To Referee 1: Thank you for reviewing this manuscript in detail. We appreciate the suggestions you gave, and think that they have led to considerable improvement of the manuscript. The comments are addressed below; green text is our response to each. Also, we made some general changes in the manuscript: 1. RHic was changed to RHib, and 2. the date rage of data used to produce the cloud populations has changed. The rationale is given below.

General changes:

- RHic (which was meant to stand for RHi climatology) has been changed to RHib (background RHi) for consistency with the terminology used in the discussion.
- In the submitted manuscript, we have used data within 1 Nov 2006 to 30 April 2014. We think that this causes an uneven sampling in season, which could affect the results of Figure 5. For this reason, in the revised version we plan to remake Figure 3-7 and Figure 9 using data within 1 Jan 2007 to 31 Dec 2013, so that the number of sampled months in each season is the same. After making this change, the largest change in percentage is 3% in Figure 3 and 1% in Figure 4. As an example, below is the new Figure 3. The resulting patterns in Figure 6, 7, and 9 are almost identical.

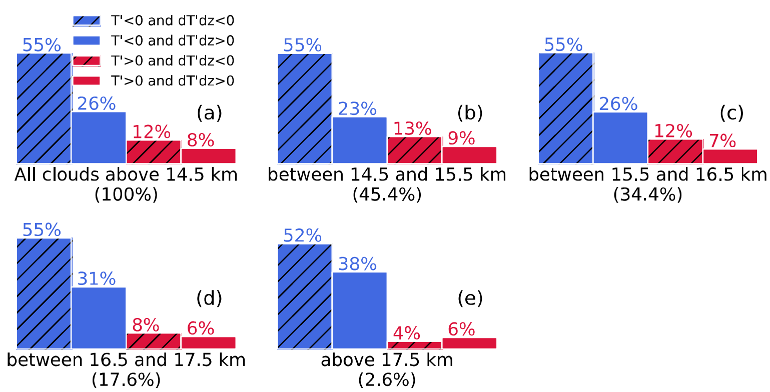

Figure 1. Same as Figure 3 in the manuscript, except made with data within 1 Jan 2007 to 31 Dec 2013.

Specific Comments (by Line Number):

45: Your narrative suggests P18 agreed with K16 in placing the highest cloud frequencies in phase 1 of the gravity wave perturbation cycle. But in P18's figure 3, large ice crystals can also be found in phase 3 (where they are sublimating). Putting aside the effect of background relative humidity, P18's abstract says "The precise location where the confinement (i.e., wave-driven localization of ice crystals) occurs … is always characterized by… a positive vertical wind anomaly." So while P18 agrees that ice is suspended in phase 1, I would mention their finding that upward vertical motions permitted the presence of TTL ice clouds within phase 3 as well.

We have changed the wording of the sentence beginning with "P18 argues …". We added a statement in the first paragraph of Section 4.1 that Phase 3's vertical motion permits the presence of ice.

76: Why is the effective vertical resolution of RO profiles ~200m when temperature estimates are given at 30m spacing? Are the temperature data too noisy, and require vertical averaging? Are you using a version of the RO temperature retrievals that removes fine scale artifacts?

It is probably more appropriate to say that the vertical 'weighting' is ~200m. At each height, the vertical weighting is more or less centered on the reported height in atmPrf, so each height has a slightly different vertical weighting even though they are only ~30m apart. In other words, the RO temperature can be interpreted as a temperature profile with a smoothing window applied, with the width of the window as the vertical weighting.

90: Explain what a "retrieval to flight ratio" means. Is that the ratio of the satellite-retrieved re to that determined from aircraft measurements? How does the +/- 20% uncertainty of retrieved re values compare with those derived from in-situ flight measurements?

Yes, the ratio is the satellite-retrieved re divided by the r_e derived from in-situ 2D stereo probe measurements. The 1.05 is the mean ratio and we have revised the paper to indicate 'mean retrieval-to-flight ratio'. It is hard to know how our choice of +/- 20% uncertainty compare with in-situ measurements. At line 280 we have noted that retrieval of re has large uncertainties especially with respect to thin cirrus., so that the reader is aware of the associated uncertainties in our results pertaining to r_e.

140: How large is the sample of CALIPSO cloud profiles corresponding to each RO profile? In terms of the spatial collocation criteria I'd like to see the results of a sensitivity analysis, where you try both a smaller- and larger-sized diameter range with respect to the RO perigee point. Depending on sample size, I might try 50 km and 200 km and see if your results are affected in any significant way. A larger diameter incorporates more data, which is statistically preferable. But as wave phase surfaces slope with increasing distance from the RO perigee point, a smaller diameter mitigates the possibility of associating the presence of clouds with the incorrect phase. Bound by these constraints, what spatial collocation criteria is optimal?

Figure 2 below shows the distribution for how many CALIPSO profiles tend to be collocated to be each RO profile. There is a sharp cutoff at around 40 profiles. Since each CALIPSO profile has 5-km horizontal resolution, 40*5km is 200 km which is consistent with the 100-km radius collocation criterion.

As you say, a smaller diameter is probably more preferable since it is less likely to be influenced by sloping wave phases. We have done some testing using 50 km and 200 km as a collocation radius. As a reminder, the radius used in the manuscript is 100 km. Please see below for the cloud population in phases obtained using these collocation radii.

For altitudes below 17.5 km, the results from using 50 km vs 100 km (Figure 3 in this document vs Figure 1), the difference is at most 1%. Above 17.5 km, there are larger differences, but the number of samples is quite small at this altitude. Overall, there is good agreement between using 100 km and 50 km, which is reassuring. Results from using 200 km (Figure 4) remains qualitatively similar, but we see 2% differences (53% vs 55% in Phase 1 in panel (a)).

It is hard to objectively assess what collocation criterion is optimal. However, here we show that the results are not particularly sensitive to the collocation diameter. The qualitative features remain largely the same. Since 100 km provides more samples than 50 km, we prefer to maintain the use of 100 km. We will mention that we tested these collocation radiuses and the results had small variations in the text.

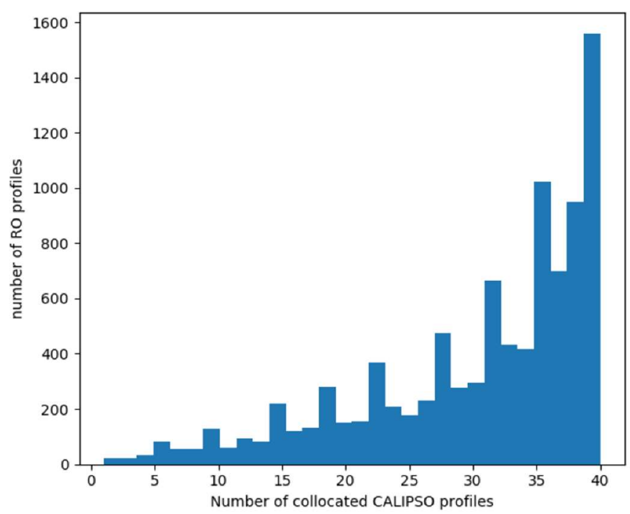

Figure 2. Distribution of the number of CALIPSO profiles collocated to RO profiles. The collocation criteria are 100-km radius and +-2 hours.

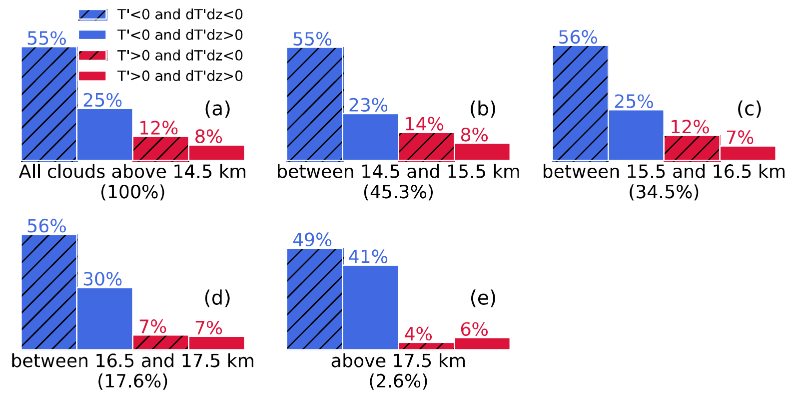

Figure 3. Same as Figure 1 except made with collocation radius of 50 km.

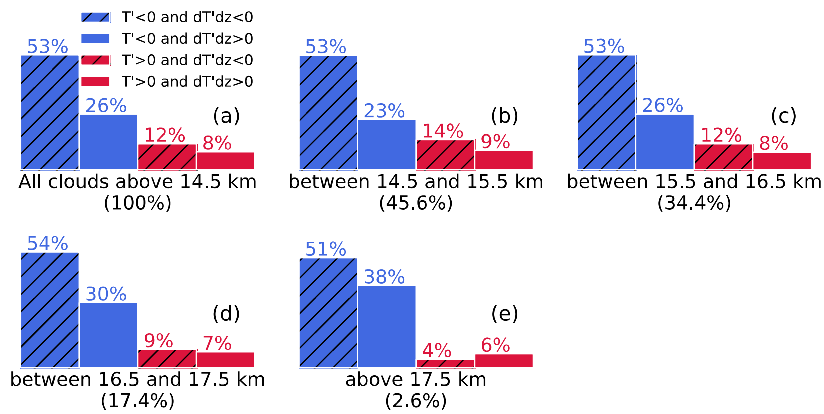

Figure 4. Same as Figure 1 except made with collocation radius of 200 km.

145: What is the mean (and standard deviation) of tropopause altitude in your RO dataset? It would follow that convective detrainment would only occur below this level; is this consistent with your finding that warm phase clouds decrease markedly above 16.5km? Or above an even lower altitude? Does this impact your comparison with K16's cloud fraction results? How so?

The mean and standard deviation of the tropopause within 20N/S of the equator is 17.14 and 0.72 km. For a detailed discussion on the differences of clouds above or below the tropopause, please see our response to the next comment. In summary, below the tropopause we find more clouds in the warm phase, consistent with the expectation that convective detrainment occurs below the tropopause level.

In K16's Figure 5, they showed that the T' of low clouds (below 16 km) associated with convection were distributed evenly across the four phases, while clouds non associated with convection were more common in Phase 1 and 2. Although we are not able to separate convective and non-convective clouds, our results here are consistent with their findings.

155: I think it would be worth re-running this analysis, having deleted altitudes that are below the mean tropopause for each geographical region and season. Does this reduce the influence of deep convection that is likely impacting your results for the western Pacific winter, or the Asian monsoon region during summer? Then either update figure 4 with new "filtered" results or else explain the difference (or lack thereof).

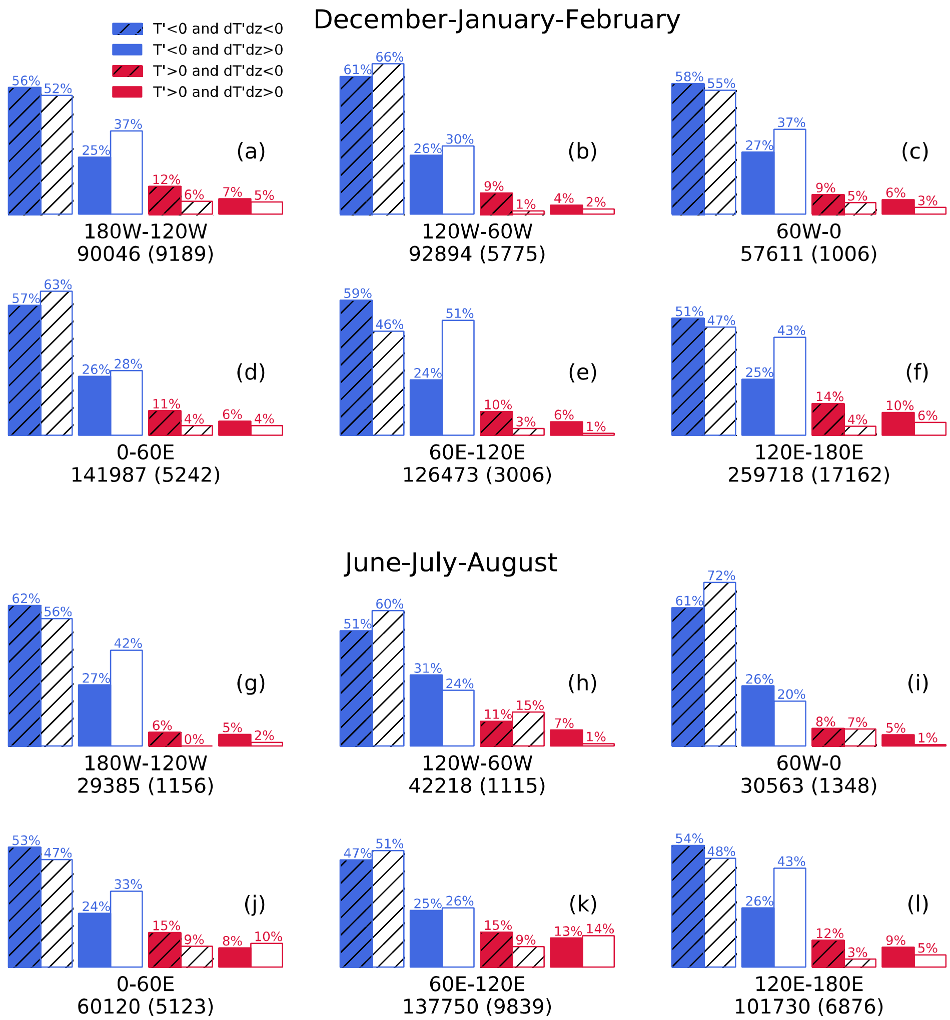

Figure 5: Figure 4 of the manuscript, updated with clouds above the tropopause. The unfilled bards represent clouds above tropopause, while the filled ones represent all TTL clouds. The number below the longitude belt label is the total number of cloud bins for all clouds and clouds above the tropopause (in parenthesis).

Please see Figure 5 for the cloud population after excluding clouds below the tropopause. The mean tropopause for each longitude band and season was calculated using RO. There are generally less clouds in the warm phases and more in the cold phases when looking only above the tropopause. In the western Pacific winter region, there is a significant reduction of clouds in the warm phase. The reduction is less apparent in the summer Asian monsoon. One possibility is that the tropopause is elevated by the upper tropospheric anticyclone (associated with the monsoon), causing outflow clouds to be present at higher altitudes. In Section 4.1 we will add discussion regarding the exclusion of clouds below the tropopause.

165: In regards to your comparison with K16's ATTREX results, I recommend re-computing your RO-derived T' values using a 30-day rather than a 7-day mean temperature. Does making this change result in a significant difference?

Please see Figure 6 below for the cloud population derived using 31-day mean temperature. Compared to the population derived from 7-day means, there tends to be slightly more clouds in the Phase 1 and less clouds in the warm phases. Unlike what K16 found, we do not see that Phase 2 having more clouds than Phase 1 in the 180W-120W band. In paragraph 3 of Section 4.1 we will add a note that we tried using the 31-day mean and explain what we found.

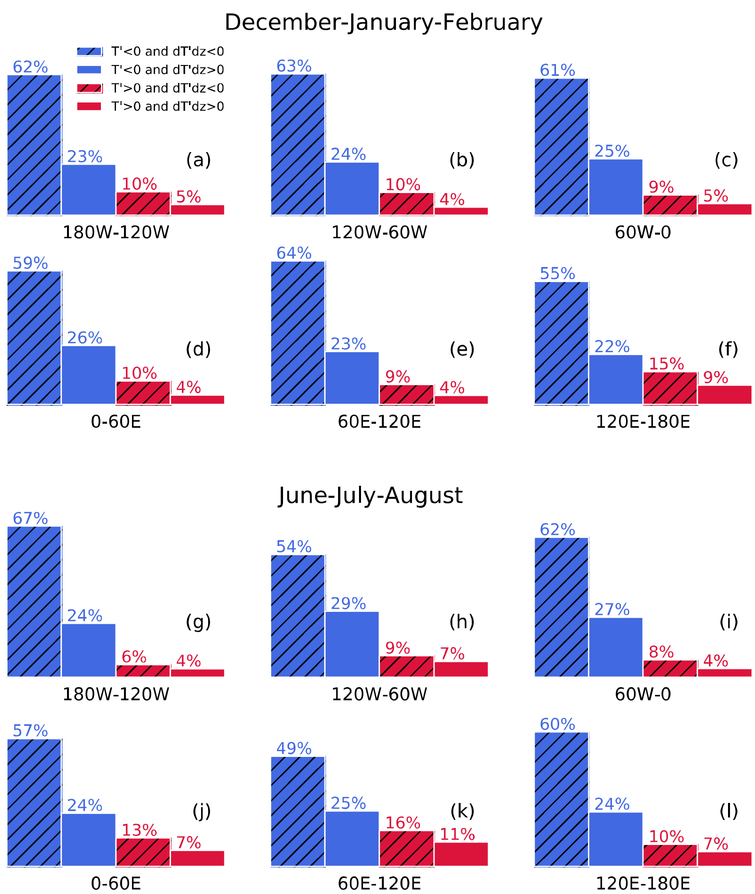

Figure 6: Same as Figure 4 in the manuscript, except derived using 31-day mean temperature as background state.

168: Re-word "line segments in Figure 1 according to each phase." Do you mean "…defined as the amount of vertical overlap between individual cloud boundaries and wave phases?" For example, what if a profile's Phase 1 region extends from 14.8 – 15.2 km, but the cloud layer product indicates a cloud only from 14.9 – 15.1 km. In that layer, is CF = 0.2km/0.4km = 0.5?

Yes, we mean that it is the amount of overlap between individual cloud boundaries and wave phases, as you stated. We have edited the sentence to state this. Your example is correct, in that case the CF is 0.5.

Table 1: I wonder if this is necessary or can you just incorporate this information into the text? If you keep the table, I would (at least) eliminate the CF=0 column, and explain why the number of cases where CF=1 far exceeds the number of cases for which 0 < CF < 1 (shown in figure 5).

We have decided to put the total count of 0<CF<1 in place of the C=0 column, and in the CF=1 column we have added a percentage showing the fraction of instances with CF=1. The new table is shown below. The discussion has been modified accordingly.

Wave Phase	0<CF<1	CF=1
Phase 1	30,519	19,410 (38.9%)
Phase 2	32,152	10,656 (24.9%)
Phase 3	12,687	9,583 (43.0%)
Phase 4	9,784	999 (9.3%)

Table 1: Table to replace Table 1 of the manuscript

As for why CF=1 tends to outnumber cases with 0<CF<1, see Figure 7. We find that the vertical thicknesses of the Phase 1 (distribution in blue) tend to be smaller than the typical cloud thickness (black). Since clouds are thicker, it is likely that the entire Phase 1 is embedded in clouds, hence the higher frequency.

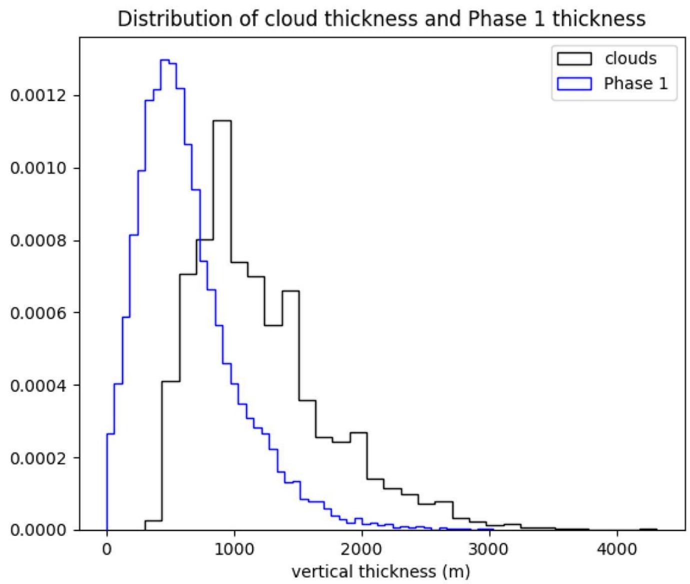

Figure 7: Distribution of Phase 1 vertical thickness (of those used to derive Figure 5 of the manuscript), and the thicknesses of clouds that overlap with Phase 1 (black)

207: Are the weak anomalies of figure 7 simply the result of variability in sampled cloud top heights diluting the overall result? Explain.

Yes, the weak anomalies are a result of cloud top height variability. Since it is common that TTL clouds are embedded in Phase 1, their cloud tops tend to be close to the T' minimum (this is depicted in Figure 6). Since we don't composite with respect to the cloud top height in Figure 7, the result is a pattern of weak anomalies.

221: How can you compute cloud fractions with 50m vertical spacing when the CALIOP 5 km cloud profile product only reports the presence of cloud every 60m?

When the CALIPSO product reports the presence of a cloud in a bin, we are assuming that the cloud occupies the entirety of the bin. Say that the CALIPSO product reports that a bin at height z has a cloud. Then we assume that the region between z0 = z + 30m and z1 = z – 30m is occupied by a cloud. Then, in the grid on which we calculate the cloud fraction (the grid has 50-m spacing), we check the overlap of the cloud boundary [z0,z1] with each bin on the grid. In this approach, the two spacings do not need to match.

Because the reported CALIPSO heights are not on a regular height grid, there isn't much advantage in using a 60-m grid since the same method would have to be used.

238-250: This explanation seems overly detailed and tedious. I would shorten this narrative, focusing on your statistical goals, and eliminate Figure 9(g)-(i), since these panels just look like averaged cloud fractions within the column, and show very minor differences between them. I would attempt to make the remaining 12 panels of figure 9 each a little bigger for visibility.

We have moved the explanation (most of lines 238-250) to the appendix, and in the main text we briefly summarize that we extract the anomalies by building a background cloud fraction. Figure 9(g)-(i) has been eliminated from Figure 9.

272-281: I'm surprised that distributions of re from 2C-ICE data are so uniform. I doubt that the difference between altitude bins is significant. I'd prefer that in figure 10, you keep the results for all TTL altitudes (panel (a)), but delete the vertical stratification presented in panels (b)-(d). Instead, I'd like to see three panels of re distributions from a more limited temporal and spatial domain (but for all altitudes combined), namely: (a) the region from 180W-120W during DJF; (b) from 120E-180E during DJF; and (c) from 60E-120E during JJA. Then as per figure 4, discuss whether convective influence might have an impact on observed ice crystal sizes.

The figure has been made as you specified (see Figure 8 below). The regions associated with frequent convection (120E-180E DJF and 60E-120E JJA) exhibit a lower mean r_e in Phase 1 and 3. The r_e distribution of these phases also tend have a higher peak compared to Phase 2 or 4. However, it is not clear whether this is a result of convection or not. One might have expected to see more large r_e in these regions due to detrainment, yet this is not observed. An interesting feature is that the distributions over 60E-120E JJA are notably narrower than those of 120E-180E DJF, but this feature seems independent of the wave phase. Since this new figure also seems to be inconclusive, we will consider some alternative ways to depict r_e, although it is likely that we will use this figure to replace the Figure 10 of the manuscript and edit the discussion accordingly.

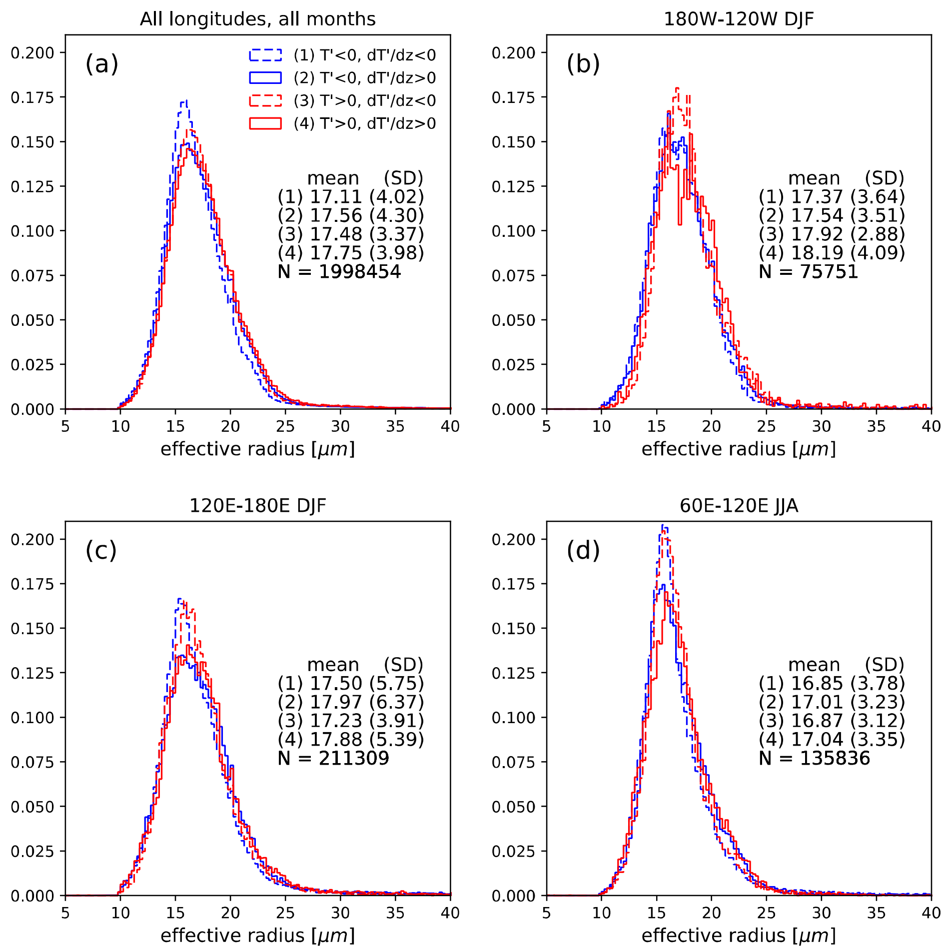

Figure 8: Revised version of Figure 10 to include different regions/seasons as requested. N is the total number of 2C-ICE r_e values in plot.

329: I would interpret this a bit differently. I think there is a clear trend in partitioning between phases 1 and 2 from low RHic values, up to 100%. Thereafter it makes little difference, except for the highest supersaturation values.

This is an interesting point. Figure 11 has been remade into a line plot to better show the cloud population as a function of RHib. RHib is now binned in 0%, 50%, 60%, 70%, …, 180%. The 0% to 50% is grouped together because of the low number of samples. As shown below (Figure 9), there is a rather apparent trend in Phase 1 from 50% to 100% as you noted. On the other hand, Phase 3 and 4 don't seem sensitive to RHib. Phase 2 also exhibits a decreasing trend as RHib increases up to 100%. In the text, the discussion has been edited for this new figure.

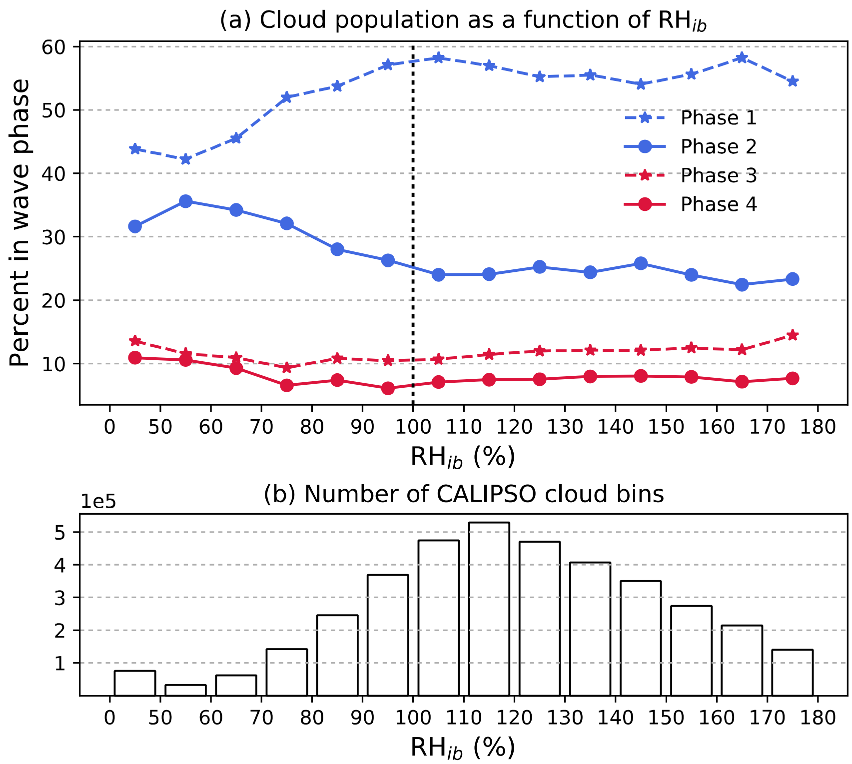

Figure 8: (a) Percentage of clouds in each wave phase a function of RHib. (b) Number of cloud bins in each RHib category. This will replace Figure 11 of the manuscript.

Technical Corrections:

Title: Suggest "Influence of a gravity wave temperature anomaly…"

Title was edited as suggested.

183: Try "-35, -33, … , -1, 1, … 33, 35"

Changed as suggested.

Figure 4: Can you include the number of profiles that contain clouds for each latitude bin?

We have added the number of CALIPSO Cloud Profile bins included in each plot. The number is shown under the longitude label. Note that this isn't the same as the number of profiles since one profile may have multiple bins with clouds.

298-300: Try something like, "We conclude that our analyses of re and RHic data is qualitatively consistent with P18's results."

Changed as suggested.

345-348: Suggest, "…as well as upward vertical motion all impact wave anomalies (…), it remains to be determined whether one has a stronger role in favoring cloud formation."

Changed as suggested.

Figure 11: Add percentage symbols (%) to each of the RHic labels and add the number of cloud encounters each panel represents.

This figure has been changed from bar plots to line plots showing cloud population as a function of RHib (see Figure 9 above). The number of cloud bins in each RHib category is also shown in the bottom plot.