#### **Response to Reviewer 2:**

My biggest concerns are about the data products used in model-data assessments. In a few instances they are using the same product for model initialisation, boundaries and model assessments (glory, tpxo, WOA), which is not an independent comparison. I do see these sorts of non-independent comparisons as a useful tool as it shows how faithful the model downscales the original product. However, on their own and without more independent assessments, we don't know whether the biases we are seeing is the model degrading the initial/boundary products or if the model is actually improving on them.

The reviewer is correct that, in some cases, we have compared the model solution in the interior of the domain to data products that were used for boundary forcing. These comparisons are nonetheless meaningful tests of the capacity of a regional model to translate boundary forcing into an interior solution that remains consistent with observations (the essential task of a regional ocean model). Unlike the data products in question, the regional model does not benefit from assimilating observations within the domain. It must explain multiple observed interior properties by dynamically extending from the specified boundaries with a single set of self-consistent explicitly specified physical and biogeochemical dynamics. Maintaining agreement with observation-based products in the domain interior thus supports the fidelity of these dynamics: skillful forcing at the boundary is a starting point, but much can go wrong between the boundary specification and the interior solution. We now more clearly explain the rationale for these comparisons in the manuscript text which reads:

Lines 287-294: "We note that several comparisons are made against gridded data products that were also used to force and initialize the NEP10k hindcast (i.e., GLORYS, TPXO, WOA23). While these comparisons are not fully independent, they are nonetheless meaningful tests of the capacity of the regional model to translate horizontal boundary and surface forcing into an interior solution that remains consistent with observations. The regional model must explain multiple observed interior properties by dynamically extending from the specified boundaries with a single set of self-consistent explicitly specified dynamics without the benefit of assimilating, or being informed by, observation from within the domain. Maintaining agreement with observation-based products in the domain interior thus supports the fidelity of these dynamics."

We do agree with the reviewer that comparison against interior observations that are not directly considered in the boundary forcing provides an even stronger test of NEP10k performance. Such comparisons assess the adequacy of both the boundary conditions and the dynamics within the model interior. We have added several observation-based comparisons to the supplement including assessment of NEP10k ability to reproduce observed tidal constituents calculated from gauges in the eastern Bering Sea and western Gulf of Alaska (Fig S13) as well as a broader comparison against the CalCOFI dataset for

temperature, salinity and biochemistry (Figs. 26-29). We feel that these, in combination with the independent data comparisons elsewhere in the text (Figs. 2a-c, 4a-c, 5a-c, 10, 11, 14, 15, 18, 19, 20, 21, 25) provide a robust understanding of the strengths and limitations of the NEP10k configuration and dynamics.

• My next but related concern is to do with the Aletutain islands. It seems likely to me that there are some very fine scale processes occurring around these that could have a noticeable effect at the 10-20km scales that you are assessing. If your model is a higher resolution than the data products then it is likely that your model is better capturing these processes than your data products so you will need some higher resolution data products to properly assess the model performance in this region. I am not too familiar with the oceanography of the region, so it could be that the authors have already considered this - but I would like to see a short discussion on how well represented they expect this region around this island chain to be represented in their model and the observational datasets.

While some of the data products that we use have resolutions that are coarser than the scale of the Aleutian Island passes, the observations upon which that are based includes the impact of these finer-scale features. That is, GLORYS may not resolve Aleutian Island throughflows well, but the Bering Sea data that it assimilates includes their effects as they are in nature, allowing the GLORYS state estimate to correct its deficiencies.

In addition, the most fisheries-critical feature in the Bering Sea/Aleutian Island system for the present intended applications of the NEP10k model is the Bering Sea cold pool, for which we have trawl surveys. Fluctuations in this feature have been shown to correlate with fisheries dynamics, including collapses, linked to billions of dollars in gains and losses (e.g., Szuwalski et al., 2024). The model was able to capture this feature, as observed with detailed trawl survey observations over nearly 3 decades, exceptionally well (Figs. 19, 20). Further exploration of the circulation around the Aleutians would indeed be useful, but is beyond the scope of this paper. We do include text in the manuscript that we would expect model performance in this region to improve at higher resolution:

Lines 1032-1038: "Increasing horizontal resolution of the NEP10k configuration may further improve representations of important regional GOA ecosystem features. For example, sea surface heights south of the Aleutian Island Chain, central to the Alaska Gyre, are lower than observed in reference datasets (Fig. 5) and could improve with better resolution of opposing horizontal flows, specifically the southwestward Alaska stream and eastward Subarctic or Aleutian Current. Higher resolution may also improve representation of transports through the Aleutian Island chain, which can significantly impact water mass properties in the Bering Sea (Stabeno et al., 1999)."

We are pursuing several different avenues of developing the NEP regional MOM6 model including a 5km resolution configuration, which we mention in the manuscript text (Lines

1053-1054). Additionally, a number of collaborators are focused on more in-depth analysis of region-specific performance of the NEP10k, with a Bering Sea-focused manuscript having been just recently submitted to this special issue of GMD (Seelanki et al., 2025).

Seelanki, V., Cheng, W., Stabeno, P. J., Hermann, A. J., Drenkard, E. J., Stock, C. A., and Hedstrom, K.: Evaluation of a coupled ocean and sea-ice model (MOM6-NEP10k) over the Bering Sea and its sensitivity to turbulence decay scales, EGUsphere [preprint], https://doi.org/10.5194/egusphere-2025-1229, 2025.

## Minor Comments

• There are quite a few acronyms which does reduce the readability, particularly in the abstract. I suggest carefully assessing which of these are really needed and which can be written out in full.

We have removed the regional initialisms (i.e., NEP, EBS, GOA, CCE) from the abstract; the only remaining abbreviation in the abstract is 'NEP10k'. We have reviewed the manuscript text to ensure it conforms with journal abbreviation conventions.

• Fig 1 is on a different projection to the rest of the figures. I am less familiar with this region, and it did make it more difficult to locate the different regions on the subsequent figures.

The projection of Figure 1 was chosen to optimize the amount of visible detail (e.g., landmask outline) and reduce the amount of unused space that results from the model domain curvature when plotted on the plate carrée projection. We have added a figure in the supplement (Fig. S1) that depicts the same information as Figure 1 but cast in the plate carrée projection that is more similar to other map projections used in this text.

# • lines 115-120. What sort of an effect does land-use changes have on this system?

We have added a few references exploring the impacts of land-use change in the region to this introductory material. These specifically reference the risks of deoxygenation and acidification mentioned in the prior sentence as they can be exacerbated by pollution, nutrient inputs and coastal engineering. The text now reads:

Lines 119-121: "These risks can be further amplified by processes resulting from changing land-use such as increased nutrient input, pollution and coastal engineering (e.g., Halpern et al., 2009; Hughes et al., 2015)."

### • Line 200: NEP domain -> NEP10km

We have made the recommended change to the manuscript text. It now reads: Lines 211-212: *"Initial and boundary conditions were regridded to the NEP10k domain using the xesmf python software package (Zhuang et al., 2023)."*  • Line 206: Suggest also including nudging timescales here.

Per the reviewer's suggestion, we have added information regarding the nudging timescales and tracer length scales to the manuscript text. It now reads:

Lines 214-218: "lateral boundary forcing also applies nudging and tracer reservoirs (the latter retains a memory of water properties exchanged with the modeling domain rather than instantaneous forcing; see Ross et al., 2023 for more details). As in Ross et al., (2023), the radiation and nudging schemes utilize 3 day inflow, 360 day outflow timescales, and both inward and outward tracer reservoir length scales were 9000 meters (Table A1). No nudging was included in the interior of the domain"

• Lines 305-306: It seems that for both of the mixed layer depth comparisons you are using the same calculation for model and observations (which is a +ve) – but I suggest rephrasing to make this clearer.

Per the reviewer's recommendation, we have made this portion of the manuscript text clearer. It now reads:

Lines 328-332: "We validated NEP10k mixed layer depth (MLD) against the 1° de Boyer Montégut (2024) monthly MLD climatology, which incorporates measurements from an assemblage of MBT, XBT, CTD casts and profiling floats, and defines the MLD as the seawater depth where potential density is 0.03 (kg/m3) greater than the density at a reference depth of 5m. From NEP10k, we used the MOM6 diagnostic variable MLD\_003, which calculates the mixed layer depth based on a user-defined reference depth (in our case, 5 meters for consistency with de Boyer Montégut). The mixed layer depth is identified as the depth where the potential density increases by 0.03 kg/m<sup>3</sup> relative to the surface reference depth"

• Line 345-350. It is hard to see where the 500 m contour is on fig 1, and it is hard to assess how many grid cells you have on the shelf. I suggest adding the 500 m contour to figure 1, and including some text to say how many grid cells wide this continental shelf region is (perhaps max and min lengths?) The shelf is quite narrow and for most of your domain, the area less than 500 m appears to be only a few grid cells wide. Is the 10km resolution model a good enough tool to describe the coastal region? If shelf conditions are a key fisheries-critical variable, then nesting into the coastal regional may be required (I am not suggesting you need to do this for the current publication). Your global observation products may also struggle to capture coastal processes. I think this paper could be complete without the coastal assessments, and a short discussion about how to approach this in future would work instead (for example, using in situ products that measure shelf scales combined with higher resolution models).

Per the reviewer's suggestion, we have added a contour indicating the 500m isobath to Figure 1. We have also added a figure to the supplement wherein we zoom in on the individual, regional shelves and report the number of gridcells, total areal extent, and approximate minima/maxima shelf distances.

The methods text now reads:

Lines 382-384: "We thus complemented the broad spatial comparisons with region-specific time series of shelf (defined as grid cells where bottom depth is less than 500 meters) conditions, where the subregions are those shown in Fig. 1 and regional shelf extents are depicted in Fig. S2."

We have also expanded the text in the discussion section to reflect the reviewers suggestions, specifically considerations for higher resolution vs. using a smaller nested grid, as well as text regarding both the need for comparison against coastal observation products and reasonable expectations for such comparisons. This portion of the manuscript text reads:

Lines 1057-1070: "For applications wherein many-fold higher resolution is necessary, it may be more practical to utilize a smaller, higher-resolution nested domain (e.g., modeling the Salish Sea in Khangaonkar et al., 2018) which can be forced by the NEP10k at the open boundaries, rather than increasing resolution for the full NEP10k domain. Continued NEP10k development will incorporate comparison against a broader array of local observation datasets similar to that of CalCOFI. Such extensive observation records are invaluable for better understanding and evaluating model performance, particularly in regions that may not be well represented in relatively coarse, gridded data products. However, it is important to approach such comparisons with realistic expectations. As shown in Fig. S28, NEP10k poorly reproduces temporal variability (i.e., low Pearson correlation coefficients) of repeated samplings of individual stations across multiple years. This is not surprising since the NEP10k hindcast does not assimilate observations and, thus, any biases in the mean locations of fronts and other features is compounded by stochastic mesoscale and submesoscale features whose precise locations and timing will not match those observed. Indeed, more coherent patterns emerge after averaging over such features (e.g., Fig. S29, Fig. 17, Fig. 20), which demonstrates that NEP10k strength and utility is in representing reasonable approximations of ecologically-important environmental conditions rather than exact reproduction of in situ observations."

### Line 394: What does CMEMS stand for?

We have clarified the manuscript text to indicate the full name Copernicus Marine Environment Monitoring Service. The text now reads:

Lines 434-435: "We assessed NEP10k's ability to generate realistic NGAO and GOADI patterns by comparing against satellite altimetry from the Copernicus Marine Environment Monitoring Service (CMEMS, 2023)."

 Section 2.5.3: This section (and the similar results section) could be of value as it shows some of the good features that help your model code run faster. However, as written, I don't think it is essential to be included as some of these results will be system specific (i.e. it is of interest to you, but not to a broader audience). If you include this section, then I would like to know more about where did your model runs. I note that you mention the computer at the end of the paper, but this needs to come earlier. I would also like to see a description of the computing system used as these results will likely vary across different computing architectures (e.g. inter-PE communication speeds will vary across different computers).

We have expanded on the description of the GAEA supercomputer that appears in the Methods section 2.2 to include information regarding the computing environment; this was added in section 2.5.3 and reads as follows:

Lines 453-459: "As mentioned in section 2.2, simulations were conducted on NOAA's Gaea High-Performance Computing system. This system consists of HPE-Cray EX 3000 nodes (2 × AMD EPYC 9654, 2.4 GHz base, 96 cores per socket), connected via HPE Slingshot 11 a high-speed interconnect designed for exascale systems. The system also features over 150 PB of shared storage using IBM Spectrum Scale parallel file systems. The model runs in a distributed-memory configuration using MPI across hundreds to thousands of cores. Additional system details can be found in the NOAA RDHPCS documentation (https://docs.rdhpcs.noaa.gov/systems/gaea\_user\_guide.html#system-overview).

As described in Section 1, the viability of the NEP10k configuration for ecosystem applications depends on its ability to not only simulate fisheries-critical features but also to run with sufficient computational economy to permit generation of the thousands of years of retrospective forecasts and projections required to provide credible uncertainty estimates (e.g., Ross et al., 2024; Koul et al., 2024; Ross et al., 2024). However, we recognize that others interested in running the NEP10k configuration may have different computing resource availability. Therefore, we report the computational performance under different NEP10k configuration options (i.e., scaling, land masking and time-step splitting) in order to provide insight into how one might optimize production on a given computing system."

### • Line 428: How did you conclude that the 400s tracer time step was best?

We did not conclude that the 400s tracer time step was best. We included sensitivity experiments wherein the tracer time step was set to 400s (as opposed to the 1200s time step used for the simulation reported in the results) to demonstrate the impact on computational efficiency and the economy granted by the timesplitting capabilities of the model. We have clarified this in the manuscript text - it now reads

Lines 481-484: "Finally, we include additional 50 x 75 PE and 50 x 100 PE simulations with the thermodynamic time step equal to the baroclinic time step (400 seconds) rather than three times the baroclinic time step (i.e., 1200 seconds) as was used in the base configuration. These last two experiments allow us to quantify and demonstrate the computational value of the flexible time stepping that MOM6 enables."

With regards to the choice of a 400s baroclinic time step - we found that this was the longest baroclinic time step with which we could integrate the model without it crashing due to numeric instability for the hindcast period (1993-2019) which is mentioned in the manuscript text:

Lines 908-914: "The NEP10k computational cost is comparable to the recently published Northwest Atlantic regional MOM6 configuration (NWA12) of Ross et al. (2023), which used a 40x40 layout (1200 PEs after land masking) to generate 1 simulation year in about 9 hours (about 10,800 PE hours per simulation year). While NWA12 was a larger domain, NEP10k required smaller baroclinic and thermodynamic time steps for stability (400 versus 600 seconds and 1200 versus 1800 seconds, respectively). The instability at longer time steps in the NEP10k configuration primarily occurred in the vicinity of the Aleutian Island chain where strong currents could be generated within tight channels."

• Line 486: Can you indicate the Aleutian island chain on Figure 1?

We have added text to the Figure 1 caption indicating the position of the Aleutian Island Chain; The caption now reads:

Lines 160-163: "These regions, from north to south, are the Bering Sea (BS), Gulf of Alaska (GoA), British Columbia (BC), Northern California Current System (NCCS), Central California Current System (CCCS), and Southern California Current System (SCCS). The southern arc of the Bering Sea polygon traces the Aleutian Island Chain"

 Line 500 – 505: I'm not convinced by this statement. It is possible that these biases represent your model improving on the tpxo dataset. In my experience (in other parts of the world), a locally produced model tends to compare better to tide gages than the tpxo model for partially enclosed areas. Is there tide gage data for this area that you can compare to? If resolution is to blame for the bias, how does the tpxo resolution compare to your model?

We have added additional supplemental comparisons of NEP10k and TPXO tide approximations against NOAA tide gauges in the eastern Bering Sea and western Gulf of Alaska (Figs. S12 and S13). We found that TPXO and NEP10k approximate the tidal M2 and K1 constituents fairly well and that TPXO generally is a good benchmark for comparison. However, in a few cases, the NEP10k-TPXO bias exaggerates the bias of NEP10k relative to the tide gauge since their respective biases are opposite signs. We have incorporated these findings into the results section of the manuscript which now reads:

Lines 561-571: "To investigate some of these biases further, we include additional, zoomed in maps of the eastern Bering Sea and Gulf of Alaska in the supplement (Fig. S12) along with comparison against several tide gauges in that region. Generally, TPXO better approximates tidal harmonic constituents than NEP10k with higher Pearson Correlation coefficients and/or lower RMSE (with exception of M2 phase). However, in cases such as the gauge in Anchorage, AK, the bias in M2 amplitude for TPXO is comparable to the bias exhibited by NEP10k. Since these biases are opposite signs, the discrepancy between the two gridded products (i.e., NEP10k-TPXO, shown in the maps in Fig. 6 and Fig. S12) exaggerates the model bias by almost a factor of two relative to the bias for the gauge. Thus, some of the more severe near-shore differences in Fig. 6 may be a reflection of how NEP10k and TPXO approximate complex coastline geometry (bottom of Fig. S13) rather than an exact indication of NEP10k performance."

• Figure 18: note here that some of the apparently poorer comparison on the shelf compared to offshore could also be because the offshore comparisons were log transformed.

This is a fair point and is now noted on Lines 784-787: "Notably, the shelf chlorophyll comparisons in Figure 18, which focus on temporal chlorophyll variability within a defined region, are not log transformed. This amplifies the discrepancies at the higher end of the observed range relative to those in full domain, which focus on the model's ability to capture order-of-magnitude cross-ecosystem differences (Fig. 10)."

• Line 734 – Suggest writing CPA in full as it is not used often, and readers will have forgotten what it stands for.

We have made the recommended change to the manuscript text. It now reads:

Lines 808-809: "The NEP10k model, however, robustly reproduces interannual variability of the cold pool area (CPA) indices, with best performance at the higher temperature thresholds (Fig. 20)"

### • Line 879: your fig 18 suggests that the coastal biases are not small

This sentence was intended to emphasize the improvement of increased near-shore chlorophyll to more realistic levels than were seen before the addition of a phytoplankton size class. We have removed reference to the biases in this sentence; it now reads:

Lines 964-965: "These benefits can be seen in the generally high coastal (relative to open ocean) chlorophyll levels along the U.S. West Coast (Fig. 10)"

### • Line 968: Quite often when you increase resolution you also need to decrease timestep. This could potentially be a lot more than 8-fold increase in computational cost!

Yes, the reviewer is correct, we were including the decreased timestep in the "eight-fold" approximation. We were specifically referring to doubling resolution only in the horizontal directions which would impart a two-fold increase in computational cost for each horizontal dimension, as well as an additional two-fold increase for shortening the timestep. This computational cost approximation for shortening the timestep is derived from balancing the CFL condition ( $u\Delta t/\Delta x$ ) in order to maintain numerical stability. Indeed, if we also doubled resolution in the vertical dimension, the computational cost would increase further. We have modified the manuscript text to reflect this clarification. The text to now reads:

Lines 1053-1057: "However, while higher resolution (i.e., ~5km) simulations are currently underway, any benefits of doubling horizontal resolution will need to be balanced against the roughly eight-fold increase in computational cost (i.e., two-fold for each horizontal dimension and an additional two-fold increase for the necessity of shortening the timestep needed to maintain Courant–Friedrichs–Lewy stability)."