Dear Editor Katsumata,

We thank you for your careful assessment of the revised manuscript and for acknowledging that its physical consistency has improved. We also thank Reviewer #1 for the additional clarification and for indicating that they would not object to publication in Ocean Science.

We agree that the four main points raised by Reviewer #1 were provided indirectly or across different sections, and we now centralize them. We appreciate that a clearer format where we present them more directly will improve the traceability of the review process and make the scientific discussion easier to follow.

We first provide direct responses to the four comments highlighted by Reviewer #1, and then address each of the remaining editorial comments in the order in which they were raised.

Reviewer comment 1 — Mean-flow lee waves

"First, the author cited the well-known diagram of internal waves generated by "tide"-topography interactions (St. Laurent and Garrett 2002, note this is incorrectly cited as St. Laurent et al., 2012) to argue that tidally-driven lee waves are not significant. However, this does not address my original concern. My comment specifically referred to the potential generation of lee waves by "mean flows", not tidal flows. My original specific comments (Fig. 1: It would be helpful to overlay the circulation pattern on the map. L113-124: are tidal flows much weaker than the mean flow?) relates to this point. Without a clear response to these points, I remain unconvinced by the authors' argument."

We thank the reviewer for clarifying that the concern was specifically about lee waves generated by mean flows over topography, rather than by tidal flow. We agree that our previous response was too focused on tide–topography interactions and did not make this distinction sufficiently explicit.

To address this point, we added the regional circulation pattern to Fig. 1a and explicitly compared the likely role of tides and mean flows. We also included a new paragraph that report some litterature review (see lines 140-143). We then performed a direct diagnostic of mean-flow/topography forcing by estimating a characteristic topographic forcing frequency as $\omega = |U| \, k_{eff}$, where $|U|$ is taken from the observed LADCP velocities and k_{eff} is derived from the effective topographic scale estimated from the local slope and detrended bathymetric roughness. This analysis is now reported in the Supplementary Material (Section 3 and Fig. S4, reported here below):

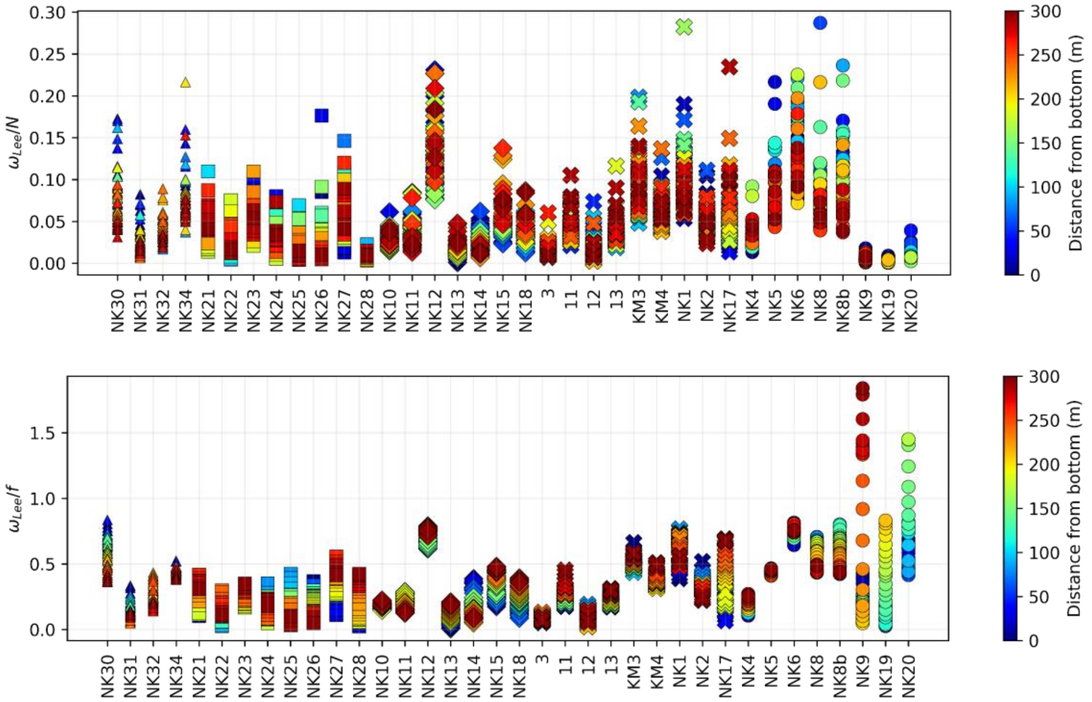

Figure S4. *Normalized characteristic topographic forcing frequency estimated from the interaction between the observed flow and bottom topography. The upper panel shows the forcing frequency normalized by stratification, while the lower panel shows the same quantity normalized by the Coriolis frequency. Values are organized by station and colored according to distance from the bottom. Most values remain well below the buoyancy frequency and largely within the subinertial regime relative to the Coriolis frequency, suggesting that the observed flow–topography interaction does not predominantly fall within a classical lee-wave generation regime.*

The results show that ω/N remains small, typically below 0.2–0.25, while ω/f is mostly below 1, with only a few localized values approaching or slightly exceeding the inertial frequency. Thus, for most stations and near-bottom bins, the mean-flow/topography interaction lies outside the classical freely propagating lee-wave regime. This does not exclude that weak or localized flow–topography interactions may occur, but it indicates that mean-flow lee waves are unlikely to be the dominant mechanism controlling the observed basin-scale variability of Rω*.

We therefore retain our interpretation that the observed patterns are more consistently explained by the modulation and redistribution of the internal-wave field by bottom slope and roughness, rather than by dominant local lee-wave generation by the mean flow.

Reviewer comment 2 — Δ-slope, β, and the single-frequency assumption

"Second, on the Δ-slope, I am convinced that the bathymetry slope does affect the calculation of Δ-slope. However, the validity of calculating β using the single-wave approximation is not addressed at all. Moreover, it remains unclear how the criticality of the bottom slope affects internal-wave frequencies (and Rω). This is related to the unanswered comment (L435-438: I think the description is not accurate. The internal-wave frequency or propagating angle doesn't change during the reflection, but wavenumber changes)."

We thank the reviewer for this clarification. We agree that the intrinsic frequency and propagation angle of an individual internal wave are conserved during reflection, while the wavenumber content may change through reflection, scattering, and redistribution of energy across vertical scales. We have corrected the manuscript accordingly (see lines 428-432).

We also agree that the classical relation between Rω, frequency, and propagation angle is exact only for a monochromatic or narrowband wave field. Our estimates are instead based on shear and strain variance integrated over finite spectral bands. We therefore now denote them Rω*, ω*, and β*, to make clear that they are broadband statistical diagnostics, not true single-wave frequency or ray-angle estimates (see lines 441-444).

Accordingly, Δ-slope is not used to infer the criticality of a unique internal-wave ray. It is used as an effective geometric index comparing the observed shear–strain variance partition with local bathymetric slope. This allows us to organize the observations statistically and assess how the internal-wave field varies with bottom geometry.

To test that this diagnostic is not an artifact of the chosen spectral bands, we added a sensitivity analysis using alternative shear and strain integration ranges, including same-band configurations (see Section 2 of the Supplementary Material). The results remain coherently organized across configurations, supporting the robustness of the observed patterns while avoiding any single-frequency interpretation.

Reviewer comment 3 — Roughness versus steepness

"Third, regarding the calculation of topographic roughness and steepness: I acknowledge the authors' clarification that the roughness derived from the GEBCO product is relatively insensitive to horizontal scale and correlates well with the topographic slope. However, roughness and steepness represent distinct physical properties and are expected to play different roles in internal wave dynamics. At present, I am confused about how each of these two parameters individually influences the observed variations of Rω. The manuscript does not provide a clear physical interpretation that distinguishes their respective effects."

We thank the reviewer for pointing out this important distinction, which we clarify more explicitly in the manuscript (see line 510-516): "*Bottom slope and roughness are therefore interpreted here as complementary, but not equivalent, descriptors of the seafloor. Slope primarily describes the large-scale geometric relation between the internal-wave field and the boundary. Roughness, instead, describes smaller-scale bathymetric variability able to scatter the wave field and redistribute variance toward shorter vertical scales. Because both quantities covary in our study region, we do not attempt to isolate their independent causal contributions; rather, we interpret roughness as a modulation of the slope-controlled geometric framework.*".

In this framework, slope provides the effective geometric control captured by Δ-slope, while roughness modulates the shear–strain partition within this framework. In particular, rougher regions tend to show lower and less variable Rω*, consistent with a more strain-dominated internal-wave field and enhanced redistribution toward smaller vertical scales.

We will also revise the Discussion to avoid formulations that could imply that slope and roughness act identically (see lines 566-577).

Reviewer comment 4 — High-wavenumber strain spectra and low-pass filtering

"Lastly, regarding the blue strain spectra at high wavenumbers: I acknowledge that the potential inclusion of double-diffusive layers, excluded from the analysis, could contaminate the strain spectra. The authors now present strain spectra restricted to the data used in the analysis; however, they apply a low-pass filter that artificially suppresses high-wavenumber signals. This approach does not fully address my concern. I would expect to see the behavior of the strain spectra at high wavenumbers without the application of a low-pass filter, in order to properly assess whether the used signals are not contaminated. For comparison, Ferron et al. (2014), which the authors argue to follow, present shear and strain spectra without applying such filtering, and their strain spectra do not exhibit a blue slope at high wavenumbers. This point needs to be explicitly addressed; otherwise, I remain unconvinced by the interpretation of the high-wavenumber strain spectra."

We thank the reviewer for this clarification. This point has been explicitly addressed in our previous response by providing a direct comparison between unfiltered and low-pass-filtered spectra (see Fig. R2 of the previous Revision Notes, reported below).

This comparison shows that the low-pass filter only affects the highest vertical wavenumbers, corresponding to wavelengths shorter than ~10 m. These wavelengths are outside the integration range used to compute Rω*, since strain variance is integrated only down to the 10 m wavelength limit. Therefore, the high-wavenumber blueing visible beyond this limit does not enter the Rω* estimates we use in the manuscript.

In the previous Revision Notes (see "Editor comment #2") we also clarify that the low-pass filtering was not used to hide or modify the variance entering the Rω* calculation. The calculation is controlled by the prescribed spectral integration limits, while the additional filtering only affects wavelengths shorter than the retained strain band. The unfiltered/filtered comparison demonstrates that the spectral structure within the retained integration bands is unchanged, and that our conclusions are not controlled by the filtered high-wavenumber part of the spectra.

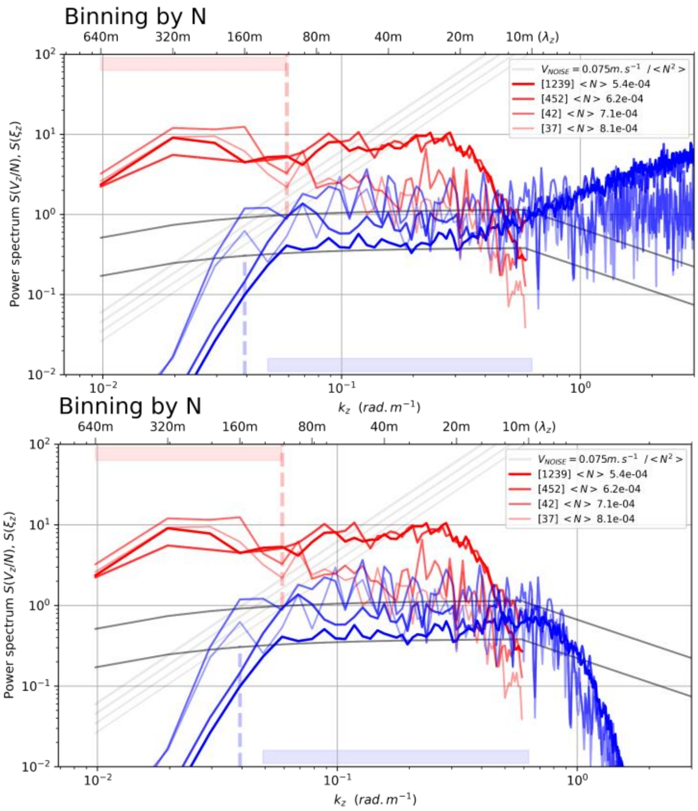

***Fig. R2.** Comparison between unfiltered (top) and low-pass filtered (bottom) [0 - 960m] from bottom spectra. The filtering only affects the highest vertical wavenumbers corresponding to wavelengths shorter than ~10 m, while the spectral ranges used for the R_{ω^*} computation remain unchanged. This demonstrates that the filtering does not alter the structure of the spectra within the integration bands used in the analysis.*

Editor comment 1 — Notation Rω*, ω*, and β*

"L.441, If the symbols with asterisk are to be used, this should be applied to Eq.(6), too. The statement is also inconsistent with what is written in L.428-431 "from which we estimate a local value of ω in the water-column, given $N(z)$ and the local f, that we assume representative of a certain regional extent in the horizontal. From ω we can estimate β perpendicular to α. The relation between $R\omega$, ω and β will be exploited later in our study to associate the beam angle with $R\omega$ value and compare it with the bathymetrical slope to infer propagation insights."

We thank the Editor for pointing out this inconsistency. We will therefore revise Eq. (6) and all related text so that the diagnostic ratio used in this study is denoted Rω*. Similarly, the quantities derived from this diagnostic mapping will be denoted ω* and β*, not ω and β, whenever they refer to the broadband statistical estimates used here. See new Equation (6) and lines 352-356.

The statement is corrected accordingly: L.428–431:

We revised this part. See lines 428-432

Editor comment 2 — Figure 3 panel references

"L.511-, Please doublecheck the panel for Fig.3 in this paragraph. E.g. Figure 3b (L.514) → Figure 3c"

We thank the Editor for noticing this typo. The panel references have been corrected.

Editor comment 3 — Use of Rω* in the finescale parameterization

"L.554, "Values for $R\omega*$ are necessary in the finescale parameterization" The value of unstarred $R\omega$ are necessary but use of $R\omega*$ should be avoided (Polzin et al., 2014, P.1398)."

We thank the Editor for this important clarification. We agree that the finescale parameterization is formally written in terms of the classical, unstarred Rω, and that our broadband diagnostic Rω* should not be presented as a direct substitute without qualification. We revise the sentence to avoid implying that Rω* itself is a required parameter of the formal finescale parameterization. We also revised the surrounding discussion to avoid suggesting that Rω* is directly used to derive dissipation rate or diapycnal diffusivity in this study.

See lines 559-563.

Editor comment 4

Supplementary material
Figure S1ab, The confidence interval for the bootstrap is not found -- 68%? 99%?

We thank the Editor for pointing this out. The bootstrap confidence interval shown in Figure S1a,b and S2a,b corresponds to a 95% confidence interval. This was mentioned in the statistical tables, but not stated clearly in the figure captions. We revised the captions of Figures S1a,b and S2a,b to specify: "The grey shaded area represents the 95% bootstrap confidence interval of the logarithmic spectral difference."; and also revised the explanatory text before Figure S1a,b to state:" The uncertainty associated with the logarithmic spectral differences is reported as the 95% bootstrap confidence interval."

Editor comment 5

"Figure S1ab, left panels, significant wave numbers are found for both B>A and A>B in one spectrum, sometimes in a neighbouring wavenumbers. It is unclear how to interpret these."

We thank the Editor for pointing this out. In Figures S1a,b, the green markers indicate wavenumbers where the difference between two spectra is statistically significant. Positive values mean that the second spectrum has more energy than the first one, while negative values mean the opposite. Therefore, green markers can appear both above and below zero when the two spectra cross, or when the difference changes sign with vertical scale. These markers should be interpreted as local significant differences, not as a single continuous physical trend. For this reason, the physical interpretation is based on the band-integrated statistics reported in Tables S1 and S2, which summarize the spectral differences over broader wavelength ranges. We now clarified in the text (see Section 1, second paragraph): "Green points may occur on both sides of zero when the two spectra cross, or when the sign of the spectral difference changes with vertical scale; they should therefore be interpreted as local significant differences, while the physical interpretation is based on the band-integrated statistics reported in Tables S1 and S2."

Editor comment 6

"Table S1, S2, "large scales 640-80m", "intermediate scales 160-80 m" -- why overlap?"

We thank the Editor for noticing this inconsistency. This overlap was a typographical error in the description. The bands used for the band-integrated statistics were intended to describe distinct scale ranges. The text has been corrected accordingly: "*large scales 640-160m"*

Editor comment 7

"L.321, "the bottommost 640m" while the 1st line of Section 2 in the supplement materila says "the bottom [0-960 m] layer". I guess you have one spectrum for a 640 m bin i.e. bin i=1 at 0-640m (from the bottom), bin i=2 at 10-650 m, bin i=3 at 20-660m, ..., bin i=33 at 320-960m. Is this correct?"

Yes, this is correct. Each spectrum is computed over a 640 m vertical window. The window is then shifted upward every 10 m from the bottom, which leads to 33 overlapping spectra in the [0-960 m] layer. We now clarify, however, that for Fig. S3 we selected spectra according to the distance from bottom of the window centre: spectra were retained when their window centre was located within 960 m above the bottom. Since each window is 640 m long, the retained windows can extend up to about 1310 m above the bottom (see first paragraph of Section 2 and Caption of Figure S3). This selection yields 1933 individual spectra which are exploited for the Spearman correlations shown in Table S4.

Editor comment 8

"Each spectrum is actually a group of 9 spectra shown in Table S3. Then you take a pair from the group, say, REF and SAME_redux. You have 33 spectra (bin i=1..33) for REF and 33 spectra (bin i=1..33) for SAME_redux. Then you calculate the spearman rank correlation for the pair. Or did you treat all spectra from all stations? Since on Figure S3, the index i runs upto about 1900. The caption says "Median of R ratio across spectral subsets" -- but which subset? Each i indicates station number or depth or both?"

We thank the Editor for pointing out this ambiguity. In Figure S3, the index i refers to the sequence of individual 640 m spectral windows used in the sensitivity test. For each station, these windows are ordered from the bottom upward, retaining only those whose centre is located within 960 m above the bottom. The sequence then continues with the next station, again from the bottom upward.

Thus, i is not a station number and not the vertical-window index of a single station. It is a pooled spectral-window index across all retained stations and depths. Each spectral-window (i) is used to compute shear and strain spectra, from which the R diagnostic is then obtained for the different bandwidth configurations listed in Table S3. This pooled set contains 1933 individual spectral windows and is used for the Spearman correlations reported in Table S4

We clarify the text in the supplementary material (see the first paragraph of Section 2): "For every single ratio R, variants of the median are diagnosed using the respective bandwidth configurations listed in Table S3, each bandwidth representing a subset. Spearman correlations will be then tested between each median series of R associated with each subset". We then clarify Figure S3 caption accordingly: "Here, i is the index of the pooled list of all valid spectra used in the test, obtained by stacking spectra from all stations and all sliding-window positions whose centre is located within 960 m above the bottom".

Editor comment 9

"Table S4, Do you really need 7 significant digits? How different is 0.541860 from 0.541861? I see a blue raw for Sh_redux2. What does this mean? I also see the same numbers in diagonal positions. This is consistent but is it necessary? I believe most interesting comparison would be between physically meaningful "SAME_redux,same1,2,3" and the parameter used in this manuscript "REF". The correlation seem smaller (0.47--0.62) than other entries. Would the use of "SAME_redux" change some of the conclusion of the manuscript?"

We thank the Editor for these comments. We agree that Table S4 was not sufficiently edited and was unnecessarily exhaustive, especially regarding the number of digits, diagonal values, and formatting. We have now simplified the table by rounding the correlations, removing redundant diagonal entries, and removing the unexplained color formatting.

The lower correlations between REF and SAME_redux indicate that imposing identical and narrower shear/strain bands changes the absolute values and local ordering of $R\omega^*$, as expected from the reduced spectral bandwidth. To assess whether this affects the physical interpretation of the results, we directly compared the REF and SAME_redux diagnostics. In Figure R3a (below), we show that SAME_redux substantially modifies the dynamic range of

Rω* and produces a noisier, more dispersed diagnostic. This is expected because this conservative same-band configuration excludes the larger vertical wavelengths that contain an important part of the shear signal, while also excluding part of the smaller-scale range contributing to the strain signal. Importantly, in Figure R3b, comparisons of Rω* against bathymetric slope, roughness, and Δ-slope show similar broad structures, with no reversal of the main trends. We therefore retain REF as the reference configuration, because it preserves the broader noise-controlled shear band while using the appropriate strain band for the CTD-derived displacement spectra. SAME_redux affects the quantitative expression of Rω*, but it does not modify the principal conclusions regarding the relationship between the internal-wave field and bottom topography.

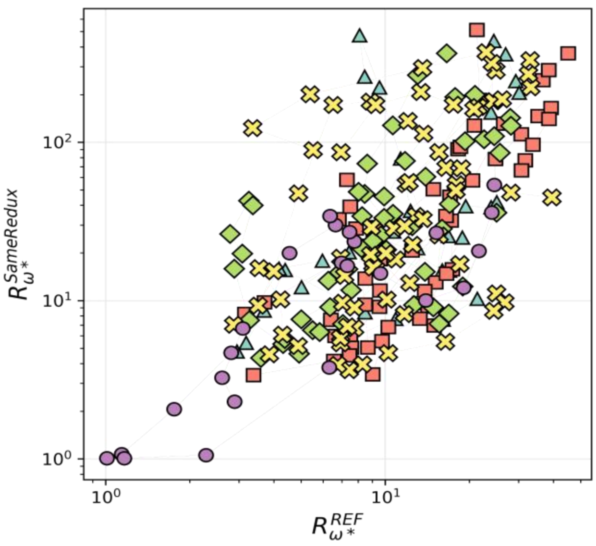

***Figure R3a**. Comparison between Rω* obtained using the reference configuration (REF) and the conservative same-band configuration (SAME_redux). SAME_redux imposes identical shear and strain wavelength ranges, resulting in a narrower spectral bandwidth than REF. This substantially modifies the absolute values and dynamic range of Rω*, reflecting the loss of part of the large-scale shear signal. Despite these quantitative differences, the positive relationship between both diagnostics indicates that the broad organization of the internal-wave field is preserved.*

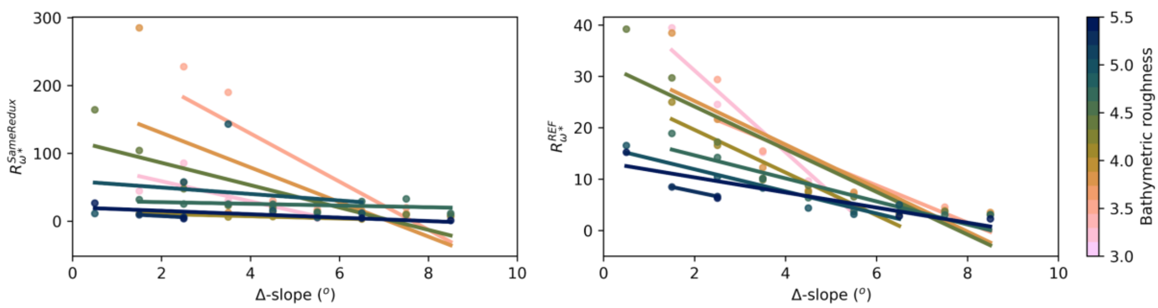

Figure R3b. Distribution of Rω* obtained with the ***SAME_redux*** configuration (left) compared to the ***REF*** configuration retained in our study (right). Ratios are binned by roughness classes of 0.3, and Δ-slope classes of 1°, then median values are calculated and a linear regression is performed for each roughness class. Note that the expected roughness ordering is not strictly reproduced in the SAME_redux Rω* scatter. This should be interpreted with caution, as this configuration relies on a reduced number of spectral components, making the shear–strain ratios less stable and less representative of the broader signal. However, the main patterns discussed in the manuscript remain visible (i.e., distinguishable overall tendency of Rω* to be scattered from high to low values with roughness increasing), supporting the robustness of the qualitative interpretation with respect to bandwidth selection.