

Supplementary Discussion

Section S1: Additional discussion of model resolution and uncertainties

Section S1.1: Linear resolution analysis

Recovery of larger scale columnar structures

A common pitfall of linear resolution analysis is to examine recovery of only one scale of test structure. Indeed, while the two resolution experiments we have discussed thus far have focused specifically on plume-like input structures, another important question is whether significantly larger scale anomalies may be erroneously “broken up” due spatial variation in sensitivity and thereby give rise to plume-like anomalies.

To examine this possibility more closely, we now turn to experiments focused on recovery of large contiguous domains of reduced Vs. Here, we designed a simple test structure, composed of a cylindrical anomaly 5000 km in diameter extending from the CMB to 1500 km depth and centered roughly 10° north of Tahiti, as might be expected if the LLSVP extended as a unit that high into the mantle. As before, the test structure has a maximum amplitude of -2%, but instead of a cosine-cap, the lateral profile of this anomaly is more like a boxcar – or, more specifically, as close to a boxcar as possible given the intrinsic smoothness of our spline basis. In Fig. S1, we show both the input test structure and the results of our resolution analysis, and we observe impressively uniform retrieval of the input anomaly, without artificially giving rise to “plume-like” anomalies similar to those imaged in SEMUCB-WM1 (either as a result of the large anomaly “breaking up” or in the form of plume-like anomalies smearing vertically off its upper surface).

We find that depth variation of amplitude recovery is very similar to that seen in our isolated-plume resolution tests (Extended Data Figs. 4 and 5, while lateral variation is also quite uniform, with the exception of a small patch at the southern-most end of the anomaly at mid-mantle depths. The latter can be seen as a local maximum in retrieved amplitude at the southern end of the S-N trending cross section in Fig. S1, as well as in the 2000 and 2500 km map views of the anomaly as a W-E trending band at approximately the same latitude as the Macdonald hotspot. Indeed, this feature is clearly an artifact of the spherical spline parameterization, as it follows *exactly* the edge of the icosahedron on which the grid of spherical spline knots is registered. Specifically, this is a known limitation of our parameterization when representing large geologically implausible constant-amplitude domains with sharp edges, such as those used in our input structure, and relates to the fact that (a) the spherical spline basis is not spectrally complete and (b) the spherical spline knot grid is locally deformed in the neighborhood surround the edges of the icosahedron. Regardless, we note that this band-like feature does not resemble any of the plumes appearing in our model – leading us again to conclude that the plume-like anomalies in SEMUCB-WM1 cannot be attributed simply to poor spatial resolution resulting in patchy recovery of broad domains of low Vs.

Finally, while the above test is focused on recovery of large-scale domains of reduced Vs that extend from the CMB through most of the height of the lower mantle, we also wish to probe the effect of a thinner, stronger anomaly concentrated closer to the CMB.

With this in mind, we present in Fig. S2 the results of a resolution analysis employing a 600 km high cylindrical anomaly at the base of the mantle, again with fairly sharp edges, a diameter of 5000 km, and centered in the same location, but now with maximum amplitude of -3%. Ideally, we would like to observe that this structure is well recovered, showing both (a) no evidence of artificial “plume-like” anomalies resembling those imaged in SEMUCB-WM1 (either embedded within or streaking / rising from the top of the retrieved structure) and (b) a depth extent that is well-confined to that of the input structure (given the often noted tradeoff between anomaly height and amplitude). Indeed, upon examining the results in Fig. S2, we confirm both of these expectations to be the case – again reinforcing the robustness of the plume-like anomalies imaged in SEMUCB-WM1, both in their morphology and their depth extent.

Section S1.2: Resampling analysis

While we have thus far focused on linear resolution analysis as a tool for probing model uncertainties, statistical resampling techniques that operate in the data space – such as the deleted jackknife⁵² – are another valuable tool for placing semi-quantitative bounds on uncertainty. Indeed, we chose to use the deleted jackknife for this very purpose in our analysis of SEMUCB-WM1⁴ and here revisit those results in the context of the plume-like features discussed in the main text (we refer the reader to [4] for a more thorough discussion of how the test was configured and our results). In Fig. S3, we show the resulting standard error estimates for the Vs component of the SEMUCB-WM1 model in map view at a range of depths, along with both the maximum and RMS uncertainty amplitudes at each depth. Importantly, we observe that the RMS uncertainty estimates at each depth are *well below* the amplitudes typically associated with the plume-like structures discussed in the main text: at least 0.5 % anomaly strength maintained throughout the mantle, with local amplitude maxima above 2 %. Indeed, these are consistent with the thresholds used in assessing plume continuity (see Extended Data Table 1).

Section S1.3: Other assessments of model quality and sources uncertainty

In [4], we provide a thorough discussion of model quality and potential sources of uncertainty, from which we briefly summarize relevant aspects here, referring the reader to [4] for a more detailed treatment (specifically, Sections 5.1 and 5.2 thereof).

- *Performance on held-out data*: waveform variance reduction for held-out waveform data (i.e. data not included in our inversion) was computed with the SEM, and it was determined that these data exhibit fits quite similar to the inversion data – indicative that we have not over-fit the inversion data (specifically, noise therein) at the expense of model correctness.
- *Performance limitations due to source errors*: SEM-based source inversion tests using the SEMUCB-WM1 model to correct for earth structure demonstrated that a subset of our events could benefit from source inversion. However, of the 20 test events considered, only one exhibited changes in body wave variance reduction in excess of 20 %, with the remainder showing improvements that are far more

- modest. This should be revisited if SEMUCB-WM1 is to be used as the starting model for an anelastic inversion.
- *Uncertainty in scaling factors*: tests using different configurations of (depth-dependent) $\ln V_p$ to $\ln V_s$ scaling factors found that our inversion is only weakly sensitive to uncertainty in this parameter, consistent with the results obtained by [53].

Section S2: Detection of plumes – Effects on the wavefield

Using the tools of linear resolution analysis, we have shown that plumes of dimension and amplitude similar to those imaged in SEMUCB-WM1 should indeed be recovered in our inversion, while we have also ruled out their being attributable simply to smearing or poor resolution of other classes of structure (e.g. large, idealized LLSVP-like anomalies). However, this approach is strictly valid only under the assumption of linearity⁴¹, and in reality only allows us to test the behavior of the last iteration of the inversion. A more rigorous test would require reproducing the complete iterative procedure employed during the development of this model with the goal of recovering a known test structure by inverting synthetic data, which would take many months. In what follows, we instead choose to illustrate the effect of plume-like structures on the seismic wavefield at periods relevant to our inversion (≥ 32 s) through synthetic forward-modelling experiments using our highly accurate modeling method of choice (SEM).

Forward modeling experiments: Setup

For simplicity, we choose a column-like anomaly extending from the CMB to the surface and of 600 km in width, quite similar to the narrower of the two test structures used in the resolution analyses above. In our earlier resolution tests, we found that a 2 % maximum amplitude plume of this dimension led to ≤ 1.5 % maximum recovered amplitude under best-case conditions (Hawaii). Since we observe ~ 2 % peak amplitudes in the lower mantle associated with the plumes in SEMUCB-WM1, we chose -3 % as the maximum amplitude for our synthetic plume-like structure (i.e. assuming a best-scenario plausible amplitude recovery ratio of 3/2). For comparison, we also prepared a second test structure with identical dimensions, but amplitude doubled to -6 %. In both cases, the simplified plume-like anomaly is placed beneath Pitcairn, owing to the primacy of this plume among those imaged in SEMUCB-WM1. The source and station geometry in our modeling test is taken directly from the SEMUCB-WM1 inversion setup (i.e. we employ a realistically distributed set of sources and receivers ringing the Pacific), and the SEM used in these synthetic forward modeling experiments is configured identically to that used in our inversion.

There are of course caveats to this approach. Chief among them is the simple geometry of our chosen anomaly – a columnar “plume” in an otherwise 1D background model (namely, the reference model for SEMUCB-WM1). While this is vastly simplified compared to using synthetic plumes based on hypothetical geodynamic models¹, such as in [54], it makes no implicit assumptions as to the veracity of the latter, and is closer to the approach taken by [55] who also used a cylindrical plume-like anomaly – albeit

stronger and sharper given their focus on the diffracted wavefield. Further, our synthetic test includes no noise. Indeed, realistic (correlated) noise would be a valuable addition if, for example, we were to use the resulting waveforms as “data” for a synthetic inversion. All of this said, the results are still quite informative – particularly as they allow us to examine certain wavefield phenomena that would otherwise be difficult to attribute unambiguously to the presence of a plume.

Forward modeling experiments: Results

In Figs. S5 – S7, we begin by examining a suite of T-component waveforms recorded at three different stations for one particular shallow test event (2 Nov. 2003, Mw 6.4 event off the west coast of the South Island of New Zealand; GCMT C110203C; 15 km depth). These particular stations (Fig. S5) were selected in order to illustrate the effect of the simplified plume model discussed above on paths that progressively transition from passing directly through the plume (GEOF.BOA), to grazing incidence (GSN.BCIP), to immediately adjacent to the plume (GSN.OTAV). Namely, at the relatively long periods of our modeling, we observe in Fig. S6 a steady progression from waveform effects dominated by focusing and very weak phase delays for paths directly traversing the plume, to defocusing and multipathing (e.g. secondary phases in the body wave coda) for paths grazing or adjacent to the -3 % amplitude plume. These effects become more pronounced in Fig. S7, where the anomaly amplitude is now doubled to -6 % in order to better highlight the amplitude variation and multipathing seen earlier at -3 %. We observe that these coda phases remain quite prominent even for paths passing within hundreds of km outside of the plume – owing to the finiteness of the zone of sensitivity at our chosen period range.

While these examples serve to illustrate these effects on direct and singly or doubly surface-reflected phases commonly recorded from shallow events, this does not represent the full breadth of the waveform dataset used in our inversion. Indeed, the deep events therein generate phases that provide complementary illumination of lower mantle structure, including long trains of multiple ScS phases (and associated depth phases), which reverberate between the surface and CMB. With this in mind, we present in Figs. S8 and S9 waveform examples for a path similar to GSN.OTAV in our earlier example (passing immediately adjacent to the plume, as opposed to passing through or grazing it) recorded from a particular deep test event (the 4 Sept. 1997, Mw 6.8 Fiji event; GCMT C090497A; 621 km depth). In Fig. S9, we once again observe pronounced waveform complexity due to multipathing along this plume-adjacent path, highlighting particularly well how the trains of multiple-ScS phases provide additional sensitivity to this plume-like anomaly.

We have seen here that both amplitude effects (focusing, defocusing) and waveform complexity due to multipathing (scattering, diffraction) provide strong evidence for the presence of a “plume” in our synthetic experiments. Thus, we have demonstrated that plumes of a dimension previously shown to be retrievable in our resolution tests, and with strength motivated by both the apparent amplitude recovery in the latter and the amplitudes of the plumes imaged in our model, *should clearly be detectable in our waveform data* (i.e. using the same period range and realistic source and station distribution). Further, arrival time delays for distinct phases in these synthetic data are

extremely weak – likely owing to finite frequency effects such as wavefront healing, which is to be expected at these relatively long periods. This latter observation implies that while the “signal” of these plumes should be readily visible to waveform inversion, which is able to make full use of the information contained in the wavefield, including waveform amplitudes and complex interacting phases, the same cannot be said of classical travel-time inversion (again, with the caveat that these are relatively long periods).

Supplementary Figures

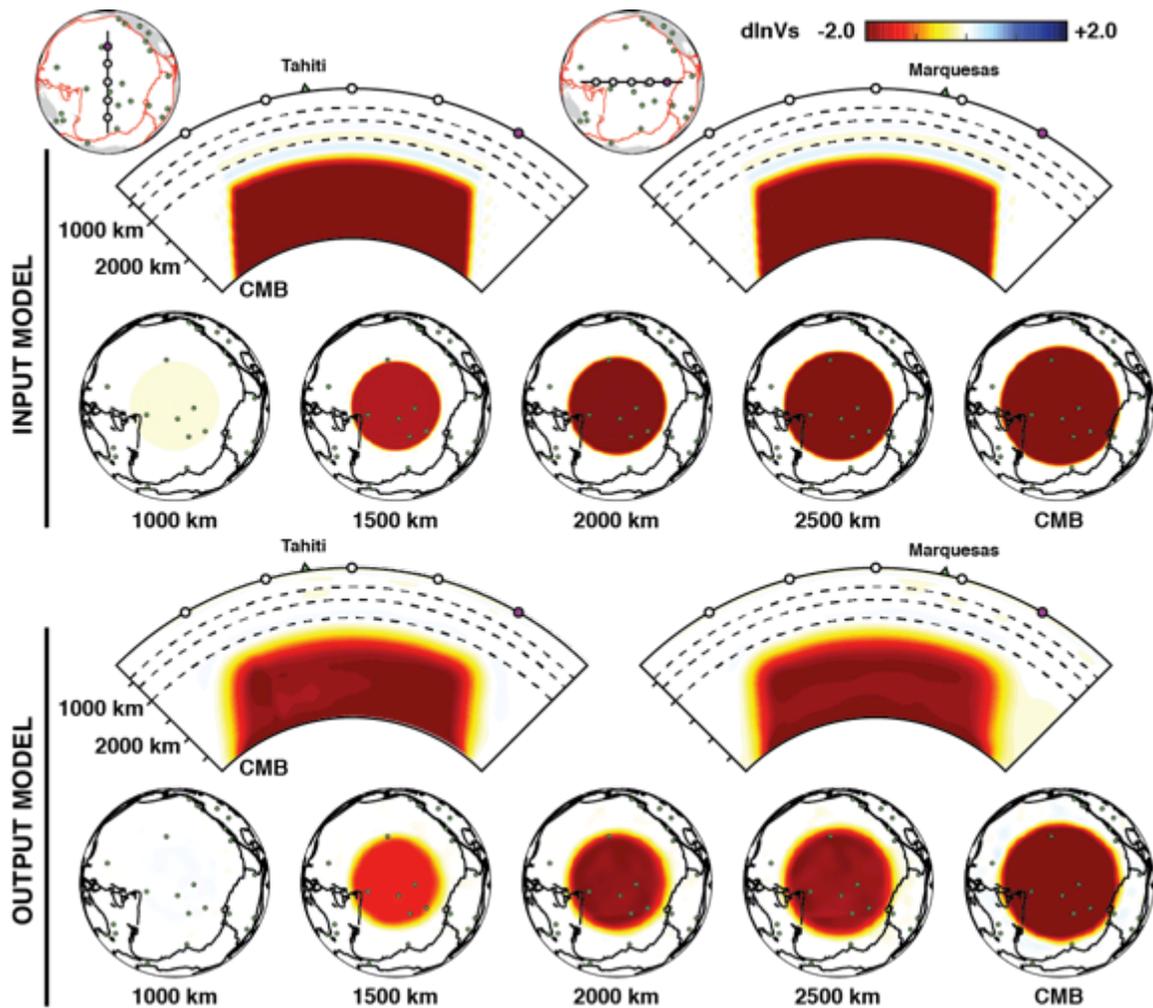


Figure S1: Linear resolution analysis focused on recovery of large contiguous domains of reduced V_s . Upper panels: Input structure in both cross-section and map view – a cylindrical anomaly of diameter 5000 km extending from the CMB to 1500 km below the surface (maximum amplitude -2%). The anomaly is centered roughly 10° north of Tahiti. Bottom panels: Output structure presented in the same manner, illustrating impressively uniform retrieval of the input anomaly, without artificially giving rise to “plume-like” anomalies similar to those imaged in SEMUCB-WM1. Depth variation of amplitude recovery is quite similar to that seen in our isolated plume resolution tests (Extended Data Figs. 3 and 4). Lateral variation in recovery is also quite uniform, as seen in both cross section and map view, with the exception of the small patch at the southern-most end of the input anomaly at mid-mantle depths. The latter can be seen in the 2000 and 2500 km map views of the anomaly as a W-E trending band at approximately the same latitude as Macdonald (it can also be seen in the S-N trending cross section). This is an artifact of the spherical spline parameterization (following exactly the edge of the icosahedron on which the grid of spline knots is based), and is a known limitation thereof when

representing large geologically implausible *constant-amplitude* domains with sharp edges, such as those used in our input structure.

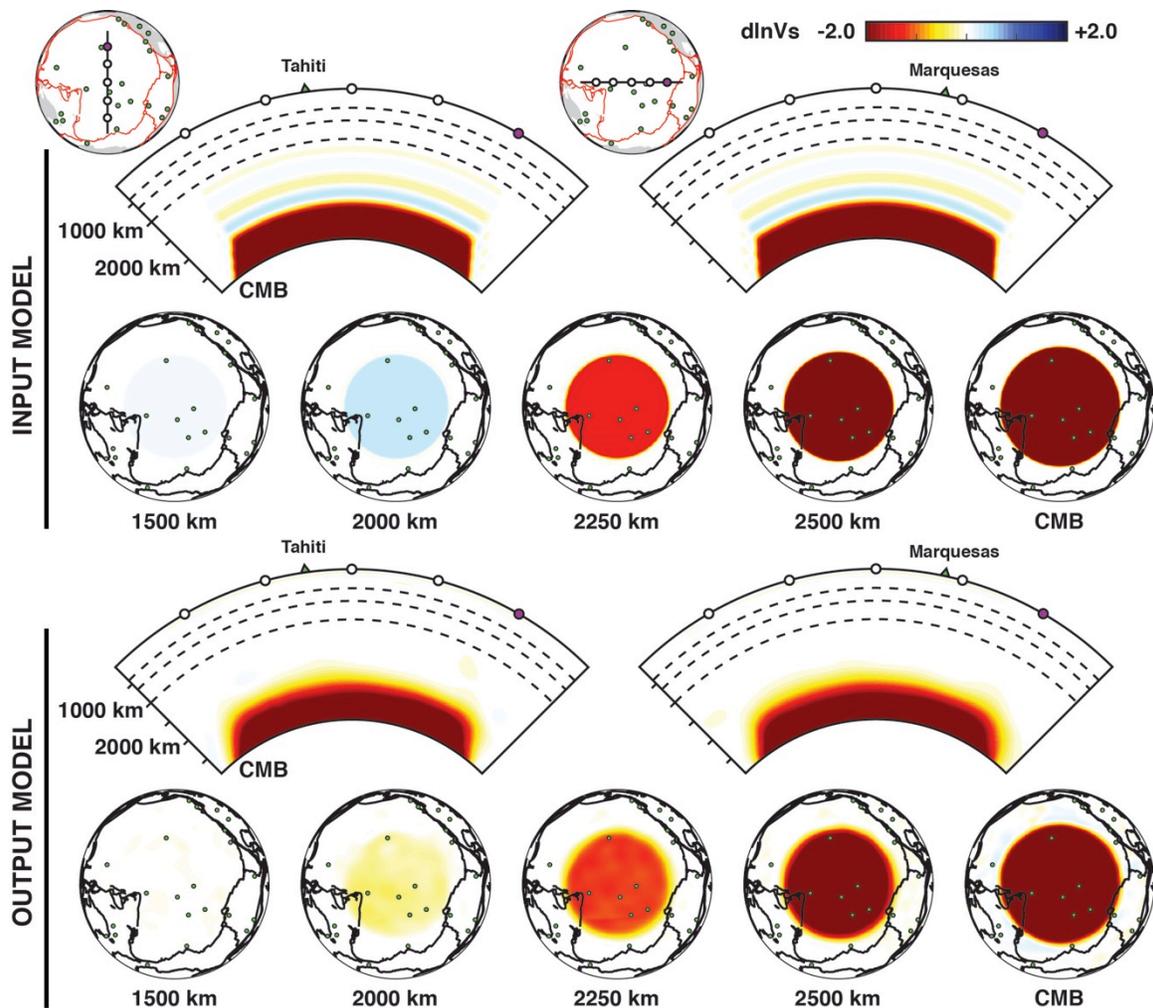


Figure S2: Linear resolution analysis focused on recovery of large contiguous domains of reduced V_s similar to Fig. S1, but now featuring a thinner, stronger anomaly. Upper panels: Input structure in both cross-section and map view – a 600 km high cylindrical anomaly of diameter 5000 km extending up from the CMB, with maximum amplitude -3%. The anomaly is again centered roughly 10° north of Tahiti. Bottom panels: Output structure presented in the same manner, again illustrating impressively uniform retrieval of the input anomaly. Note that although the input structure saturates at -3%, we intentionally over-saturate these plots at -2% in order to better capture low-amplitude artifacts, should they exist. Despite this, we see no evidence of artificial “plume-like” anomalies resembling those imaged in SEMUCB-WM1, either embedded within or rising from the top of the input structure. At the same time, we do note some degree of Gibbs-like “ringing” above our input structure from aliasing of its sharp upper boundary by our cubic b-spline basis. However, this weak ($< 0.4\%$ peak amplitude) short-wavelength parameterization artifact is not retrieved in our output structure – largely owing to our chosen model covariance operator, which was designed to prevent artifacts not only of uneven data coverage, but also Gibbs-like phenomena due to our parameterization.

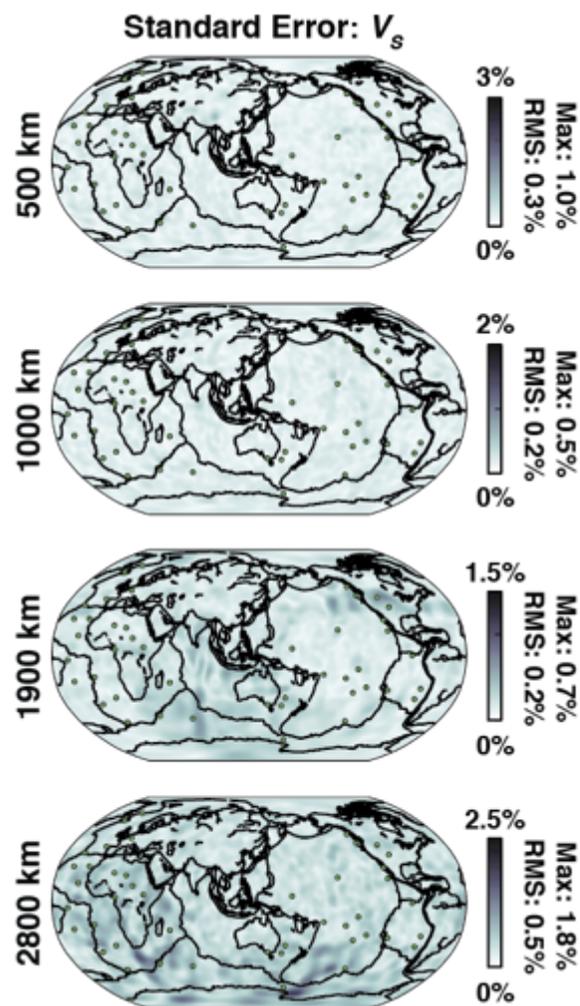


Figure S3: Standard error estimates of the V_s component of the model SEMUCB-WM1 obtained using jackknife resampling (modified from [4]). See suppl. text section S1.2 for details. At each depth, we show standard error estimates in map view (utilizing the same color scale bounds as used to present the model structure in [4]), along with the maximum and RMS uncertainty amplitudes at each depth. We note that the RMS uncertainty estimates at each depth shown here are well below the amplitudes typically observed for the plume-like structures imaged in this work (maxima often above 2% anomaly strength, in general above 0.5% throughout the mantle), as well as the thresholds used in assessing plume continuity (see Extended Data Table 1).

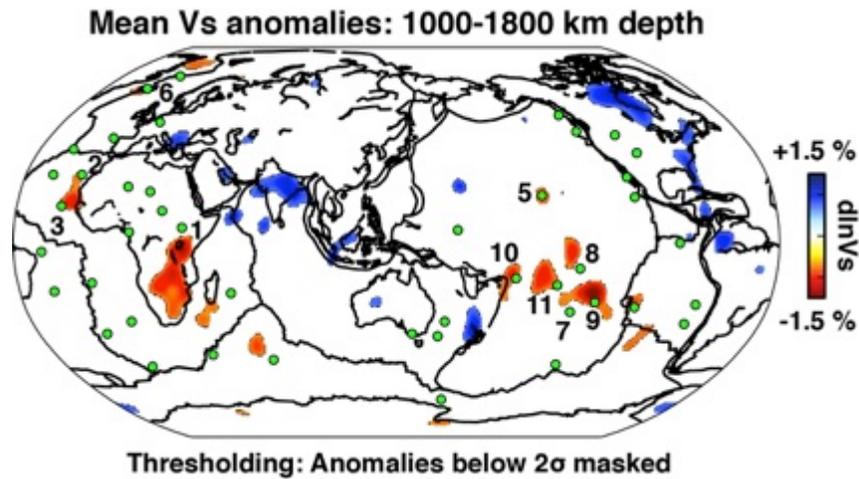


Figure S4. Map view of relative Vs anomalies in SEMUCB-WM1 averaged between 1000 and 1800 km depth, where anomalies weaker than the 2σ level (computed based on lateral variation over this mid-mantle depth-averaged map) have been masked (i.e. only *outlier* features with respect to the average structure of the mid-mantle are shown). We find that all of the "primary" plumes of Fig. 4 in the main text correspond well with some of the most-significant low-velocity anomalies in this mid-mantle average (and have been numbered accordingly). Green circles denote hotspot locations²⁷.

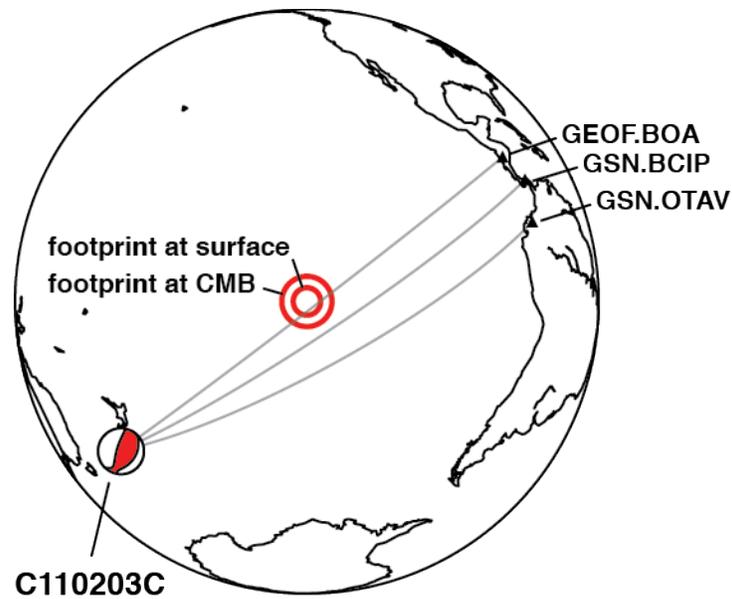


Figure S5: Geometry for the first synthetic forward-modeling experiment discussed in Suppl. Section S2, showing the source location and mechanism (depth: 15 km), receiver locations, and synthetic plume location. The plume has a constant diameter of 600 km, meaning that its angular width changes as a function of depth. Thus, we show its lateral extent projected at both the surface and the CMB.

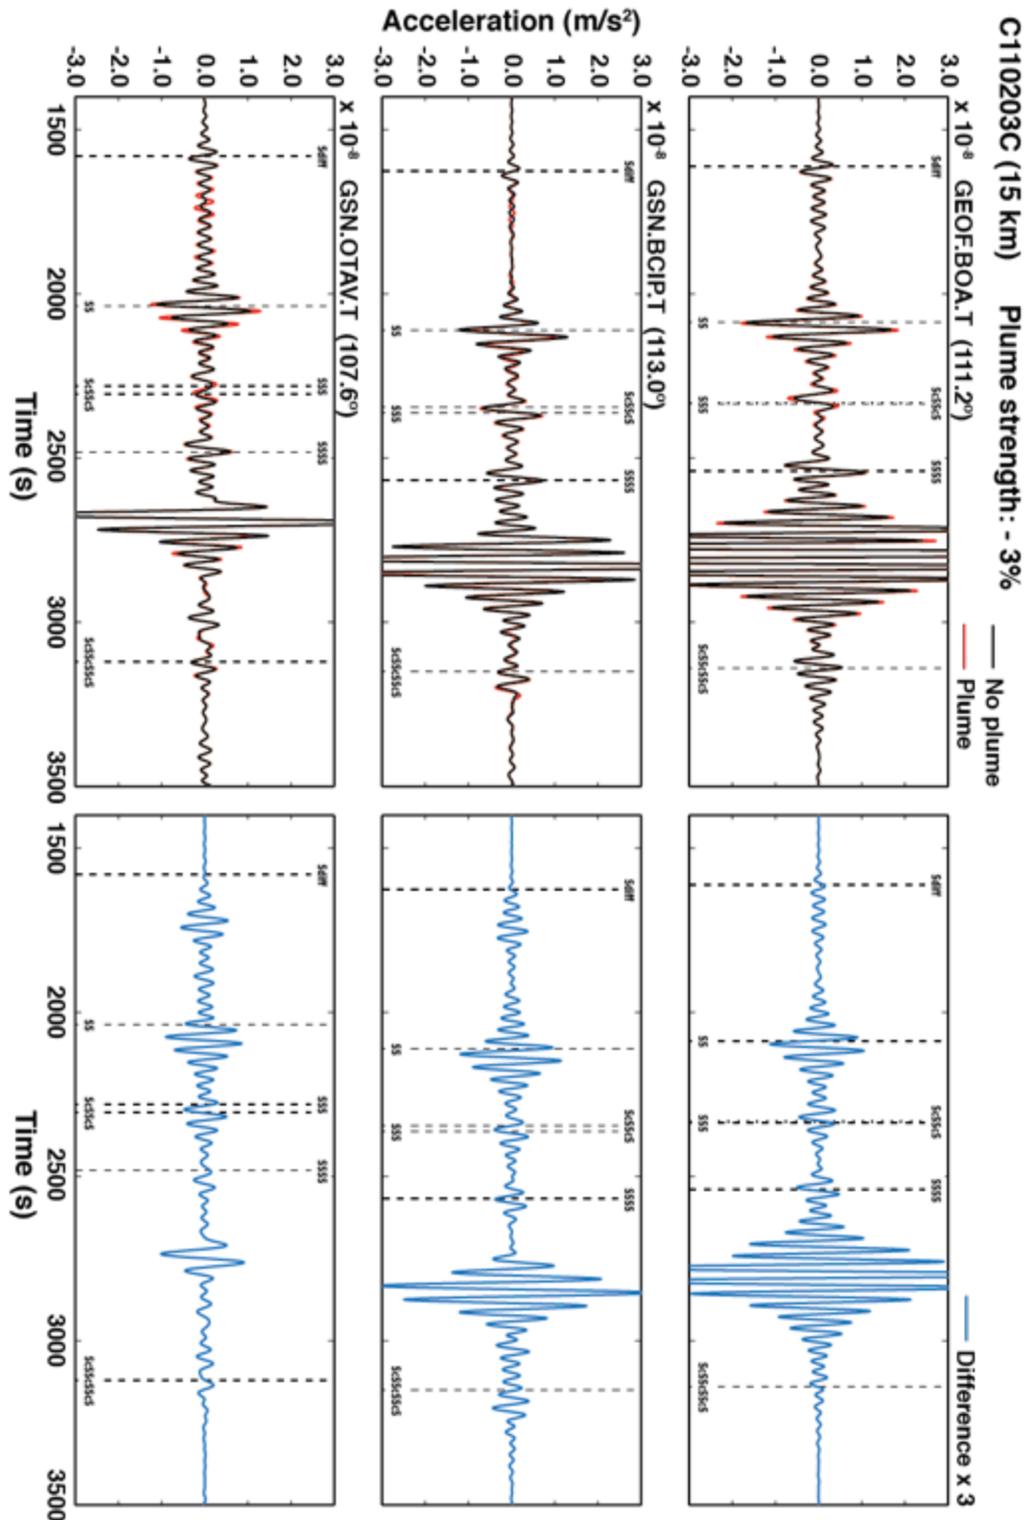


Figure S6: T-component waveform comparisons demonstrating the effect of a plume of width 600 km and -3% maximum amplitude at periods relevant to our inversion. See Suppl. Section S2 for details. Left panels: SEM synthetics with (red) and without (black) the presence of a plume, recorded at the three stations shown in Fig. S5. Right panels:

difference waveforms between plume and plume-free synthetics at left amplified 3x. As noted in the suppl. text, we observe a steady progression from waveform effects dominated by focusing and phase delays for paths directly traversing the plume, to defocusing and multipathing (e.g. secondary phases in the body wave coda) for paths grazing or adjacent to it. Amplitude effects and waveform complexity provide key constraints on earth structure that are readily useable in waveform inversion. Further, we find that the anomalous phases appearing in the residual waveforms shown at right are typically within a factor of only 2-3x smaller in amplitude relative to the direct phases with which they are associated – implying that they are unlikely to be systematically buried in noise. Note also that the time axis is compressed in order to include both body waves and surface waves, making it difficult to see phase delays in the left panels.

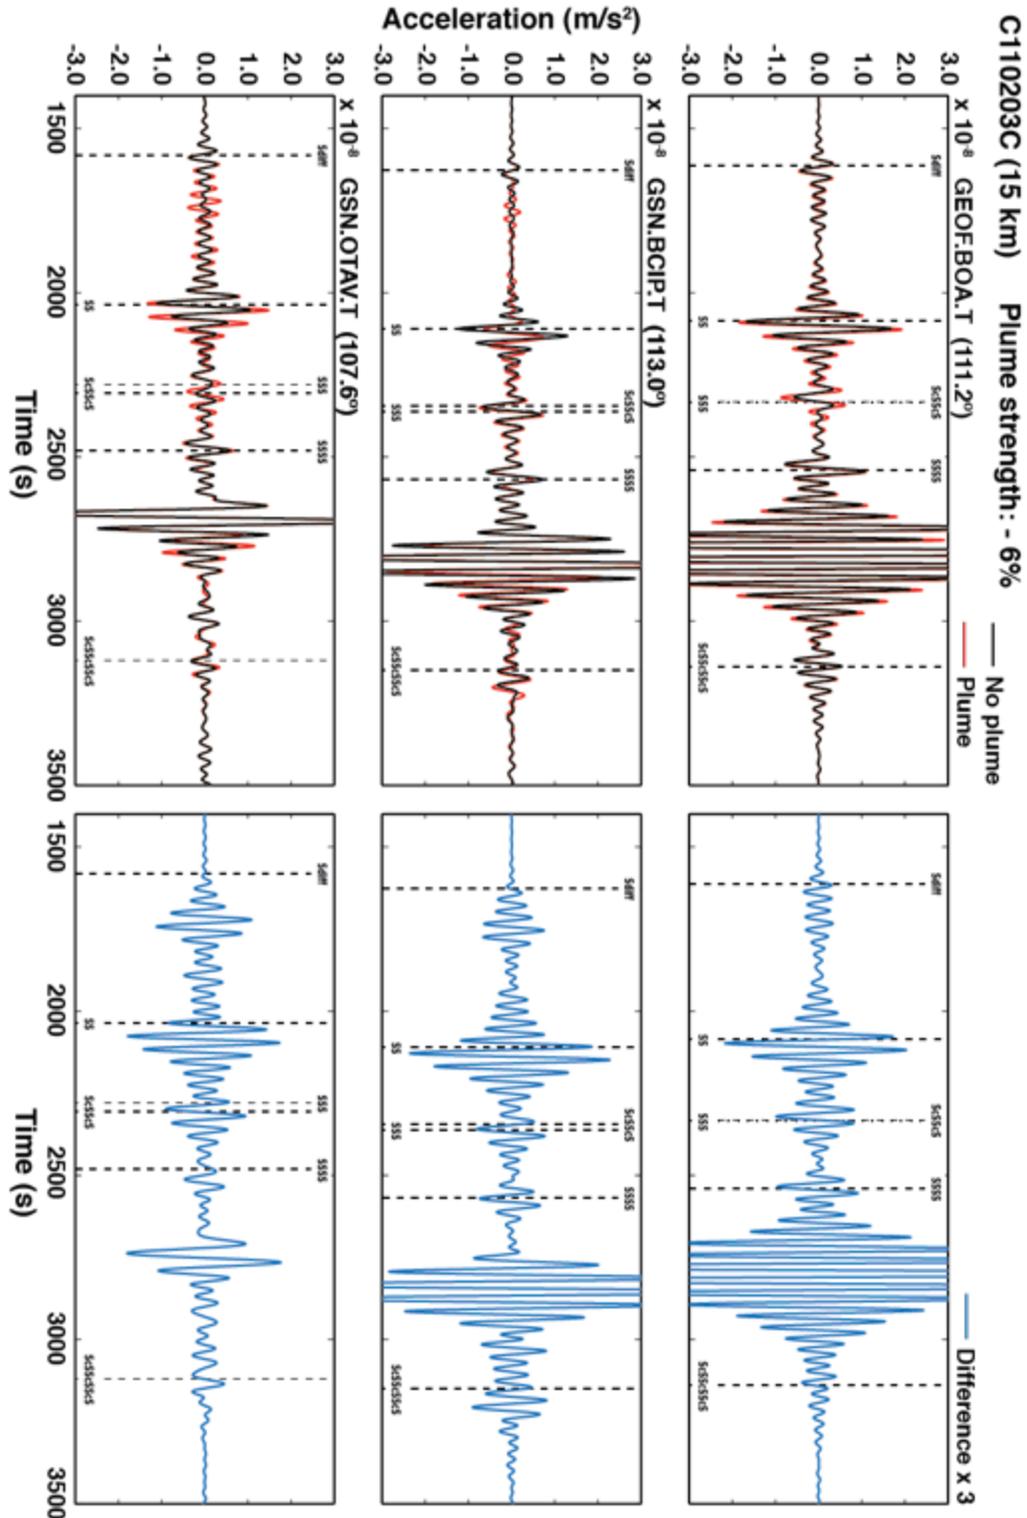


Figure S7: T-component waveform comparisons demonstrating the effect of a plume of width 600 km width and -6% maximum amplitude at periods relevant to our inversion, presented in a similar manner to Fig. S6. See Suppl. Section S2 for details. We again observe a progression from focusing and phase delays for paths directly traversing the plume, to defocusing and multipathing for paths grazing or adjacent to it.

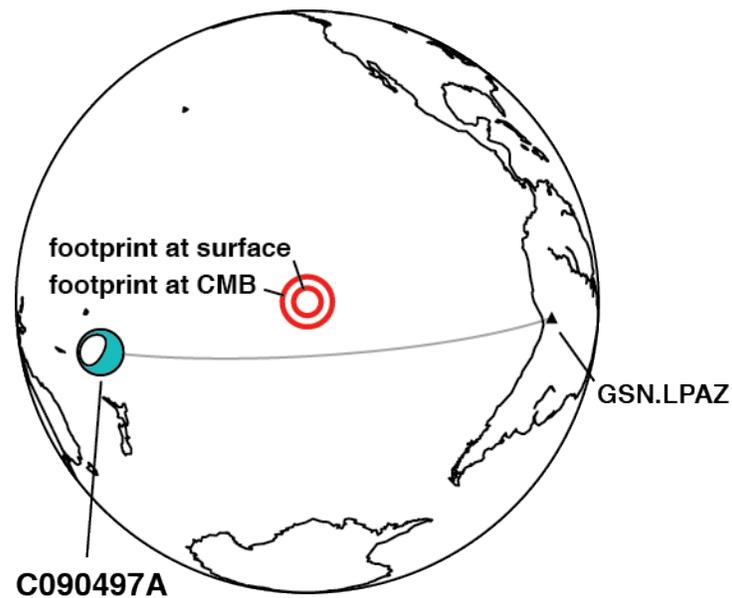


Figure S8: Geometry for the second synthetic forward-modeling experiment discussed in Suppl. Section S2, showing the source location and mechanism (depth: 621 km), receiver location, and synthetic plume location. As in Fig. S5, we show its lateral extent projected at both the surface and the CMB.

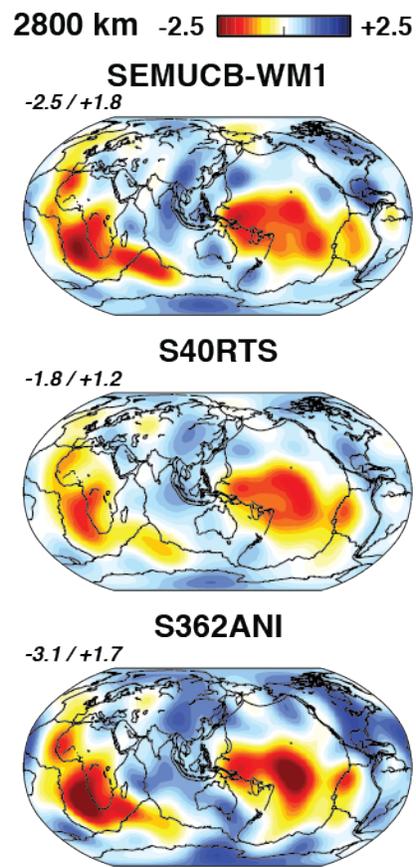


Figure S10. Comparison of shear velocity maps at 2800 km depth, filtered to long wavelengths (max degree 12 in a spherical harmonics expansion of the models). Top: SEMUCB-WM1, middle: S40RTS⁴⁶, bottom: S362ANI⁵³. Maximum and amplitude variation in each filtered model inset. The long wavelength features in SEMUCB-WM1 are in good agreement with those of models developed using more classical approaches.

Supplementary References

52. Efron, B. & Stein, C. The jack knife estimate of variance, *Ann. Stat.* **9(3)**, 586–596 (1981).
53. Kustowski, B., Ekström, G. & Dziewonski, A. Anisotropic shear-wave velocity structure of the Earth's mantle: a global model, *J. geophys. Res.* **113**, 303–318 (2008).
54. Hwang, Y.K., Ritsema, J., van Keken, P. E., Goes, S. & Styles, E. Wavefront healing renders deep plumes seismically invisible. *Geophys. J. Int.* **187(1)**, 273–277 (2011).
55. Rickers, F., Fichtner, A. & Trampert, J. Imaging mantle plumes with instantaneous phase measurements of diffracted waves. *Geophys. J. Int.* **190**, 650–664 (2012).