We would like first to thank both referees for their valuable inputs. Many of their remarks proved pertinent, and overall contributed to make the manuscript better.

General comments:

The authors have made an interesting analysis of semidiurnal tides in the HYCOM model using the time-lagged Eulerian and Lagrangian autocovariances of vertical isotherm displacement. They compared output from a 32-day-long HYCOM simulation with Argo park mode data and moored thermistor data, and found that the "total internal tide" variance in HYCOM is too small, especially in the far Southern Ocean. They also found that the Eulerian and Lagrangian estimates of the tidal variance at near-zero lag agree very well using HYCOM data, which bolsters this analysis and their previous analysis of Argo data. Finally, they used Caspar-Cohen's technique to estimate the "intrinsic" and "apparent" decorrelation times of the internal tide, and found that the observed tides decorrelate faster than the HYCOM tides. The authors make the interesting observation that the mean (stationary) IT in HYCOM is too large compared to altimetry, but the total IT (stationary + nonstationary) is too small compared to Argo and moorings.

They discussed some reasons for the discrepancies between HYCOM and the observations, but it was unclear if any of their suggestions could explain the quantitative differences. With regard to the too stationary tides in HYCOM, they did not mention the possible roles of missing small-scale mesoscale or submesoscale variability in HYCOM or the deficit of high-frequency wind forcing.

Overall, this is a nice piece of work which I think will be of interest to many readers of Ocean Science. I have many small comments, listed below. While I would say I have no major concerns, my comments could justify some new analyses or revisions of results presented, so I recommend Major Revision.

Comments:

l1: Is "total" needed? Why not omit or say "tidal"? Throughout the abstract, "total" is used, but it is not contrasted with "partial" or another "non-total" quantity to understand what distinction is implied by "total".
Ans.: Reworked abstract. Suppressed “total”.

l11: "beams" -> "waves" or "beams of waves"
Ans.: Corrected

l17: Omit "at any given position", since later in the sentence you state that you are referring specifically to "their generation site".
Ans.: The sentence is correct: the phase difference accounts for the propagation of the waves from the generation site to any given position, but it is constant in time.

l19: "causes" -> "cause"
Ans.: Corrected

l23-l27: I think I understand what the authors' are getting at, but I found the first three sentences confusing. When we look at the plot of an autocovariance, such as is suggested by the first phrase, we would see the envelope of autocovariance decay, and the coherent fraction of the signal will dominate the autocovariance at long lag. It seems like this paragraph may be muddling the ideas of what happens to the autocovariance as a function of increasing lag, versus what happens to the wave energy as a function of increasing propagation distance. I would suggest re-thinking the purpose of this paragraph and re-writing it to more-clearly articulate the point you wish to make.

Ans.: Reworked paragraph, re-centering it on the stationary/nonstationary wave field.

l42: Once again, "total" is used without distinguishing it properly. It seems like it should be clearly defined above, when the ideas of the coherent and incoherent signals are defined.
Ans.: "total" was defined earlier in the text (l.25). Added italic font.

l52: Finally, "the autocovariance at short time lags", is identified with the "total variance". Some sort of explanation needs to be provided earlier. But how is tidal variability distinguished from noise and high-frequency ocean variability when looking at the "short time lags"?
Ans.: "total" was defined earlier in the text (l.25). Provided short explanation on how noise is filtered out.

l52: "On the other hand" -- I am not clear what is the other part of the contrast. Omit this phrase?
Ans.: Contrasts the Lagrangian decorrelation downside with the fine time resolution upside.

l55: I am not sure "intrinsic" is the right word. An intrinsic quality ought to be one which is unaffected by extrinsic factors. But, the decorrelation is entirely caused by interactions with the propagation medium. Perhaps it is best to stick with the Eulerian vs Lagrangian distinction, and when the autocovariance is discussed, it seems like you need to be clear whether you are discussing an Eulerian or Lagrangian autocorrelation.
Ans.: Agreed. Replaced "intrinsic decorrelation" by "decorrelation" or "decorrelation of the IT", and further "apparent decorrelation" by "Lagrangian decorrelation", everywhere.

l63: omit "the strength of"
Ans.: Corrected

l73: "can vary" by how much?
Ans.: Replaced by "The sampling period of the park phase can occasionally vary by more than a few seconds." The vast majority of the park phases we use have a sampling period of 1h. Very rarely do park phases have sampling periods significantly shorter (and even more rarely longer) than one hour.

l81: Did you use exactly the same dataset as in Geoffroy and Nycander (2022)? I would be interested to know how many 32-day records there are, from how many individual drifters. Also, can you remind us exactly what the "data" consist of? Is it time series of isopycnal displacement, inferred from temperature measurements during the park phase, using temperature profiles from the start and end?
Ans.: Added information at the beginning of section 2, and in section 2 and 3 (isotherms displacement is properly defined in section 3).

l111: Is dT/dP in the numerator the same as dTbar/dP?
Ans.: No: Tbar in the numerator is not a function of z, and they are calculated independently.

l114: "obtained" -> "estimated"
Ans.: Corrected

l123: Sorry if I misunderstand what is meant by "unbiased" here, but isn't this a biased estimator when the expected value is taken for fixed N?
Ans.: This estimator is unbiased: for increasingly large N, the estimator converges to the true value.

l140: I do not understand why the sine component is included. The autocovariance is an even function, so any projection onto the sine must be noise, right ? [...]

Ans.: Indeed, the acov is an even function. However, when two tidal constituents close in frequency are present, the autocovariance of the resulting beating can be expressed in terms of a cosine and a sine component (short derivation below).

Consider a tidal variability as a sum of two tidal constituents:

$$h(t) = h_0(t) + h_1(t) = A_0\cos(\omega_0 t + \phi(t)) + A_1\cos(\omega_1 t + \phi(t)).$$

Here ω_i and A_i are the angular frequency, and the amplitude of the constituent i, respectively. $\phi(t)$ is an AR1 process representing random phase modulations (for simplicity, taken to be identical for both constituents). One can show that the autocovariance of h is the sum of the autocovariance of h_0 and h_1 (i.e. the crosscovariance of h_0 with h_1 is 0). Denoting $R(\tau)$ the autocovariance, following Geoffroy and Nycander, 2022, we have:

$$R_h(\tau) = R_{h_0}(\tau) + R_{h_1}(\tau) = \exp(-\sigma_\phi^2 + R_\phi(\tau))\big(\frac{A_i^2}{2}\cos(\omega_i\tau) + \frac{A_j^2}{2}\cos(\omega_j\tau)\big).$$

By noting that $A + B = A\cos(\frac{a+b+a-b}{2}) + B\cos(\frac{b+a+b-a}{2})$, expanding and recollecting the terms, one can rewrite the expression for $R_h(\tau)$ as:

$$R_h(\tau) = \exp(-\sigma_\phi^2 + R_\phi(\tau))\bigg(\frac{A_i^2 + A_j^2}{2}\cos(\frac{\omega_i+\omega_j}{2}\tau)\cos(\frac{\omega_i-\omega_j}{2}\tau) - \frac{A_i^2 - A_j^2}{2}\sin(\frac{\omega_i+\omega_j}{2}\tau)\sin(\frac{\omega_i-\omega_j}{2}\tau)\bigg),$$

here $\frac{\omega_i+\omega_j}{2}$ is the mean frequency and $\frac{\omega_i-\omega_j}{2}$ the beat envelope.

[…] Likewise, I don't understand the total error defined in equation (4). And why would a robust estimator (median for \tilde{SEM}) be combined with a non-robust estimator (Var A)?
Ans.: The definition of the total error in equation (4) was not rigorous. Referee #2 also pointed at the wrong assumption of Gaussian statistics when computing the confidence interval of the complex demodulate. For both these reasons, we now use a Monte Carlo method to estimate the confidence interval of the complex demodulate.

l157: "not significantly different" -- Well, I agree that they do fall within each others' standard errors, but they look significantly different to me. What is the probability of the offset over so many different lags; how many d.o.f. do you think are in these estimates?
Ans: Modified the sentence: "Apart from the first couple of demodulates, the HYCOM demodulate series appears consistently smaller than the Argo one." The number of d.o.f. is not straightforward to estimate, mainly because one first has to estimate how many independent values our 767-h records contain. Eq. (24) in Awe's 1964 paper "Errors in correlation between time series" gives a way to compute such an estimate. However, since we do not know the underlying true autocovariance function, we cannot precisely evaluate this. Using the local mean autocovariance computed from Argo data for the local example shown in Fig. 2, we estimate that values separated by L~12 h may be considered independent of one another. Hence, when computing the mean autocovariance at a given \tau, we have roughly N_p(N-\tau)/L d.o.f., with the number of 32-day records N_p=8 and the number of values in each record N=767. For small \tau, we have approximately 8*767/12 ~ 500 d.o.f. For the corresponding Lagrangian data from HYCOM we get ~250 d.o.f. (here L~39 h). Thus, they have largely enough d.o.f. to be considered different. The new confidence interval estimates reflect that conclusion.*

l170: Can you explain why you estimated the Eulerian autocovariance along the Lagrangian trajectory? Are you trying to account for the geographic variability of the Eulerian autocovariance?

Ans.: Precisely. That way we are not introducing any discrepancy due to the (random) Lagrangian spatial sampling. Added explanations in the text.

Fig 3: This is related to the above question: Why are the Eulerian errorbars so small compared to the Lagrangian? Maybe you could spend a little more time explaining how this plot relates to Fig 2. Are the Lagrangian HYCOM curves in Fig 2 identical to those in Fig 3?
Ans.: The Eulerian mean autocovariance uses many more d.o.f. There are about 60 times more 32-day segments that are used to compute the Eulerian mean autocovariance resulting in about 4600 d.o.f. (L is also larger, i.e. the data less independent). Added explanations in the main text. Added in the caption of Fig. 3 that the Lagrangian HYCOM curves in Fig. 2 and 3 are identical.

Fig 4: Maybe use the same color for the HYCOM curves in each plot? Is the red curve in Fig 3b the same as the black curve in Fig 4b? They seem to both be labelled as demodulates of the HYCOM Eulerian autocovariance, but they seem to have different numeric values (R(740hr) < 10m^2 in Fig 3b, but R(740hr) > 10m^2 in Fig 4b).
Ans.: The color choice we made is to have the HYCOM Lagrangian data plotted in black, and the in situ data in red whenever possible. The red curve in Fig 3b should indeed be the same as the black curve in Fig 4b (Added text in the caption of Fig.4). There was an error in the plotting script.

l190: "for each particle" -> "for each HYCOM particle" ? But if this paragraph applies to HYCOM, how can there be outliers?
Ans.: Replaced "for each particle" by "for each HYCOM particle". \eta_1000 values can be unstable when facing small temperature gradients in the denominator of Eq. (1) and (5), leading to unrealistically large variance. For consistency with the Argo results, we use the same quality checks on the variance of \eta_1000 computed using HYCOM data. Added explanations in the text.

l193: Are the HYCOM Eulerian autocovariances computed all along the trajectory? This would seem to use so many more degrees of freedom compared to the Lagrangian estimates. I am not sure why this is done or why it would be justified. Paragraph at line 195: This is a very good comparison and a little surprising, to me.
Ans.: Yes, the HYCOM Eulerian autocovariances are computed all along the trajectory subsampled every 12 h. As written before, the HYCOM Eulerian autocovariance estimates do use many more degrees of freedom. As a result the formal error is much smaller.

Fig 6: Why are the maps drawn so small?
Fig 7 + 9 + 13 + 14: Please enlarge the maps and panels.
Ans.: Enlarged Fig. 6, 7, 8, 9, 13, 14.

l215: The "fairly constant" ratio is not apparent to me in Fig 7c. Should it be?
Ans.: The ratio itself is not plotted (added as a comment in the main text), however we write its mean value and standard deviation in the text.

l225: Would it be fair to guess that there are also many Argo trajectories that were excluded in the Southern Ocean due to the 0.1m/s drift speed cutoff criterion? I wonder if you would see the difference in HYCOM vs Argo if you made a Fig 8 based on the drift speed?
Ans.: This criterion removes roughly 3000 Argo segments (~15%), but the final maps only get ~1% less coverage. This mostly affects the Argo data density in the regions east of Drake's passage, east of Algulas, and in the Equatorial Pacific. Below is a figure similar to Fig. 8 but based on the mean drift speed: the Argo mean autocovariance remains more or less the same while the Lagrangian HYCOM autocovariance is significantly smaller for drift speeds smaller than 0.033 m s-1. Coincidentally, the HYCOM bins with a mean speed smaller than 0.033 m s-1 are mostly poleward of 50 deg S.

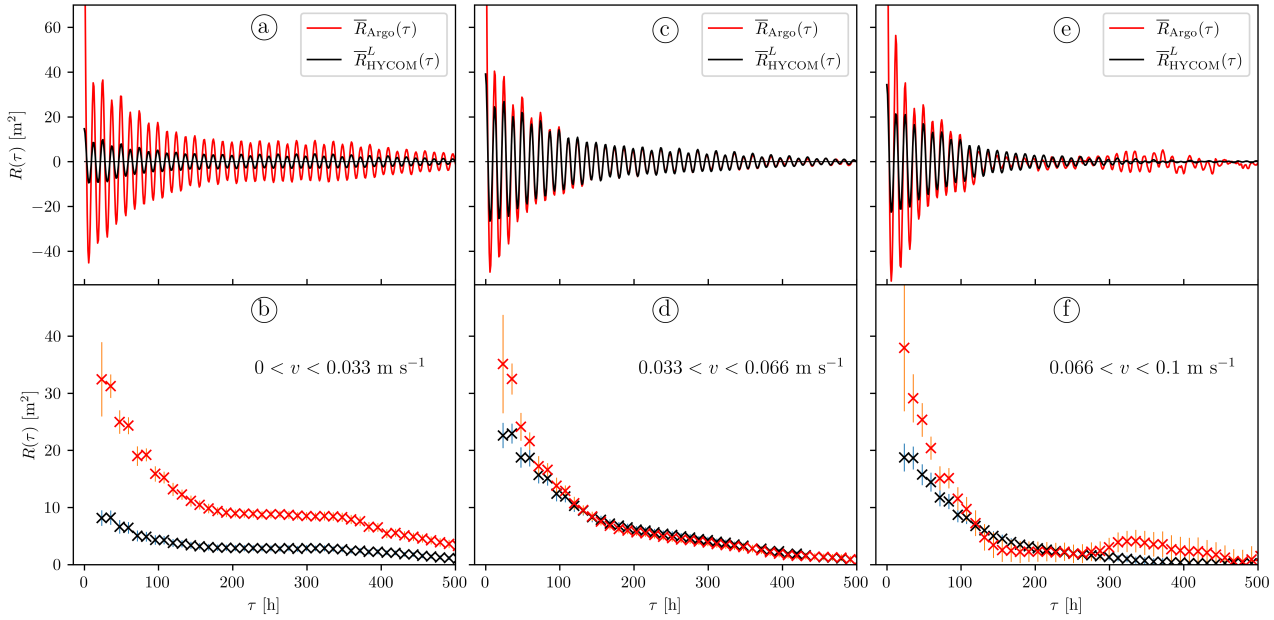

l227: Once again, "intrinsic" does not seem to be the right word.
Ans.: Deleted "intrinsic".

l235: Is this an expected property of the Rayleigh distribution for the pdf of the modulated wave amplitude? You might want to look into this in the acoustics or optics literature. I don't believe this has been observed previously for narrowband ocean internal waves.
Ans.: As pointed at by referee #2, the complex demodulates are not Gaussian distributed. Following, our definition of the complex demodulate, if the fitted C and S are Gaussian random variables with a same standard deviation, then A = sqrt(C^2+S^2) follows a Rice distribution. We did not give more thought about the expected distribution of the local mean autocovariance.

l247: Why are you comparing the SCVF_15 statistic? This is a ratio of sample statistics and likely to be very noisy. I don't really know what to make of Fig 9a. With such a small dataset, I would like to see the ratio of total variance (the demodulate amplitude at tau=48hr), instead.
Ans.: Indeed the SCVF_15 statistic from local mean autocovariance was very noisy. We now only compute it from larger populations of sample autocovariances. Fig 9 is now showing the variance instead of SCVF_15.

l260: Previously (in the Argo vs HYCOM comparison) you used the ratio of the demodulate amplitude at 48hr lag. Why not use that same quantity for comparison? Oh -- I see it in Table 1.
Ans.: -

l262: I am not sure what the "discrepancy" refers to.
Ans.: Poor writing, rewrote the sentence.

l265: SCVF^{15} -> SCVF_{15}
Ans.: Corrected

l273: I don't think "no impact" is the correct way to characterize the previous results. There is considerable scatter in Fig 5, and Fig 3b shows that the estimates differ. Also, it is unclear to me why you don't try to make the estimates more consistent by extrapolating the demodulate amplitude to zero lag.

Ans.: Replaced “no impact” by “no significant impact”. We do not know what is the true autocovariance function and how it behaves close to 0 time lag. The first demodulate is a conservative, simple, and robust estimate of the IT variance. It is also consistent within this work.

l275: "mecanism" -> "mechanism"
Ans.: Corrected

l276: Why not just call it "Lagrangian decorrelation" instead of "apparent decorrelation"? If I had been a reviewer on Gaspar-Cohen, I would have made the same suggestion.
Ans.: Replaced "apparent decorrelation" by "Lagrangian decorrelation", everywhere.

l306: "sinusoide" -> "sinusoid"
Ans.: Corrected

Table 2: It is interesting that the \omega_{AM} frequency corresponds to M2-S2 beating, but the amplitude (\sigma^2_{AM}) does not.
Ans.: The M2-S2 beating, although probably dominating the semidiurnal amplitude modulation, is definitely not the only contribution to this amplitude modulation (other constituents close to M2 play a role).

l330-358: Modes discussion.
Ans.: -

l364-377: Bathymetry discussion. Surely the importance of the errors depends on the horizontal spatial scale of the errors. While this is interesting discussion, it ought to consider the wavenumber spectrum of the error.
Ans.: We do not have quantitative information on this error spectrum.

l386-401: Stratification. None of these discussions really deal with the overall quantitative difference of HYCOM vs obs which is about 0.74 (HYCOM/Argo) or 0.51 (HYCOM/Mooring) equatorward of 50 deg. Both datasets have problems in terms of their spatial coverage, but the Argo comparison seems much more meaningful. I am unclear which of the authors' proposed sources of bias could account for the 26% deficit compared to Argo.
Ans.: Added summary sentence at the end of Sect. 4.

l403: "run with" -> "run of"
Ans.: Corrected

l416: Shouldn't the factor of 1.5 mentioned here equal the reciprocal of the 0.74 value at l216? Have I misunderstood this?
Ans.: This factor of 1.5 was already mentioned l.225. Contrarily to the 0.74 value at l.216, it does include latitudes north of 50 deg N. Rewrote with factors consistent throughout the section. Added table 1 to summarize the statistics for the different groups.

l422: "stationnary" -> "stationary"; also, I think the "big O" notation should be reserved for asymptotics, and here it is better to say "about" or "approximately". Finally, I don't think "becomes stationary" is an appropriate descriptor; this would be like saying a time series "becomes its mean".
Ans.: Corrected. Replaced “the IT becomes stationary” by “the IT autocovariance reaches its stationary limit“.

l429: "unaffected" -> "unaffected in the mean"
Ans.: Corrected

l339: "supposedly account" -> "supposedly accounts"
Ans.: Corrected

l452: latex formatting needs help in the URL.
Ans.: Corrected

l454: Zaron's URL has changed to https://ingria.ceoas.oregonstate.edu/~zarone/downloads.html
References: inconsistent capitalization is used in article titles
Ans.: Corrected

l557: "and contributors, T. P."
Ans.: Corrected