Caveats on tomographic images: Supplementary Material

Gillian R. Foulger, Giuliano F. Panza, Irina M. Artemieva, Ian D. Bastow,
Fabio Cammarano, John R. Evans, Warren B. Hamilton, Bruce R. Julian,
Michele Lustrino, Hans Thybo, Tatiana B. Yanovskaya

Tomography Perils and Pitfalls

Many issues arise when interpreting a tomographic model. It is a highly imperfect "lens" through which to view Earth. The extent and causes of this imperfection are understandably vague for geologists, geochemists, and others outside the specialty, but even many currently active tomographers seem to be unaware of many of its perils and pitfalls or to misjudge their extent.

Here we restrict ourselves to a terse description of issues arising in teleseismic tomography, the use of the terminal portions of *P*, *S*, and core phases to estimate regional wave speed structure (or *Q*) beneath a seismograph array. It is the earliest, simplest, and most linear of current tomography methods. In this Supplement we do not address local-earthquake or global tomography, which are highly nonlinear problems due to the presence of ray turning points inside the model volume. The presence of turning points leads to first-order interactions between wave speed and raypaths, requiring a conservative approach to the inversion process and serious risk of falling into a "local minimum", *i.e.*, a solution that is far from broadly optimal. Neither do we address surface-wave tomography or methods used in refraction or reflection seismology, though analogous problems will be present, generally turned 90° in space from those discussed here.

Background: What is tomography?

The data used are variations from some norm (some standard Earth velocity model) in the arrival times of *P*, *S*, or some other phase from distant earthquakes, as received at an array of seismographs that overlies the "target volume"; these are called "traveltime residuals". For teleseismic tomography, the target volume comprises some portion of the crust and upper mantle for which one wishes to know the wave-speed variations ("anomalies"). In practice one uses "relative" travel-time residuals (typically by subtracting the mean residual). This reduction is permissible because teleseismic tomography cannot recover absolute velocities in the target volume (therefore cannot predict absolute traveltimes). The reduction also is desirable because it improves the accuracy of the measured relative traveltime residuals.

One next creates some description of an initial wave-speed structure inside the target volume. That structure inevitably simplifies the volume and its structure, commonly into a set of constant-wave-speed blocks or nodes, each given initial wave speeds that are reasonable for that part of the world at that depth (to put the rays in the right places). Those initial wave speeds typically are for a simple layered Earth model. This rather profound simplification of Earth wave-speed structure into blocks or nodes within a layered Earth is called "parameterization" and imposes substantial restrictions and assumptions on the structure one eventually computes, imperfections that are not explicit in the subsequent mathematical solution.

The computation itself is a fitting process, equivalent to fitting a line to a scatter plot. (Indeed, deciding *a priori* that the scattered data must express some simple line is the equivalent of parameterizing the tomographic model and, of course, can be anywhere from a

good approximation to nonsensical. Regardless, any functional form is bound to be simpler than reality in a system as complex as the Earth.) The tomographic fitting process then extracts the best fitting wave-speed variations for each block or node in the parameterized target volume, based upon the traveltime data and where the associated rays travel within that volume.

Up to the present, the mathematical fitting methods used, "the inversions", typically use a least-squares fit with additional constraints beyond the observed data (typically by "damping" either the amplitude or smoothness of the model). Those additional constraints are required to select a particular result from an infinite number of possible ones (infinite for various mathematical and physical reasons, as showed for example by Backus and Gilbert, 1967, 1968, 1970). The selection of that additional constraint also can cause and obscure uncertainties in the resulting wave-speed model. So in addition to formal measures of fitting accuracy yielded by the inversion procedure, or by tests with synthetic data sets, one has to account for the effects of parameterization, selecting the additional constraint, and the inability of teleseismic tomography to extract absolute velocity. However, none of these appear in the numerical result in any clear manner, if at all.

Resolution Perils and Pitfalls — General

The resolution capabilities of a given data set (set of timed rays at various locations, backazimuths, and inclinations) must be tested, evaluated, and understood in detail. However, at present such matters typically are only tested in a cursory, optimistic manner. It is never sufficient to evaluate only the diagonal elements of **R** (the resolution matrix), or equivalent techniques for inversion methods not generating an explicit **R**; nor is it sufficient to do a "checkerboard test". Many of these issues are tested numerically and illustrated in the following section.

The checkerboard test is a widely used technique for estimating the resolving capabilities of a teleseismic tomography data set. One creates a synthetic data set of traveltime residuals for a checkerboard wave-speed structure (horizontally alternating blocks of high and low wave speed). Unfortunately, all damped teleseismic models have resolution kernels (equivalent to a column of **R** or the result for a synthetic model containing a single anomalous block) that oscillate along the horizontal direction. That is, they have negative values adjacent to the block of interest as well as smearing positively above and below the block (Fig. S1). Because of the shape of such resolution kernels, teleseismic tomography is highly compatible with horizontal-checkerboard structures. Put bluntly, checkerboard tests are the most optimistic possible evaluation of resolution. Actual resolution is inevitably worse, often much worse, for real Earth structures.

Some of the inherent drawbacks of the use of checkerboard tests and other simplifications may be overcome by synthetic modeling of realistic features. Test the degree to which the ray set can recover a plausible synthetic model of the main features of the region (*e.g.*, Alinaghi *et al.*, 2007). If the synthetic model is realistic, this approach allows the interpreter to test whether specific features may be fully or partly recovered from the real data, and how anomalies might smear together or into distal parts of the model.

At a minimum, proper evaluation of **R** requires one to understand the single-block "impulse" response (resolution kernels) resulting from their data set and to do this separately for all major regions of the model (Fig. S2). Resolution behavior generally can be assumed to be roughly similar within each of these regions and to vary more or less smoothly between them. Note especially that **R** and its kernels always will be poor below a depth approximately equal

to the array aperture, and will also be poor in the periphery of the model, beyond the array bounds. In these areas ray crossfire inherently is limited, and nothing can improve the situation. With a good data set one may be able to infer that some wave-speed anomaly lies along the dominant ray bundle at some distance anywhere from the surface to well beyond the model boundary, but there is no way to determine where the wave-speed anomaly is along the ray bundle. Furthermore, systematic timing errors, real anomalies outside well-resolved areas, and even strong anomalies within the well-resolved area commonly smear into these poorly resolved regions to satisfy the inversion's damping term by roughly accounting for those features of the traveltimes with weak, distributed features far from the true causal anomaly. Features in these areas are, therefore, always artifacts or indistinguishable from artifacts.

There are also widespread misconceptions of what constitutes a "good" set of observed rays (or of first-Fresnel zones, not distinguished here because they simply affect a smoothing operation so mainly benefit second-order features near the model's noise level). An optimal teleseismic tomography data set must everywhere contain roughly equal numbers of rays in all ray orientations and locations. That is, the number of rays is less important than their distribution of orientations at any given location in the model.

A common, related trap results from inhomogeneous source-event distribution, that is, from using data sets with most rays in bundles that have only a subset of possible orientations. For example, in the western U.S., nearly all viable rays originate along the Pacific rim (northwest, southwest, and southeast *P* and *S*) or at distances above ~140° (near-vertical core phases). In contrast, the northeast azimuth has very few earthquakes and inevitably has very few observed rays of good quality and fewer per event even when observable. Additionally, rays from the SW mostly come from deep sources in the south Pacific yielding impulsive waveforms that can be timed more accurately than the shallow-source phases available from the NE. This situation is equivalent to a properly "reversed" NW-to-SE-directed refraction line but an "unreversed" SW-to-NE-directed line. Features on the SW–NE plane always will be poorly resolved and strongly smeared along *P* and *S* ray bundles from the SW as well as, sometimes, vertically along the core phase bundle. Subsetting the SW ray set leaves that bundle observed with lower-noise data than NE body waves, which are lower in timing accuracy. This quality asymmetry leads to lesser but similar bias to the imbalance in ray count.

The effects of misconceptions and inadequate analysis of **R** turn up regularly in the literature. In particular, recent inferences of deep (deeper than the array aperture) low wave-speed anomalies beneath Yellowstone and Hawaii are likely either artifacts or indistinguishable from artifacts. Claiming these features as evidence of plumes from the core-mantle boundary, or other deep sources for "hot spots", is unsupported by the data. The only viable evidence we have seen for melt or compositional variations beneath "hot spots" is above the upper-mantle discontinuities. Furthermore, there are several interpretations that imply dipping structures in regions where rays are dominated by one bundle of *P* or *S* — such "dipping" also is an artifact or indistinguishable from artifact.

Resolution Perils and Pitfalls—Numerical Examples

All tomographic solutions are inherently smoothed. It is impossible to obtain a precise local solution everywhere because not all points of the medium are crossed by raypaths or are not crossed by rays that are widely distributed in orientation. In both cases, information about wave speed in such areas is weak, absent, or spatially distorted. The degrees of smoothing

and distortion at a given location determine useful resolution and the limits of meaningful interpretation.

Even in the limit where the resolution matrix **R** is the identity (**R** = **I**; all parameters perfectly resolved), wave speed anomalies within blocks (or surrounding and contributing to nodes) are at least smoothed within those block volumes, and the degree of that smoothing is determined by block size (or node spacing), which is determined by ray coverage. In reality it is impossible to achieve such "perfect" resolution (Backus and Gilbert, 1967, 1968, 1970) and difficult even to approach that ideal. Thus, an estimate of the best possible resolving power at any location is the volume and shape of the actual averaged region at each location and those "impulse response" volumes will vary in size and shape across the model depending on local ray density and how evenly those rays are distributed in position and orientation.

In Figs. S3(*d*) and (*g*), a simple synthetic model where the ray-orientation distribution is almost uniform and from all sides, the averaging area can be approximated by a circle or sphere — the "volume-of-averaging area". When the rays have an uneven distribution of orientations (the usual case around the edges of any teleseismic tomography model and below a depth about equal to the aperture of the observing array), then the volumes resolved will be stretched along the predominant ray direction. The simplest approximation of the "resolved" volume in this case is an ovoid (or in two-dimensions an ellipse). Resolution will be least along the major axis of the ovoid and best along the minor axis. The aspect ratio of such ovoids can be quite large.

This variation in the "resolvable ovoid" and corresponding diag(**R**) are illustrated by two-dimensional synthetic examples in Figs. S3(*a*)–(*n*). Sources and receivers are present on all sides of that model area; in teleseismic tomography the stations are at the surface and sources effectively along the bottom of the model, making the situation worse. Starting models of wave speed are shown in (*c*), (*f*), (*i*), and (*l*) for various arrangements of low- and high-wave speed anomalies, including checkerboard and anti-checkerboard models. Traveltimes were calculated for the two sets of raypaths in (*a*) and (*b*), respectively good and poor ray distributions. Both ray distributions are poor at the periphery. No random noise was added to the synthetic traveltimes. Had it been added, useful resolution would be worse.

In Figs. S3(*a*) and (*b*), green areas are the resolvable ellipses (Yanovskaya, 1997) centered on several of the true (circular) anomalies. For the good set of raypaths (*a*), these averaging areas in the central area are approximately circular and centered at the original anomalies (the azimuthal distribution of paths is most uniform here). However, near the periphery, especially in the corners, the ovals are elongated parallel to dominant ray directions (effectively radial). Thus, the tomographic reconstruction of wave speed anomalies near the centers of (*d*), (*g*), (*j*), and (*m*) are faithful, but for the poor ray set used in (*e*), (*h*), (*k*), and (*n*), yield badly distorted features even near the center. All are stretched in the direction of the most numerous rays. In particular, the ray set in (*b*) has semi-parallel rays along the "southwest–northeast" diagonal, and the resolvable ellipses are elongated along the diagonal and are large. Thus, the wave speed anomalies in (*e*), (*h*), (*k*), and (*n*) are both badly smeared along the diagonal and weakened relative to the good ray set.

A number of the original anomalies effectively vanish or have lowered estimated wave speed perturbations. Even with the good ray set, estimated perturbations between anomalies of the same sign (anti-checkerboards of (*f*) and (*l*)) bias toward the adjacent values, with individual features effectively vanishing in (*h*) and (*n*) and merging with like-sign neighbors and weakening in (*e*), (*k*), and (*m*) and even in the corners of (*j*). Wave speed estimates midway between anomalies fall to zero (rather to the layer mean wave speed) where those anomalies

are of opposite sign, as in (*d*) and the north-south direction of (*g*). The midpoints between like-sign anomalies are biased to the same sign, as in (*g*), (*j*), and especially (*m*).

The dense-anomaly results in (*j*) and (*m*) illustrate checkerboard optimism. The wave speed sign alterations in (*j*) are more-or-less recovered at the center of the model but those in (*m*) are recovered only in the north-south direction and merge into single anomalies east-west. The well-observed, effectively isolated features in (*d*) and (*g*) are recovered fairly well regardless of sign. Even so, points midway between like-sign anomalies bias toward the sign of the neighboring anomaly. These isolated features simulate impulse responses for each feature with only the low-amplitude edges of kernels merging.

The distributions of true anomalies in the Earth thus have a first-order effect on the ability of the data set to resolve those anomalies. Checkerboard tests do not illuminate such effects. On the contrary, they tend to mask them. An example of the problems for a real experiment is illustrated in Fig. S4. This shows the results of inverting a synthetic dataset similar to that collected in the Ethiopian rift by Bastow *et al.* (2008), to test the ability of the experiment to retrieve a structure comprising a pair of high-wave-speed rift flanks. A strong zone of low wave-speed is imaged between the flanks, that did not exist in the original structure. This entirely fictitious body would be readily interpretable in terms of geology.

Cryptic Perils and Pitfalls

In addition to the well (if not widely) understood problems with **R** described above, we note the following:

(1) **R** truly is singular. It is a "lens" that cannot be corrected for or "seen through" any better than is described formally by **R**, and often far less well than common tests of **R** suggest.

(2) Long sums of small numbers (small off-diagonal values in any kernel of **R**) can cause artifacts that many do not recognize or consider. In the extreme, these off-diagonal elements sum to large positive values that effectively smear all evidence of absolute wave speeds across the entire depth range of the model, rendering them invisible. The same phenomenon is present when modeling broad lateral variations in wave speed, such as a horizontal low-wave speed "lens" embedded in an otherwise featureless background (as may be present near the Moho beneath the Yellowstone caldera and certainly is present as low-speed breccia, tuff, hydrothermally altered rock, and other materials filling surface calderas). The edges of the horizontal lens will be located correctly and characterized by a doublet anomaly in the horizontal direction, where the peak-to-peak speed range of the doublet is approximately equal to the (damped) perturbation within the real lens. However, absolute wave speeds vanish, such that positive and negative peaks of the doublet have similar magnitudes. This is so, again, because damped inversions have negative side lobes in the kernels of **R** in horizontally adjacent blocks. Furthermore, the center of the lens is eviscerated and spread into a tall, very weak diamond-in-X pattern of apparent perturbation, as discussed by Evans and Achauer (1993, their Fig. 13.7, upper-right panel).

(3) Because it is impossible to recover absolute wave speeds with teleseismic tomography, all "high" and "low" wave speeds calculated are relative to some unknown, unknowable wave speed. For example, it is incorrect to characterize areas surrounding a thermo-compositional low-wave speed anomaly as "fast". These peripheral areas may, in fact, be near or even below typical Earth model wave speeds despite their apparent elevation.

(4) More insidious still are the cryptic effects of parameterizing the wave speed model and the various other approximations and assumptions required to make the inversion mathematically realizable and unique. These effects have been known since the beginning of tomography, illustrated for example, by the inversion for a single-block anomaly. The structure is rendered well when it matches the location of a model block, but vanishes when it straddles adjacent model blocks (Ellsworth, 1977). There simply is no escaping the fact that Earth does not comprise layers of blocks or nodes, nor is it damped, nor can the resulting model be uniquely selected from the infinite set of possible solutions. While some authors appear to believe that **R** is truth, full and precise, it is not, and the differences can be substantial, invisible, and insidious.

(5) Cross-correlation phase-picking methods have come into nearly universal use, commonly using the method of VanDecar and Crosson (1990). However, these methods inherently disregard physics and known Earth structure. In particular, the first 10 to 20 s of *P*, and a similar range before and after the first arrival of *S*, are contaminated by non-random signals from crustal and upper mantle converted phases and multiples. If they were not thus contaminated, receiver function analysis would be impossible. Furthermore, diffracted phases passing around strong anomalies are present before or after the direct ray and may mislead the analyst or subsume additional systematic errors into cross-correlation picks (inversions generally assume that the direct phase is measured and that phase generally is the largest signal). While such errors are unavoidable, they are smallest during the largest early half cycle of the direct phase (generally the first or second large half cycle). Diffracted and converted phases are commonly an order of magnitude smaller during this peak-signal interval, yet when summed over a wider correlation window become significant (again, the sum of many small numbers). Worse yet, because these errors are non-random (deterministic) they must map into the inversion result somewhere and will not average out like random noise. Unfortunately, where they will map is essentially unknowable, although they may affect poorly-resolved regions more because in well-sampled areas redundant sampling of anomalies may suppress expression of these errors. Thus, no travel-time picking method should be used that does not, at least in its final step, correlate only the first or second strong half cycle of the direct wave (as is easily done with the method of Lou *et al.*, 2013).

Lou *et al.* (2013) describe a revised correlation picking method and associated software that makes central the use of a stacking the target-phase records, thus improving the analyst's ability to judge the quality of the data and time of strongest signal. In addition to relative traveltimes, their method (1) returns robust absolute first-arrival estimates, which may be used to help constrain the mean structure of the region, and (2) makes it feasible for the analyst to restrict the correlation window to the first large peak or trough of the direct phase in question (*e.g.*, their Figs. 2(*a*) and 3(*b*)). We consider this restriction essential for accurate data analysis and all that follows. Because correlation squares the signal, the high amplitudes of the direct-phase first arrival and relative absence of the lower amplitudes of contaminating phases within this narrow window makes these picks as accurate as can currently be done. Though largely automatic, the method of Lou *et al.* (2013) requires the analyst to look at every event and apply their judgment rather than blindly trusting an algorithm to determine picks and their qualities. In our experience, far too many things happen in fielded instruments and the Earth for an algorithm to be reliably used as a black box. The software provided by Lou *et al.* (2013) makes implementation of their method straightforward, and we recommend its use.

(6) One corollary to (5) is that a human being must verify all arrival-time measurements by inspection and skilled judgment, removing or correcting those that fall prey to “cycle

skipping", seismograph problems such as polarity reversals, crossing phases, local noise bursts, and so forth. An independent test of data quality can be made by plotting all observed relative residuals at each station onto a lower-hemisphere plot according to the ray approach direction beneath the station. In such a plot, strong, isolated traveltime anomalies are candidates requiring additional manual review to confirm or remove them from the data set. Obviously, the seismograms take precedence to such "polar plots", or one risks "cooking" the data.

(7) Damping may be too low (resulting in the same checkerboard or patchy pattern discussed above) or too high (resulting in underestimated anomaly amplitudes). Furthermore, there is now good evidence that damped inversions always significantly depress the amplitude of recovered anomalies even when damping is chosen optimally (Christoffersson and Husebye, 2011). The significance of this effect can be assessed using model-driven resolution tests.

(8) Scaling of the wave-speed scale, particularly when the most anomalous regions saturate, is suspect. Authors should at the least present a range of scaling options in e-supplements to any tomography paper, and should in the body of the paper show at least some unsaturated plots to demonstrate the relative strengths of various claimed features.

Implications

The effects of these perils and pitfalls upon discussions, in particular of "hot-spot" origins, has caused widespread confusion and resulted in interpretations not supported by the data. What is clear from teleseismic tomography and other methods is that most "hot spots" are underlain by relative low wave-speed anomalies that may appear to largely or fully terminate at ~200–600 km depth, commonly near an upper-mantle discontinuity or the bottom of the low-velocity zone (LVZ).

It is not clear, however, whether various proposed features below this ~200–600 km upper-mantle depth are meaningful or are simply artifacts of the method, the particular ray sets used, and the strength of shallower anomalies. Indeed, well-established highly anomalous crustal and upper-mantle structures are certain to smear outward into these poorly resolved parts of models and to be partly or entirely responsible for deeper apparent anomalies. Many similar features were noted and disregarded by the most experienced tomographers decades ago, simply because these deeper features were recognized as uninterpretable, either artifacts or indistinguishable from artifacts.

We consider all proposed deeper anomalies at "hot spots" to date to be unconvincing, and see no good evidence of wave-speed anomalies below a few hundred kilometers depth. Our shallow-only inference is most easily explained either as a near-solidus bottom-up source in the upper mantle, or as a surface-down penetration to no deeper than about the upper-mantle discontinuities. Given similarly strong reservations about global-tomography models, in part because of their inconsistent results and very weak perturbations, leaves only one known test for deep, weak, low-wave speed conduits in the lower mantle beneath "hot spots" — by looking for guided waves traveling within the conduit (Julian and Evans, 2010). Several suitable datasets exist for this method, for Yellowstone, Iceland, and Hawaii.

Figures for Supplementary Material

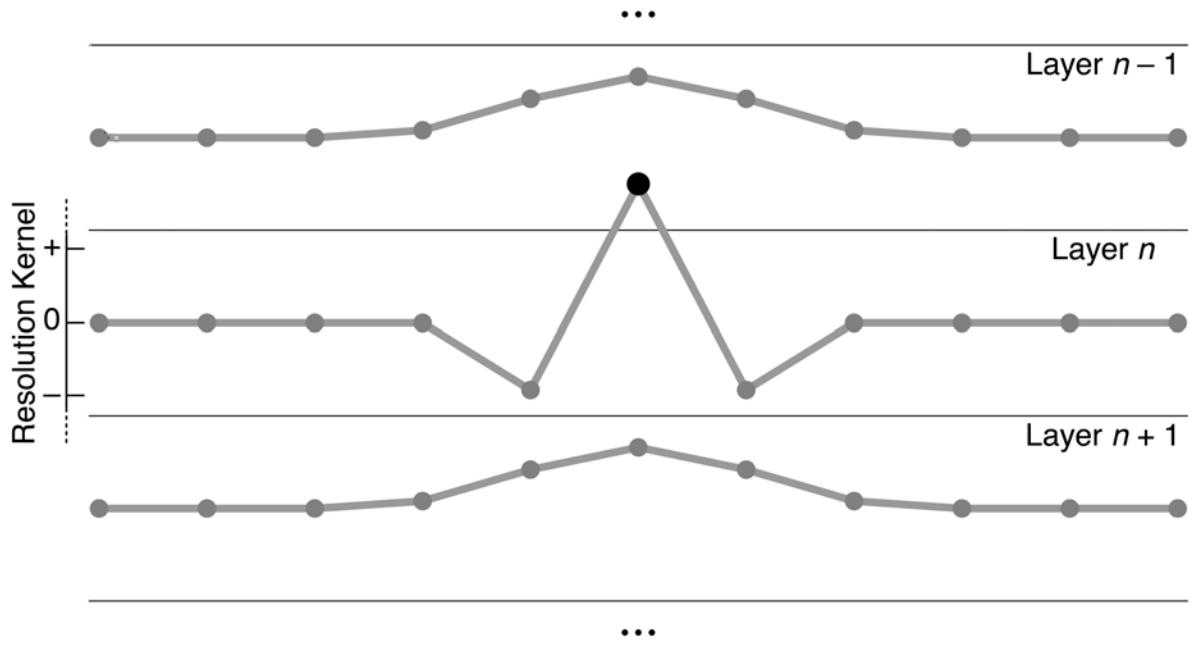

Fig. S1 Schematic of a portion of one typical kernel (column) of **R** for a block or node at the center of layer *n* for a typical teleseismic tomography study (gray lines and dots) showing the relative magnitudes of diagonal (black dot) and off-diagonal (gray dots) elements of **R**. These dots are (horizontally) at block centers and vertically proportional to the amplitude of **R** in that element (relative scale at left). Such representations can be thought of as the "impulse response" to the presence of a single-block anomaly at the center of layer *n*, for the filter comprised of a particular dataset and (to a lesser degree) the particular inversion method used.

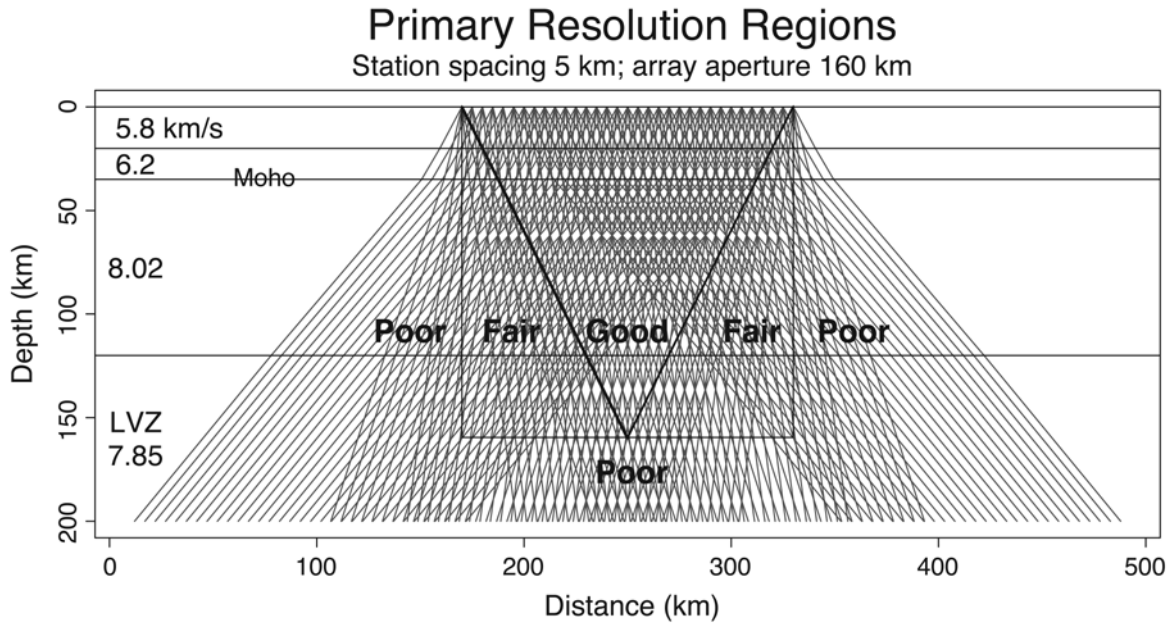

Fig. S2 Typical but idealized two-dimensional ray set and primary resolution regions for teleseismic tomography. Arbitrary scaling; no vertical exaggeration (after Evans and Achauer, 1993).

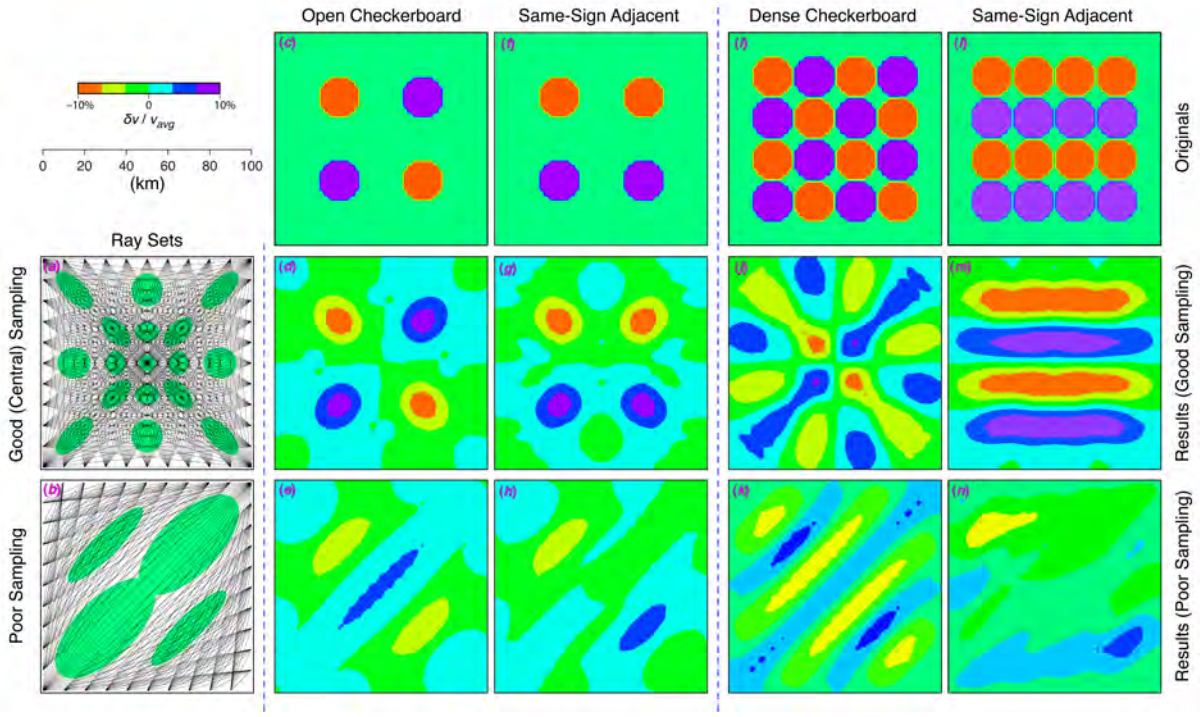

Fig. S3 Two-dimensional synthetic-data tests of resolution for (*a*) a best-case and (*b*) more typical ray sets, with similar numbers of rays in each. The models used to create synthetic traveltimes are in (*c*), (*f*), (*i*), and (*l*). Inversion results using the best-case ray set are in (*d*), (*g*), (*j*), and (*m*) and those for the more typical ray set are in (*e*), (*h*), (*k*), and (*n*). Green ovals in (*a*) and (*b*) are estimated minimum-resolvable objects in various locations (from Yanovskaya, 1997).

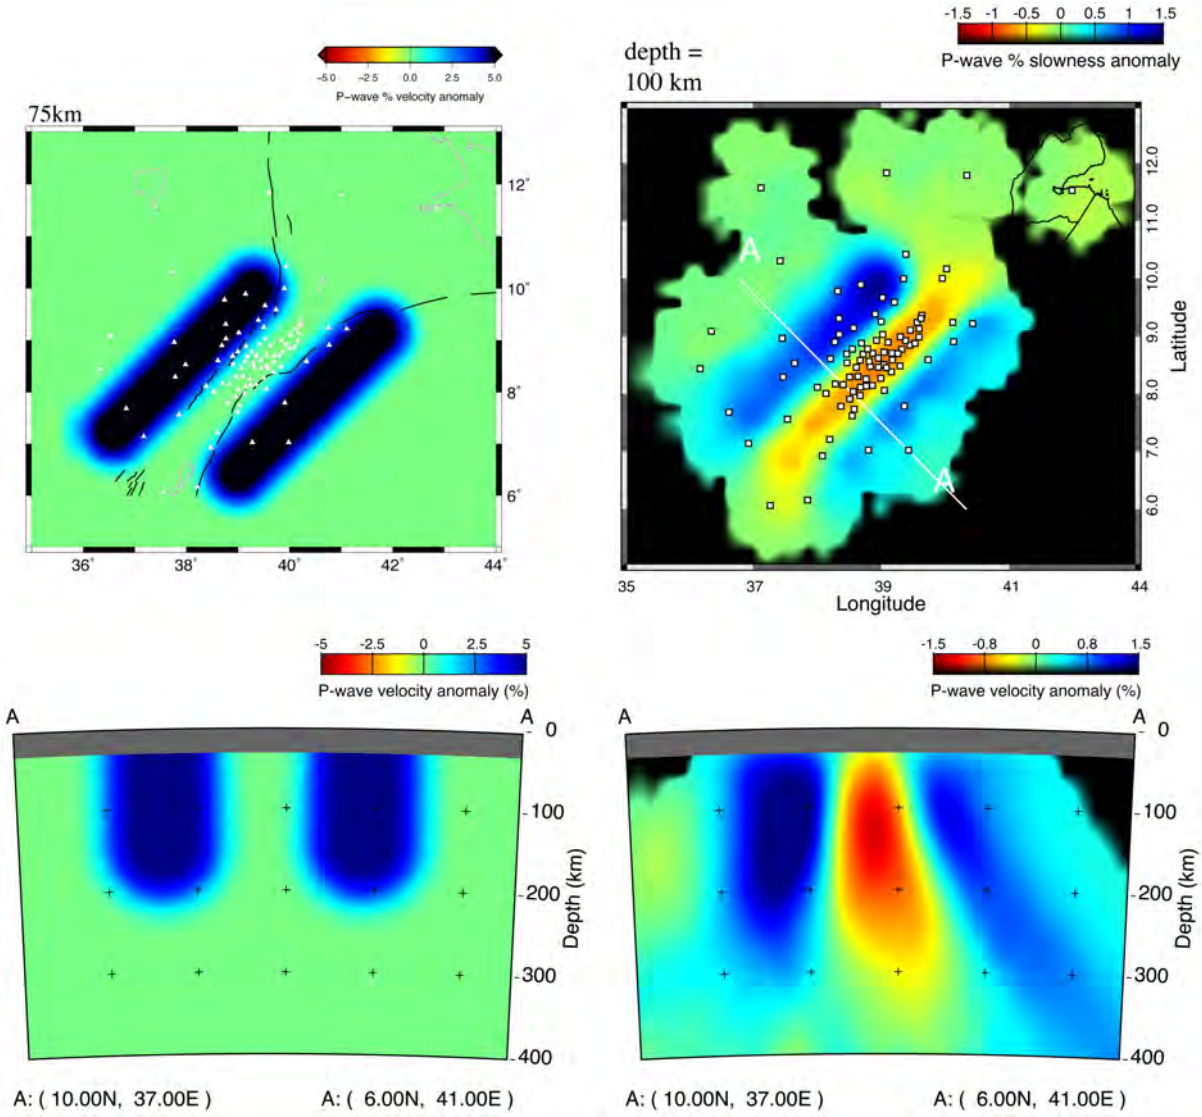

Fig. S4 Synthetic wave-speed model for the Ethiopian rift (left) consisting of high-wave-speed (Δv_P =5%) rift flanks. A relative-arrival-time dataset is computed for the same station-earthquake pairs used in the study of Bastow *et al.* (2008). The resulting tomographic model is characterized not only by fast wave-speed flanks but also by a low-wave-speed zone beneath the rift valley. In reality, in Ethiopia, *P*- and *S*-wave arrival times are ubiquitously late compared to the global mean, with the implication that the "high" wave-speed (blue) regions are low wave-speed compared to the global mean (Bastow, 2012; Bastow *et al.*, 2008). A model result with high- and low-wave-speed structure is the inevitable consequence of relative arrival-time inversions.

References for Supplementary Material

Alinaghi, A., Koulakov, I. and Thybo, H., 2007. Seismic tomographic imaging of P- and S-waves velocity perturbations in the upper mantle beneath Iran. *Geophys. J. Int.*, **169**, 1089–1102.

Backus, G.E. and Gilbert, F., 1967. Numerical applications of a formalism for geophysical inverse problems. *Geophys. J. Roy. Astron. Soc.*, **13**, 247– 276.

Backus, G.E. and Gilbert, F., 1968. The resolving power of gross earth data. *Geophys. J. Roy. Astron. Soc.*, **16**, 169–205.

Backus, G. and Gilbert, F., 1970. Uniqueness in the inversion of gross Earth data. *Phil. Trans. Royal Soc. London*, **A266**, 123–192.

Bastow, I.D., 2012. Relative arrival-time upper-mantle tomography and the elusive background mean. *Geophys. J. Int.*, **190**, 1271–1278.

Bastow, I.D., Nyblade, A.A., Stuart, G.W., Rooney, T. and Benoit, M.H., 2008. Upper Mantle Seismic Structure Beneath the Ethiopian Hotspot: Rifting at the Edge of the African Low Velocity Anomaly. *Geochem. Geophys. Geosyst.*, **9**, Q12022, doi:10.1029/2008GC002107.

Christoffersson, A. and Husebye, E.S., 2011. Seismic tomographic mapping of the Earth's interior — Back to basics revisiting the ACH inversion. *Earth Sci. Rev.*, **106**, 293–306.

Ellsworth, W.L., 1977. *Three-dimensional Structure of the Crust and Mantle Beneath the Island of Hawaii*. Ph.D., Massachusetts Institute of Technology, 327 p.

Evans, J.R. and Achauer, U., 1993. Teleseismic velocity tomography using the ACH method: theory and application to continental-scale studies. In: *Seismic Tomography: Theory and Applications* (H.M. Iyer and K. Hirahara, eds). Chapman and Hall, London.

Julian, B.R. and Evans, J.R., 2010. On possible plume-guided seismic waves. *Bull. Seismol. Soc. Am.*, **100**, 497–508.

Lou, X., van der Lee, S. and Lloyd, S., 2013. AIMBAT: A Python/Matplotlib Tool for Measuring Teleseismic Arrival Times. *Seismol. Res. Lett.*, **84**, 85–93.

VanDecar, J.C. and Crosson, R.S., 1990. Determination of teleseismic relative phase arrival times using multi-channel cross-correlation and least squares. *Bull. Seismol. Soc. Am.*, **80**, 150–169.

Yanovskaya, T.B., 1997. Resolution estimation in the problems of seismic ray tomography. *Izvestiya Phys. Solid Earth*, **33**, 762–765.