

Chapter 2

The Thermal Properties of Iron-bearing Silicate Perovskite and the Implications for Lower Mantle Structures

Abstract

The high-temperature and -pressure equations of state of 0% and 13% Fe-bearing Mg-silicate perovskite are measured using powder diffraction in the diamond anvil cell. The derived equation of state properties are broadly consistent with previous work, though the Mg-endmember shows slightly larger volumes at high pressure than previous studies. The derived equations of state are incorporated into an ideal mixing model used to probe the density stability of perovskite material at a range of temperatures and compositions. This stability analysis shows that the metastable dome theory for Large Low Shear Velocity Provinces (LLSVPs) is possible but unlikely, with a probability of less than roughly 10%, given the narrow range of compositions values consistent with metastability.

2.1 Introduction

The Earth's lower mantle is thought to be composed of primarily aluminous $(\text{Mg, Fe})\text{SiO}_3$ perovskite coexisting with $(\text{Mg, Fe})\text{O}$ ferropericlase. While the exact phase proportions depend on the assumed compositional model for the lower mantle, e.g., pyrolytic vs. picritic, iron-bearing silicate perovskite is thought to dominate making it the most common mineral in the silicate Earth (*Kesson et al.* 1998, *Mattern et al.* 2005). As the primary lower mantle phase, it plays a major role in determining the physical properties of lower mantle material, thereby exerting large control over mantle dynamics and evolution. Characterizing its temperature-dependent equation of state is therefore vital to both interpreting seismic observations and providing accurate input to geodynamics calculations.

Most previous high-pressure experimental studies of silicate perovskite have either focused on the Mg-endmember and neglected the role of Fe (*Fiquet et al.*, 2000), or they have studied a range of Fe compositions but were restricted to ambient temperature (*Lundin et al.*, 2008; *Dorfman et al.*, 2013). There are also numerous studies that have examined natural compositions, including ferric iron and aluminum, which can also have an important impact on the equation of state behavior of perovskite (*Catalli et al.*, 2011). While clearly crucial to our understanding of the deep Earth, such natural compositions add significant complication to interpreting results, and thus it is important to couple them with studies the focus on simpler systems where it is easier to gain an understanding of systematic behavior. The majority of these previously mentioned studies have also used strongly non-hydrostatic pressure media, including NaCl and Ar. Recent work, however, has shown that non-hydrostatic stress states may alter a mineral's compression behavior (e.g., *Takemura,*

2007; You Shu-Jie and Chang-Qing, 2009; Iizuka *et al.*, 2010), rendering such experiments perhaps less representative of conditions within the Earth. In this study, we seek to help provide a more complete picture of perovskite by studying a simplified chemical system in great detail, examining iron-bearing and iron-free synthetic samples at both high pressures and temperatures.

In this work, we report the temperature-dependent equation of state for polycrystalline perovskite-structured $(\text{Mg}_{1-x}, \text{Fe}_x)\text{SiO}_3$ samples with two different iron compositions relevant to the lower mantle ($x = 0\%$ and 13%). These samples were probed with x-rays under a wide range of conditions between 30 and 130 GPa and room temperature up to 2500K. These results are then compared to assess the effect of iron composition on silicate perovskite's high-temperature compression behavior. Finally, we assess the plausibility of different competing models to explain perovskite-rich lower mantle structures.

2.2 Methods

Our perovskite samples were made from synthetic $(\text{Mg}_{1-x}, \text{Fe}_x)\text{SiO}_3$ pyroxene starting materials with iron compositions of $x = 0\%$ and 13% . These compositions were verified using micro-probe analysis, and Mössbauer spectroscopy on the Fe-bearing sample showed it to contain undetectable levels of ferric iron, constraining it to less than $3\% \text{Fe}^{3+}$ (Jackson *et al.*, 2009; Zhang *et al.*, 2011). The iron-free MgSiO_3 glass starting sample was mixed with roughly 10 wt% Au, which served as an infrared absorber, ensuring good coupling with the laser. To minimize contamination of the sample, the grinding and mixing was performed between two pristine tungsten-carbide cubes. Since their Fe content is sufficient

to couple with the laser, the iron-bearing samples were not mixed with gold, avoiding any potential Fe–Au alloy formation which might affect experimental results; All samples were then loaded into symmetric diamond anvil cells using pre-indented Re gaskets and 250 micron culets. The samples were placed on top of ~ 10 micron thick insulating dehydrated NaCl plates, which are crucial for reducing thermal gradients within the sample during laser heating. Each sample was also loaded with a few synthetic ruby spheres for offline pressure determination using the ruby fluorescence method (e.g., *Jacobsen et al.*, 2008; *Silvera et al.*, 2007; *Mao et al.*, 1986). The cells were loaded with a Ne pressure medium using the GSECARS gas-loading system at the Advanced Photon Source (APS) of Argonne National Laboratory (*Rivers et al.*, 2008). The samples were pressurized to ~ 30 GPa where they were then laser annealed within the stability field of perovskite.

Powder diffraction measurements were taken at the Sector 13-ID-D GSE-CARS beamline at the APS. Diffraction spectra were taken in roughly 2 to 4 GPa steps between 35 and 110 GPa for each sample. The pressure for each spectrum was determined using the Ne pressure medium as the primary pressure marker, *Dewaele et al.* (2008), together with the offline ruby fluorescence measurements. As discussed in Section 2.3.2, great care is taken in the analysis of the neon diffraction lines to assess the effect of deviatoric stress on the measured pressure. In roughly 5 to 10 GPa steps, the samples were laser heated in stages up to a maximum temperature of 2500 K. These temperatures were determined from thermal emission spectra by fitting them to a blackbody curve (*Heinz and Jeanloz*, 1987; *Shen et al.*, 2001).

The raw powder diffraction images were converted to one-dimensional spectra using

a suite of routines written in MATLAB. In these routines, the observing geometry is first determined from calibration spectra using an automated statistical method. This has the advantage over the standard ‘click-based’ method employed in FIT2D that it requires little user input and generates a reproducible result that is determined by the maximum likelihood configuration given the data. Using the derived geometric calibration, it is straightforward to calculate the observed diffraction angle (2θ) corresponding to each pixel on the CCD. The diffraction angle can then be converted into the more physically relevant inverse d-spacing ($1/d$), which describes the spacing between a set of coherently reflecting crystal planes, by simply applying Bragg’s law for first-order reflections, $1/d = 2 \sin(2\theta/2)/\lambda$, where λ is the wavelength of the monochromatic x-rays used to probe the sample. Each raw image is then ‘integrated’ assuming Poisson statistics for the uncertainties in number of photons hitting each pixel. For some of the diffraction images, which contain over-exposed pixels, a further preprocessing step is required to produce accurate 1D spectra (for more details on our data reduction pipeline see 2.A). The final step in the integration process is to subtract off a decent guess of the background spectrum so that data sets can be easily examined and fit. We employ the Bayesian background identification and subtraction method presented in *David and Sivia* (2001), which automatically determines a reasonable polynomial background curve assuming the potential presence of large positive deviations due to the presence of as-yet unmodeled diffraction peaks.

After these processing steps, we obtain a set of 1D powder diffraction spectra for a range of pressures and temperatures for both the 13% Fe-bearing perovskite sample and the pure Mg-endmember sample. To gain a good overview of the room temperature data,

we construct interpolated cold-compression maps shown in Figures 2.1 and 2.2. To best highlight all aspects of the data, we plot the intensity using a geometric color-scale¹, rendering every line visible to the eye even though their amplitudes vary over many orders of magnitude. In each figure, we also overplot the fitted line positions for each phase, showing that dozens of perovskite peaks are visible in the spectra as well as the high-intensity peaks from neon which are used as in situ pressure markers, as described in Section 2.3.2.

2.3 Analysis

The analysis of the 1D integrated spectra is broken into four parts, each of which are each discussed below. First is the process of peak-fitting, which transforms the continuous spectrum into a discrete set of measured peak positions and associated uncertainties. The next task is to determine the pressure–temperature conditions inside the diamond anvil cell. In this section, particular care is taken to determine the deviatoric stress contribution to the pressure. Thirdly, the set of peak positions corresponding to perovskite in each spectrum are transformed into estimates of the unit cell parameters along with uncertainties. Finally, the estimated volumes, temperatures and pressures are all combined to retrieve the confidence regions for the overall equation of state of each sample.

¹We use the inverse hyperbolic sin transformation to visualize the background subtracted intensities, which generates a roughly geometric color-scale. This transformation is well suited to values that vary over orders of magnitude, since it reduces to a simple log-transformation for large values, while still allowing for negative values that result from background subtraction (*Burbidge et al.*, 1988). Note that this transform is only used for visualization and plays no role in the actual data fitting.




Figure 2.1: Cold compression data map for 13% Fe-bearing perovskite sample showing intensity as a function of pressure and inverse d-spacing. The upper panel shows the raw interpolated intensities, while the lower panel overplots the fitted line positions for the entire data set. Diffraction intensities are shown using a geometric (or roughly logarithmic) scale, allowing even lines of very small amplitude to be visible. The phases with lines apparent in at least some spectra are the perovskite sample (pv), neon(ne), gold (au), the rhenium gasket (re), stishovite (stish), and the post-stishovite CaCl_2 structure of silica (sio2).




Figure 2.2: Cold compression data map for pure endmember Mg-perovskite sample showing intensity as a function of pressure and inverse d-spacing. The upper panel shows the raw interpolated intensities, while the lower panel overplots the fitted line positions for the entire data set. Diffraction intensities are shown using a geometric (or roughly logarithmic) scale, allowing even lines of very small amplitude to be visible. The phases with lines apparent in at least some spectra are the perovskite sample (pv), neon(ne), gold (au), B2-structured NaCl (nacl), the rhenium gasket (re), stishovite (stish), and the post-stishovite CaCl_2 structure of silica (sio2).

2.3.1 Peak Fitting

The process of obtaining unit cell parameters from a set of integrated 1D diffraction spectra usually follows one of two main strategies: whole pattern refinement or individual peak-fitting. The first method, officially called Rietveld refinement, is carried out by modeling the entire spectrum as a whole, and is thus often referred to as whole pattern refinement. Using this method, every crystal phase that is present in the spectrum is assumed to have a predetermined crystal structure, which dictates the relative line amplitudes. This global model is simultaneously fitted to the entire spectrum, where the free parameters are taken to be the unit cell parameters for each phase (which determine the line positions), the apparent phase fraction determining the relative amplitude of the lines from different phases, and a set of parameters that control the line widths and shapes as a function of diffraction angle. The fitting method follows the standard weighted least-squares, or chi-square minimization approach, which is appropriate to independent normally distributed data as derived using the Bayesian framework in Chapter 3. There exist numerous codes that perform this type of fitting and have graphical user interfaces (GUIs) improving ease of use, including the program GSAS (General Structure Analysis System) with its EXPGUI interface (*Toby, 2001*).

While whole pattern refinement can be quite convenient, it has a number of limitations. As it is typically applied, the user constructs a complete model of the crystal system, which allows the calculation of the entire diffraction pattern including the relative amplitudes of each line in the profile, which are tied to both the overall structure and the relative atomic positions. This is an excellent approach when the data are well represented by an ideal pow-

der diffraction model. In many high-pressure diamond anvil cell experiments, however, the diffraction pattern is seen to evolve with increasing pressure and temperature, causing the relative amplitudes of each line to shift away from their default values expected at ambient conditions. There are many possible explanations for this behavior, including: incomplete sampling of randomly oriented crystallites resulting from small number statistics for crystallites in the diffraction condition; crystal preferred orientation (due perhaps to deviatoric stress in the sample chamber); or perhaps systematic shifts of the atomic positions with increasing pressure. While all of these effects can potentially be directly modeled, they require large amounts of very precise—usually single crystal—diffraction data, making it impractical for the vast majority of experiments.

In the face of this large set of unmodeled physical processes, the whole pattern refinement method is often incapable of accurately representing the relative line intensities in high-pressure experiments. We may choose to ignore such systematic errors in our model, dismissing them as minor effects, but we do so at the risk of inducing strong biases in our results. Therefore, when applying the whole pattern refinement method to high-pressure data, one must be very careful to ensure that the relative amplitudes are well described by the model to within the observational Poisson noise. If they are not, then the best-fit model curves, which may look generally reasonable to the eye, are likely to be heavily weighted toward only a few diffraction peaks within spectrum. In particular, the fit will be depend almost exclusively on only the smallest subset of intense diffraction lines that can constrain the unit cell parameters (representing a range of hkl values). All of the information contained in the clearly visible but low intensity lines can easily be ignored in favor of small

yet meaningless improvements to the fit quality for the most intense peaks.

One of the typical ways to deal with this issues is to employ the LeBail fitting method, as implemented in GSAS (*Toby, 2001*). This is an iterative algorithm that allows the relative peak intensities to adjust away from their expected ideal values. It accomplishes this by approximately apportioning the observed intensities amongst the diffraction peaks present at each portion of the spectrum. The problem with this method is that it can often lead to unstable or “runaway solutions”, where the fitting procedure must be artificially halted early on, prior to convergence, otherwise the line intensities change so drastically that all of the model parameters, including the dimensions of the unit cell, begin to undertake unrealistically large excursions from reasonable values. This reflects a general property of data fitting in situations where the model parameters are under-constrained—when too many degrees of freedom are given to the model, the fitting algorithm will take advantage of those parameters, adjusting them to reduce the misfit, even if those adjustments are not sensible and lead to unphysical results. The LeBail method is therefore a good choice when the data quality is extremely high and thus there is little potential for bias resulting from overfitting. Unfortunately, this is often not the case for multiphase high-pressure diffraction spectra, especially when the spectra include low symmetry phases like perovskite.

It is for these reasons that we favor the second method of individual peak fitting. In this approach, the spectrum is split up into different sections that contain diffraction lines of interest. These lines are then fit individually, or in small clustered groups where the peaks are seen to overlap one another. Using this approach, it is simple to leave the relative peak intensities unconstrained, and thus avoid the issues outlined above. This advantage comes

with the extra cost of needing to identify each peak in the spectrum to be modeled, which can be fairly time consuming. Of course obtaining reasonably good fits is also correspondingly easier when breaking the spectrum up into smaller sections as compared to fitting the spectrum as a whole. This is typically done by assuming a reasonable baseline model in terms of the unit cell parameters for each spectrum and then locating the diffraction lines that appear to best line up with the predicted peak positions. The major downside to this approach is that some amount of line misidentification is inevitable when observed peaks in the spectrum happen to occur in the expected region of diffraction angle space. In Chapter 3, we discuss in great detail the process of obtaining good estimates of unit cell parameters while avoiding the systematically biasing effects of misidentified peaks.

After obtaining estimates and uncertainties for the crystal dimensions from each spectrum (see Chapter 3), we carry out a careful study of the observed neon diffraction peak positions to determine the pressure inside the diamond anvil cell, account for deviatoric stress effects. Combining the estimates of pressure, temperature, and crystal volume, we estimate the equation of state parameters by fitting the derived unit cell dimensions for each spectrum. Rather than merely reporting the ‘best-fit’ parameter values as is quite common throughout the literature, we present the correlated uncertainties in our model parameters and demonstrate how important these correlations are crucial to interpreting the difference between different equations of state.

2.3.2 Determining P-T Conditions

Determination of the externally imposed environmental conditions, pressure and temperature, experienced during diamond anvil cell experiments is obviously one of the most important tasks for obtaining accurate and unbiased material properties measurements. Unfortunately, this task is significantly more complicated than one might naively guess. In this section, we discuss how we employ careful statistical methods together with detailed physical models to retrieve the temperature and pressure conditions for each measurement.

In situ temperatures during laser-heating are obtained from measurements of the thermal emission of the sample. The laser heating system at the Sector 13 IDD beamline of GSECARS is equipped with a set of mirrors that simultaneously focus a laser-heating spot on the sample, while carrying the thermally radiated light from the sample back to two independently calibrated spectrometers. The sample's thermal radiation spectrum is then fit at the beamline assuming a gray-body spectrum, enabling estimation of the temperature for both the upstream and downstream sides of the sample. Since the spectrometers on either side are independently calibrated, they provide two independent measurements of the same quantity, assuming small axial thermal gradients. One issue that should be addressed, however, is that the calibration process necessarily leads to some degree of systematic disagreement between the upstream and downstream measurements; as seen in Figures 2.3 and 2.4, the raw temperature measurements for the upstream and downstream sides are systematically shifted from the one-to-one line. This calibration error is readily corrected by plotting the histogram of the difference between the upstream and downstream measurements, which yields a distribution that is reasonably approximated as normal, as shown in

the insets in Figures 2.3 and 2.4. The mean and standard deviation of the best-fit normal distribution therefore give the systematic temperature bias and approximate temperature error for each dataset. We subsequently correct the upstream and downstream measurements toward more realistic values by shifting them each by half the typical measurement difference, thereby accounting for any systematic biases in the spectral temperature calibration.

Improving determination of in situ pressures inside the diamond anvil cell has received considerable attention over the past decade, as many workers have focused on increasing the internal consistency amongst the equations of state for numerous pressure markers as well as increasing the absolute accuracy by accounting for deviatoric stresses. Deviatoric stress, which describes the non-hydrostatic component of the stress tensor, arises primarily in response to the uniaxial nature of the compression imposed by the standard diamond anvil cell design, which squeezes the sample only along the axial beam-path direction. The pressure medium surrounding the sample, in this case neon, is therefore relied upon together with the rhenium gasket that forms the walls of the sample chamber, to translate this entirely uniaxial imposed stress into a nearly hydrostatic stress field appropriate to understanding the state of minerals in the deep mantle. The effectiveness of this transformation depends on strength of the pressure medium, hence the preference for ultra-soft media like neon to other harder materials like NaCl.

Based on the uniaxial compression design, it is generally assumed for diamond anvil cell experiments that the dominant part of the deviatoric stress tensor is an excess normal component along the axial direction relative to the two radial directions (*Singh, 1993*), and all other non-hydrostatic contributions are neglected as small high-order corrections. The

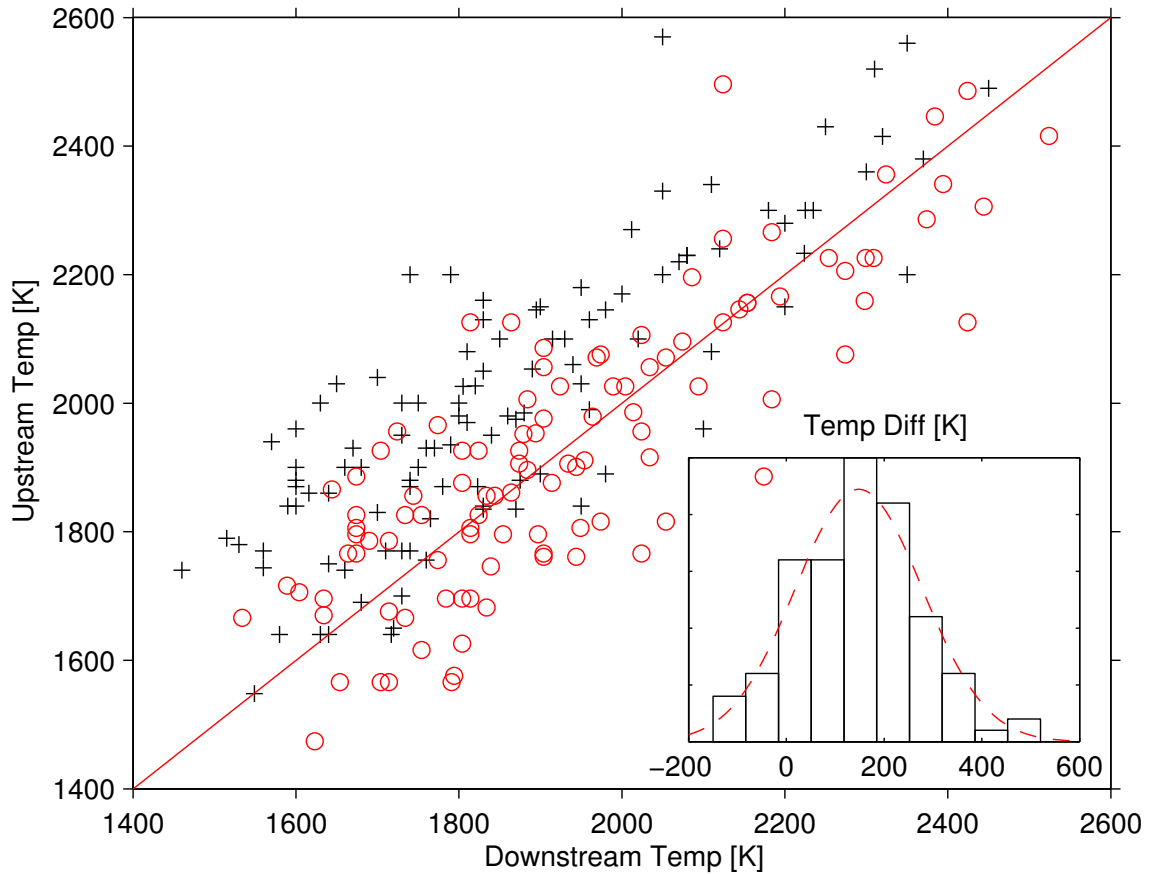


Figure 2.3: Measured spectral temperatures for Fe-bearing sample during laser heating. The main figure plots upstream and downstream measurements against each other together with the one-to-one line in red, with the raw measurements in black crosses and the corrected measurements in red circles. The inset figure shows the distribution of measured temperature differences, consistent with an absolute calibration difference of 150 K between the upstream and downstream spectrometers and a random measurement error for each spectrum of 130 K. This calibration difference is used to correct the raw measurements to their most likely value, shifting the red corrected points toward the one-to-one line.




Figure 2.4: Measured spectral temperatures for Mg-endmember sample during laser heating. The main figure plots upstream and downstream measurements against each other together with the one-to-one line in red, with the raw measurements in black crosses and the corrected measurements in red circles. The inset figure shows the distribution of measured temperature differences, consistent with an absolute calibration difference of 125 K between the upstream and downstream spectrometers and a random measurement error for each spectrum of 190 K. This calibration difference is used to correct the raw measurements to their most likely value, shifting the red corrected points toward the one-to-one line.

deviatoric stress experienced inside the sample chamber is bounded above by the yield strength of the pressure medium, which marks the point at which the material will begin to plastically deform and flow, generally allowing the sample chamber to thin along the axial direction, attempting to equalize the stresses along the axial direction (imposed by the diamond anvil surfaces) and along the radial direction (imposed by the edges of the gasket hole).

In this study, we rely primarily on the diffraction lines of the neon pressure medium to determine in situ pressures. While multiple pressure markers were placed inside the sample chamber, including NaCl and Au, the majority of the powder diffraction spectra show relatively weak signals from those markers, making their pressure determinations highly uncertain. In contrast, the diffraction lines from neon represent the most intense reflections in every spectrum.

Using the peak-fitting procedure described above and in Chapter 3, we are able to accurately retrieve peak positions for both the neon 111 and 200 lines in every diffraction spectrum, as well as the neon 220 line in most of the Mg-endmember sample spectra. For a cubic crystal under hydrostatic stress, the crystal plane spacing is given simply as

$$d_P(hkl)^{-1} = a_P^{-1} \sqrt{h^2 + k^2 + l^2} \quad (2.1)$$

where $d_P(hkl)$ is the d-spacing under purely hydrostatic stress and $a_P = V^{\frac{1}{3}}$ is the unit cell length. An estimate for the in situ pressure can then be determined by obtaining the apparent unit cell length for each diffraction line and converting the corresponding volume into a pressure using the well-determined equation of state for the pressure standard

in question. In our case, we use the neon equation of state reported in *Dewaele et al.* (2008), which provides a carefully constrained Debye-Grüneisen equation of state using high-cadence room-pressure compression data up to 200 GPa and precise resistive heating high-temperature measurements between 300K and 1000K. We find that the measured neon line positions imply apparent unit cell sizes, and corresponding pressure values, that are systematically shifted relative to one another, reflecting the presence of non-negligible deviatoric strain. By assuming that excess uniaxial stress dominates the deviatoric stress tensor, *Singh* (2009) provide the equations that describe the effect that deviatoric strain has on the measured lattice plane spacings from powder diffraction experiments:

$$d_m(hkl) = d_P(hkl)[1 + (1 - 3 \sin^2 \theta)Q(hkl)] \quad (2.2)$$

$$Q(hkl) = \frac{t}{6G^X(hkl)}$$

where $d_m(hkl)$ is the measured d-spacing of the hkl line and $Q(hkl)$ is a scaled ratio of the deviatoric stress t to the “diffraction” shear modulus, $G^X(hkl)$.

The diffraction shear modulus expresses the average value for a polycrystalline aggregate averaged over the subset of crystallites that contribute to the measured diffraction intensity. Since only crystallites in certain orientations satisfy the diffraction condition and contribute to the measured diffraction peaks, the diffraction shear modulus varies depending on the set of crystal planes, and is therefore a function of the hkl values. Following the typical approach, the average aggregate elasticity is assumed to be a weighted harmonic mean of the Voigt limit (derived assuming stress continuity along grain boundaries) and the Reuss limit (derived assuming strain continuity along grain boundaries). Accordingly, the

diffraction shear modulus is assumed to follow the empirical form:

$$[G^X(hkl)]^{-1} = \alpha[G_R^X(hkl)]^{-1} + (1 - \alpha)G_V^{-1} \quad (2.3)$$

where $G_R^X(hkl)$ is the Reuss limit for the diffraction shear modulus, G_V is the Voigt limit for the shear modulus, and α is the weighting factor typically assumed to fall between 0.5 (equal weighting of Voigt and Reuss bounds) and 1 (the Reuss limit). Note that $\alpha = 0.5$ is not the Voigt-Reuss-Hill average, which is simply the arithmetic mean of the Voigt and Reuss bounds, since Equation (2.3) uses a harmonic mean. Nevertheless, *Hill* (1952) points out that the choice of averaging scheme is arbitrary and should be chosen empirically. For diffraction measurements of shear moduli, the harmonic mean has been shown to effectively capture the elastic state of polycrystalline aggregates (*Singh*, 2009). The expressions of these limits for a cubic material are given by *Singh* (2009) as:

$$G_V = [C_{11} - C_{12} + 3C_{44}]/5$$

$$[G_R^X(hkl)]^{-1} = 2[S_{11} - S_{12} - 3S\Gamma(hkl)] \quad (2.4)$$

$$\Gamma(hkl) = (h^2k^2 + k^2l^2 + l^2h^2)/(h^2 + k^2 + l^2)^2$$

where $S \equiv S_{11} - S_{12} - S_{44}/2$ is the elastic anisotropy, S_{ij} are the elastic compliances, and C_{ij} are the elastic stiffnesses, which are just related by a matrix inversion ($\mathbf{C} = \mathbf{S}^{-1}$). When written in this form, it is clear how stress continuity implies simple additivity of stiffnesses, which have units of pressure, whereas strain continuity implies simple additivity of compliances.

In order to evaluate the diffraction shear modulus, we must be able to estimate the elastic moduli for each spectrum. A theoretical study of the elastic properties of noble gas solids by *Tsuchiya and Kawamura* (2002) showed that the normalized elastic constants, $C'_{ij} = C_{ij}/K$, of Ne, Ar, Kr, and Xe are all nearly constant over a wide pressure range between 0 and 200 GPa. *Shimizu et al.* (2005) performed Brillouin spectroscopy measurements of solid neon between 5 and 7 GPa and verified that the normalized elastic constants are approximately independent of pressure, though the measured values are systematically shifted from the theoretical calculations. As discussed by *Tsuchiya and Kawamura* (2002), this is likely the result of inadequately accounting for many-body interactions at finite temperatures. We therefore estimate the elastic properties of neon by assuming that the relative elastic constants are independent of pressure and given by the values measured by *Shimizu et al.* (2005): $C'_{11} = 1.13$, $C'_{12} = 0.93$, and $C'_{44} = 0.40$. It is important to note that while vibrational measurements are sensitive to the adiabatic elastic constants, due to the high frequencies involved, deviatoric strain measurements are sensitive to the isothermal elastic constants. Therefore, we can obtain the appropriate unnormalized elastic constants using $C_{ij} = C'_{ij}K_T$ and $K_T = (C_{11} + 2C_{12})/3$.

The approach adopted by most workers—including *Singh* (1993); *Takemura and Dewaele* (2008); *Dorfman et al.* (2012)—is to make a number of approximations in order to render Equations (2.2) to (2.4) in a convenient linear form. These involve dropping what are assumed to be small correction terms. We note, however, that such approximations are potentially undesirable, as well as unnecessary. Instead, we recognize that Bragg's law gives the direct relation between the diffraction angle and inverse d-spacing for each line,

$\sin \theta_m = 0.5\lambda d_m^{-1}$. Substituting this into Equation (2.2) and algebraic manipulation results in a simple cubic equation for the modeled inverse d-spacing, $d_{\mathcal{M}}^{-1}$, of each line:

$$0 = \frac{\lambda^2 t}{8G^X(hkl)} d_{\mathcal{M}}^{-3} + \left(1 + \frac{t}{6G^X(hkl)}\right) d_{\mathcal{M}}^{-1} - \frac{\sqrt{h^2 + k^2 + l^2}}{a_p} \quad (2.5)$$

recalling that $G^X(hkl)$ depends on the value of α and is given by Equations (2.3) and (2.4). Since this is just a simple cubic equation, it can be solved trivially using any standard polynomial root-finding method, yielding the modeled inverse d-spacing $d_{\mathcal{M}}^{-1}$ as a function of the unknown axial deviatoric stress t and unit cell length a_P .²

It is now just a simple matter of fitting the deviatoric stress and neon volumes to each set of measured lines. Since each spectrum contains only 2 or 3 neon lines, fitting these values without imposing further constraints leads to a rather large scatter amongst the determined deviatoric stress values, due to the large correlations between the fitted deviatoric stress and unit cell size. Figure 2.5 shows the unconstrained coupled determinations of deviatoric stress and pressure (evaluated using the determined unit cell size together with the equation of state from *Dewaele et al. (2008)*). While there is large apparent scatter, it is clear that the deviatoric stress is confined to low values of less than a few GPa. Further constraints can be imposed on the stress state retrieval by noting that the relative deviatoric stress fraction, stress over pressure, shows a clearly preferred value of around 0.5%, as indicated in the inset histogram in Figure 2.5. Assuming such a model where the axial deviatoric stress is linearly proportional to pressure is also reasonable given that deviatoric stress

²It should be noted that while there are 3 potential solutions to any cubic equation, only one of them gives a reasonable value for $d_{\mathcal{M}}^{-1}$, while the other two solutions are unphysical having either extremely large, negative, or imaginary values and are thus easily identified.

typically scales roughly with bulk modulus, which is itself roughly linear in pressure. We therefore employ this model as a prior constraint on the deviatoric stress and refit the set of neon line positions to obtain the best constant for the deviatoric stress fraction for each dataset, yielding values of 0.5% and 0.6% for the Mg-endmember and Fe-bearing samples, respectively, in agreement with the central peaks of the apparent deviatoric stress fraction histograms. Through this analysis, we obtain final determinations of the pressure and its error, accounting for deviatoric stress.

2.3.3 Extracting Crystal Volumes and Cell Dimensions

After obtaining line position estimates and errors from the powder diffraction spectra, the next step is to fit these positions with a crystal reflection model to estimate unit cell dimensions and volumes. We use the Pbnm crystal structure determination for Mg-perovskite from *Sugahara et al.* (2006) to model the line positions for our sample. From this structural refinement, we obtain a list of the hkl values for the reflections expected to be present in the perovskite spectrum³. For an orthorhombic crystal symmetry, it is then trivial to calculate line positions from the reflection list using the following equation:

$$\left(\frac{1}{d}\right)^2 = \left(\frac{h}{a}\right)^2 + \left(\frac{k}{b}\right)^2 + \left(\frac{l}{c}\right)^2 \quad (2.6)$$

where hkl are the reflection order parameters, abc are the crystal unit cell parameters, and d is the atomic plane spacing.

By applying the Bayesian mixture model presented in Chapter 3, we obtain reason-

³We also obtain a list of relative peak intensities at ambient pressure that can be used as a rough order-of-magnitude guide for the high-pressure data, which is helpful in line identification

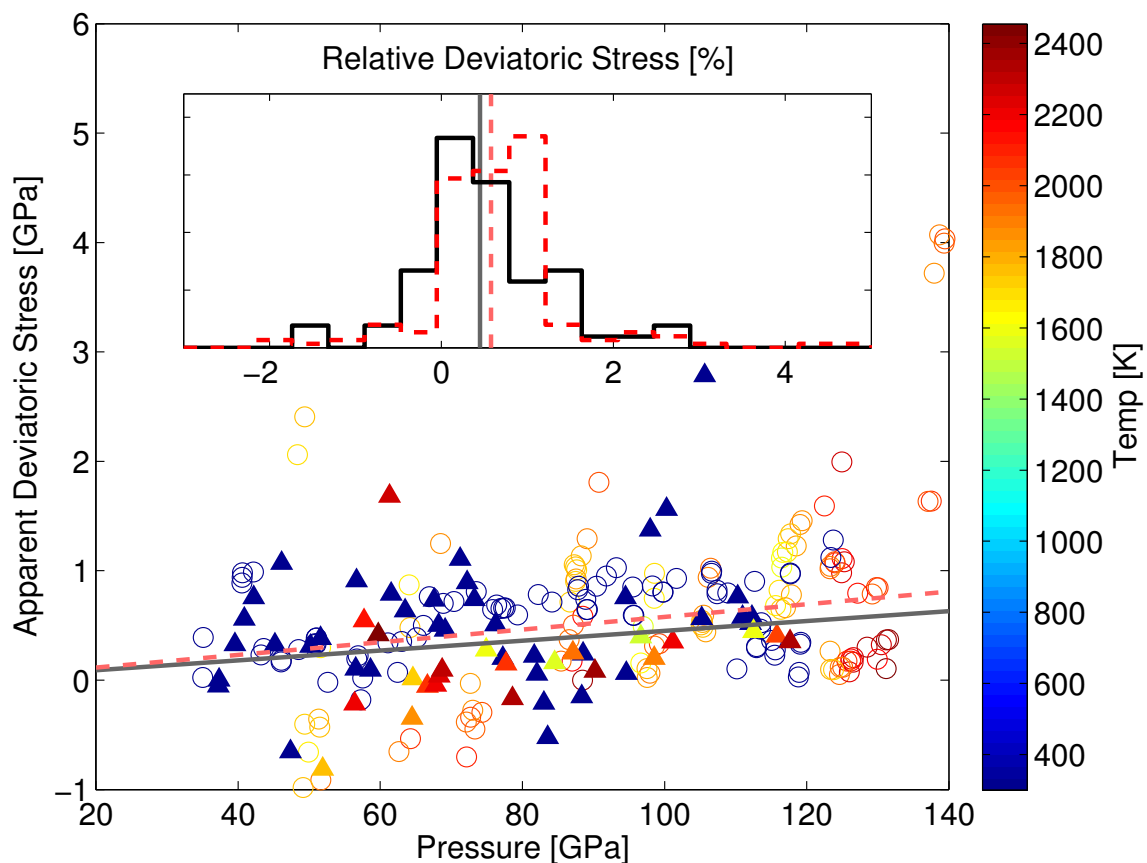


Figure 2.5: The apparent deviatoric stress is plotted against the neon pressure and color coded by temperature, obtained by analyzing the 111, 200, and 220 lines of neon. By fitting the two or three neon lines apparent in each spectrum, we obtain unconstrained simultaneous estimates of the axial deviatoric stress and the neon unit cell size, which are converted into pressure estimates using the neon equation of state from *Dewaele et al.* (2008). While there is large apparent scatter, the relative deviatoric stress, as a fraction of total pressure, shows a preferred value of roughly 0.5%, as seen in the inset histogram. The Mg-endmember sample is shown in solid black and filled symbols, while the 13% Fe sample is shown in dashed red and open symbols.

able estimates for the unit cell dimensions of the perovskite sample for each spectrum. As mentioned in Chapter 3, the choice of parameterization for any statistical analysis problem is important in order to ensure reasonable retrievals, since the parameterization is linked to the prior probabilities. The parameters we choose to fit are the volume of the unit cell $V = abc$, and the normalized axial ratios $(c/a)^* = c/\sqrt{2}a$ and $(b/a)^* = b/a$. Given this parameterization, it is simple to extract the pressure-volume curve as well as track the change in the axial ratios with compression, which are limited to the few percent level. Using the standard uncertainty estimation method, also described in 3, we can obtain reasonable estimates for the measurement errors on volume and the axial ratios.

2.3.4 Obtaining the Equation of State from P-V-T Measurements

Having determined the sample volume for a range of pressure conditions both at ambient and elevated temperatures, we are now prepared to evaluate the high-temperature equation of state by fitting the P-V-T data. This process is performed in two steps: first we obtain the cold portion of the equation of state using only ambient temperature measurements. We then fit the remaining laser-heated data to estimate the uncertainty bounds on the hot parameters. In the following section, we present the method we use to fit these data and obtain rigorous uncertainty bounds on the parameters.

In order to convert measured volumes sampled at a discrete set of pressures and temperatures into a form useful for interpolation and extrapolation, we must have a theory-based description of the equation of state. We use the simple thermal pressure approximation, which says that the free energy of a material can be divided into cold and hot compo-

nents, which lead to separate cold and hot contributions to the pressure since $P(V, T) = \partial F(V, T)/\partial V_T$. For convenience, the cold contribution to the pressure is often defined with reference to ambient temperature conditions, $T_0 = 300K$, rather than absolute zero. This requires that we subtract off the nonzero thermal contribution at the reference temperature, as indicated below:

$$P(V, T) = P_c(V) + P_{th}(V, T) - P_{th}(V, 300K) \quad (2.7)$$

where P_c is the “cold” portion of the total pressure defined at 300K, described below, and P_{th} is an expression for the thermal contribution, for which we use a Mie-Grüneisen-Debye formulation.

At ambient temperature conditions, and in the absence of phase transitions, most solid materials are well described by a Vinet equation of state (Vinet *et al.*, 1989). In particular, for data spanning a large pressure range, Cohen *et al.* (2000) showed that the Vinet equation of state yields more accurate extrapolation behavior over a large compression range as compared to the more commonly used third-order Birch-Murnaghan equation of state. The Vinet equation is derived from an approximate universal representation of the shape of the inter-particle energy well, which yields the following expression for the “cold” contribution to the pressure as a function of volume:

$$P_c(x) = 3K_0(1 - x)x^{-2} \exp[\nu(1 - x)] \quad (2.8)$$

where $x = (V/V_0)^{\frac{1}{3}}$ and $\nu = \frac{3}{2}(K'_0 - 1)$

where x is the average axial strain, V_0 is the zero-pressure volume, K_0 is the zero-pressure

bulk modulus, and K'_0 is the pressure derivative of the bulk modulus at zero pressure ($K' \equiv \partial K / \partial P$).

Somewhat more complex is the representation of the thermal component of pressure, which relies upon the simplified Debye crystal model to approximate the energetic contribution of thermal vibrations in a crystalline solid. The Mie-Grüneisen-Debye expression for the thermal pressure is given below:

$$P_{th}(V, T) = \frac{C_{vm}\gamma}{V}TD\left(\frac{\Theta}{T}\right) \quad (2.9)$$

where $C_{vm} = 3RN_{\text{cell}}$ is the Dulong-Petit high-temperature limit for the volumetric heat capacity of the unit cell, γ is the thermodynamic Grüneisen parameter, and Θ is the Debye temperature. The function $D(x)$ is the Debye integral, which together with the Debye temperature determines how rapidly the heat capacity approaches its high-temperature limit, and is defined as

$$D(x) = \frac{3}{x^3} \int_0^x \frac{y^3 dy}{e^y - 1} \quad (2.10)$$

where this integral must be evaluated numerically. What remains is an expression for the Grüneisen parameter, as well as a corresponding expression for the Debye temperature. A common expression for the Grüneisen parameter is in the form of a power law in volume, reflecting the fact that it is roughly independent of temperature, sensitive primarily to

compactness of the crystal lattice:

$$\begin{aligned}\gamma(V) &\equiv -\frac{d \ln \Theta}{d \ln V} = \gamma_0 (V/V_0)^q \\ \Theta(V) &= \Theta(V_0) \exp \left[- \int_{V_0}^V \frac{\gamma(\nu)}{\nu} d\nu \right] \\ &= \Theta_0 \exp \left[\frac{\gamma_0}{q} \left(1 - \left(\frac{V}{V_0} \right)^q \right) \right]\end{aligned}\tag{2.11}$$

where γ_0 is the Gruneisen parameter at zero pressure, Θ_0 is the Debye temperature at zero pressure, and q is the power law exponent.

Given the complete equation of state formulation described above, we can determine the parameter values by fitting the measured perovskite volumes for each spectrum. It should be noted that the errors associated with this fitting procedure appear on both dependent and independent variables P , T , and V . This situation only adds a minor complication to the fitting process, since we can use the “effective variance method” to estimate effective error bars appropriate for use in the standard least-squares method (*Orear*, 1982). This method is based on the assumptions that the model being fit to data behaves reasonably linearly with respect to the model parameters over the “high probability” region of parameter space (i.e., a first-order Taylor expansion of the model about the best fit is a good local approximation). The result of this method (which is derived using a Bayesian maximum likelihood estimation procedure) is a simple and intuitive expression for the effective variance, which can be used in place of the measurement error in a standard least-squares analysis:

$$\sigma_{\text{Tot}}^2 \approx \sigma_P^2 + \left(\frac{\partial P}{\partial T} \right)^2 \sigma_T^2 + \left(\frac{\partial P}{\partial V} \right)^2 \sigma_V^2\tag{2.12}$$

where $\partial P/\partial T$ and $\partial P/\partial V$ are simply model derivatives evaluated at the location of the best fit. For the effective variance method, we have the freedom to choose for the dependent variable whichever variable is most convenient, since all possible choices result in expressions that are identical to first order. We chose to consider P as the dependent variable, as it is much more convenient given that the EOS models are written with pressure as functions of volume and temperature, and these functions cannot be analytically inverted. The form of the effective variance is also quite attractive, as it is identical to the expression for the propagation of independent error sources, a rather intuitive way of viewing the combination of these different error contributions.

We combine this procedure for determining total effective errors, with the analysis method described in Chapter 3 to obtain the best-fit equation of state parameter values, along with the associated confidence regions. The best-fit values of the parameters, together with one-dimensional marginalized uncertainties are found in Table 2.1—while the precise values presented in that Table correspond to the global fit to all ambient and heated data, the hot data contain little extra information on the cold parameters and therefore alter their best-fit values by less than the $1\text{-}\sigma$ uncertainties. The cold compression data are plotted together with the 68% confidence interval on the equations of state in Figure 2.6. From that figure, it is clear from the narrower confidence bounds on the compression curve that the Fe-bearing data has reasonably higher accuracy than the Fe-free data. Nevertheless, these two compression curves are extremely similar to one another, with the Fe-bearing curve showing only slightly greater compressibility, which is reflected in the slightly smaller values of both K_0 and K'_0 . It is important not to overstress this difference, as the data clouds

Table 2.1: Vinet and Mie-Grüneisen-Debye Equation of State Parameters

X_{Fe}	0%	13%
V_0	162.51(19)	163.23(18)
K_0	259.0(56)	254.0(38)
K'_0	4.28(17)	4.060(85)
γ_0	3.19(25)	1.579(93)
q	2.16(43)	0.63(27)

We fix the zero-pressure Debye temperature to the value from *Anderson* (1998), $\Theta_0 = 1100K$, since our high-temperature data all significantly exceed this value and are thus largely insensitive to it.

strongly overlap one another, reflecting only small changes to the “cold” equation of state parameters. Discussion of the comparison with previous equation of state measurements, as shown in the inset figure, is deferred to the Discussion section.

The thermal equation of state parameters are obtained by fitting the laser-heated measurements including the thermal pressure contribution, given by the Mie-Grüneisen formulation in Equation 2.9. This method introduces three additional “hot” EOS parameters, Θ_0 , γ_0 , and q . However, the sample behavior is insensitive to the zero-pressure Debye temperature, Θ_0 , for very high temperatures well above its value, where the heat capacity is well approximated by the Dulong-Petit high-temperature limit. Since all laser-heated measurements exceed its nominal value of $\Theta_0 = 1100K$ (*Anderson*, 1998), we fix it to this value throughout this investigation. Note that this is fundamentally different from fixing data sensitive parameters like K'_0 , since they have a clear effect on the equation of state over the region of interest, whereas Θ_0 has almost no effect over the temperature ranges of either the data or the lower mantle, and therefore negligible correlation with the other equation of state parameters.

The proper method for incorporating the results from the cold parameter fit uses the




Figure 2.6: The cold compression curves for 0% and 13% Fe-bearing perovskite samples in black and red, respectively. The pressure and volume error bars reflect the $1\text{-}\sigma$ uncertainties as determined from the measurement errors. The pale shaded regions indicate the 68% confidence bounds on the equation of state, showing the somewhat higher confidence in the Fe-bearing equation of state. The inset figure compares the highly correlated 68% confidence bounds on K_0 and K'_0 with past investigations. The 0%, 9%, and 13-15% Fe-contents are shown in black, orange, and red. The uncertainties from this work are shown in heavy solid lines, (Lundin *et al.*, 2008) fits are shown in dashed lines, and (Fiquet *et al.*, 2000) in dot-dashed lines. (It should be noted that this comparison mixes third-order Birch-Murnaghan and Vinet equations of state. Also the confidence regions for previous work are estimated assuming similar parameter correlations across studies).

Table 2.2: Equation of State Parameter Correlations

0%					
	V_0	K_0	K'_0	γ_0	q
V_0	+1.00	-0.72	+0.43	-0.02	-0.07
K_0	-0.72	+1.00	-0.93	-0.24	-0.14
K'_0	+0.43	-0.93	+1.00	+0.34	+0.26
γ_0	-0.02	-0.24	+0.34	+1.00	+0.96
q	-0.07	-0.14	+0.26	+0.96	+1.00
13%					
	V_0	K_0	K'_0	γ_0	q
V_0	+1.00	-0.92	+0.75	+0.06	+0.06
K_0	-0.92	+1.00	-0.94	-0.27	-0.27
K'_0	+0.75	-0.94	+1.00	+0.42	+0.45
γ_0	+0.06	-0.27	+0.42	+1.00	+0.97
q	+0.06	-0.27	+0.45	+0.97	+1.00

For those unfamiliar, the correlation matrix is just a convenient scaled form of the covariance matrix which enables easy interpretation, where Σ is the covariance matrix and the correlation coefficient is simply $\rho_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\Sigma_{jj}}}$

Bayesian concept of a prior, which is discussed thoroughly in Chapter 3. The prior simply expresses the probability distribution for the parameters prior to any particular step of the analysis. Thus, we can use the output probabilities from fitting the ambient temperature data as the input to the heated data modeling step. Note that many investigators will fix the cold parameters to their best-fit values when fitting the hot parameters. This is both unnecessary and undesirable, since it underestimates the uncertainties in the hot parameters by neglecting potential covariances between the cold and hot parameters. Instead, we fit the heated data points with a full thermal model, recognizing that nearly all of the information about the cold parameters comes from the cold parameter posterior (converted into a prior for the hot data analysis). As far as the prior on the hot data is concerned, we choose to use uninformed flat priors since we have significantly more simultaneous high-pressure

high-temperature data than any previous studies *Fiquet et al.* (2000), and all previous studies chose to fix one or more significant parameters, biasing their uncertainty estimates to unrealistically small values.

The resulting best-fit parameters, with their marginal $1-\sigma$ uncertainties, are reported in Table 2.1. To summarize all the information for the posterior, we also report the correlation matrix for both samples in Table 2.2. The residual volumes due to thermal expansion are shown for each sample in Figure 2.7, where the best-fit cold equation of state has been subtracted off for the purpose of visual inspection, leaving only the residual thermal expansion signal. Data are shown with effective error bars in volume, allowing for simple visual inspection, and are determined using the effective variance method with volume as the dependent variable. The volumetric measurement error is determined using the Bayesian method discussed in Ch. 3 and the pressure and temperature contributions to the effective error are derived in Section 2.3.2. Since both panels in Figure 2.7 are shown with the same axis limits, it is plainly visible that the effect of iron is to dramatically reduce the thermal expansion at high pressures. This observation is directly reflected in the best-fit estimates for the thermal parameters γ_0 and q , which relate to the overall scale for the thermal expansion and its compression (or Volume) dependence. Comparing the figure with the retrieved values, we can see that the Fe-free sample has both a greater average thermal expansion (larger γ_0), yielding wider spacing between isotherms, and a greater compression dependence (larger q), resulting in steeper isotherm slopes. We can therefore see that while the addition of Fe has only a minor effect on the cold equation of state parameters, it has a huge impact on the thermal properties of silicate perovskite.

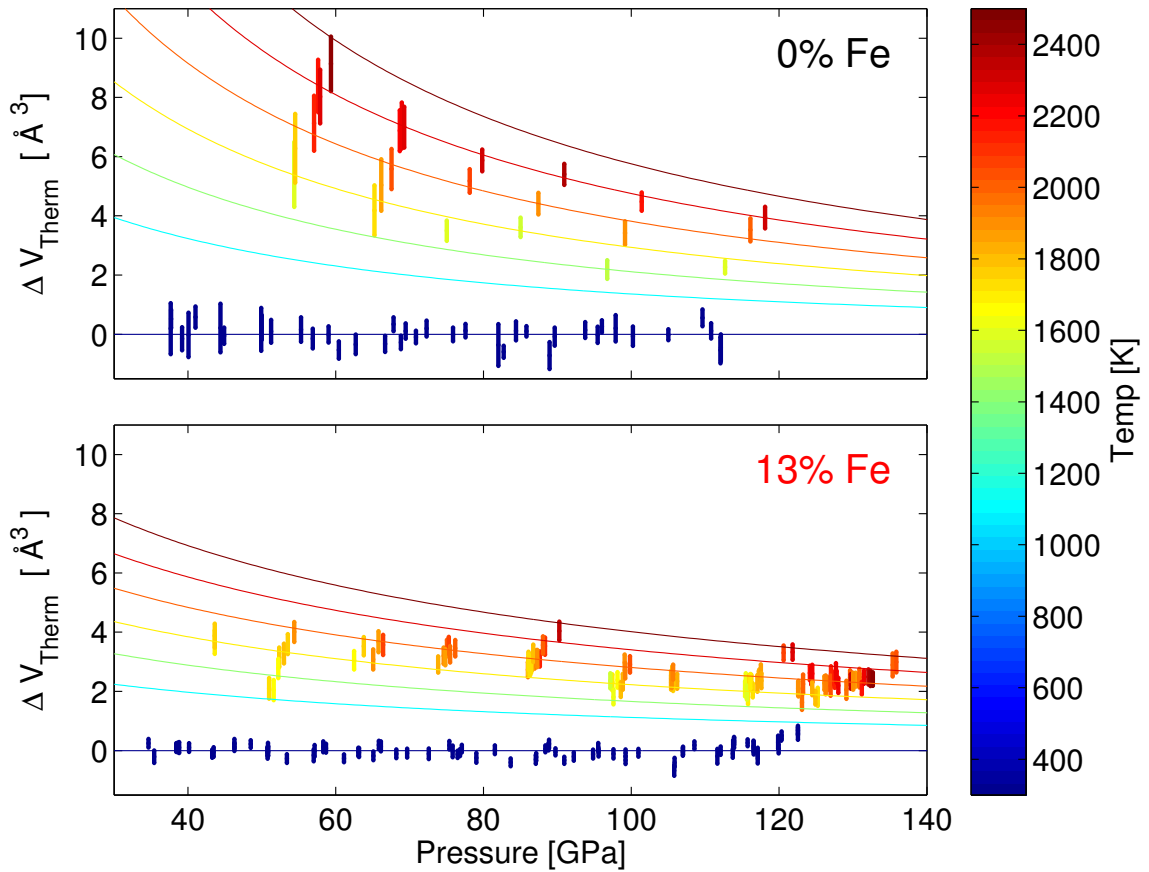


Figure 2.7: Excess volumes due to thermal expansion are shown for the 0% and 13% Fe-bearing samples color coded by temperature, together with the best-fit isotherms. The strong decrease in the apparent thermal expansion due to the addition of iron is clearly visible in the noticeably smaller envelope for the elevated temperature isotherms.

2.4 Discussion

Given the posteriors estimated for both Fe-free and 13% Fe-bearing perovskite, we are now ready to compare our results with previous investigations. Most importantly, we evaluate how apparent differences across studies compare with the correlated uncertainties associated with each study. After that, we then use our posterior estimates to determine the variation of important physical quantities like density, bulk modulus, and thermal expansion for a range of pressures and temperatures. Finally, we incorporate the equation of state determinations into an ideal mixing model in order to evaluate the plausibility of different perovskite-rich compositional models for deep mantle structures.

2.4.1 Equation of State Comparison and Uncertainties

Using the general Bayesian approach of estimating covariance matrices from the curvature of the posterior in log-space, we obtain the 1D uncertainty bounds for each parameter together with the correlation matrix presented in Tables 2.1 and 2.2. These results can be directly visualized using a “stair-step” plot, which is presented in Chapter 3, Figure 3.5, which shows the correlated uncertainty bounds on each pair of parameters by plotting every parameter against every other parameter. The stair-step plot is just a way of visualizing the information conveyed by the best-fit and covariance matrix. Looking at the elements of the correlation matrix in Table 2.2, we can get much of the same information, noting in particular the highest correlation of any two parameters exists between the thermal parameters γ_0 and q , reflecting how difficult it is to independently constrain their values. In spite of that difficulty, the large quantity of thermal data for the Fe-bearing sample has very

effectively constrained the thermal parameters to a correlated but small region of parameter space, as plainly shown in the lower-right-hand corner of Figure 3.5;

These results should then be compared directly to previous investigations of the perovskite equation of state. Unfortunately, this task is made rather difficult for two reasons. Firstly, none of the previous investigations report the full covariance matrix (or equivalent correlation matrix) for their equation of state parameters. Second, and more important, is the fact that all previous investigations fixed a number of important equation of state parameters in order to improve the apparent constraints on the remaining parameters. The issues with this common but incorrect practices are discussed thoroughly in Chapter 3. Therefore, we must settle for the rough comparison shown in the inset of Figure 2.6 where the correlated 68% confidence regions on the bulk modulus K_0 and its pressure derivative K'_0 are shown for this study in solid lines as well as *Fiquet et al.* (2000) in dot-dashed lines and *Dewaele et al.* (2008) in dashed lines. All of the perovskite samples are color coded according to composition as before, with 0% Fe in black, 9% Fe in orange, and 13-15% Fe in red. As the figure indicates, all studies are consistent with bulk moduli near 255 GPa for a wide range of iron contents up to 15%. To give the previous studies reasonable correlations between the parameters (they do not report covariance matrices), we simply apply the average correlation coefficient from this study. Unfortunately, the effect of previous studies fixing parameters remains, where the ellipse sizes are likely significantly smaller than they should be. As a result of the large correlated uncertainties in previous studies, trends of K'_0 with iron content are obscured. In this study, we do find a discernible decrease in K'_0 with increasing iron, but the magnitude of the change is rather small, as the cold compression

data still largely overlap one another.

While considerable attention is always given to comparing EOS parameters, in reality, we are most interested in the behavior at mantle-relevant pressures, rather than the ambient-pressure properties expressed by the equation of state parameters. For this reason, we propagate our posterior uncertainties in order to predict the behavior at lower mantle conditions, by drawing randomly from the posterior and calculating properties at elevated pressure states. In Figure 2.8, we plot the 68% correlated confidence regions for this study, along with rough estimates of the same for previous studies. In that figure, we can see how previous studies have not uncovered any strong evidence for systematics that may exist for perovskite as a function of iron content. If anything, they find hints that the addition of iron causes the high pressure volume of perovskite to increase (shifting the curve upward), but has no discernible effect on bulk modulus—in contrast, we find the opposite, that adding iron can slightly reduce high pressure volumes, though volumes at ambient pressure increase. It should also be noted, however, that true confidence ellipses are larger than depicted for these other studies, implying that their true scatter could be larger than appears.

While these results were obtained using reasonable priors on the zero-pressure volume data, it is important to realize that they inherently rely on the assumption that metastable volume measurements of perovskite are relevant to its behavior within its stability region. It has not been well established how similar the behavior of perovskite is inside and outside its stability range, but there are some indications that the axial ratios may undergo anomalous deviations near the edge of stability. It is also common for minerals to undergo softening as

they approach a stability limit, and therefore it is reasonable to expect that metastable measurements provide, at best, only a rough idea of the compression behavior. For this reason, we also refit the non-laser-heated data in order to test how the apparent core-mantle boundary properties of perovskite depend on the usage of metastable data. In the thick dashed lines in Figure 2.8, we plot the 68% confidence regions resulting from uniformed priors on V_0 , thereby removing all influence of metastable data. Under this alternate hypothesis, we can clearly see a 5% to 10% increase in the apparent bulk modulus of perovskite at the core-mantle boundary, though there is little change in the volume. Additionally, there is a large inflation of the uncertainties, resulting from the reduced constraint on the equation of state model, resulting in a reduction of the CMB differences between iron-bearing and iron-free samples allowing the confidence regions to partially overlap. This result matches our intuition from Figure 2.6, in which the two data clouds are seen to largely overlap one another. Unlike for the cold parameters discussed here, however, the usage or absence of the metastable data has little effect on the thermal equation of state parameters.

As investigated by other authors, we also compare the evolution of the crystal axial ratios with increasing pressure. Past work (*Lundin et al.*, 2008; *Dorfman et al.*, 2013) have found that the addition of iron causes a noticeable change in the normalized axial ratios, which are generally observed to grow roughly linearly with pressure. The normalized unit cell parameters derived by *Andrault et al.* (2007) are: $a^* = a(V/\sqrt{2})^{-1/3}$, $b^* = b(V/\sqrt{2})^{-1/3}$, and $c^* = c(2V)^{-1/3}$. They are constructed to yield values of exactly one for an ideal cubic perovskite, and deviate progressively with increasing distortion. As discussed in Chapter 3, the use of average crystal length scale in place of pressure for plots

like these greatly improves them, by enabling the simultaneous consideration of a range of temperatures, as well as providing a space that is often more closely linear.

In the upper panel of Figure 2.9, we show the normalized axial ratios measured at ambient temperature for both perovskite samples over the entire pressure range. The Fe-bearing sample is shown in solid symbols and the Fe-free sample is shown in open symbols. Furthermore, these axial ratios can be converted into an estimate of the octahedral tilt angle of the corner-sharing silica octahedra within the perovskite structure. It is the progressive tilting of these octahedra with increasing pressure that accommodates the gradual distortion of the unit cell induced by compression. From *O’Keeffe et al.* (1979), we can calculate the octahedral tilt angle as a function of the unit cell parameters:

$$\phi = \cos^{-1} \left(\frac{\sqrt{2}a^2}{cb} \right) = \cos^{-1} \left(\frac{1}{(c/a)^*(b/a)^*} \right) \quad (2.13)$$

where $(c/a)^*$ and $(b/a)^*$ are the normalized axial ratios. In the lower panel of Figure 2.9, we show the linear evolution of the octahedral tilt angle with compression for the two samples. For both the normalized axial ratios and the octahedral tilt angles, there is such a strong similarity between the two samples that we find the past reported resolvable differences as a function of iron content somewhat doubtful, at least for our samples.

2.4.2 Physical Property Profiles

The most scientifically interesting questions that can be addressed using the perovskite equation of state depend on the pressure profiles of various physical properties like thermal expansion and bulk modulus. In order to evaluate these profiles, we must obtain the expres-




Figure 2.8: Comparison of 68% confidence bounds on local properties at core-mantle boundary pressure conditions are shown for Fe-bearing and Fe-free perovskite. All lines are color coded with Fe-free samples in black, intermediate $\sim 9\%$ Fe samples in orange, and $\sim 13\text{--}15\%$ Fe-rich samples in red. The thick lines show the results from this study: the solid lines display the results when metastable ambient-pressure measurements are used to constrain V_0 , whereas the dashed lines represent the results assuming the metastable behavior of perovskite is a poor predictor high-pressure properties. The thin dashed lines show the three compositions studied by *Lundin et al.* (2008) and the thin dash-dotted line shows the results from *Fiquet et al.* (1998).




Figure 2.9: Normalized axial ratios and octahedral tilt angles shown for the 0% and 13% Fe-bearing samples. The 13% Fe-bearing sample is plotted in circles and the Fe-free sample is plotted in triangles.

sions for the relevant thermodynamic parameters. Below we derive the expressions for the isothermal bulk modulus by taking derivatives of the pressure, with both a ‘‘cold’’ contribution from the Vinet equation of state and a ‘‘hot’’ contribution from the Mie-Grüneisen-Debye thermal correction:

$$\begin{aligned}
 K_T(V, T) &= K_c(V) + K_{th}(V, T) - K_{th}(V, 300) \\
 K_c(V) &= -K_0 x^{-2} \{x[1 - \nu(1 - x)] - 2\} \exp[\nu(1 - x)] \\
 K_{th}(V, T) &= \frac{C_{vm} \gamma T}{V} \left[\frac{3\gamma \Theta/T}{e^{\Theta/T} - 1} - D \left(\frac{\Theta}{T} \right) (q - 1 + 3\gamma) \right]
 \end{aligned} \tag{2.14}$$

where $\nu = 3/2(K'_0 - 1)$ and $x = (V/V_0)^{1/3}$ is the average axial strain, and recalling that the Debye temperature Θ and the Grüneisen parameter γ are given by Equation (2.11). Combined with the expression for the heat capacity of a Debye solid, a thermodynamic identity can be used to obtain the thermal expansion parameter at any desired volume and temperature:

$$\begin{aligned}
 C_V(V, T) &= C_{vm} \left[4D \left(\frac{\Theta}{T} \right) - \frac{3\Theta/T}{e^{\Theta/T} - 1} \right] \\
 \alpha(V, T) &= \frac{\gamma(V) C_V(V, T)}{V K_T(V, T)}
 \end{aligned} \tag{2.15}$$

Using these expressions, we determine the isothermal profiles of thermal expansion, bulk modulus and volume for the two perovskite samples. In Figure 2.10, we show the propagated 68% confidence intervals for these material properties evaluated at 300K, shown in the light shaded regions, and at 3000K, shown in the dark shaded regions. From these profiles, it is clear that the largest difference between the two perovskite compositions is

the sizable affect that the addition of iron has on the thermal expansion. By adding 13% Fe to perovskite, we see a roughly 50% drop in the thermal expansion at the core-mantle boundary, as seen in the top panel. Additionally, we can see that because of this drastic difference in thermal behavior, the bulk modulus and volume see a much reduced change upon increasing temperature for the iron-bearing sample as compared to the Mg-endmember.

It should be noted that we obtain a significantly different thermal behavior for the Fe-free endmember sample than *Fiquet et al.* (2000). A previous investigation by *Mosenfelder et al.* (2009), however, pointed out the none of the previous work on the thermal expansion of MgSiO_3 perovskite is consistent to better than a factor of about two at mantle conditions, and therefore we do not find this level of disagreement particularly surprising. Furthermore, the difference in thermal expansion between the two samples is plainly visible to the eye in the raw data shown in Figure 2.7, further bolstering the robustness of this result.

2.4.3 Perovskite-Dominated Chemical Piles

With a well characterized equation of state model in hand for both Fe-free and 13% Fe-bearing perovskite, we are able to quantitatively assess the behavior of perovskite-dominated lower mantle structures. In this section, we develop a thermodynamic model of Mg-Fe perovskites at arbitrary compositions using an ideal mixing model for combining the equation of state results from our two perovskite samples. We then apply this simple perovskite model to determine the physical properties of perovskite at core-mantle boundary pressure-temperature conditions. Comparison with seismic estimates of bulk mantle physical properties allows us to compare the relative merits of different possible explanations




Figure 2.10: Shaded regions show 68% confidence intervals for 300K (light shading) and 3000K (dark shading) isothermal properties of Fe-free perovskite in black and 13% Fe-containing perovskite in red. The effect of additional iron is rather small at room temperatures, with nearly identical volumes and thermal expansions and only slight differences in bulk modulus. The distinctions becomes quite pronounced at elevated temperatures, however, due to the significant difference in thermal expansion. The low-pressure region for the Fe-free sample is rather poorly constrained by the data, but this appears to have little effect on the high-pressure behavior, which retains narrow uncertainty bounds.

for seismically observed core-mantle boundary structures.

Large Low Shear Velocity Provinces (LLSVs) represent one of the most important first-order features of the lower mantle and may represent the oldest structures in the Earth. It has been proposed that the seismically distinct LLSVs are likely to represent both chemically and thermally distinct structures on the core-mantle boundary (*Tackley, 2011*). The two LLSVs, beneath Africa and the Pacific, present a particular challenge for explanation as they appear to have sharp and often steep-walled boundaries and stand roughly 1000 km high off the CMB. Residing at the base of a convecting system, it is a challenge to understand how such structures might remain for geologic time without mixing away through the process of entrainment. The general interpretation of the sharp seismic reflector at the top of the structure is that these features cannot be explained solely by thermal anomalies (*Tackley, 2011*), which would tend to have more diffuse edges, though there is still some disagreement in the literature on this point (*Davies et al., 2012*). The viewpoint of a chemically distinct pile is further bolstered by an apparent anticorrelation between shear-wave velocity anomalies and both bulk sound velocity and density anomalies within the structures relative to bulk mantle, contrary to the general trends of most heated material (*Ishii and Tromp, 1999*).

Under the umbrella of chemically distinct explanations for these lower mantle structures, there are two broad endmember theories that account for their large topographic relief relative to the CMB, as discussed in *Tan and Gurnis (2007)*. The domes can either be explained as passive chemically dense piles that are dynamically propped up by external convective stresses or as free-standing and internally convecting metastable piles. The pas-

sive pile explanation is the more intuitive of the two, in which the piles reside at the base of the mantle since their constituent material is denser than bulk mantle material. Fighting their tendency to spread out and pool as thin shallow layers on the CMB, some external force must be invoked to sweep them into domed piles, which is often attributed to cold dense plates descending to the CMB and pinching the sides of these structures in order to dynamically prop them up. This story of course assumes that a sufficient flux and geometry of down-going plates is available at the core-mantle boundary in order to achieve the requisite force to counteract the negative chemical buoyancy of the piles. Under the competing scenario, no external force is required, but rather the chemically distinct piles are made of a material that is less dense than the surrounding mantle at the base, but experiences a density crossover, or height of neutral buoyancy, near the top of the pile about 1000 km above the CMB. Under this explanation, the pile is undergoing slow internal convection, where light material at the base rises to the neutral buoyancy point and then cools and falls back to the CMB, thereby enabling the pile to prop itself up without the help of external stresses. To achieve this behavior, the pile material must be less dense than bulk mantle (on account of elevated temperatures that result from the thermal boundary layer separating the pile from the rest of the mantle) as well as being very stiff compared to the average mantle, enabling a density crossover point. These properties are a function of both the unknown composition and poorly constrained excess temperature of the pile material. Nevertheless, they are a direct consequence of the equation of state properties of the material, and can therefore be directly probed given a thermodynamic model.

While we know very little about the detailed composition of these distinct structures

in the deep Earth, one possible model for such structures is in terms of an iron-enriched and perovskite-dominated pile. As perovskite is the primary phase in the lower mantle and given the sizes of the structures, it is also likely the primary component of the LLSVPs. With the equation of state information obtained in this study, we are perfectly positioned to investigate such a possibility. In order to probe a range of iron compositions, we construct an ideal mixing model of the 0% and 13% Fe composition perovskite samples. The ideal mixing model is based simply on the assumption that at constant environmental conditions, temperature and pressure, the behavior of a mixture of two components is given simply assuming that volumes and internal energies mix in proportion to the relative amount of each component, and also that the entropies are also additive together with an additional term that describes the ideal entropy of mixing. Under these assumptions, it is straightforward to derive the combination of thermodynamic parameters that can be simply averaged together in proportion to the quantity of each component. Using this method, we are able to take known or assumed equations of state for two different compositions and mix them under the assumptions of ideality, to obtain the approximate physical properties of a mixture of any desired composition. Since it relies on the assumption of linear additivity, there are clearly limits to this method. We therefore restrict ourselves to examining compositions for the mantle structures that are less than twice the difference between the two mixing members, and thus we consider composition up to 25 mole % Fe.

The results of this mixture model comparison are given in Figure 2.11. In this exercise, use the above described mixture model to calculate the material properties of the mixture at the core-mantle boundary. Imagining that both the iron composition and the excess tem-

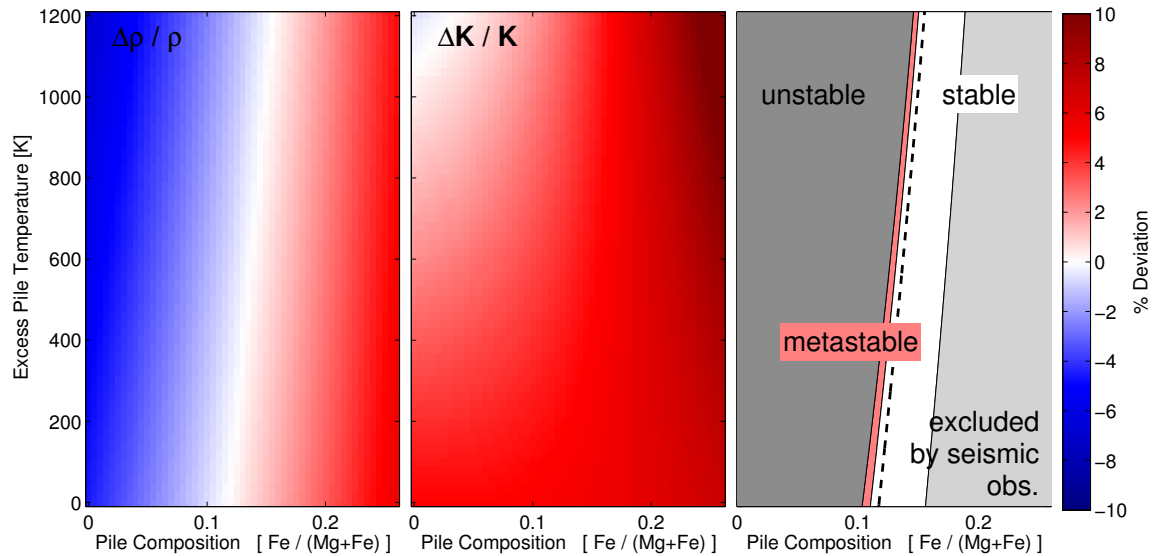


Figure 2.11: The ideal mixture model is used to sample the behavior of perovskite under a range of temperatures and compositions, corresponding to the potential conditions in a perovskite dominated LLSVP. The densities and bulk moduli are compared to the core-mantle boundary values from PREM. This yields the percent contrast maps between the perovskite chemical pile and the bulk mantle shown in the left-hand and central panels. The differences in the local properties are combined in a local linear expansion to estimate the height of the density crossover, which is transformed into the pile stability map shown in the right-hand panel. Densities greater than PREM yield stable passive piles, but excess densities greater than about 1.5% are generally excluded as they would be visible in seismic observations. A small region of space allows for metastable domes between 600 and 1200 km tall, and everything taller is regarded as initially in conflict with observations, and for low enough densities, entirely unstable at the core-mantle boundary.

perature of the LLSVP are unknown, we search a grid over those variables to determine the possible densities and bulk moduli of the mixture. In the left-hand and central panels of the figure, we show the percentage contrast between the mixture model (as a function of composition and temperature) and the bulk mantle as represented by PREM (*Dziewonski and Anderson, 1981*). Using a simple linear expansion of the density and bulk modulus difference at the core-mantle boundary, we can obtain rough estimates for the height of the density crossover. Comparing those predicted heights to a generous range of values corresponding to reasonable estimations for the observed heights of the actual LLSVPs, lying somewhere between 600 and 1200 km above the CMB, we can map out the parameter space by noting the relative stability of the structure as a function of composition and temperature. This result is shown in the right-hand panel of Figure 2.11. In the figure, the dark dashed line represents the equal density line, where to the right of that line, the chemical pile material is denser than ambient mantle and thus stable. At very large iron compositions, however, the excess density becomes so high that it should be plainly visible in seismic observations, especially normal mode inversions like those done by *Ishii and Tromp (1999)*. We therefore restrict the allowable compositions to those with excess density contrasts of less than 1.5%, which is a somewhat conservative estimate. Immediately to the left of the equal density line is material that is slightly buoyant and thus the red shaded sliver represents material with a height of neutral buoyancy lying within the reasonable limits of 600 and 1200 km above the CMB, representing solutions consistent with a metastable dome explanation for the LLSVPs. For lower iron compositions, the pile material becomes even less dense until it is entirely unstable escaping from the base of the mantle. As is ob-

vious from the figure, there is only a tiny sliver of allowable phase space that corresponds to the delicate balance required by the metastable dome hypothesis. In contrast, passive piles are extremely insensitive to composition and temperature. We can thus assess the relative plausibility of the passive pile and metastable dome explanations by calculating the fraction of allowable parameter space occupied by material behaviors consistent with the two theories—this is given simply by the relative area of the red shaded region as compared to the white region. This fraction is just below 10%, indicating that while metastable domes are possibly consistent with our current understanding of the perovskite equation of state, they imply very constraints on the temperature-dependent composition of the LLSVP material and are thus not very likely.

2.5 Conclusion

A series of high-pressure and -temperature laser-heated diamond anvil cell experiments were performed in a nearly hydrostatic neon pressure medium to determine the equations of state of 0% and 13% Fe-bearing Mg-perovskite. Using neon as the primary pressure indicator, we develop a new statistical technique for estimating deviatoric stresses in the diamond anvil cell. These stresses are confined throughout the experiments to values of about half a percent of the pressure. Using a series of Bayesian data analysis techniques that are detailed in Chapter 3, we obtain estimates of the unit cell volumes which are then fit with thermal equations of state.

The equation of state parameters found for the Fe-bearing and Fe-free samples are generally consistent with previous measurements perovskite measurements, the Fe-free sample

exhibits slightly elevated volumes at high pressure as compared to previous values. The derived equation of state parameters for the Fe-bearing sample are significantly more precise than any previous perovskite measurements, providing tight constraints on both the ambient-temperature and high-temperature properties. The most dramatic results are the large thermal expansions measured for the Mg-endmember sample, corresponding to the high values determined for the thermal equation of state properties.

Finally, we incorporate the pair of equation of state determinations into an ideal mixture model enabling the calculation of approximate equation of state properties for any arbitrary iron composition. Using this mixture model, we examine the range of plausible values in temperature-composition space to explain the presence of LLSVPs. Using plausibility arguments, we show that metastable perovskite domes are possible, given our knowledge of the equation of state, but represent a very precise balance between the iron content and temperature of the chemical piles on the CMB, and are not therefore particularly likely. We instead find the passive pile explanation more compelling, as it does not require that the Earth occupy a special narrow region of parameter space merely by chance.

— Appendices —

2.A Data Reduction Pipeline

Our data pipeline, written in MATLAB, transforms raw powder diffraction images into properly normalized one-dimensional spectra. This data processing procedure involves first obtaining the observation geometry from a diffraction image of a known calibration

standard. If necessary, the data image is corrected for saturation effects due to CCD pixel blooming. Intensity weighting factors, determined by the geometry, are applied to the data transforming them to idealized detector measurements. These transformed data are integrated producing a 1-D spectrum with error bars resulting from counting statistics. Finally, the background signal is estimated and subtracted from each spectrum, preparing it for convenient spectrum comparison or fitting.

We determine the detector geometry using a method similar to that reported in (*Hinrichsen, 2006*). After determining an approximate beam-center location, the diffraction data for a known calibration standard—such as CeO_2 or LaB_6 —is sliced radially using bi-cubic interpolation to obtain a set of 1-D radial spectra. Peak fitting with pseudo-Voigt profiles is then used to extract the 2θ locations of every line in each radial slice. These calibration line positions (with associated uncertainties) are then fit by varying the experimental geometry parameters controlling the orientation and position of the detector: detector distance, beam-center location, and detector tilt and rotation. The geometric calibration parameters are chosen as the values with the maximum likelihood given the measured line positions using the standard least-squares method (as is appropriate to simple data fitting with Gaussian uncertainties and in the absence of strong prior information).

For diffraction images containing both strong and weak x-ray scatterers, it is often impossible to obtain high quality spectra that do not suffer from over-exposure in certain regions of the image. Due to the basic properties of CCD detectors, exposure of a pixel beyond its full-well depth causes 'blooming' in the final image, where electrons spill over into neighboring pixel wells causing full-intensity streaks to emanate from the excessively

bright points in the image. The typical approach to this problem is to adjust exposure times to limit its occurrence. Unfortunately, in many cases this method is either ineffective (e.g., in the presence of very weak scatterers) or impractical (e.g., during high-temperature measurements). It would be highly advantageous to be able to use these data while minimizing the impact of the erroneous intensities caused by blooming—such an approach is made possible by the conservation of electrons within the affected region. Since the total number of electrons, equal to the number of photons registered by the CCD, remains constant as electrons spill into neighboring pixel bins, simple summation in an overexposed region will give an approximate total intensity for that region. In order to determine how to reasonably distribute this total amongst the affected pixels, we leverage the angular symmetry inherent to powder diffraction by setting the relative intensity of each pixel to the values from integrated 1-D spectrum. Iterating this procedure then results in reasonable pixel intensities which will no longer induce wild bias into the final integrated spectrum.

The final data reductions step is the integration procedure itself, which operates on the powder diffraction images (pre-corrected for saturation if necessary) using the geometric parameters derived from the calibration image. Using the equations presented in *Hinrichsen et al. (2008)*, we determine the proper intensity weighting factors required to transform each measured pixel value into an equivalent ideal pixel intensity for a perpendicular and hemispherical detector. This factor is combined with the standard 2-D Lorentz and polarization correction factors (see, e.g., *Hinrichsen et al. (2008)*), forming a single overall weighting factor for each pixel on the detector. Since this intensity weighting map is independent of the data collected, it need only be calculated once for each detector configura-

tion. To obtain a reasonably smooth one-dimensional spectrum, as remarked by *Hammersley et al.* (1996), measured pixel intensities are divided amongst sub-pixel regions according to standard bi-cubic interpolation. These sub-pixels are then sorted into bins according to their diffraction angle 2θ . The sub-pixel intensities are combined together as weighted observations of a Poisson process, using the intensity weight map described above, resulting in a one-dimensional spectrum of intensity as a function of diffraction angle. As a final post-processing step, the method derived by *David and Sivia* (2001) is used to estimate and subtract a robust Chebyshev polynomial background from the spectrum. This acts as a good initial guess of the background, which can later be refined, and provides a simple flat spectrum ready for analysis and visualization.