Referee 1:

This manuscript tackles a pivotal issue in seismology: enhancing the accuracy of core-phase travel-time reconstruction through coda wave interferometry, a topic of paramount importance for advancing the imaging of Earth's deep interior. The authors employ a perturbation analysis to quantify travel-time deviations and establish a critical angular threshold, constituting a valuable theoretical contribution that addresses persistent uncertainties regarding the physical correspondence between reconstructed coda-derived phases and true inter-station core phases. However, several substantive issues pertaining to literature integration, theoretical rigor, completeness in data analysis, and contextualization with prior research must be resolved to bolster the manuscript's robustness and impact. Based on the following assessment, a major revision is recommended.

We thank Referee 1 for the thorough review and positive assessment of our work. We appreciate the recognition of our study's contribution and the constructive identification of areas for improvement.

We have carefully reviewed all the specific comments. In the revised version, we will address each point. We sincerely thank you for your time and valuable feedback, which are essential for improving this work.

Major revisions:

1, Insufficient Quantitative Analysis of Data Volume, Distribution, and Deviation Angle Dependence on Reconstruction Stability:

A central conclusion of this study is that "a sufficiently large number of global earthquakes" enables reliable core-phase reconstruction. However, this claim lacks the necessary quantitative foundation and fails to systematically assess how the stability of travel-time retrieval depends on key data characteristics. Specific shortcomings include:

a) Undocumented Processing Parameter: The manuscript does not specify the coda time window (e.g., start and end times relative to the origin time) used for the correlation analysis. This critical parameter directly influences the extracted waveforms and must be reported to ensure reproducibility.

In the revised manuscript, we have specified the coda time window.

Revision (lines 205–207):

“For each earthquake–station pair, we extracted the late coda time window from 10,000 to 40,000 s after the origin time. This window This window contains dominant multiply scattered waves that have sampled deep Earth structures.

Unquantified Impact of Event Count: While 205 earthquakes were used, the study provides no analysis of how reconstruction stability degrades with smaller datasets. To objectively define data sufficiency, a sensitivity analysis is required. For instance: How does the standard deviation of the reconstructed travel times for key phases (e.g., ScS, PKIKP^2) evolve as the number of events used in the stack is progressively reduced from 205 to, for example, 150, 100, or 50 (e.g., via bootstrap resampling)? Does a minimum event count threshold exist, below which the stability deteriorates significantly or the proposed cubic scaling relationship between accuracy and the angular threshold breaks down?

Response:

Thank you for this important suggestion. Following your recommendation, we have performed a **bootstrap sensitivity analysis** to assess how reconstruction stability degrades with smaller earthquake datasets. This analysis provides an objective basis for evaluating reconstruction reliability.

Revision (lines 213–226):

In the stacked correlograms, prominent deep Earth phases --- specifically PcP, ScS, and PKIKP^2 --- are clearly distinguishable (Fig. 7). To quantitatively evaluate whether the illumination condition has been satisfied, we performed a bootstrap analysis by progressively increasing the number of earthquakes in the stack from 10 to 120, in increments of 10. The stabilization of travel times with an increasing event count serves as a diagnostic for meeting the illumination condition. For each subset size, we performed 250 random realizations to compute the mean travel time and standard deviation for the extracted ScS and PKIKP^2 phases

Our results reveal two key findings (Fig. 8). First, for both phases, the mean travel time stabilizes and the standard deviation decreases as the number of earthquakes increases, indicating clear convergence toward a stable value. Second, significant travel-time deviations---exceeding 1 s ---are observed when fewer than approximately 20 earthquakes are used, particularly for the ScS phase. This indicates that a minimum event count is required to achieve locally uniform illumination. However, we emphasize that this threshold is not solely dependent on the event count, but also on the azimuthal distribution of sources, their focal mechanisms and the number of correlation traces stacked per distance bin. The cubic scaling relationship derived in Eq. (18) assumes locally uniform illumination; when this condition is violated—as seen in the smaller earthquake subsets—the scaling relationship breaks down, resulting in the observed high standard deviations.

New Figure 8 has been added showing the bootstrap results, with mean travel time and standard deviation plotted as functions of earthquake count for both ScS and PKIKP².

c) Incomplete Investigation of Deviation Angle Dependence: While Figure 10 presents a valuable analysis of convergence with increasing maximum deviation angle (φ), the current approach of using cumulative ranges (e.g., 0-10°, 0-20°, etc.) limits its power to validate the theoretical framework. A more rigorous, binned analysis is needed. The data should be stacked and analyzed within discrete, non-overlapping ranges of the deviation angle φ (e.g., 0-10°, 10-20°, 20-30°, etc.).

This binned comparison is critical for directly testing the theoretical prediction of whether travel times remain accurate and unbiased across all directions of wave incidence under the local uniformity assumption.

Such an analysis would reveal: 1) If travel-time biases exist in specific ranges of the deviation angle, and 2) The magnitude of any such biases as a function of φ.

A finding of consistent travel times across all discrete angular bins would strongly corroborate the theoretical model. Conversely, identifying systematic biases in certain angular ranges would provide invaluable insights into the limits of the local uniformity condition and guide future data selection.

Response:

Thank you for this constructive suggestion. We agree that a **discrete, non-overlapping binned analysis** provides a more rigorous test of the theoretical framework. In the revised manuscript, we have adopted your recommended approach.

Revision (lines 227–234 and Figure 9,10):

According to our theoretical framework, when the locally uniform illumination condition is satisfied and the Γ value is large (as it is for both ScS and PKIKP^2), correlation signals across different deviation planes should converge toward the true inter-station arrival time. To verify this convergence, we partitioned the correlograms based on specific ranges of the deviation

angle φ. The resulting stacked correlograms exhibit clear signals in both the ScS and PKIKP^2 time windows; however, as predicted, the PKIKP^2 signals appear more focused than those of ScS (Fig. 9). This trend is further illustrated by comparing the azimuthal deviation (φ) ranges of (0°, 20°) and (40°, 60°) (Fig. 10a, b): while travel-time deviations for ScS reach up to 3 s, they remain nearly identical for PKIKP 2 (Fig. 10c, d).

These results provide direct validation of our theoretical framework: phases with larger Γ values are robust to the correlation at diverse deviation planes.

***Figure 9 and 10** has been completely revised to show the discrete binned analysis, with travel-time picks for each bin and quantitative comparison between phases.*

2, Missed Opportunity for Theoretical Validation through Comparative Analysis of ScS and PKIKP^2 Phases

The manuscript estimates a critical angle of 18° for the ScS phase and applies this framework in the real-data analysis. However, it does not fully leverage the contrasting behaviors of ScS and PKIKP^2 phases observed in Figure 10 to rigorously test and validate the underlying theory. Specifically, the ScS travel time shows a clear dependence on the deviation angle φ, while the PKIKP^2 travel time remains stable. This striking discrepancy represents a critical opportunity to strengthen the study's conclusions, yet it remains largely unexplained.

Response: Thank you for this incisive observation. We agree that the contrasting behavior of ScS and PKIKP² provides an opportunity to validate our theoretical framework. In the revised manuscript, we have provided a consistent explanation for this contrast.

To transform this observation into a powerful validation of the theoretical framework, the authors should:

a) Perform Phase-Specific Theoretical Calculations: The critical angle and the expected travel-time deviation are functions of the wave period (T) and the phase-specific travel time (t(p)). The manuscript must present the theoretically predicted critical angle and the scaling of travel-time accuracy specifically for the PKIKP² phase, rather than implicitly assuming the ScS-derived value applies.

Revision (lines 175-183):

Core phases retrieved from late coda correlations are characterized by steep incidence angles, which typically yield relatively large values of Γ. To illustrate this effect, we compare two representative core phases at an inter-station distance of 10.0°:

(i) ScS wave: Travel time ~ 1000 s, threshold angle θ_0 ~ 18.0°, incidence angle 3.0°, yielding $\Gamma \sim 6.0$;

(ii) PKIKP^2 wave: Travel time ~ 2500 s, threshold angle θ_0 ~ 11.0°, incidence angle 1.0°, yielding $\Gamma \sim 11.0$.

The significantly higher Γ value for PKIKP^2 indicates that its reconstructed waveform exhibits greater convergence across diverse deviation planes compared to the ScS phase. This comparison demonstrates that for phases with steep incidence angles, contributions from earthquakes at all deviation angles must be accounted for. In such cases, even sources located far from the great-circle path can contribute constructively to the correlation signal, facilitating the robust reconstruction of deep-earth phases.

b) Provide a Physically Consistent Explanation for the Contrast: The fundamental difference in the sensitivity of ScS and PKIKP^2 to the deviation angle φ must be explained within the proposed theoretical framework. This discussion should explicitly link the distinct ray paths of the two phases (e.g., ScS reflecting off the core-mantle boundary versus PKIKP^2 traversing the inner core) to the potential magnitude of the deviation function $\delta(\theta, p)$ and its higher-order derivatives in Eq. 19. For instance:

Does the more complex path of PKIKP² through the inner core lead to a different "smoothness" of δ(θ, p) near θ=0, resulting in smaller higher-order terms and thus greater robustness to a limited deviation angle range?

Conversely, does the ScS path make it more susceptible to structural heterogeneity near the core-mantle boundary, amplifying the higher-order derivatives and making its reconstruction more sensitive to the angular distribution of sources?

By quantitatively calculating the phase-specific theoretical parameters and then using them to explain the empirically observed difference in stability between ScS and PKIKP², the authors can demonstrate that their model not only predicts general behavior but also accurately captures the specific physics governing different core phases. This would significantly elevate the impact of the study by providing a unified and predictive theoretical explanation for the key observational result in Figure 10.

Revisions (line 169-173):

When $\theta_0 > i$ (i.e., $\Gamma > 1$), even coda waves radiated from earthquakes in a plane perpendicular to the inter-station plane fall within the stable angular range defined by θ_0 (see Fig. 3b). In this regime, a larger Γ implies a smaller i relative to θ_0, indicating that the late coda waves radiated in these deviated planes align more closely with the target inter-station ray path. Consequently, for a given core phase, a larger Γ value leads to a tighter convergence of correlation signals across different deviation planes within the selected time window.

In our theoretical analysis, because the Γ value for $PKIKP^2$ is significantly larger than that for ScS, the correlation signals in the $PKIKP^2$ window exhibit greater convergence across different φ ranges. We refrain from extending the analysis to higher-order terms of $\delta(\theta, p)$, first because the smoothness of $\delta(\theta, p)$ affects travel time reconstruction accuracy but does not influence the stability of reconstructed waves. Second, constraining the behavior of $\delta(\theta, p)$ after coda waves undergo multiple reflections is difficult, and extending the analysis would introduce greater theoretical uncertainty.

3, Unaddressed Discrepancies in Figures 8 and 9: Both figures indicate that most core phases exhibit obvious waveform differences as the inter-station distance approaches 0°. The manuscript does not investigate the origin of this systematic pattern, which could stem from physical phenomena (e.g., near-field effects, 3D structural complexities) or methodological artifacts (e.g., inadequate azimuthal coverage at very short distances). Explaining this observation is vital for affirming the method's reliability across the entire distance range.

Thank you for this important observation regarding the systematic waveform differences at near-zero inter-station distances. We now provide a physical explanation grounded in our theoretical framework.

Revision (lines 250-257):

In Fig. 7, the reconstructed core phases at near-zero inter-station distances (<4°) exhibit systematically lower signal-to-noise ratios (SNR) than at larger distances. This pattern can be explained within our theoretical framework. The constructive interference required to build a specific phase (e.g., PKIKP^2) depends on coda waves with incidence angles nearly identical to that of the target phase. For late coda waves with near-vertical incidence, the propagation path from the earthquake source to the station is substantially longer, resulting in stronger geometric attenuation and thereby reducing the amount of correlated late coda energy available for constructive interference. In addition, as shown in Fig. 6c, the number of correlation traces stacked in near-zero distance bins is relatively low compared to bins at moderate distances. The reduced stacking fold further limits the enhancement of coherent signals.

4, Need for a Unifying Theoretical Discussion on I2* versus True Phase Reconstruction for PKIKP²:

The manuscript reports stable PKIKP² travel times across a wide range of deviation angles (Fig. 10). This finding appears to contradict a body of prior work (e.g., Wang & Tkalčić, 2019, 2020; Costa de Lima et al., 2022) which argues that coda correlations typically retrieve an I2* wavefield—a modified Green's function whose travel times exhibit a dependence on the distribution of seismic sources (e.g., varying with deviation angle). The authors have a critical opportunity to use their theoretical framework to explain and reconcile these differing observations, thereby making a seminal contribution to the debate on what is physically extracted from coda correlations.

Thank you for this insightful suggestion. Following your recommendation, we have synthesized these points into a generalized criterion for true travel-time extraction that reconciles previous I2* observations with our stable PKIKP² results.

To achieve this, the authors must:

a) Explicitly Discuss the Discrepancy within Their Theoretical Context: The discussion should directly engage with the findings of the aforementioned I2* studies. The core argument should posit that the critical difference lies in whether the condition of "local uniform wave incidence" (quantified by the critical angle θ_0) stem from datasets or phases where this local uniformity condition was not satisfied. In contrast, the current study, potentially by leveraging a massive global dataset for PKIKP^2, may have met this condition, thus successfully retrieving the true inter-station travel time.

Revision (lines 279-285):

This unified interpretation suggests that the I2* phenomenon and true PKIKP^2 phase reconstruction are not mutually exclusive, but rather endpoints on a continuum defined by the degree to which locally uniform illumination and structural smoothness are achieved. In previous studies where the I2* effect was dominant, the illumination condition was likely not satisfied due to limited event counts or restricted azimuthal coverage — a scenario mirrored in our bootstrap results when using fewer than 20 earthquakes (Fig. 8). In contrast, the stable PKIKP^2 travel times reported here, achieved using a globally distributed datasets panning a decade, satisfy these criteria. Consequently, our theoretical frame work reconciles these seemingly contradictory findings by providing a generalized criterion for the extraction of true travel times from the late coda.

b) Provide a Physical Mechanism for PKIKP^2 Stability Based on the Perturbation Analysis: The authors must use their theoretical framework to explain why the PKIKP^2 phase in their study is robust. The key lies in Eq. (19): the travel-time error scales with θ_0^3 and the higher-order derivatives of the deviation function $\delta(\theta,p)$.

The authors should argue that the specific ray path of PKIKP^2 (traversing the inner core) results in a "smoother" $\delta(\theta,p)$ near $\theta=0$ (i.e., very small higher-order derivatives). Combined with its specific period-to-travel-time ratio yielding a small θ_0, this makes the phase inherently robust to variations in the deviation angle φ once a basic illumination threshold is crossed.

This provides a physical mechanism for why their method, under the right conditions, avoids the deviation-angle-dependent biases characteristic of I2* retrieval.

Revision (lines 258-260):

The near-vertical incidence of PKIKP^2 results in a high Γ value, which in turn facilitates an extremely tight convergence of correlation signals across different deviation planes (Fig. 9 and 10). This geometric advantage ensures that arrival times remain nearly identical regardless of the source deviation angle φ.

In our analysis, as $\delta(\theta, p)$ is difficult to characterize analytically after multiple reflections, we have not further investigated its impact on the stability of PKIKP^2.

c) Propose a Generalized Criteria for True Travel-Time Extraction: The manuscript should synthesize these points into a clear proposition: The transition from retrieving a biased I2 to the true PKIKP^2 travel time occurs when the angular range of incident waves meets or exceeds the phase-specific critical angle θ_0 and the structural setting leads to a sufficiently smooth deviation function. This would powerfully contextualize their results, suggesting that the previous I2* observations and their own stable result are not fundamental contradictions but are explained by the degree to which the conditions of their unified theory are met.

Revision (lines 260-278):

A long-standing debate in coda correlation seismology centers on whether extracted phases represent true inter-station PKIKP^2 arrivals (the Green's function) or a modified wavefield (I2*) whose travel times exhibit a systematic dependence on source distribution (Wang & Tkalcic, 2020a, b; Costa de Lima et al., 2022). Based on our theoretical analysis and empirical results, we propose that the transition from a biased I2* measurement to a true $PKIKP^2$ phase travel time occurs when the following two conditions are jointly satisfied:

(i) Illumination condition for stable reconstruction: The angular range of incident waves meets or exceeds the phase-specific critical angle θ_0, ensuring that the stationary phase zone is adequately sampled. This condition depends on the number of earthquakes, focal mechanisms, their azimuthal distribution relative to the inter-station path, and the distribution of inter-station directions stacked per distance bin. The degree to which this condition is satisfied can be evaluated quantitatively through bootstrap convergence analysis (Fig. 8), where stabilization of travel times with increasing event count indicates that the illumination condition has been achieved. It can also be assessed by examining the convergence of the correlation signals across different deviation planes.

(ii) Smoothness condition for accurate travel time recovery: The structural setting along the ray path must yield a sufficiently smooth deviation function $\delta(\theta,p)$, rendering higher-order terms in Eq. (19) negligible. Seismic ray theory assumes a smooth velocity structure (Chapman, 2004), which produces a smooth wavefront and thus smooth travel time variations from the source to the receiver and its surroundings — which in turn ensures a smooth $\delta(\theta,p)$. Our analysis therefore remains valid within the ray-theoretical framework. Although this smoothness condition cannot yet be independently verified in the real Earth, it is a necessary prerequisite for interpreting stable measurements as true arrivals, pending future validation against earthquake-derived empirical travel times.

Minor comments:

1, The Introduction's discussion of the research gap concerning travel-time deviations for coda-based core phases requires sharper focus and better integration with key literature. Notable studies on core-phase extraction (e.g., Wang & Tkalčić, 2019, JGR-Solid Earth; Poli et al., 2017, Earth and Planetary Science Letters; Phạm & Tkalčić, 2022, Nature Communications) should be cited to delineate the knowledge gap and underscore the novelty of this work.

We have revised the Introduction to integrate these studies.

2, To strengthen the practical motivation, the Introduction should explicitly mention applications of coda-based core phases in imaging Earth's interior. Citing relevant studies (e.g., Wang, Song, & Xia, 2014, Nature Geoscience; Tkalčić & Phạm, 2018, Science; Wang & Tkalčić, 2022, Nature Astronomy) would highlight the significance of accurate travel-time reconstruction.

We have added these applications.

3, Several conclusions would benefit from citations to post-2020 research to enhance timeliness. For instance, referencing recent comparative studies on core-phase travel-time accuracy (e.g., Costa de Lima et al., 2022, JGR-Solid Earth) would help contextualize the findings within the latest advancements.

We have updated relevant sections with post-2020 citations, including Costa de Lima et al. (2022) and Phạm & Tkalčić (2022).

4, The terminology for the angle of wave arrival at a station is inconsistent, alternating between "incident angle" and "incidence angle." The standard seismological term "incidence angle" should be used consistently throughout the manuscript.

We have standardized to "incidence angle" .

5, The manuscript uses "azimuth" to describe the orientation of the earthquake-station plane relative to the inter-station plane. This usage is not explicitly distinguished from the standard seismological definition of azimuth (the angle from true north to the source-station great-circle path). To prevent confusion, the authors should clearly define their custom azimuth (or use “deviation angle”) and clarify its relationship to the conventional term.

We now use "deviation angle " consistently and define it clearly in the text (**lines 166-167**):

where ϕ represents the angle between the earthquake-station plane and the two-station plane, representing the deviation of the source from the great-circle path (see Fig. 3a)

6, “microseisms (1–50 period)” should be “microseisms (periods 1–50 s)”

corrected

7,“205 large earthquakes (≥M 6.8)” should be “205 large earthquakes (M ≥ 6.8)”

corrected

Referee 2:

The manuscript presents a simplified geometrical analysis of coda wave correlation retrieval of body wave phases that sample the Earth's core. The key topic it seeks to address is the reliability of the travel time information when sources are not near equally distributed. Analytic perturbation analysis, numerical experiments, and an observational example from a decade of data across much of North America provide multiple useful perspectives. Varying truncation angles are shown to influence the recovered travel time information for ScS and PKIKP2, with the small threshold angles of PKIKP2 leading to more stable results across different deviation angle ranges.

We thank Referee 2 for the thoughtful review and for highlighting the important distinction between measurement stability and true arrival time equivalence. We appreciate the constructive feedback that has helped strengthen the manuscript.

I think the results provide practical insights into core phase interferometry. However, in some case stability of results is too readily accepted as an indication that these measurements can be considered true empirical arrival times that are physically interpretable like measurements from earthquake sources. The analysis does represent a step forward on a difficult problem and I think it could be a valuable contribution to literature with minor revisions. More specific feedback to motivate revisions is provided below.

Thank you for pointing out this important distinction. We agree that stability alone is not sufficient to equate our measurements with direct earthquake-derived arrival times. To address the issue of accurate travel-time reconstruction, we have synthesized the necessary conditions as follows:

Revision (lines 260–278, as detailed in response to Referee 1, Point 4c):

A long-standing debate in coda correlation seismology centers on whether extracted phases represent true inter-station PKIKP^2 arrivals (the Green's function) or a modified wavefield (I2∗) whose travel times exhibit a systematic dependence on source distribution (Wang & Tkalcic, 2020a, b; Costa de Lima et al., 2022). Based on our theoretical analysis and empirical results, we propose that the transition from a biased I2∗ measurement to a true PKIKP^2 phase travel time occurs when the following two conditions are jointly satisfied:

(i) Illumination condition for stable reconstruction: The angular range of incident waves meets or exceeds the phase-specific critical angle θ_0, ensuring that the stationary phase zone is adequately sampled. This condition depends on the number of earthquakes, focal mechanisms, their azimuthal distribution relative to the inter-station path, and the distribution of inter-station directions stacked per distance bin. The degree to which this condition is satisfied can be evaluated quantitatively through bootstrap convergence analysis (Fig. 8), where stabilization of travel times with increasing event count indicates that the illumination condition has been achieved. It can also be assessed by examining the convergence of the correlation signals across different deviation planes.

(ii) Smoothness condition for accurate travel time recovery: The structural setting along the ray path must yield a sufficiently smooth deviation function δ(θ,p), rendering higher-order terms in Eq. (19) negligible. Seismic ray theory assumes a smooth velocity structure (Chapman, 2004), which produces a smooth wavefront and thus smooth travel time variations from the source to the receiver and its surroundings — which in turn ensures a smooth δ(θ,p). Our analysis therefore remains valid within the ray-theoretical framework. Although this smoothness condition cannot yet be independently verified in the real Earth, it is a necessary prerequisite for interpreting stable measurements as true arrivals, pending future validation against earthquake-derived empirical travel times.

60-65. What is the justification for assuming that all reflection wavefields for each source are uncorrelated? Does this require a minimum separation distance depending on the period? It may seem like a detail, but it is potentially relevant to the observational case in which seismicity tends to be clustered along plate boundaries.

Thank you for this important question. We have clarified this assumption in the revised manuscript.

Revision (lines 70–72):

In practical coda correlation analysis, researchers compute cross-correlations of late coda waves generated by individual earthquake events and subsequently stack the resulting CCFs for each station pair to enhance coherent arrivals.

This procedure does not account for the separation distances between event pairs.

176-180. At least some basic processing steps should be given for the observational correlograms as it is needed to facilitate reproducibility.

Revision (lines 202–212):

We selected 205 large earthquakes (M ≥ 6.8) from 2010 to 2020 with global distribution. For each event, we downloaded broadband waveforms from stations in the USArray Transportable Array. The spatial distribution of the earthquakes and stations is shown in Figs. 6a and 6b. For the late coda radiated by each earthquake, instrument responses were removed, and the data were bandpass filtered between 0.02 and 0.07 Hz (15–50 s period) to retain core-sensitive body wave energy. For each earthquake–station pair, we extracted the late coda time window from 10,000 to 40,000 s after the origin time. This window contains dominant multiply scattered waves that have sampled deep Earth structures. For each station pair and each earthquake, we computed the cross-correlation of the coda waveforms following the procedure of Bensen et al. (2007). We then calculated the deviation angle φ between the earthquake–station plane and the two-station plane. Correlograms were stacked within selected φ ranges and inter-station distance bins (bin width 1°). Arrival times of target phases were picked from the stacked correlograms using automated peak detection with manual verification. The number of cross-correlation functions stacked in each bin is statistically summarized in Fig. 6c.

181-184. the similarity illustrates stability, not necessarily convergence to an accurate Green’s function

Thank you for your thoughtful comment. We agree that similarity alone does not guarantee convergence to the accurate Green's function. In the revision, we have clarified that accurate reconstruction of inter-station travel times additionally requires δ(θ, p) to be smooth—an assumption that is justified within the ray-theoretical framework.

Revision (Lines 272-278):

(ii) Smoothness condition for accurate travel time recovery: The structural setting along the ray path must yield a sufficiently smooth deviation function δ(θ,p), rendering higher-order terms in Eq. (19) negligible. Seismic ray theory assumes a smooth velocity structure (Chapman, 2004), which produces a smooth wavefront and thus smooth travel time variations from the source to the receiver and its surroundings — which in turn ensures a smooth δ(θ,p). Our analysis therefore remains valid within the ray-theoretical framework. Although this smoothness condition cannot yet be independently verified in the real Earth, it is a necessary prerequisite for interpreting stable measurements as true arrivals, pending future validation against earthquake-derived empirical travel times.

Figure 7. I'd suggest adding theoretical travel time curves for the phases discussed and/or add a panel with 1-D Green's functions. Either or both

would help highlight that some phases agree well with the Green’s function and others do not.

We have added theoretical travel time curves to Figure 7, calculated using the Taup software with the ak135 reference model.

Can the perturbation analysis help predict any specific artifacts found in observational correlograms?

This is an insightful question. The stationary phase analysis of Kennett & Phạm (2018) demonstrates that certain artifacts in coda correlations arise from correlating different wave types that share the same slowness. Our perturbation analysis focuses on the recovery accuracy of individual phases within a general medium structure. It does not directly address correlations between different wave types and therefore cannot predict the formation of such artifacts.

Figure 9. The correlograms look similar but another quantitative metric would be helpful, such as the standard deviation of travel time shifts between the 0-10 and 0-90 degree signals. Some core studies look for very deviations (although often with shorter period data), so it would be good context.

Thank you for this practical suggestion. We have added quantitative metrics in **Figure 10**.

The text could be clearer about whether the difference in lags time between ScS and PKIKP2 in Figure 10 can be quantitatively predicted or is just broadly consistent with the expectations.

We have clarified this point in the revised text. The observed difference in lag-time sensitivity between ScS and PKIKP^2 in Figure 10 is broadly consistent with the theoretical expectations from our Γ-based analysis. While our framework provides a quantitative prediction of the relative stability, we cannot yet provide a precise quantitative prediction of the absolute travel-time bias for each phase.