climate*, **28**, 2635–2649, doi:10.1175/JCLI-D-14-00511.1.

——, A. Sobel, and I. Held, 2005: Diagnosis of subtropical humidity dynamics using tracers of last saturation. *J. Atmos. Sci.*, **62**, 3353–3367, doi:10.1175/JAS3533.1.

——, C. Rella, Z. Sharp, K. Samuels, and D. Ward, 2011: Surface measurements of upper tropospheric water vapor isotopic

composition on the Chajnantor Plateau, Chile. *Geophys. Res. Lett.*, **38**, L17803, doi:10.1029/2011GL048557.

Gettelman, A., E. J. Fetzer, A. Eldering, and F. W. Irion, 2006: The global distribution of supersaturation in the upper troposphere from the Atmospheric Infrared Sounder. *J. Climate*, **19**, 6089–6103, doi:10.1175/JCLI3955.1.

Held, I., and B. Soden, 2000: Water vapor feedback and global warming. *Annu. Rev. Energy Environ.*, **25**, 441–475, doi:10.1146/annurev.energy.25.1.441.

Hurley, J. V., and J. Galewsky, 2010a: A last saturation analysis of ENSO humidity variability in the subtropical Pacific. *J. Climate*, **23**, 918–931, doi:10.1175/2009JCLI3193.1.

——, and ——, 2010b: A last-saturation diagnosis of subtropical water vapor response to global warming. *Geophys. Res. Lett.*, **37**, L06702, doi:10.1029/2009GL042316.

——, ——, J. Worden, and D. Noone, 2012: A test of the advection-condensation model for subtropical water vapor using stable isotopologue observations from Mauna Loa Observatory, Hawaii. *J. Geophys. Res.*, **117**, D19118, doi:10.1029/2012JD018029.

Jouzel, J., and L. Merlivat, 1984: Deuterium and oxygen 18 in precipitation: Modeling of the isotopic effects during snow formation. *J. Geophys. Res.*, **89**, 11 749–11 757, doi:10.1029/JD089iD07p11749.

Karcher, B., 2012: Supersaturation fluctuations in cirrus clouds driven by colored noise. *J. Atmos. Sci.*, **69**, 435–443, doi:10.1175/JAS-D-11-0151.1.

Kennedy, J., 1995: Particle swarm optimization. *Proc. IEEE Int. Conf. on Neural Networks*, Perth, WA, Australia, IEEE, 1942–1948, doi:10.1109/ICNN.1995.488968.

Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi, 1983: Optimization by simulated annealing. *Science*, **220**, 671–680, doi:10.1126/science.220.4598.671.

Noone, D., J. Galewsky, Z. D. Sharp, and J. Worden, 2011: Properties of air mass mixing and humidity in the subtropics from measurements of the D/H isotope ratio of water vapor at the Mauna Loa Observatory. *J. Geophys. Res.*, **116**, D22113, doi:10.1029/2011JD015773.

Peacock, J. A., 1983: Two-dimensional goodness-of-fit testing in astronomy. *Mon. Not. Roy. Astron. Soc.*, **202**, 615–627, doi:10.1093/mnras/202.3.615.

Pierrehumbert, R., 1998: Lateral mixing as a source of subtropical water vapor. *Geophys. Res. Lett.*, **25**, 151–154, doi:10.1029/97GL03563.

——, H. Brogniez, and R. Roca, 2006: On the relative humidity of the atmosphere. *The Global Circulation of the Atmosphere*, T. Schneider and A. Sobel, Eds., Princeton University Press, 143–185.

Rémillard, B., and O. Scaillet, 2009: Testing for equality between two copulas. *J. Multivar. Anal.*, **100**, 377–386, doi:10.1016/j.jmva.2008.05.004.

Samuels-Crow, K. E., J. Galewsky, Z. D. Sharp, and K. J. Dennis, 2014: Deuterium excess in subtropical free troposphere water vapor: Continuous measurements from the Chajnantor Plateau, northern Chile. *Geophys. Res. Lett.*, **41**, 8652–8659, doi:10.1002/2014GL062302.

Schneider, T., K. Smith, P. O'Gorman, and C. Walker, 2006: A climatology of tropospheric zonal-mean water vapor fields and fluxes in isentropic coordinates. *J. Climate*, **19**, 5918–5933, doi:10.1175/JCLI3931.1.

Soden, B. J., and I. M. Held, 2006: An assessment of climate feedbacks in coupled ocean–atmosphere models. *J. Climate*, **19**, 3354–3360, doi:10.1175/JCLI3799.1.

Sun, D.-Z., and R. Lindzen, 1993: Distribution of tropical tropospheric water vapor. *J. Atmos. Sci.*, **50**, 1643–1660, doi:10.1175/1520-0469(1993)050<1643:DOTTWV>2.0.CO;2.

Wallace, J. M., G.-H. Lim, and M. L. Blackmon, 1988: Relationship between cyclone tracks, anticyclone tracks and baroclinic waveguides. *J. Atmos. Sci.*, **45**, 439–462, doi:10.1175/1520-0469(1988)045<0439:RBCTAT>2.0.CO;2.

Wright, J. S., A. H. Sobel, and G. A. Schmidt, 2009: Influence of condensate evaporation on water vapor and its stable isotopes in a GCM. *Geophys. Res. Lett.*, **36**, L12804, doi:10.1029/2009GL038091.

——, A. Sobel, and J. Galewsky, 2010: Diagnosis of zonal mean relative humidity changes in a warmer climate. *J. Climate*, **23**, 4556–4569, doi:10.1175/2010JCLI3488.1.