Physics of the Earth and Planetary Interiors 210-211 (2012) 21–35

Contents lists available at SciVerse ScienceDirect

Physics of the Earth and Planetary Interiors

journal homepage: www.elsevier.com/locate/pepi

On the likelihood of post-perovskite near the core–mantle boundary: A statistical interpretation of seismic observations

Laura Cobden [a,*], Ilaria Mosca [a,1], Jeannot Trampert [a], Jeroen Ritsema [b]

[a] Department of Earth Sciences, Utrecht University, Utrecht, The Netherlands
[b] Department of Geological Sciences, University of Michigan, Ann Arbor, MI, USA

ARTICLE INFO

Article history:
Received 6 January 2012
Received in revised form 23 July 2012
Accepted 20 August 2012
Available online 21 September 2012
Edited by George Helffrich

Keywords:
Post-perovskite
Core–mantle boundary
Seismology
Mineral physics

ABSTRACT

Recent experimental studies indicate that perovskite, the dominant lower mantle mineral, undergoes a phase change to post-perovskite at high pressures. However, it has been unclear whether this transition occurs within the Earth's mantle, due to uncertainties in both the thermochemical state of the lowermost mantle and the pressure–temperature conditions of the phase boundary. In this study we compare the relative fit to global seismic data of mantle models which do and do not contain post-perovskite, following a statistical approach. Our data comprise more than 10,000 P_{diff} and S_{diff} travel-times, global in coverage, from which we extract the global distributions of $d\ln V_S$ and $d\ln V_P$ near the core–mantle boundary (CMB). These distributions are sensitive to the underlying lateral variations in mineralogy and temperature even after seismic uncertainties are taken into account, and are ideally suited for investigating the likelihood of the presence of post-perovskite. A post-perovskite-bearing CMB region provides a significantly closer fit to the seismic data than a post-perovskite-free CMB region on both a global and regional scale. These results complement previous local seismic reflection studies, which have shown a consistency between seismic observations and the physical properties of post-perovskite inside the deep Earth.

1. Introduction

The lowermost 150–300 km of the mantle (i.e. the D″ region) is one of the most enigmatic and seismically complex regions of the Earth. Postulated as the source region for mantle plumes (Williams et al., 1998), a graveyard for subducted slabs (Van der Voo et al., 1999), and potentially a primitive, melt-bearing layer (Fiquet et al., 2010), it represents a zone in which strong thermal, chemical and structural heterogeneity may be expected. Nonetheless, the 2004 discovery of a phase transition in $(Mg,Fe)SiO_3$ perovskite (Pv) to post-perovskite (pPv) at pressures approaching the Earth's lowermost mantle (Murakami et al., 2004; Oganov and Ono, 2004; Shim et al., 2004) opened up the possibility to explain much of the region's observed seismic behaviour in terms of this transition.

Theoretical calculations indicate that pPv is seismically distinct from Pv: it has a ~1.5–2% higher *S*-wave velocity (Stackhouse and Brodholt, 2007; Wentzcovitch et al., 2006) but a reduced bulk sound velocity (Nishio-Hamane and Yagi, 2009; Hustoft et al., 2008), which may be manifested as little or no change (±0.5%) in *P*-wave velocity. Thus, the Pv to pPv phase transition may be responsible for the so-called D″ discontinuity observed about 150–300 km above the core–mantle boundary (CMB) (Lay et al., 2005), in which a 2–3% increase in *S*-velocity is accompanied by a generally smaller or absent increase in *P*-velocity (Wysession et al., 1998). Lateral variations in post-perovskite content may further explain the intermittent anti-correlation in bulk sound and *S*-wave velocities at the CMB (Wookey et al., 2005). However, the depth and width of the Pv to pPv phase transition is strongly temperature- and composition-dependent (Grocholski et al., 2012; Catalli et al., 2009; Wookey et al., 2005). Given the likely thermochemical heterogeneity in D″ and the large experimental uncertainties on the position of the phase boundary (Hirose, 2007), it is still uncertain if, or where, pPv exists within the mantle. A number of studies (e.g., Chaloner et al., 2009; Hutko et al., 2008; Kito et al., 2007; van der Hilst et al., 2007; Lay et al., 2006; Wookey et al., 2005) have presented recordings of seismic waveforms in the lowermost mantle that are consistent with the presence of a Pv-to-pPv transition. This consistency is primarily based on agreement between the observed structure of the D″ discontinuity and the predicted seismic structure of the perovskite to post-perovskite phase transition. Yet, it is possible that the intermittent deep mantle reflections with variable amplitudes (Wysession et al., 1998) are due to structural complexity associated with subducted slabs,

* Corresponding author. Present address: Institute for Geophysics, University of Münster, Corrensstraße 24, 48149 Münster, Germany. Tel.: +49 2518334726; fax: +49 2518336100.
E-mail address: laura.cobden@uni-muenster.de (L. Cobden).
[1] Present address: Department of Earth and Environmental Sciences, Munich University, Munich, Germany.

0031-9201/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.pepi.2012.08.007

changes in bulk chemistry, or anisotropic effects (Lay and Garnero, 2007; Thomas et al., 2004a,b) in a pPv-free mantle.

To complement existing local studies, we will present a robust statistical study of the fit of different mantle models (Table 1) to seismic wave speed variations on a global to regional scale, considering both pPv-free and pPv-bearing mineral assemblages. Our approach resembles a hypothesis test: is post-perovskite present, or absent, at the core–mantle boundary? Rather than focussing on the structure of the Pv-to-pPv transitions and attendant seismic discontinuities, as in previous studies, we focus on the large-scale S and P wave speed structure of the lowermost mantle using both seismic and mineral physics constraints, and a fitting procedure based on the Metropolis algorithm (Mosegaard and Tarantola, 1995).

Lateral seismic wave speed variations are usually expressed as relative perturbations of P-wave velocity (V_P) and S-wave velocity (V_S) with respect to a 1D reference model, and referred to as $d\ln V_P$ and $d\ln V_S$, respectively. They are normally calculated via a tomographic inversion (e.g., Masters et al., 2000). However, the inverse problem is ill-posed, and the required regularisation makes it very difficult to assess the corresponding resolution and uncertainty of the wave speed variations. In order to make a quantitative and self-consistent interpretation of the seismic models in terms of thermochemical parameters, and hence pPv content, it is essential that the mineral physics parameters are averaged over the resolving length prescribed by the seismic models and that uncertainties are treated rigorously. Thus, in this work, we use a simplified wave propagation theory, the path-average approximation (Mosca and Trampert, 2009), which fixes the resolution, and converts seismic travel-time measurements directly into $d\ln V_S$ and $d\ln V_P$. This approximation has already been validated on seismic data (Mosca and Trampert, 2009) and has the advantage of allowing a precise quantification of all uncertainties without the need for a regularised inversion. Our approach is reminiscent of a Backus–Gilbert hypothesis test (e.g., Trampert and van Heijst, 2002), which solves an inference problem by fixing the desired resolution *a priori*.

The path-average approximation is accurate for spatial scale lengths of 2000–3000 km laterally, i.e. spherical harmonic degree 8. In this study we use only core diffracted phases, which give rise to a specific path-average sensitivity kernel that fixes the depth resolution to about 50 km near the CMB, with a corresponding known uncertainty. Following this approach allows us to map lateral P and S wave speed variations with the same spatial resolution, albeit with different – but quantifiable – uncertainties. It is essential to have such coherency between P- and S-wave structure in order to be able to distinguish thermal from chemical variations (Cobden et al., 2009; Hernlund and Houser, 2008; Deschamps and Trampert, 2003). Moreover, it is imperative to include uncertainties – often neglected or undefinable in tomographic studies – in the analysis to ensure a reliable, quantitative physical interpretation of the seismic data.

Another important difference in our approach is that instead of interpreting geographic maps of $d\ln V_S$ and $d\ln V_P$, or single values of the velocity perturbation at specific locations, we derive histograms of $d\ln V_S$ and $d\ln V_P$, and also the seismic parameter $R = d\ln V_S / d\ln V_P$ near the CMB. The width and shape of the histograms reflect the magnitude, frequency and nature of lateral variations in temperature and composition near the CMB (e.g., Hernlund and Houser, 2008; Deschamps and Trampert, 2003). By compiling histograms, we lose the 3D information on where, geographically, a particular velocity perturbation has come from. However, for our purpose this is not relevant: we wish simply to test whether the morphology of the histograms is better fit by a CMB region where pPv is present, or where it is absent. It is necessary to consider R in addition to $d\ln V_S$ and $d\ln V_P$ because globally-compiled histograms of $d\ln V_S$ and $d\ln V_P$ do not fully represent the correlations between V_P and V_S variations, which are important to be able to distinguish between different compositions. R is produced by combining two variables ($d\ln V_S$ and $d\ln V_P$) into one, and when used in isolation, can be interpreted non-uniquely (Deschamps and Trampert, 2003). R must therefore be analysed in conjunction with the separate $d\ln V_S$ and $d\ln V_P$ distributions. Our seismic approximation is particularly suited to the calculation of R, since $d\ln V_S$ and $d\ln V_P$ have the same spatial resolution.

We calculate the seismic properties ($d\ln V_P$, $d\ln V_S$ and R) of different hypothetical thermochemical structures at the CMB (Table 1) via thermodynamic computations based on the most recent mineral physics data. For each structure, we make no *a priori* assumptions on the bulk chemical composition or temperature at the CMB. Instead, we create thousands of models in which temperature and composition are allowed to vary freely within extremely broad ranges. From these thousands of models we select the subset whose properties match the seismic data within uncertainties, and perform a statistical analysis of the level of fit to determine the likelihood of a pPv-free, Pv-free and (Pv + pPv)-bearing CMB region.

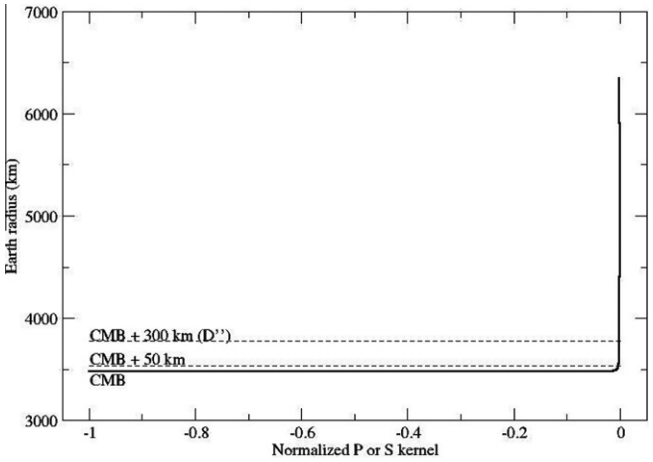

Fig. 1. Sensitivity kernel of the path-average approximation for P_{diff} and S_{diff} waves.

Table 1
Summary of different thermochemical scenarios tested in this study.

	Name	Description
Case 1	pv + ppv	Both Pv and pPv are present and lateral velocity variations arise from lateral variations in both temperature (T) and composition (X)
Case 2	no ppv	No pPv is present, but there are lateral variations in both T and X
Case 3	no pv	No Pv is present, but there are lateral variations in both T and X
Case 4	ppv:pv = 70:30	Both Pv and pPv are present in a fixed ratio of 70:30; T and X vary laterally
Case 5	Fixed T[a]	Both Pv and pPv are present at fixed T and variable X
Case 6	Fixed X[b]	Both Pv and pPv are present at fixed X and variable T

[a] Six sets of models, each with different (but fixed) temperature, were calculated, ranging between 3500 and 4000 K.
[b] Hundred sets of models each with different (but fixed) composition were calculated.

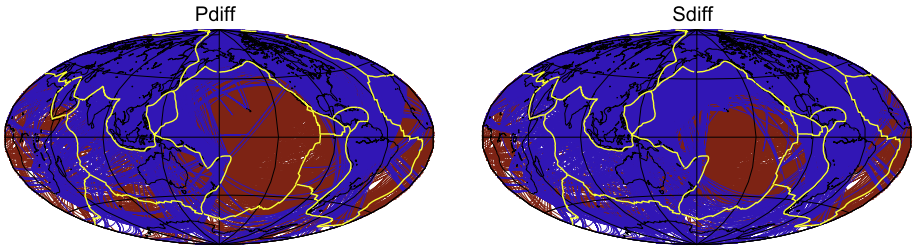

Fig. 2. Ray coverage for P_{diff} and S_{diff}; Blue indicates rays for which $\delta T < 0$; red indicates rays for which $\delta T > 0$. Yellow lines are plate boundaries. Blue ray paths are plotted on top of red paths, such that many of the regions appearing blue in this plot are also traversed by red ray paths. Ray paths have been selected to have the same global coverage and ray density for *P* and *S*. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

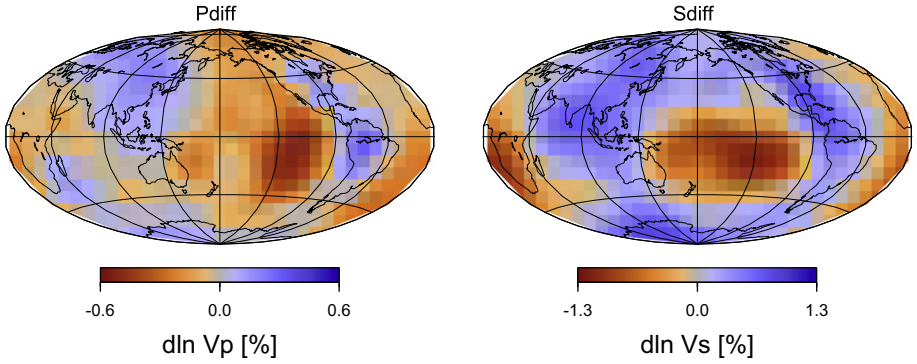

Fig. 3. Maps of inferred $d\ln V_S$ and $d\ln V_P$ near the CMB using P_{diff} and S_{diff} and the path-average approximation.

2. Methods

2.1. The path-average approximation

The path-average approximation is often used in long-period waveform inversions (e.g., Kustowski et al., 2008), and is still the method of choice in surface wave tomography (e.g., Ekström, 2011). Mosca and Trampert (2009) proposed to use it for modelling body wave travel-time residuals and benchmarked it for forward and inverse problems. According to the path-average approximation, the travel-time residual, δT, measured between the arrival-time of a wave propagating in the real Earth and that calculated in some reference model along a given ray path, is equal to the average two-way vertical travel-time perturbation, $\delta\tau$, from the Earth's surface to the bottom of the ray, along the source–receiver locus Δ:

$$\delta T(p) = \frac{1}{\Delta}\int_0^{\Delta} \delta\tau(p,\theta,\phi)d\Delta \quad (1)$$

$$\delta\tau(p,\theta,\phi) = -2\int_{r_{bot}}^{a} K(r)\delta\ln V(r,\theta,\phi)dr \quad (2)$$

where $K(r)$ is a sensitivity kernel which depends only on radius r; p is the ray parameter, $\delta\ln V$ is a lateral relative velocity perturbation with respect to the reference model, θ and ϕ are latitude and longitude, respectively, r_{bot} is turning radius of the ray and a is the radius of the Earth. The kernel $K(r)$, which resembles a delta function (Fig. 1), has a sharp peak at the ray turning depth, and near-zero values at other depths. Thus, the two-way vertical travel-time perturbation is mainly sensitive to the local velocity perturbation near the depth at which the ray bottoms. This is of course a zero-order, high-frequency approximation. In practice, the sensitivity also depends on the seismic wavelength. For the frequencies considered

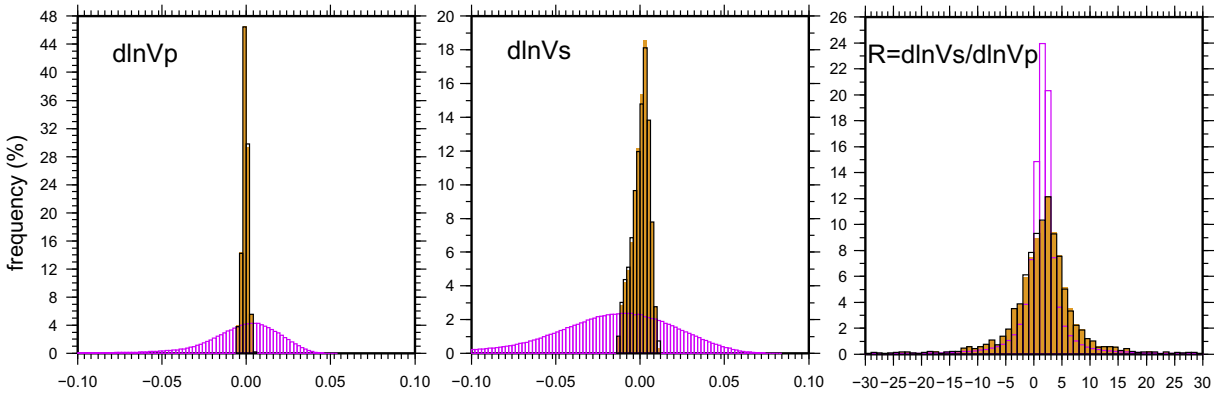

Fig. 4. Histograms showing global distributions of $d\ln V_P$, $d\ln V_S$ and R near the CMB (orange bars), compiled from the data mapped in Fig. 3. Pink lines: example of the $d\ln V_P$, $d\ln V_S$ and R distributions for 800,000 mineral physics models at the CMB (Case '1' described in Table 1). Black lines: $d\ln V_P$, $d\ln V_S$ and R distributions for the subset of these models remaining after application of the Metropolis algorithm. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2
Thermal and chemical ranges of models tested in this study.

Parameter	Lower bound	Upper bound
Temperature (K)	2300	4800
(A). Vol% pv + ppv + mwus in mineralogical assemblage[a]	85	100
(B)% (pv + ppv) within (A)	60	100
(C)% $FeSiO_3$ within pv + ppv	0	20
(D) Partition coefficient Fe-Mg between (p)pv and mwus	0.0001	2
(E)% ppv within (pv + ppv)[b]	0	100
(F)% Al_2O_3 within $CaSiO_3$ + Al_2O_3 + SiO_2 component	0	100
(G)% SiO_2 within $CaSiO_3$ + Al_2O_3 + SiO_2 component	0	100

pv = $(Mg,Fe)SiO_3$ perovskite, ppv = $(Mg,Fe)SiO_3$ post-perovskite, mwus = (Mg,Fe)O magnesiowustite.

[a] Other minerals in the assemblage are $CaSiO_3$ perovskite, Al_2O_3 perovskite + postperovskite, and SiO_2 seifertite; their total volume is equal to (100-A)%. % Al_2O_3 which is present as post-perovskite rather than perovskite is set to be equal to (E).

[b] For Cases (1),(5),(6) only (refer to Table 1 for description of the six tested CMB cases).

here, the delta-like sensitivity kernel is sub-wavelength, and the error due to this modelling assumption is incorporated as a formal uncertainty compared to the more correct ray or finite-frequency theory (Mosca and Trampert, 2009).

The path-average approximation simply allows us to convert measured travel-time residuals for the Earth into maps of lateral velocity variations at particular depths assuming that we know the ray-theoretical take-off angles (ray parameters) and the turning depths of the rays. Selecting data with the same ray parameter, Eq. (1) is used to convert the measurements into maps of two-way travel-times, similar to surface wave tomography. According to Eq. (2), these maps are then local depth integrals of the wave speed with the corresponding kernel $K(r)$. P_{diff} and S_{diff} waves – i.e., waves which diffract along the core–mantle boundary interface – are uniquely suitable for the procedure, as they each have a single, known ray parameter and a fixed bottoming depth at the CMB, 2891 km. If we assume the wave speed to be constant over the narrow sensitivity kernel (Fig. 1), we can divide the two-way travel-time maps by the radial integral of the kernel. This will give us a seismic model with the depth resolution fixed to the width of the kernel. Thus, the obtained path-average model has a fine vertical resolution but a large uncertainty. In standard tomography based on a more accurate wave propagation theory, the situation would be reversed. A correct characterization of uncertainty is important for our hypothesis testing. Using a more advanced description, such as full ray or finite frequency theory, would require us to set up an inverse problem involving the whole mantle, with the disadvantage of quantifying resolution and subjective uncertainties depending mostly on regularisation. Since we are interested in modelling the histograms of the seismic heterogeneities, a better route is a simple benchmarked theory with physically meaningful uncertainties.

2.2. Seismic data preparation and uncertainty

Our study uses a set of 10,011 travel-time residuals for both P_{diff} and S_{diff} (Fig. 2). These residuals are measured relative to PREM (Dziewonski and Anderson, 1981) from the maximum cross-correlation between the observed and PREM-predicted waveforms (Ritsema and van Heijst, 2002). Only measurements for which there is a high correlation between the observed and predicted *P*- and *S*-waveforms are used. The data are adjusted for crustal, ellipticity and source relocation corrections as described in Ritsema and van Heijst (2002). Measurements are made on the vertical and transverse components for *P* and *S*, respectively, at a frequency of 20–25 s, corresponding to a wavelength at the CMB of approximately 275–340 km for P_{diff} and 145–180 km for S_{diff}.

From the 10,011 ray paths, maps of the two-way vertical travel-time perturbation $\delta\tau$ near the CMB are calculated using the path-average approximation (Eq. (1)). These maps, for P_{diff} and for S_{diff}, are expanded into spherical harmonics up to degree and order L:

$$\delta\tau(p,\theta,\phi) = \sum_{l=0}^{L}\sum_{m=-l}^{+l} c_{l,m}\psi_{l,m}(\theta,\phi) \qquad (3)$$

where θ and ϕ are latitude and longitude, respectively; p is the ray parameter; l and m are the angular and azimuthal order; L is the maximum spherical harmonic degree; $\psi_{l,m}$ are the spherical

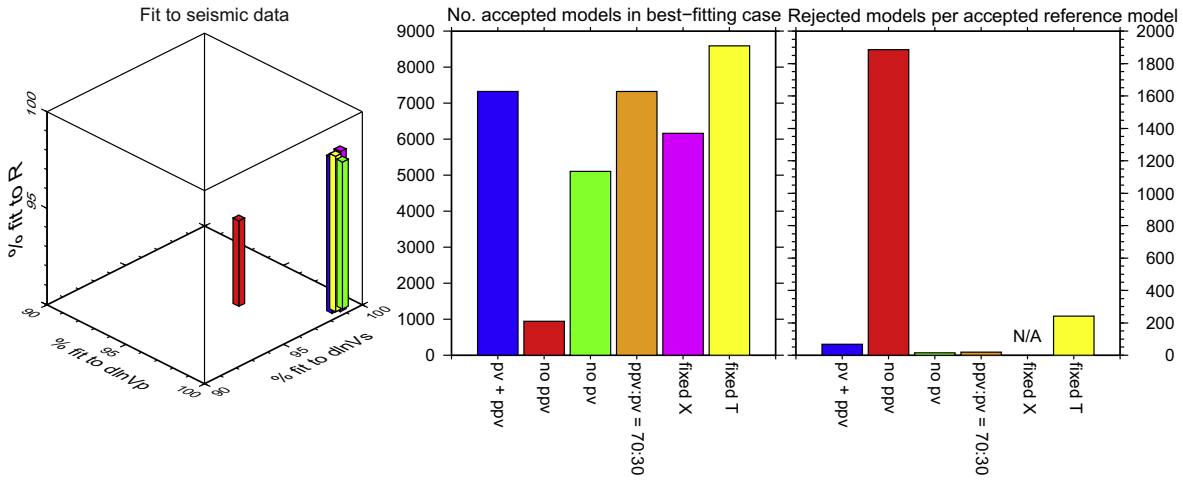

Fig. 5. Summary of fit to seismic data of six different thermochemical scenarios near the CMB. Left: % fit of each scenario to global $d\ln V_P$, $d\ln V_S$ and R distributions (orange bars, Fig. 4) following application of the Metropolis algorithm. Colour scheme the same as in other panels; note that the "ppv:pv = 70:30" case is not visible because it is almost identical to the "fixed *X*" case which plots on top of it. Centre: number of remaining models from 800,000 initial models for each scenario, after application of Metropolis algorithm. Right: number of iterations required to find a starting (reference) model whose density and velocities lie within 1% of PREM. Compositions for the fixed *X* cases were defined using the first 100 reference models of Case 1 (Table 1), and the data shown in this figure are from one of those 100 models selected at random. Six sets of models were calculated at fixed *T*; in this figure we plot the data from a simulation where the temperature was fixed to 3700 K. Although only single data points are plotted for fixed *X* and fixed *T*, the averages of the 100 and 6 sets of models, respectively are very similar to the randomly-selected models shown here. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

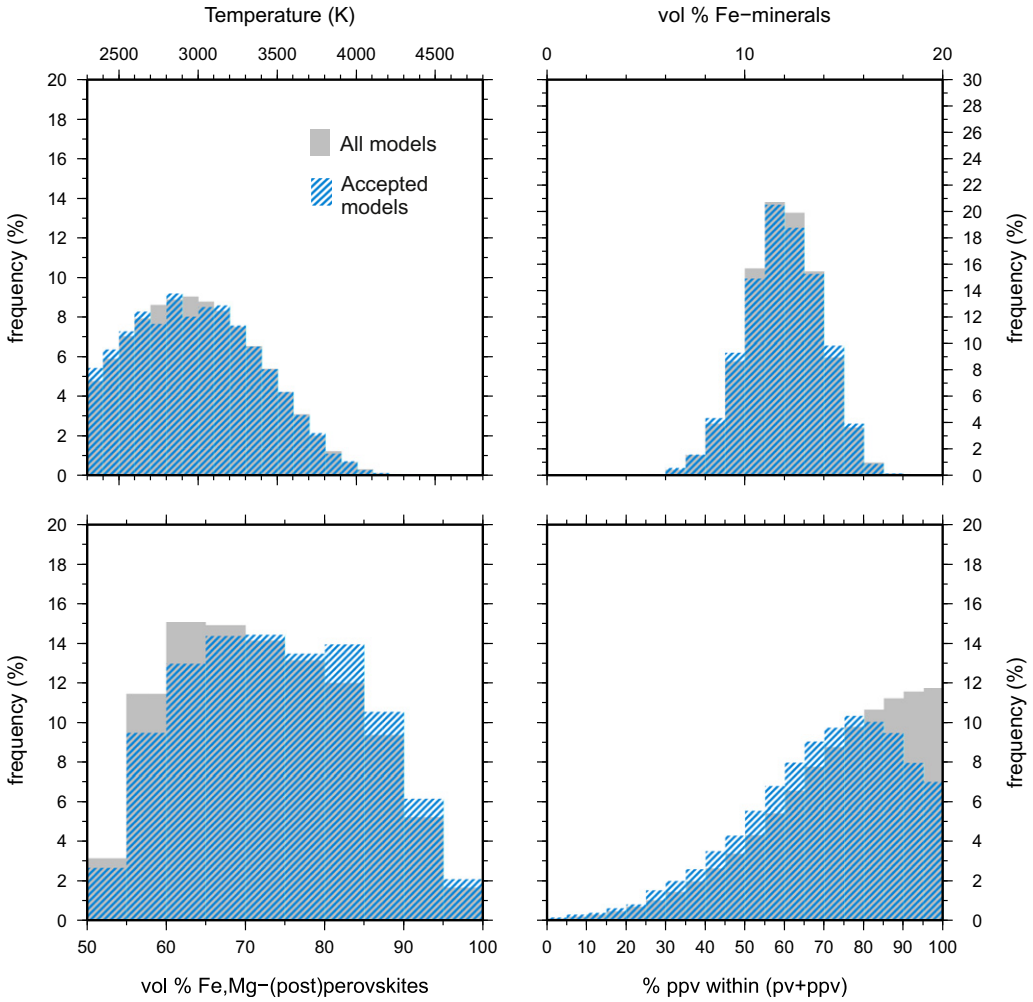

Fig. 6. Thermochemical ranges of starting (reference) models, i.e. physical 'reference' models relative to which $d\ln V_S$ and $d\ln V_P$ are calculated, from Case 1 (refer to Table 1). Grey – all 800,000 reference models; blue – subset of accepted reference models following application of the Metropolis algorithm. Although initial reference models are selected at random from a uniform distribution of thermochemical properties (Table 2), reference models are rejected if their velocity and density do not lie within 1% of PREM (see Section 2.4). This leads to asymmetric ranges in the non-rejected models' thermochemical properties, as shown above. Note (1) the close correspondence between the blue and grey histograms, indicating a consistency between the seismic and mineral physics reference models, and (2) the high percentages of pPv in most reference models. Vol% Fe-minerals is the total sum of Fe-perovskite, Fe-postperovskite and FeO wustite in a given model. Vol% Fe,Mg-(post)perovksites is the total sum of (Fe,Mg) perosvkite plus (Fe,Mg)post-perovskite. Histograms are normalised to 100% for each dataset, i.e. the frequency percentages for the blue histogram are relative to the subset of accepted models only, and not the initial 800,000 models. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

harmonic functions, and $c_{l,m}$ are the harmonic coefficients. The path average approximation works best for long wavelength structure, so we consider only the first 8 degrees. This is no real restriction, however, since diffracted waves naturally average small-scale lateral variations (e.g., Lay, 2007). To ensure that the structure up to degree 8 is not distorted by spectral leakage from higher order terms (Trampert and Snieder, 1996), we initially made a spherical harmonic expansion up to degree 20 with a lateral smoothing constraint. Due to the good ray coverage, we find that the first 8 degrees are independent of regularisation and spectral leakage.

The spherical harmonic expansion (Eq. (3)) allows us to infer lateral wave speed variations without an explicit inversion. The kernel K in Eq. (2) is almost zero until the point of maximum curvature (about 50 km above the CMB, Fig. 1) and therefore only seismic velocity variations in the lowermost 50 km of the mantle contribute to the travel-time anomalies. The vertical average of $\delta\ln V$ within the lowermost 50 km of the mantle may be calculated simply by dividing $\delta\tau$ by the radial integral of the kernel, i.e. −816.62 for V_P and for −1543.94 V_S. This straightforward procedure provides maps of relative variations of wave speeds near the CMB with a lateral resolution of ~2000–3000 km (Fig. 3). Because our dataset provide an excellent global coverage (Fig. 2) there is no geographic bias towards a particular region.

The uncertainty in the $\delta\tau$ values has contributions from (1) the imperfect fit of the spherical harmonic expansion to Eq. (1), u_{sh}; and (2) the path-average approximation, u_{pava}. The first term, u_{sh}, is equal to the variance of the misfit between the measured travel-time residuals and those given by our chosen degree 8 map. To address the second term, we followed Mosca and Trampert (2009), and estimated the uncertainty of the path average approximation relative to ray theory by comparing deterministic as well as random seismic models of the Earth's mantle expanded up to spherical harmonic degree 8.

$$u_{\mathrm{pava}} = \sum_{i=1}^{N} \frac{[dT_{\mathrm{ray}_i} - dT_{\mathrm{pava}_i}]^2}{N} \Big/ \sum_{i=1}^{N} \frac{dT^2_{\mathrm{ray}_i}}{N} \tag{4}$$

where dT_{ray} is the travel time residual predicted by ray theory, dT_{pava} is the residual predicted by the path average approximation, and N is the total number of observations through our synthetic seismic

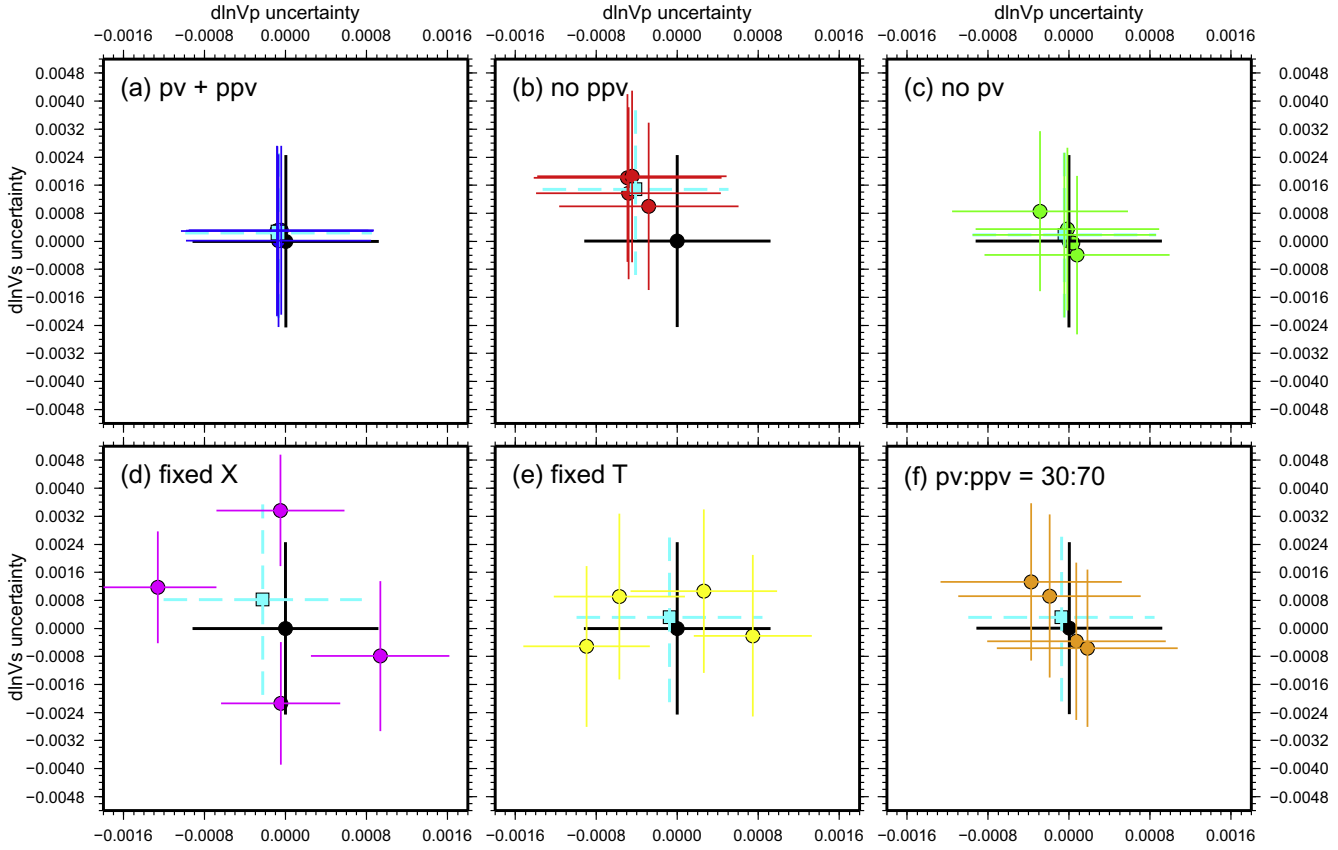

Fig. 7. Uncertainty distributions of accepted models for different CMB scenarios. Black circle: mean $d\ln V_P$ uncertainty and $d\ln V_S$ uncertainty of all 800,000 models (i.e. both zero); black lines: standard deviation of $d\ln V_P$ and $d\ln V_S$ uncertainties for all models; blue squares: mean $d\ln V_P$ uncertainty and $d\ln V_S$ uncertainty for the accepted models in each scenario; should ideally plot at (0,0) if uncertainties are unbiased. After cluster analysis is applied to divide the accepted models into four groups (clusters), the coloured circles and lines indicate the means and standard deviations of uncertainties for each cluster. Again ideally these should all be centred at (0,0). Note in cases (b) and (d) the significant offset of the blue squares from zero, and differing degrees of scatter of the four clusters about the blue squares in each case. Plot (a) exhibits the least amount of scatter. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 3
Alternative representation of the data shown in Fig. 7. We normalised the histograms of $d\ln V_P$ and $d\ln V_S$ uncertainties for the accepted models in each of the six Cases, and calculated their overlap with the normalised uncertainty distributions of the 800,000 initial models. The product of the overlap in $d\ln V_P$ with the overlap in $d\ln V_S$ (third column) is a measure of the degree of bias in the accepted models' uncertainties.

	%Overlap $d\ln V_P$ uncertainties	%Overlap $d\ln V_S$ uncertainties	P (no bias)
Case 1 (pv + ppv)	96.71	95.44	0.92
Case 2 (no ppv)	82.53	75.83	0.63
Case 3 (no pv)	97.08	95.00	0.92
Case 4 (ppv:pv = 70:30)	96.24	94.67	0.91
Case 5 (fixed T)	96.26	94.38	0.91
Case 6 (fixed X)	89.59	84.00	0.75

models. We find that u_{pava} is dominant and estimate that the total uncertainty, i.e. the square root of $(u_{\mathrm{sh}} + u_{\mathrm{pava}})$, is 0.75 s for P_{diff} and 3.8 s for S_{diff}. This is quite large compared to uncertainties in the cross-correlation measurements and is entirely due to our approximation for wave propagation. Although earthquake relocation corrections have been applied (Ritsema and van Heijst, 2002), source uncertainties remain. However they average out in the construction of the maps because many different earthquakes have been used without a particular geographical bias. The corresponding uncertainties in the velocity perturbations are thus 0.75/816.62 for $d\ln V_P$ and 3.8/1543.94 for $d\ln V_S$.

Using the maps of $d\ln V_S$ and $d\ln V_P$ near the CMB, we construct histograms of $d\ln V_S$ and $d\ln V_P$ together with the seismic parameter R (Fig. 4, orange bars) from 4926 equal-area blocks.

The residuals used in this study have been specifically selected to have identical geographic coverage and density for P_{diff} and S_{diff}, i.e. the same source, receiver and epicentral distance for each ray path. Horizontal resolution is determined by the maximum degree L of our spherical harmonic expansion (Eq. (3)) and vertical resolution is fixed by the path average kernel (Fig. 1; Section 2.1). This ensures that the resolution of our P and S velocity anomalies are identical, which is important when using the seismic parameter R. Computation and analysis of R is only appropriate when $d\ln V_P$ and $d\ln V_S$ are obtained independently, i.e. from separate seismic datasets, with the same geographic ray coverage (Deschamps and Trampert, 2003), as is implicit in our method. Further, the seismic reference model relative to which they are each measured should represent the same laterally-weighted average of the 3D velocity

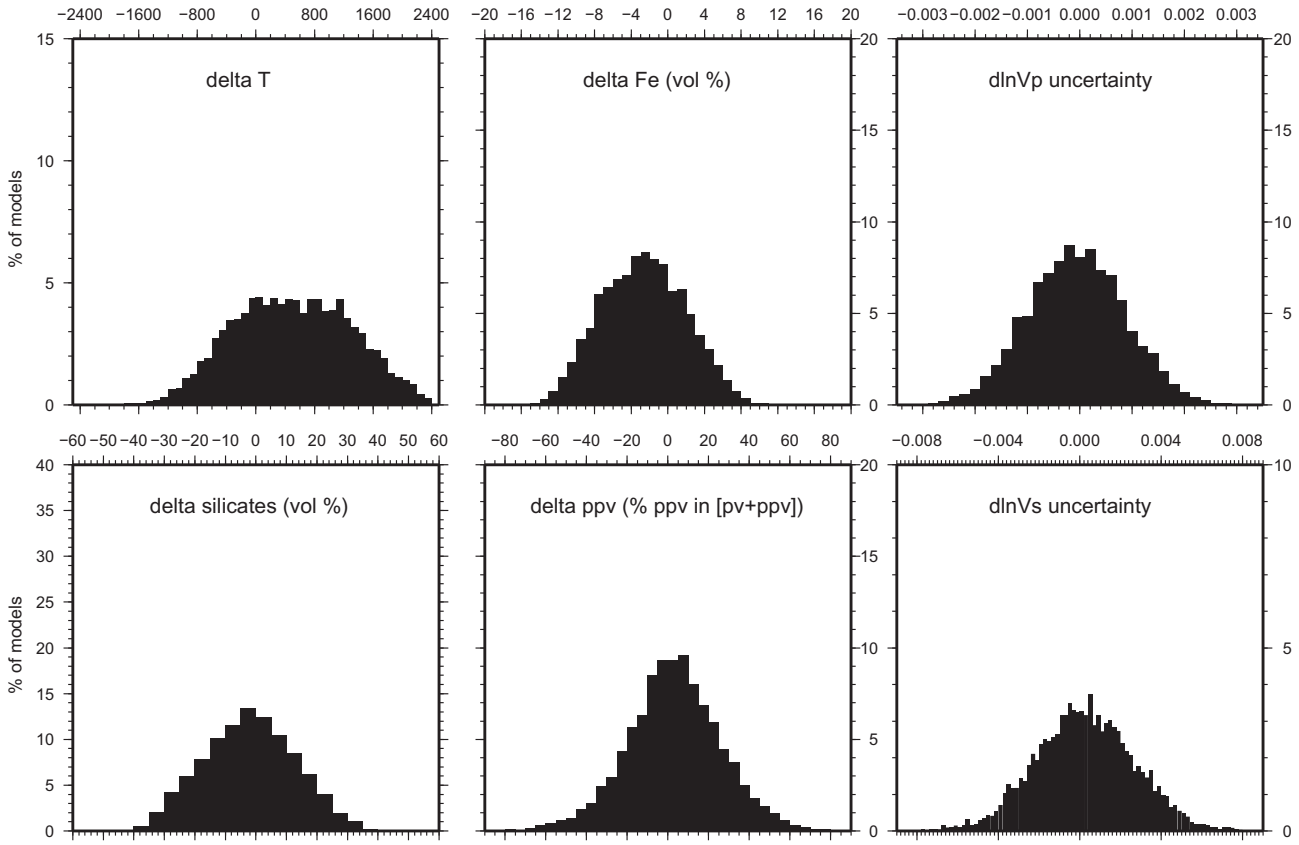

Fig. 8. Ranges of thermal and chemical variations, plus seismic uncertainties, in accepted models of Case 1 (refer to Table 1). The thermochemical ranges are not always symmetrical about zero. This is due to the imposed constraint that reference models must fit PREM to within 1%, which gives a suite of reference models whose average is not at the centre of the ranges defined in Table 2 (Fig. 6). The broad ranges here should not be interpreted as representative of the real CMB: they arise due to trade-offs between different parameters which can fit the seismic data non-uniquely.

structure. Our data satisfy this second criterion, given the close correlation between lateral changes in $d\ln V_P$ and in $d\ln V_S$ (Fig. 3).

2.3. Mineral physics calculations

The bulk modulus K, shear modulus G, and density ρ for a given mineral at the CMB pressure of 135.75 GPa, and variable temperature, are calculated using a third-order finite strain Birch–Murnagan equation of state, coupled with a Mie–Gruneisen thermal pressure correction for the temperature, as described in Stixrude and Lithgow-Bertelloni (2005). Mineral elastic parameters are taken from Xu et al. (2008), with the exception of the perovskite and post-perovskite parameters which have been modified to take into account more recent experimental results (Stixrude and Lithgow-Bertelloni, 2011). These elastic parameters have been generated using a large compilation of experimental and theoretical mineral physics data within a rigorous thermodynamic framework and they are fully consistent with the equation of state. We include 10 end-member mineralogical phases in our calculations: $MgSiO_3$ perovskite and post-perovskite, $FeSiO_3$ perovskite and post-perovskite, FeO wustite, MgO periclase, Al_2O_3 perovskite and post-perovskite, $CaSiO_3$ perovskite, and SiO_2 seifertite. Seifertite is a high pressure polymorph of SiO_2 with α-PbO_2 structure, which is expected to occur near the CMB (e.g., Murakami et al., 2003).

For a chosen mineralogical assemblage, the overall bulk and shear moduli are assumed to be the Voigt–Reuss–Hill average of the constituent minerals. The density is the average of the densities of all the minerals present, weighted according to their volumetric proportions. Seismic velocities V_P and V_S are calculated using these averaged values for K, G and ρ. Velocities are not corrected for anelasticity since the effects of anelasticity have been shown to be negligibly small at the base of the mantle (Brodholt et al., 2007), where the effect of high pressure dominates over high temperature.

To address the effect of mineral physics uncertainties, we allowed each mineral elastic parameter in our dataset to vary within the uncertainty bounds published in Stixrude and Lithgow-Bertelloni (2011) and Xu et al. (2008). We found that the effect on the seismic properties of a given thermochemical structure was extremely small, especially in comparison to the seismic data uncertainties, and did not modify our results relative to simulations in which no mineral physics uncertainties were included. This small effect arises partially because we are working with velocity perturbations (Section 2.4) rather than absolute velocities. We evaluated the mineral elastic parameters only at the CMB, and not over the seismological sensitivity kernel $K(r)$, because the effect of elastic parameter uncertainties is larger than the amount by which the parameters would be modified by integrating over the kernel.

2.4. Combining seismic observations with mineral physics calculations

In a mineralogical sense, $d\ln V_P$ and $d\ln V_S$ may be described as a sum of partial derivatives:

$$d\ln V_i = \sum_M \frac{\partial \ln V_i}{\partial M} dM \tag{5}$$

where M is some physical variable, e.g., temperature, % Fe, % perovskite, etc., and V_i is V_P or V_S. Thus, the width and shape of the $d\ln V_P$

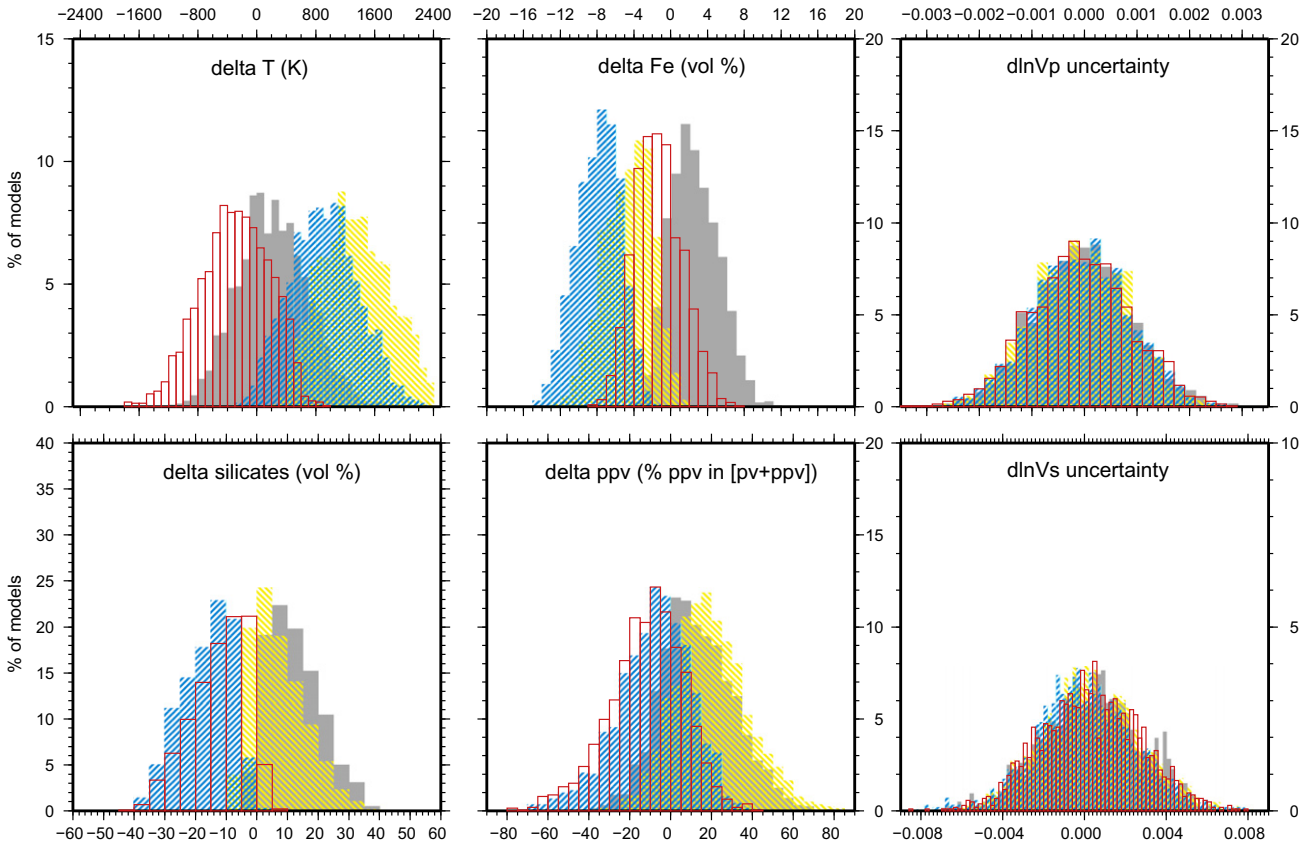

Fig. 9. The thermochemical variations and uncertainty distributions of the accepted models for Case 1, as shown in Fig. 8, now divided into four clusters (details in Section 3.2), to illustrate correlations between different parameters. Each cluster is plotted as a different colour histogram. 'delta' refers to the difference in T or X between the reference and perturbed composition for a given model. The four clusters for the seismic uncertainties are almost identical, because they are not correlated with thermochemical variations, and plot on top of each other. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

and $d\ln V_S$ distributions in Fig. 4 may reflect the ranges and magnitudes of the underlying thermochemical variations near the CMB – provided that the uncertainties in the seismic data do not exceed the sensitivity of $d\ln V_P$ and $d\ln V_S$ to such variations. We therefore attempt to construct assemblages of thermochemical variations which replicate the histograms (orange bars) in Fig. 4, including a correction for the effect of seismic uncertainty in V_i.

Six different mantle scenarios are considered (Table 1). We model $d\ln V_P$ and $d\ln V_S$ as the relative finite difference between a starting (reference)- and end (perturbed)-thermochemical model, i.e.:

$$d\ln V_i = \frac{V_{i(\text{perturbed})} - V_{i(\text{reference})}}{V_{i(\text{reference})}} \quad (6)$$

This formulation is used because the partial derivatives in Eq. (5) vary significantly and non-linearly with temperature and composition. Ideally a "reference model" should be the physical structure represented by PREM (Dziewonski and Anderson, 1981), the 1D seismic reference model relative to which the observed seismic velocity perturbations are measured. However, it is unclear if a unique physical reference structure even exists, and at best it remains poorly constrained (Cobden et al., 2009; Deschamps and Trampert, 2004). Therefore we restrict our reference models to only those in which V_P, V_S and density fall within 1% of PREM. This restriction maintains a consistency between mineral physics and seismic data, but still allows for a wide thermochemical range of reference models that encompasses uncertainty in the properties of any real physical reference model.

For each CMB scenario (Table 1), we generate a set of 800,000 reference- and perturbed-model pairs, and calculate their associated $d\ln V_P$, $d\ln V_S$ and R. Every reference model and perturbed model has a different temperature and composition, selected at random in a Monte-Carlo procedure from the broad ranges listed in Table 2, assuming a uniform distribution of each parameter between its limits. The limits are deliberately broad in order to include all feasible compositions at the CMB and to compensate for uncertainties in the phase equilibria, which cannot be quantified explicitly. We recognise that some of the selected models may, in reality, be thermodynamically unstable, but at present, phase equilibria (for the perovskite to post-perovskite transition) near the CMB are too poorly defined to be incorporated as an additional constraint. We could also, in principle, place tighter restrictions on the bulk chemical composition of each model using geodynamic or geochemical constraints, but we have chosen not to because we are interested in what the seismic data can resolve independently of such external constraints.

Reference models that do not fit the PREM constraint are rejected and recalculated until a 1% fit is obtained. Notably, many more iterations of the Monte Carlo procedure are required for ppv-free mineral assemblages than for ppv-bearing assemblages, before an acceptable reference model is found (Fig. 5). Further, for simulations in which the ppv content is allowed to vary, a large ppv content is present in most reference models (Fig. 6).

According to our mineral physics calculations, the average CMB temperature lies between 3000 ± 700 K (Fig. 6). Mineral physics experiments and ab initio simulations predict a temperature of between 2400 and 4200 K on the mantle side of the CMB (Asanuma

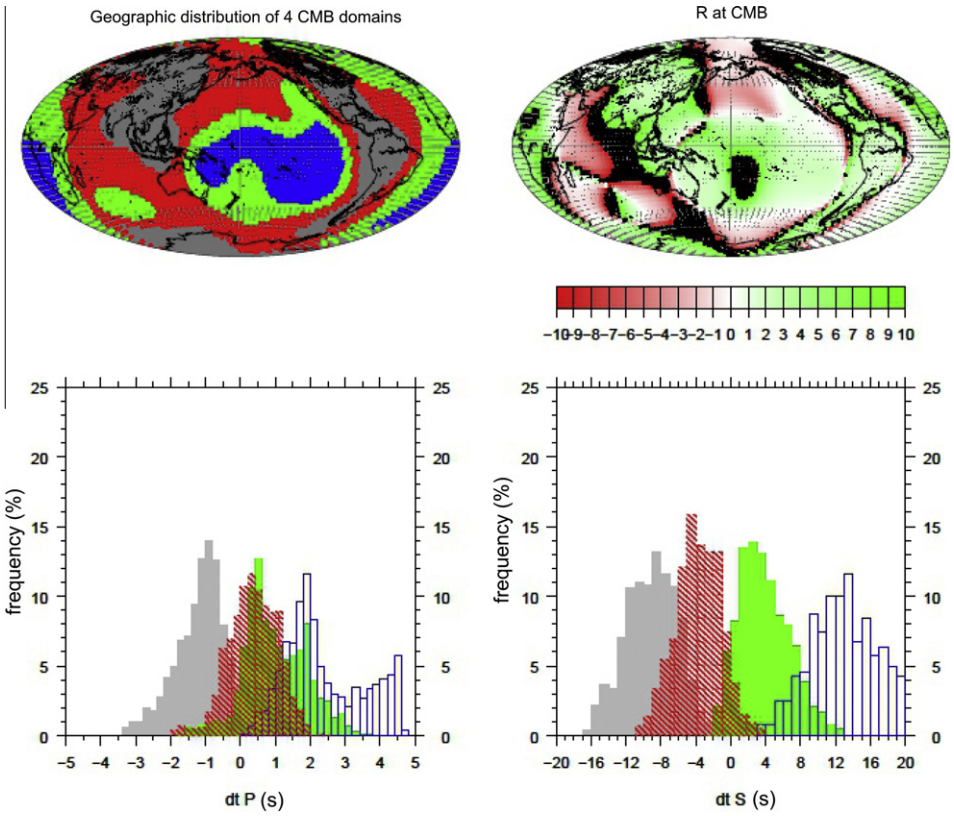

Fig. 10. Division of CMB into four domains (red, green, blue, grey) on the basis of the local travel time residuals. Bottom: P_{diff} and S_{diff} travel-time residuals for the four domains using cluster analysis. Top left: geographic locations of each cluster. Top right: R values near the CMB. Black indicates regions where the colour scale is saturated, i.e. $|R|$ is greater than 10. We see that the red cluster is associated with negative or high magnitude values of R (positive and negative). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

et al., 2010; Kamada et al., 2010; Campbell et al., 2007; Oganov et al., 2002), based on experiments on the melting temperature of iron or iron-alloys in the core, so the temperatures of our reference models are consistent with this range. It is unsurprising that the models favour the lower end of the experimental range, given that there may be high vertical thermal gradients in the bottom few km of the CMB, and an adiabatic geotherm extrapolated from the Earth's surface predicts a temperature of ~2400–2600 K by CMB depths (e.g., Cobden et al., 2009).

To account for seismic uncertainties, a correction is applied to each calculated $d\ln V_i$, following the same distribution as the uncertainty of the real seismic data (see Section 2.2). For a given set of 800,000 reference- and perturbed-model pairs, the corrections are normally distributed with a standard deviation of 0.00092 for $d\ln V_P$ and 0.00246 for $d\ln V_S$. From here onwards, the term 'model' refers to a reference- plus perturbed-model pair.

3. Statistical analysis and interpretation

3.1. Metropolis algorithm

We find that for every set of 800,000 models, the corresponding histograms of $d\ln V_P$, $d\ln V_S$ and R have very different morphologies from those of the seismic data (e.g., compare the pink lines with the orange bars, Fig. 4). We therefore apply a Metropolis algorithm (Mosegaard and Tarantola, 1995) which iteratively rejects or accepts models, until the subset of accepted models produces histograms with a close fit to the seismic data (black lines, Fig. 4). If we run the algorithm multiple times for the same set of 800,000 models, the thermochemical properties of the accepted models remain the same even though different individual models are accepted or rejected on each run. Our version of the Metropolis algorithm simultaneously fits all three of $d\ln V_P$, $d\ln V_S$ and R.

In all six of the thermochemical scenarios which we test, it is possible to extract a subset of models, using the Metropolis algorithm, which fit the seismic data histograms to more than 90%, and in many cases more than 98% (Fig. 5). Although the fit is slightly higher in cases in which pPv is included than in the case where pPv is excluded, this factor alone is not conclusive evidence for the presence of pPv, given the relatively good fit across all situations.

Our analysis illustrates the importance of incorporating seismic uncertainties into the physical interpretation of seismic data. For example, without adjusting for uncertainties, the R value of any fixed-composition model is always 2.1. Adding the correction for seismic uncertainty allows a set of models to be created which fit the broad R distribution of the seismic data to within 98% (Fig. 5). Without further investigation this would imply that the possibility of a CMB containing only temperature variations is viable.

We therefore study the distributions of the seismic uncertainties within the accepted models. Ideally a subset of models drawn from the initial 800,000 (via the Metropolis algorithm) should fit the $d\ln V_P$, $d\ln V_S$ and R histograms due to its underlying thermochemical variations. In this case, the uncertainties of the accepted models would follow the same normal distribution as the 800,000 initial models, with a mean of zero. If it is not possible to fit the $d\ln V_P$, $d\ln V_S$ and R histograms with the available thermochemical variations, then this may be compensated by systematically accepting models with certain uncertainty values biased away from the original normal distribution, thus changing the shape of

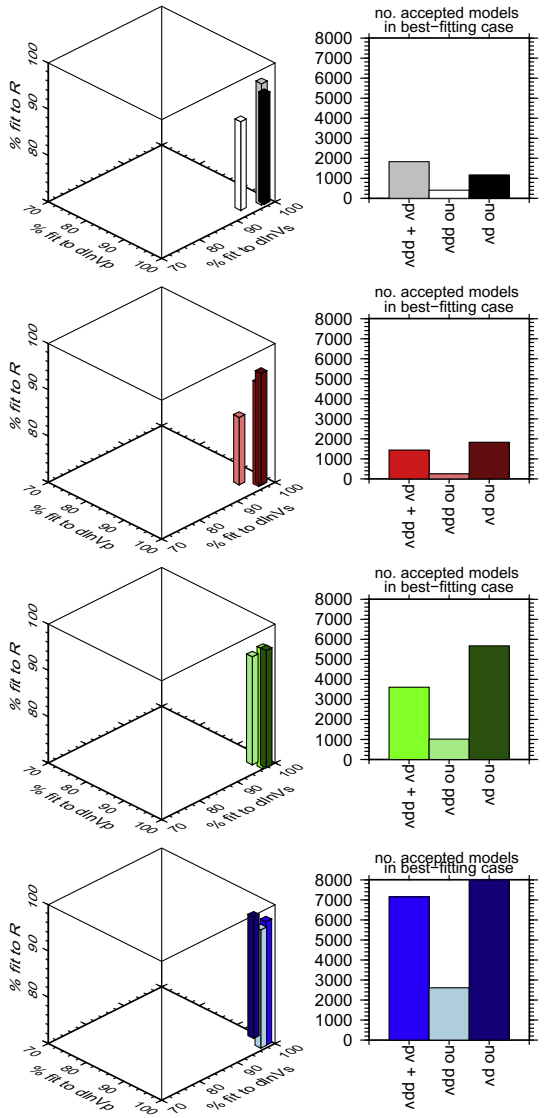

Fig. 11. Fit of Cases 1–3 (Table 1) to seismic data for each of the four CMB domains. Refer to Fig. 10 for colour legend. Right column: number of accepted models from 800,000 initial models after Metropolis algorithm. Left column: percent fit of each case to the regional distributions of $d\ln V_P$, $d\ln V_S$ and R; colour scheme as in the right column. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the $d\ln V_P$, $d\ln V_S$ and R distributions so as to match the seismic data. Fig. 7 shows the means of the seismic uncertainties for the accepted models in each CMB scenario (blue squares). In most cases the means are zero for both $d\ln V_P$ and $d\ln V_S$. However in two cases we observe systematic biases in the uncertainties: a CMB in which no-pPv is present, and a CMB in which pPv is present but only temperature variations occur laterally (i.e. composition is fixed). The uncertainties are significantly skewed towards negative values in $d\ln V_P$, and towards positive values in $d\ln V_S$ (e.g., for the no-pPv scenario, the uncertainty mean is shifted from zero by 45% of a standard deviation for $d\ln V_P$ and 60% for $d\ln V_S$). This misfit is also indicated in Table 3, where we show the degree of overlap between the full uncertainty distributions of the accepted and initial models. This analysis strongly indicates that lateral variations in chemical composition, together with some amount of pPv, should be present near the CMB. In this respect, our results differ from earlier studies (e.g., Kawai and Tsuchiya, 2009) which conclude that lateral velocity heterogeneity is predominantly generated by temperature variations.

3.2. K-means cluster analysis

When we examine the ranges of the thermochemical variations within the accepted "best-fit" model set for each CMB scenario, we find that an extremely broad range of each parameter is still present (Fig. 8). These broad ranges arise from a combination of uncertainty in the thermochemical state of the reference (starting) model (Fig. 6), and trade-offs between temperature and composition, or between different compositional variables, that are non-uniquely compatible with the seismic data. Seismic velocity perturbations as such do not allow us to place tight constraints on the magnitude of lateral thermal or compositional variations near the CMB.

Direct correlations between pairs of thermochemical parameters are not immediately obvious when studied in isolation. Therefore we apply a K-means cluster analysis (Sparks, 1973) to the accepted models. This algorithm divides and sorts multi-dimensional data into groups, or clusters, on the basis of shared physical properties. By trial and error we find that having four clusters illustrates the correlation between parameters most clearly. The most seismically-relevant physical properties which we input as separate parameters in the analysis are: temperature variation, iron variation, silica variation and post-perovskite variation. We also include the seismic uncertainties in $d\ln V_P$ and $d\ln V_S$.

Iron variation is defined as the variation in the total volume of all Fe-bearing minerals (Fe-Pv, Fe-pPv and FeO); silica variation is defined as the variation in total volume of all Si-bearing perovskites and post-perovskites plus free SiO_2; post-perovskite variation (where applicable) is defined as the change in percentage of pPv within the total (Pv plus pPv) content. We do not include calcium and aluminium variations as separate parameters in the cluster analysis, because we found that, for our mineral physics parameters, neither exhibits significant correlations with other thermochemical parameters.

We illustrate a cluster analysis of Case 1, the CMB containing both Pv and pPv (laterally-varying), in Fig. 9. Where there is a correlation between two or more parameters, the clusters will occupy different regions of the total ranges for each parameter, with distinct means. For example, the clusters show a significant anti-correlation between temperature and iron content: positive temperature variations (i.e., $T_{\mathrm{pertubedmodel}}$ is greater than $T_{\mathrm{referencemodel}}$) are associated with negative iron variations (e.g., compare the positions of the grey and yellow clusters in Fig. 9), and vice versa. The magnitude of this anti-correlation depends in turn on the sign (positive or negative) of the silica content variations (e.g., compare the positions of the grey and red clusters). The correlation with pPv-content is complex, depending on all three of temperature, iron and silica content.

If a particular variable has no correlation with the parameters included in the analysis, then the mean and range of each cluster for that variable should be the same for all clusters, and equal to the mean and range of all the models prior to clustering. We see this happening with the seismic uncertainty distributions in Fig. 9.

In Fig. 7, we show the trade-off between seismic uncertainties and thermochemical parameters for each of the six CMB models. While we observe almost zero trade-off for Case 1, a CMB containing lateral variations in both Pv and pPv, variable degrees of trade-off occur for the other five CMB cases. This suggests that the morphology of the observed seismic histograms (Fig. 4) is more likely to arise from a CMB containing lateral variations in the pPv/(Pv + pPv) fraction than one in which the fraction is kept fixed. However the difference in misfit between Case 1 and the other models in which pPv is present, is not as significant as the overall misfit of the pPv-free CMB. We do not think that the enhanced fit of Case 1 is caused intrinsically by its having an extra degree of freedom compared to the other five cases, but rather because it has

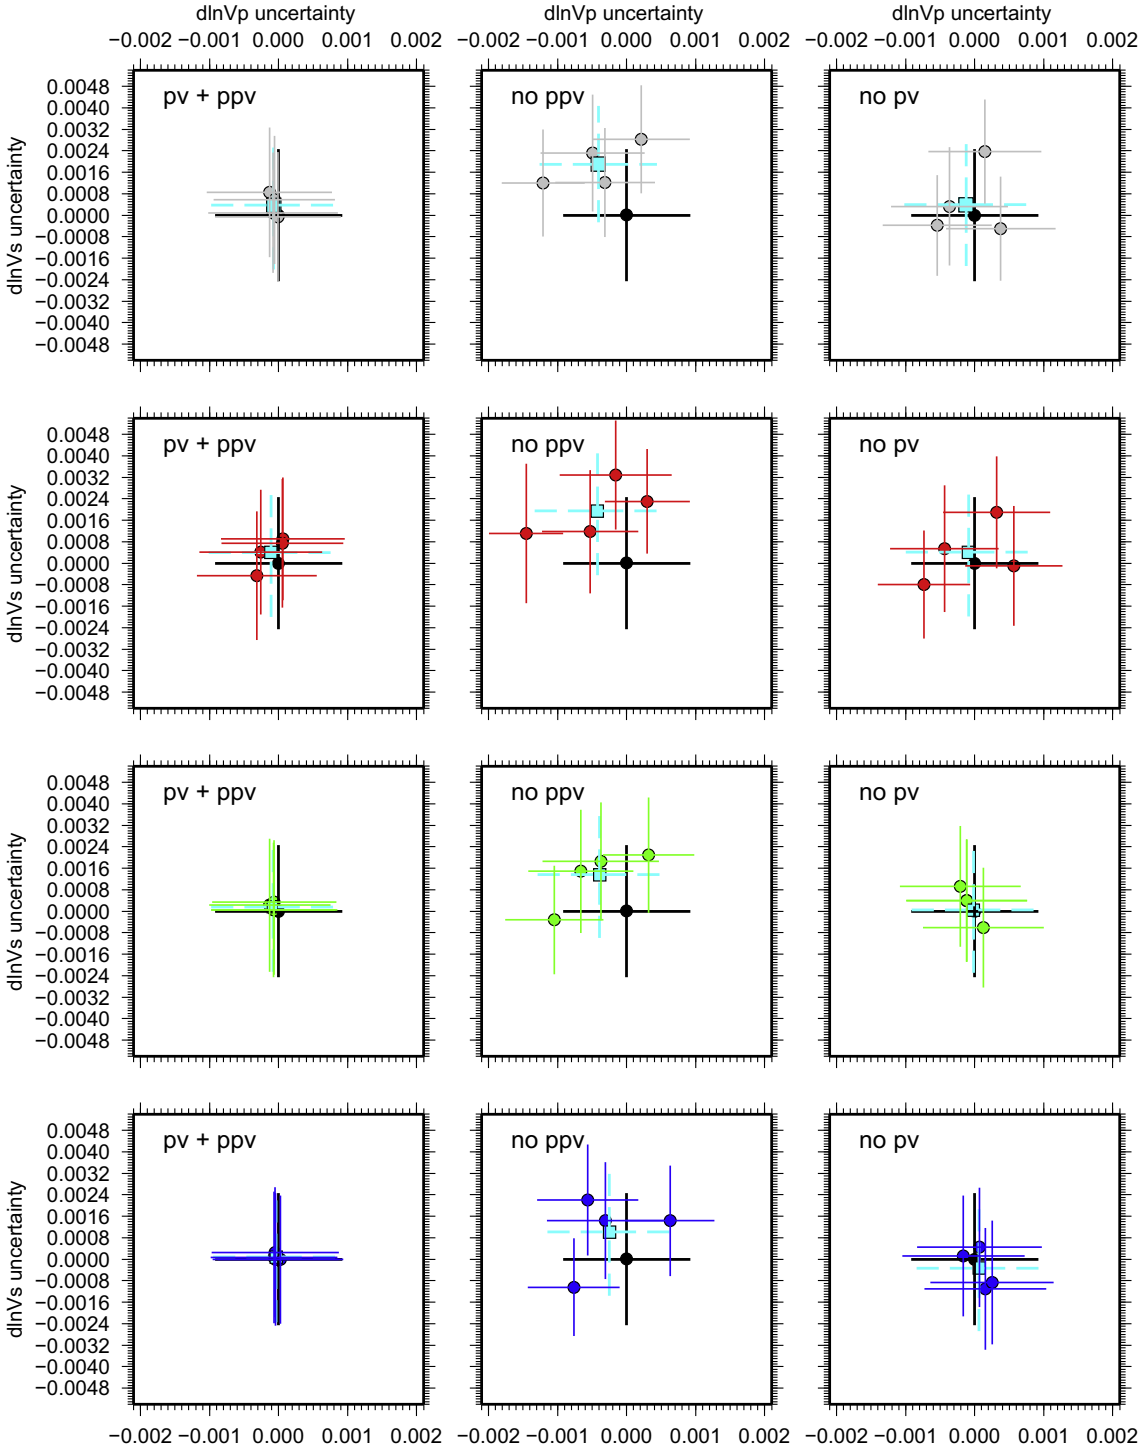

Fig. 12. Distributions of uncertainties for the four regional CMB domains. Colour scheme as in Fig. 10. Black circle and lines: mean and standard deviation of all 800,000 models. Light blue: mean and standard deviation of all accepted models. Other colours: means and standard deviation of 4 clusters into which accepted models are divided. As seen in the global dataset, there remains a significant misfit for the CMB containing no pPv, for all four regions. The greatest scatter is seen for the "red" region, which is associated with negative and high-magnitude R values (Fig. 10). We attribute this scatter to a larger uncertainty on the velocity perturbations: near-zero $d\ln V_P$ values can give rise to extremely large magnitude R values which can rapidly change from positive to negative. Further, the contribution of seismic uncertainty to the velocity perturbation becomes proportionately greater as either $d\ln V_S$ or $d\ln V_P$ tend towards zero. From a mineral physics standpoint, there are also fewer thermochemical changes that can produce negative R values than positive ones, so a smaller proportion of the 800,000 initial models have seismic properties compatible with the distributions of the seismic data for the red region, leaving fewer models after the Metropolis algorithm (less than 1000), which in turn makes any subsequent statistical analysis less reliable. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

more of a particular degree of freedom – namely lateral variations in pPv content – than the other cases. Cases 3–5 do also allow lateral variations in pPv content (due to lateral variations in chemistry, even though the Pv:pPv ratio may be fixed), and have the same number of degrees of freedom as Case 2, but Case 1 allows the greatest flexibility in the size of the lateral variations in post-perovskite and this is probably what gives it a better fit. We also wish to emphasise that the most significant result in Fig. 7 is the relative misfit of a pPv-free CMB rather than the minor differences between the various pPv-bearing models.

Table 4
Alternative representation of the data shown in Fig. 12. We normalised the histograms of $d\ln V_P$ and $d\ln V_S$ uncertainties for the accepted models, and calculated their overlap with the normalised uncertainty distributions of the 800,000 initial models. The product of the overlap in $d\ln V_P$ with the overlap in $d\ln V_S$ (third column) is a measure of the degree of bias in the accepted models' uncertainties. Colour scheme as shown in Fig. 10.

	%Overlap $d\ln V_P$ uncertainties	%Overlap $d\ln V_S$ uncertainties	P (no bias)
Blue			
pv + ppv	97.71	97.24	0.95
no ppv	87.97	84.10	0.74
no pv	96.37	94.05	0.91
Green			
pv + ppv	95.83	95.29	0.91
no ppv	82.57	77.80	0.64
no pv	97.68	96.27	0.94
Red			
pv + ppv	94.25	91.97	0.87
no ppv	79.54	67.60	0.54
no pv	94.64	92.60	0.88
Grey			
pv + ppv	95.12	92.39	0.88
no ppv	79.89	67.95	0.54
no pv	92.75	91.99	0.85

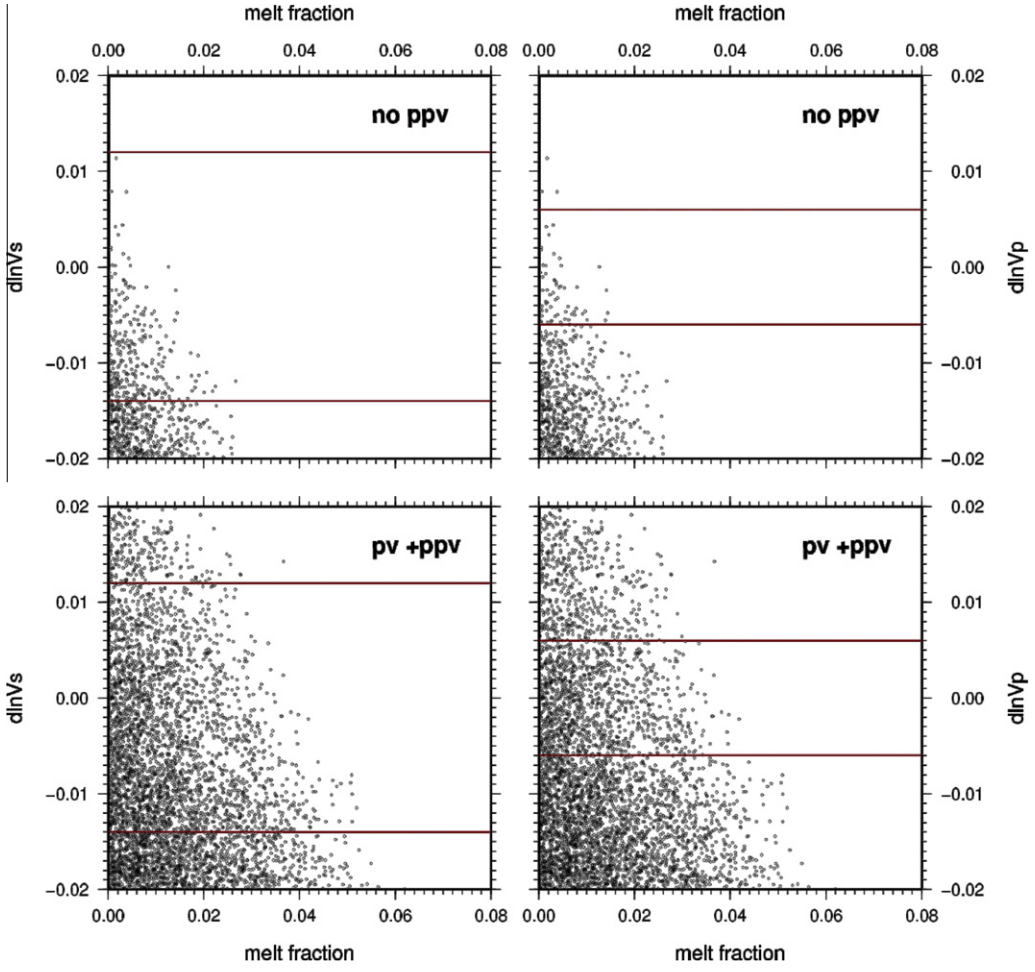

Fig. 13. $d\ln V_S$ and $d\ln V_P$ versus melt fraction for 800,000 thermochemical models, i.e. all models before the metropolis algorithm. Melt fraction varied randomly between 0 and 0.3 for all models in which the perturbed (final) temperature exceeded 4000 K. (1) Upper panel: a CMB containing no pPv (Case 2); (2) Lower panel: a CMB containing both Pv and pPv (Case 1). Red lines show max and min $d\ln V_S$ and $d\ln V_P$ of the seismic data, and indicate that at a horizontal resolution of 2000–3000 km, in a pPv-free CMB the melt fraction cannot exceed 1.5%, whilst for a pPv-bearing CMB, the melt fraction cannot be greater than 4%. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

If seismic uncertainties are excluded from the cluster analysis (not plotted here), then the means of each seismic uncertainty cluster are identical, with the exception of the fixed-composition CMB. This indicates that the correlations between different thermochemical parameters are, in most cases, stronger than correlations between thermochemical parameters and seismic

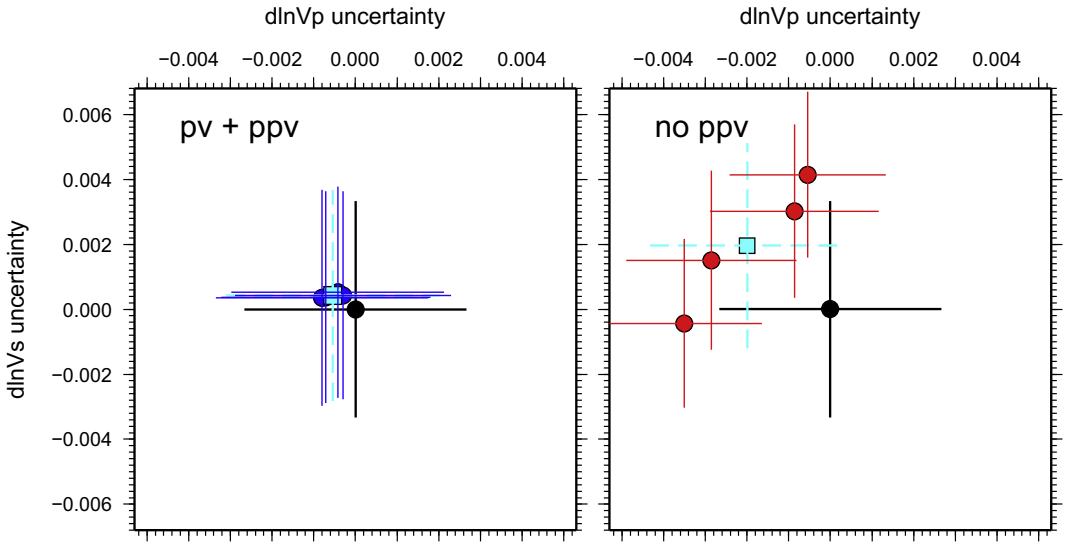

Fig. 14. Uncertainty distributions of accepted models for a CMB containing both Pv and pPv (blue, left) and containing no pPv (red, right), when uncertainty ranges of all 800,000 initial models are increased to match the maximum anisotropy in the model of Panning and Romanowicz (2006). Black circle and lines: mean and standard deviation of 800,000 initial models. Light blue square and dashed lines: mean and standard deviation of accepted models. Red and dark blue circles and lines: means and standard deviations of four clusters of the accepted models, after performing cluster analysis (see text). Compared with Fig. 7, the misfit to the seismic data is larger in both the above cases. However, the misfit of a pPv-free CMB is greater than half a standard deviation for both $d\ln V_S$ and $d\ln V_P$, and, once again, significantly larger than the misfit of a pPv-bearing CMB. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

uncertainties. A large correlation with the seismic uncertainties remains in the fixed-composition case, because the only physical variable is temperature, which cannot trade-off with chemical variations in order to fit the seismic data.

4. Discussion and conclusions

Previous seismic studies have suggested the existence of a "double-crossing" of the Pv–pPv phase transition (Hernlund et al., 2005), due to strong vertical temperature gradients above the core, whereby Pv converts to pPv at the top of D″ but reverts back to Pv at some point above the CMB. In this study we observe a strong misfit to the seismic data when the CMB region is pPv-free. This implies that, at shallower depths than the depth interval investigated here, and at the length scales of this study, not all pPv in the mantle should have already back-transformed to Pv. It does not imply that back-transitions do not happen at all. While our data cannot resolve the "post-perovskite lenses" observed on local scales (e.g., van den Berg et al., 2010; Kito et al., 2007; van der Hilst et al., 2007; Lay et al., 2006; Hernlund et al., 2005) they do support the possibility of lateral variations in pPv content near the CMB, and are thus consistent with such observations.

In addition, lateral variations in post-perovskite content may be associated with particular geodynamic domains (Tackley et al., 2007). For example, pPv could be restricted to colder regions associated with accumulated subducted material, and/or regions such as the large, low-shear velocity provinces (LLSVPs) underneath the Pacific and Africa (Lay et al., 2006), where a distinct chemistry (e.g., Fe-enrichment; Trampert et al., 2004; Ishii and Tromp, 1999) expands the stability field of pPv (e.g., Grocholski et al., 2012). To test this, we divided the CMB into four sub-domains (Fig. 10) on the basis of their seismic properties, using K-means cluster analysis. The only two parameters input to the cluster analysis were the two-way vertical travel time residuals δT (Eq. (3)) for P_{diff} and S_{diff}. The analysis was able to correlate seismically slow regions beneath the Pacific and Africa, and distinguish them from seismically-fast regions beneath the Americas, Eurasia and Antarctica. At the transition between these two seismic end-members, where δT changes from positive to negative for P_{diff} and S_{diff}, negative or high magnitude values (positive and negative) of R occur.

We observed that the biases in the seismic uncertainties seen on a global scale for a pPv-free mantle remain in all four regions (Figs. 11, 12 and Table 4), with no significant differences between, for example, the LLSVPs beneath Africa and the Pacific, and the faster, sub-continental regions. Thus the presence of pPv does not appear to be restricted to particular regions of the CMB based on our subdivision into four domains.

Our findings are robust within the limits of the currently-available mineral physics data, and provided that no other physical cause, most notably partial melting, could explain the velocity distributions in Fig. 4. The possibility that the seismic data may be explained by partial melting of a pPv-free CMB is not substantiated by this study. Assuming that melting occurs at temperatures above 4000 K near the CMB (e.g., Fiquet et al., 2010), we adjusted the $d\ln V_P$ and $d\ln V_S$ values in Cases 1 and 2 (see Table 1) for every model in which the perturbed (end) temperature exceeded 4000 K, since $d\ln V_P$ and $d\ln V_S$ should decrease in the presence of partial melting. $d\ln V_P$ and $d\ln V_S$ were reduced following the parameterisation of Berryman (2000) and Williams and Garnero (1996), with the melt fraction initially varying between 0 and 0.3 (a random value within these limits was selected for each model). At these high melt fractions, the reduction in $d\ln V_P$ and $d\ln V_S$ is dramatic, such that the $d\ln V_P$ and $d\ln V_S$ values both lie outside the range of the seismic data for almost all melt-bearing models (Fig. 13), and none of the accepted models (i.e. models accepted after application of the Metropolis algorithm) contained any melt. We then repeated the calculation but (based on Fig. 13) restricted the maximum melt fraction to 4% for Case 1 and 1.5% for Case 2. This time, a small percentage (around 3%) of the accepted models did contain melt, but the low melt fraction combined with the low number of models did not improve the fit of a pPv-free mantle to the seismic data (or worsen it for a pPv-bearing mantle), and thus, the seismic data histograms in Fig. 4 cannot be explained by partial melting effects. Note that while a large amount of melting near the CMB is not consistent with our global seismic dataset, this does not preclude the possibility that melting occurs locally on small length scales which are below the resolution of our data.

Furthermore, we cannot rule out the possibility that the seismic observations may be explained by an as-yet unknown mineral or physical structure at the base of the mantle whose seismic properties are similar to post-perovskite.

We do not think that our inferences are biased by anisotropy. We have very good azimuthal ray coverage (Fig. 2) similar to surface wave studies. The studies of azimuthal anisotropy using surface waves have shown that the isotropic part is unbiased by anisotropy (Trampert and Woodhouse, 2003). Because our modelling is similar that of to the great circle approximation in surface waves, we conclude that the vertical travel-time maps provide a good representation of the azimuthal average, and therefore contain no bias due to azimuthal anisotropy. There have also been many reports of radial anisotropy in D″ using body waves, (e.g., Panning and Romanowicz, 2004, 2006), which indicate that horizontally-polarised shear waves, V_{SH}, may be 1% faster, on average, than vertically polarised shear waves, V_{SV}. Our S_{diff} measurements are taken from V_{SH} only, whilst mineral physics calculations provide isotropic velocities. Assuming that there exists 1% radial anisotropy globally near the CMB, an approximate conversion from V_{SH} to V_{SV} can be done by adding 0.01 to each $d\ln V_S$ value in the seismic data, and recalculating the corresponding R values. On performing this calculation, we found that the higher misfit of a pPv-free CMB relative to a pPv-bearing CMB remained. Therefore even with the assumption that there exists maximum anisotropy in our V_{SH} dataset, our conclusions do not change. In reality, the anisotropic bias in our V_{SH} measurements may be less than 1%: normal mode evidence actually suggests that significant radial anisotropy in D″ is unlikely (Beghein et al., 2006). This range of 1% anisotropy which we tested for is also representative of the possible apparent SH_{diff}–SV_{diff} splitting reported by Komatitsch et al. (2010). We also tested a laterally varying model of transverse anisotropy (Panning and Romanowicz, 2006) by increasing the standard deviation of the seismic uncertainties added to the thermochemical models to match the range of anisotropy reported by Panning and Romanowicz (2006), namely 0.00267 for $d\ln V_P$ and 0.00333 for $d\ln V_S$. We found that this has no effect on our conclusions. The fit to seismic data generally decreases without changing the trend of our observations (Fig. 14).

Given the long wavelength of our P_{diff} and S_{diff} measurements (~250 km), combined with the spherical harmonic expansion of the dataset (which gives a horizontal resolution of the order of 2000–3000 km), we also do not expect small-scale features such as ULVZs to be resolved by, or map significantly into, the $d\ln V_P$ and $d\ln V_S$ observations. ULVZs are thought to exist locally along the edges of large, low shear velocity provinces (Rost et al., 2010) and have horizontal dimensions of up to a few hundred kilometres, (e.g., Ni and Helmberger, 2001).

A last concern could be the effect of CMB topography on our measurements. Using SPECFEM3D_GLOBE (e.g., Tromp et al., 2010)), we generated synthetic seismograms for the 3D mantle model S362ANI (Kustowski et al., 2008), with crustal model CRUST2.0 (Bassin et al., 2000) on top. We then ran a second simulation in which we also included a spherical harmonic degree 6 model of CMB topography (Mosca, 2010), assuming an amplitude of ±10 km for the peak topography. This amplitude is extreme, relative to seismic observations, which suggest that the maximum amplitude of CMB topography is actually around ±2 km (Tanaka, 2010). We found that at a period of 20 s, the effects of topography are small and the seismic waveforms remain virtually unchanged. We saw at most a 0.2 s difference for P waves and a 1 s difference for S waves, with a slight amplitude change for S_{diff} only. This observation suggests that diffracted waves are not a useful tool for studying the CMB topography.

We conclude that pPv presents a detectable seismic signature in both the 1D average velocity structure, and global long-wavelength distributions of 3D velocity variations, near the CMB. Contemporaneously, a pPv-free CMB is statistically unlikely (Fig. 5) and requires systematic biases in seismic uncertainties (Fig. 7) to fit the seismic data. Confirmation of a significantly enhanced fit of a pPv-bearing CMB to the seismic data relative to a pPv-free CMB provides the strongest evidence to date for the existence of pPv within the mantle. We further favour a CMB in which both Pv and pPv are present in laterally-varying proportions, although the difference in fit between this scenario and a CMB containing zero Pv, is small, and the possibility of a CMB in which 100% of the Mg,Fe,Al-silicates exist as the pPv phase cannot be ruled out.

Acknowledgements

We thank an anonymous reviewer whose comments helped to improve the manuscript, and Lars Stixrude and Carolina Lithgow-Bertelloni for helpful discussion of an earlier version of the manuscript. Part of this work was funded by the Dutch National Science Foundation under Grant number NWO: VICI1865.03.007 and by the European Commission under Grant number NERIES INFRAST-2.1-026130. J.R. was funded by NSF Grant EAR-0944167.

References

Asanuma, H., Ohtani, E., Sakai, T., Terasaki, H., Kamada, S., Kondo, T., Kikegawa, T., 2010. Melting of iron–silicon alloy up to the core–mantle boundary pressure: implications to the thermal structure of the Earth's core. Phys. Chem. Miner. 37, 353–359.

Bassin, C., Laske, G., Masters, G., 2000. The current limits of resolution for surface wave tomography in North America. EOS Trans. AGU 81, F897.

Beghein, C., Trampert, J., van Heijst, H.J., 2006. Radial anisotropy in seismic reference models of the mantle. J. Geophys. Res. Solid Earth 111, B02303.

Berryman, J.G., 2000. Seismic velocity decrement ratios for regions of partial melt in the lower mantle. Geophys. Res. Lett. 27, 421–424.

Brodholt, J.P., Helffrich, G., Trampert, J., 2007. Chemical versus thermal heterogeneity in the lower mantle: the most likely role of anelasticity. Earth Planet. Sci. Lett. 262, 429–437.

Campbell, A.J., Seagle, C.T., Heinz, D.L., Shen, G., Prakapenka, V.B., 2007. Partial melting in the iron–sulfur system at high pressure: a synchrotron X-ray diffraction study. Phys. Earth Planet. Inter. 162, 119–128.

Catalli, K., Shim, S., Prakapenka, V., 2009. Thickness and Clapeyron slope of the post-perovskite boundary. Nature 462, 782-U101.

Chaloner, J.W., Thomas, C., Rietbrock, A., 2009. P- and S-wave reflectors in D′ beneath southeast Asia. Geophys. J. Int. 179, 1080–1092.

Cobden, L., Goes, S., Ravenna, M., Styles, E., Cammarano, F., Gallagher, K., Connolly, J.A.D., 2009. Thermochemical interpretation of 1-D seismic data for the lower mantle: the significance of nonadiabatic thermal gradients and compositional heterogeneity. J. Geophys. Res. Solid Earth 114, B11309.

Deschamps, F., Trampert, J., 2004. Towards a lower mantle reference temperature and composition. Earth Planet. Sci. Lett. 222, 161–175.

Deschamps, F., Trampert, J., 2003. Mantle tomography and its relation to temperature and composition. Phys. Earth Planet. Inter. 140, 277–291.

Dziewonski, A.M., Anderson, D.L., 1981. Preliminary Reference Earth Model. Phys. Earth Planet. Inter. 25, 297–356.

Ekström, G., 2011. A global model of Love and Rayleigh surface wave dispersion and anisotropy, 25–250 s. Geophys. J. Int. 187, 1668–1686.

Fiquet, G., Auzende, A.L., Siebert, J., Corgne, A., Bureau, H., Ozawa, H., Garbarino, G., 2010. Melting of peridotite to 140 Gigapascals. Science 329, 1516–1518.

Grocholski, B., Catalli, K., Shim, S.H., Prakapenka, V., 2012. Mineralogical effects on the detectability of the postperovskite boundary. Proc. Natl. Acad. Sci. USA 109, 2275–2279.

Hernlund, J.W., Thomas, C., Tackley, P.J., 2005. A doubling of the post-perovskite phase boundary and structure of the Earth's lowermost mantle. Nature 434, 882–886.

Hernlund, J.W., Houser, C., 2008. The statistical distribution of seismic velocities in Earth's deep mantle. Earth Planet. Sci. Lett. 265, 423–437.

Hirose, K., 2007. Discovery of post-perovskite phase transition and the nature of D″ layer. In: Hirose, K., Brodholt, J., Lay, T., Yuen, D.A. (Eds.), Post-perovksite: The Last Mantle Phase Transition. American Geophysical Union, Washington, DC, pp. 19–35.

Hustoft, J., Catalli, K., Shim, S., Kubo, A., Prakapenka, V.B., Kunz, M., 2008. Equation of state of NaMgF(3) postperovskite: implication for the seismic velocity changes in the D″ region. Geophys. Res. Lett. 35, L10309.

Hutko, A.R., Lay, T., Revenaugh, J., Garnero, E.J., 2008. Anticorrelated seismic velocity anomalies from post-perovskite in the lowermost mantle. Science 320, 1070–1074.

Ishii, M., Tromp, J., 1999. Normal-mode and free-air gravity constraints on lateral variations in velocity and density of Earth's mantle. Science 285, 1231–1236.

Kamada, S., Terasaki, H., Ohtani, E., Sakai, T., Kikegawa, T., Ohishi, Y., Hirao, N., Sata, N., Kondo, T., 2010. Phase relationships of the Fe–FeS system in conditions up to the Earth's outer core. Earth Planet. Sci. Lett. 294, 94–100.

Kawai, K., Tsuchiya, T., 2009. Temperature profile in the lowermost mantle from seismological and mineral physics joint modeling. Proc. Natl. Acad. Sci. USA 106, 22119–22123.

Kito, T., Rost, S., Thomas, C., Garnero, E.J., 2007. New insights into the P- and S-wave velocity structure of the D″ discontinuity beneath the Cocos plate. Geophys. J. Int. 169, 631–645.

Komatitsch, D., Vinnik, L.P., Chevrot, S., 2010. SH_{diff}–SV_{diff} splitting in an isotropic Earth. J. Geophys. Res. Solid Earth 115, B07312.

Kustowski, B., Ekstrom, G., Dziewonski, A.M., 2008. Anisotropic shear-wave velocity structure of the Earth's mantle: a global model. J. Geophys. Res. Solid Earth 113, B06306.

Lay, T., 2007. Deep Earth structure: lower mantle and D″. In: Romanowicz, B., Dziewonski, A.M. (Eds.), Seismology and Structure of the Earth. Elsevier, Amsterdam, pp. 619–654.

Lay, T., Garnero, E.J., 2007. Reconciling the post-perovskite phase with seismological observations of lowermost mantle structure. In: Hirose, K., Brodholt, J., Lay, T., Yuen, D. (Eds.), Post-Perovskite: The Last Mantle Phase Transition. American Geophysical Union, Washington, DC, pp. 129–153.

Lay, T., Hernlund, J., Garnero, E.J., Thorne, M.S., 2006. A post-perovskite lens and D″ heat flux beneath the central Pacific. Science 314, 1272–1276.

Lay, T., Heinz, D.L., Ishii, M., Shim, S., Tsuchiya, J., Tsuchiya, T., Wentzcovitch, R.M., Yuen, D.A., 2005. Multidisciplinary impact of the deep mantle phase transition in perovskite structure. EOS Trans. AGU 86, 1–5.

Masters, G., Laske, G., Bolton, H., Dziewonski, A.M., 2000. The relative behaviour of shear velocity, bulk sound speed, and compressional velocity in the mantle: implications for chemical and thermal structure. In: Karato, S., Forte, A., Liebermann, R., Masters, G., Stixrude, L. (Eds.), Earth's Deep Interior: Mineral Physics and Tomography from the Atomic to the Global Scale. American Geophysical Union, Washington, DC, pp. 63–87.

Mosca, I., 2010. Probabilistic Tomography using Body Wave, Normal Mode and Surface Wave Data. PhD thesis, Utrecht University.

Mosca, I., Trampert, J., 2009. Path-average kernels for long wavelength traveltime tomography. Geophys. J. Int. 177, 639–650.

Mosegaard, K., Tarantola, A., 1995. Monte-Carlo sampling of solutions to inverse problems. J. Geophys. Res. Solid Earth 100, 12431–12447.

Murakami, M., Hirose, K., Ono, S., Ohishi, Y., 2003. Stability of $CaCl_2$-type and α-PbO_2-type SiO_2 at high pressure and temperature determined by in-situ X-ray measurements. Geophys. Res. Lett. 30 (5), 1207.

Murakami, M., Hirose, K., Kawamura, K., Sata, N., Ohishi, Y., 2004. Post-perovskite phase transition in $MgSiO_3$. Science 304, 855–858.

Ni, S., Helmberger, D.V., 2001. Probing an ultra-low velocity zone at the core mantle boundary with P and S waves. Geophys. Res. Lett. 28, 2345–2348.

Nishio-Hamane, D., Yagi, T., 2009. Equations of state for postperovskite phases in the MgSiO(3)–FeSiO(3)–FeAlO(3) system. Phys. Earth Planet. Inter. 175, 145–150.

Oganov, A.R., Ono, S., 2004. Theoretical and experimental evidence for a post-perovskite phase of $MgSiO_3$ in Earth's D″ layer. Nature 430, 445–448.

Oganov, A.R., Brodholt, J.P., Price, G.D., 2002. Ab initio theory of phase transitions and thermoelasticity of minerals. In: Gramaccioli, C.M. (Ed.), Energy Modelling in Minerals. EMU Notes in Mineralogy, vol. 4. Eötvös University Press, Hungary, pp. 83–170.

Panning, M., Romanowicz, B., 2004. Inferences on flow at the base of Earth's mantle based on seismic anisotropy. Science 303, 351–353.

Panning, M., Romanowicz, B., 2006. A three-dimensional radially anisotropic model of shear velocity in the whole mantle. Geophys. J. Int. 167, 361–379.

Ritsema, J., van Heijst, H.J., 2002. Constraints on the correlation of P- and S-wave velocity heterogeneity in the mantle from P, PP, PPP and PKPab traveltimes. Geophys. J. Int. 149, 482–489.

Rost, S., Garnero, E.J., Stefan, W., 2010. Thin and intermittent ultralow-velocity zones. J. Geophys. Res. Solid Earth 115, B06312.

Shim, S.H., Duffy, T.S., Jeanloz, R., Shen, G., 2004. Stability and crystal structure of $MgSiO_3$ perovskite to the core–mantle boundary. Geophys. Res. Lett. 31, L10603.

Sparks, D.N., 1973. Euclidean cluster analysis. J. R. Stat. Soc. Ser. C: Appl. Stat. 22 (1), 126–130.

Stackhouse, S., Brodholt, J.P., 2007. High-temperature elasticity of $MgSiO_3$ post-perovskite. In: Hirose, K., Brodholt, J.P., Lay, T., Yuen, D. (Eds.), Post-perovskite: The Last Mantle Phase Transition. American Geophysical Union, Washington, DC, pp. 99–113.

Stixrude, L., Lithgow-Bertelloni, C., 2005. Thermodynamics of mantle minerals – I. Physical properties. Geophys. J. Int. 162, 610–632.

Stixrude, L., Lithgow-Bertelloni, C., 2011. Thermodynamics of mantle minerals – II. Phase equilibria. Geophys. J. Int. 184, 1180–1213.

Tackley, P.J., Nakagawa, T., Hernlund, J.W., 2007. Influence of the post-perovskite transition on thermal and chemical mantle convection. In: Hirose, K., Brodholt, J.P., Lay, T., Yuen, D. (Eds.), Post-Perovskite: The Last Mantle Phase Transition. American Geophysical Union, Washington, DC, pp. 229–247.

Tanaka, S., 2010. Constraints on the core–mantle boundary topography from P4KP-PcP differential travel times. J. Geophys. Res. Solid Earth 115, B04310.

Thomas, C., Garnero, E.J., Lay, T., 2004a. High-resolution imaging of lowermost mantle structure under the Cocos plate. J. Geophys. Res. Solid Earth 109, B08307.

Thomas, C., Kendall, J.M., Lowman, J., 2004b. Lower-mantle seismic discontinuities and the thermal morphology of subducted slabs. Earth Planet. Sci. Lett. 225, 105–113.

Trampert, J., Deschamps, F., Resovsky, J., Yuen, D., 2004. Probabilistic tomography maps chemical heterogeneities throughout the lower mantle. Science 306, 853–856.

Trampert, J., Woodhouse, J.H., 2003. Global anisotropic phase velocity maps for fundamental mode surface waves between 40 and 150 s. Geophys. J. Int. 154, 154–165.

Trampert, J., Snieder, R., 1996. Model estimations biased by truncated expansions: possible artifacts in seismic tomography. Science 271, 1257–1260.

Trampert, J., van Heijst, H.J., 2002. Global azimuthal anisotropy in the transition zone. Science 296, 1297–1299.

Tromp, J., Komatitsch, D., Hjoerleifsdottir, V., Liu, Q., Zhu, H., Peter, D., Bozdag, E., McRitchie, D., Friberg, P., Trabant, C., Hutko, A., 2010. Near real-time simulations of global CMT earthquakes. Geophys. J. Int. 183, 381–389.

van den Berg, A.P., De Hoop, M.V., Yuen, D.A., Duchkov, A., van der Hilst, R.D., Jacobs, M.H.G., 2010. Geodynamical modeling and multiscale seismic expression of thermo-chemical heterogeneity and phase transitions in the lowermost mantle. Phys. Earth Planet. Inter. 180, 244–257.

van der Hilst, R.D., de Hoop, M.V., Wang, P., Shim, S.-H., Ma, P., Tenorio, L., 2007. Seismostratigraphy and thermal structure of Earth's core–mantle boundary region. Science 315, 1813–1817.

Van der Voo, R., Spakman, W., Bijwaard, H., 1999. Mesozoic subducted slabs under Siberia. Nature 397, 246–249.

Wentzcovitch, R.M., Tsuchiya, T., Tsuchiya, J., 2006. MgSiO(3) postperovskite at D″ conditions. Proc. Natl. Acad. Sci. USA 103, 543–546.

Williams, Q., Revenaugh, J., Garnero, E., 1998. A correlation between ultra-low basal velocities in the mantle and hot spots. Science 281, 546–549.

Williams, Q., Garnero, E.J., 1996. Seismic evidence for partial melt at the base of Earth's mantle. Science 273, 1528–1530.

Wookey, J., Stackhouse, S., Kendall, J.M., Brodholt, J., Price, G.D., 2005. Efficacy of the post-perovskite phase as an explanation for lowermost-mantle seismic properties. Nature 438, 1004–1007.

Wysession, M.E., Lay, T., Revenaugh, J., Williams, Q., Garnero, E.J., Jeanloz, R., Kellog, L.H., 1998. The D″ discontinuity and its implications. In: Gurnis, M., Wysession, M.E., Knittle, E., Buffet, B.A. (Eds.), The Core–Mantle Boundary Region. American Geophysical Union, Washington, DC, pp. 273–298.

Xu, W., Lithgow-Bertelloni, C., Stixrude, L., Ritsema, J., 2008. The effect of bulk composition and temperature on mantle seismic structure. Earth Planet. Sci. Lett. 275, 70–79.