Reviewer A:

I reviewed the revised contribution titled The Distribution of Interseismic Locking on the Cascadia Subduction Zone Using Bayesian Methods by Pearson et al., submitted to Seismica. The authors have done a nice job revising the manuscript in response to some of the concerns, although there are still a few points I think should be improved.

My main concern is still with the apparent roughness of the results, in comparison to the 'uncertainty' derived from the Bayesian sampling, which I think must be underestimated. While the model roughness has actually decreased compared to the original version, the transect A-A' in Figure 4(c) suggests that the authors have better than 95% confidence (2-sigmas) in the existence of a well-resolved patch of low to moderate coupling around 70km from the trench, in between two highly coupled zones. This entire region is offshore, and from my own experiments I am having a hard time believing that this can really be a well-resolved pattern in the coupling. Instead, I think there are a few possibilities that could explain the combination of a rough mean model and low standard deviation:

Possibility (1): The formal uncertainties of GNSS velocities are typically under-estimated for use in modeling; did the authors use the reported uncertainties (e.g. 0.05 mm/yr is common for cGPS)? These values are probably too low; in the literature it is common to impose a minimum uncertainty of 0.25 or even 0.5 mm/yr to avoid giving too much weight to individual stations that may simply be located on unstable ground (for example). Alternatively, there could be other reasons why the model uncertainty is underestimated.

Possibility (2): The Bayesian sampler may perform poorly when there is no smoothness prior. Fundamentally, the Bayesian sampler is still doing a similar thing to a linear least squares inversion, which is to seek the lowest-misfit point in model space. This is particularly true if Point (1) is the case, in which case the posterior PDF will be very narrow and the model will report a mean model very close to the linear solution. Regardless of which inverse method is used, the existence of data noise means this lowest-misfit point is very likely to be at a location in model space that represents a rough model; this is the reason for using Laplacian smoothing in linear inversions, and this is not addressed by the use of a Bayesian sampler.

However, it seems that the choice to forego regularization is a relatively common one in the Bayesian sampling literature; such un-physically rough models have been reported elsewhere (e.g. Nocquet, 2018). If the uncertainties are accurately characterized (Point 1) this may be acceptable, though I feel there is possibly still something else missing here. Short of duplicating the results myself, I cannot come to a final conclusion about why the results are so rough, but I encourage the authors to think more deeply about the possible reasons, and if they cannot solve the problem, at a minimum to expand on these reasons in the discussion as much as they are able to.

We thank the reviewer for their detailed comments regarding the apparent roughness of the mean locking ratio results and the corresponding uncertainty derived from the Bayesian sampling.

The uncertainties used in our study are calculated following the methodology of Newton et al. (2021), some of which are low due to very long (20+ year time series) with a strongly linear trend. We rescale the data covariance matrix (see Section 2.3) with additional factors based on glacial isostatic adjustment and viscoelasticity, which does increase uncertainties. Only 6 out of 1812 uncertainties (1812 values includes uncertainty in E/N/U directions for GNSS, as well as leveling values) are then less than 0.25 mm/yr. Additionally, the process of ensuring the data covariance matrix is singular positive definite also acts to scale up the diagonal of the data covariance matrix. We don't think that underestimated uncertainties are the leading cause of the roughness.

We feel that possibility 2 in the reviewer's comments is the more appropriate explanation. In practice, it is not straightforward to reduce this roughness with our method. The Bayesian inversion framework we employ does not include explicit smoothing, and the data covariance matrix must be constructed to be singular positive definite for stability, as discussed in Section 2.3. Within this framework, the occurrence of relatively rough mean models with locally low uncertainty is consistent with previous studies in the literature. As was stated, Nocquet (2018) (as well as others—such as Minson et al., 2014) report similar patterns in Bayesian sampling results without smoothing, highlighting that such behavior is a common feature of this approach. We have expanded on this in the manuscript discussion (see Line 500) to ensure that readers understand this aspect of the model.

Finally, I still disagree with the authors that 'locking' is the correct term for their model parameters; many of the papers they cite in justification of this choice were written many years ago, and I believe that the community's understanding of this terminology has now shifted such that 'kinematic coupling' is the more accepted term. I refer the authors to Sato et al. (2025) which does an excellent job distinguishing between the two terms.

- We acknowledge that terminology is continuing to evolve. Our goal is to take the practical approach of using terminology that can more easily be compared with previous publications. We note that there are still very recent papers that use the term locking (e.g. Barra et al., 2024; DeSanto et al., 2025; Li & Chen, 2024; Pollitz, F.F., 2025; Sharma et al., 2023). We agree that it is important to clarify, given that both terms are used in the literature (sometimes both terms within a single paper).

We followed your guidance in the first review regarding ensuring that we defined the term, see Line 56.

Reviewer B:

Yes, I think your statement is correct: "we expect that the mean result from this probabilistic model should not necessarily look smooth". There is no prior in this inversion that would favor smooth means, but I am still surprised that the mean solution is rough AND uncertainties appear to be quite small. This implies the inversion can resolve short wavelength variations in coupling, but of course that is not the case, because the elastic Green's functions are low-pass and surface displacements are insensitive to very short-wavelength slip variations. To me, this hints at either (1) poor mixing and undersampled posterior distribution, but I can't know that without seeing some analyses that demonstrate sampling is working (trace plots, correlation plots, etc), or (2) the likelihood is quite "sharp", meaning perhaps underestimated data uncertainties (I could not find what data uncertainties you are using and the data set is not yet available on Zenodo), or some of the data have very small uncertainties, or as you have discussed, data uncertainties are not sufficiently correlated.

I don't want to dwell on this – the uncertainties are not particularly important to the results of this paper – the main conclusions hold, regardless. I just spend a lot of time worrying about uncertainties and I am curious about why your locking uncertainties (Figure 4) are so small.

We thank the reviewer for their expanded thoughts on this topic. As discussed previously, we don't see evidence that undersampling is the problem. Additionally, since the uncertainties are scaled up within the data covariance matrix (see response to Reviewer #1), we don't think that underestimated data uncertainty is driving the width of posterior distributions. We think that this is something that's worth exploring in future Bayesian inversions.

Okay, I see now what you are saying. Yes, it makes sense to me that the data covariance structure could influence the roughness of the solution, just as you say. Your illustration of the influence of off-diagonal terms is interesting. This is a hint that the small uncertainties (previous response discussion) may have more to do with the data errors than the sampling(?).

See response to Reviewer #1 regarding roughness and uncertainties. We do agree that issues with correlations between data are likely a greater issue.

But truncated Gaussians are not flat-topped, and samples drawn should not be flat-topped. So, I still don't understand, but perhaps this is a minor point.

The appearance is due to a combination of truncation and resampling, which can lead to some appearing as though they are flat-topped.

1. Reviewer B noted that they could not verify the data uncertainties because the dataset is not yet available on Zenodo. The manuscript is also currently inconsistent on this point: Section 2 states that the refined velocity dataset "is archived on Zenodo," while the Data and Code Availability statement still reads "(will upload to Zenodo after paper acceptance)." In line with Seismica's open data policy, please upload the dataset (combined GNSS and leveling data, Green's functions, data covariance matrix, and posterior distributions) to Zenodo now and replace the placeholder with the DOI. I will verify the content of your Zenodo repository being in agreement with Seismica's requirements.

- Dataset has been uploaded to Zenodo. https://doi.org/10.5281/zenodo.21479776

2. Both reviewers raised the combination of rough mean models and locally low posterior uncertainties, and Reviewer B mentioned that convergence diagnostics (e.g., trace plots or correlation plots) would be to rule out poor mixing. Your response states that undersampling is not the cause. Please add one or two such figures to the supplement to document this and address the concern.

- We have added figures to the supplement (see Fig. S5).

3. If you retain the term "locking," please consider citing Sato et al. (2025), which Reviewer A recommended, alongside your kinematic definition in the introduction, to acknowledge the distinction the reviewer raised.

- We have added a sentence citing that paper and acknowledging the distinction (Line 56).