

Synthetic inversions for density using seismic and gravity data

Nienke Blom,¹ Christian Boehm² and Andreas Fichtner²

¹*Department of Earth Sciences, Universiteit Utrecht, P.O. Box 80021, 3508 TA Utrecht, The Netherlands. E-mail: n.a.blom@uu.nl*

²*Institute of Geophysics, Department of Earth Sciences, ETH Zürich, Sonneggstrasse 5, 8092 Zürich, Switzerland*

Accepted 2017 February 27. Received 2017 February 6; in original form 2016 September 9

SUMMARY

Density variations drive mass transport in the Earth from plate tectonics to convection in the mantle and core. Nevertheless, density remains poorly known because most geophysical measurements used to probe the Earth's interior either have little sensitivity to density, suffer from trade-offs or from non-uniqueness. With the ongoing expansion of computational power, it has become possible to accurately model complete seismic wavefields in a 3-D heterogeneous Earth, and to develop waveform inversion techniques that account for complicated wavefield effects. This may help to improve resolution of density. Here, we present a pilot study where we explore the extent to which waveform inversion may be used to better recover density as a separate, independent parameter. We perform numerical simulations in 2-D to investigate under which conditions, and to what extent density anomalies may be recovered in the Earth's mantle. We conclude that density can indeed be constrained by seismic waveforms, mainly as a result of scattering effects at density contrasts. As a consequence, the low-frequency part of the wavefield is the most important for constraining the actual extent of anomalies. While the impact of density heterogeneities on the wavefield is small compared to the effects of velocity variations, it is likely to be detectable in modern regional- to global-scale measurements. We also conclude that the use of gravity data as additional information does not help to further improve the recovery of density anomalies unless strong *a priori* constraints on the geometry of density variations are applied. This is a result of the inherent physical non-uniqueness of potential-field inverse problems. Finally, in the limited numerical setup that we employ, we find that the initially supplied anomalies in *S*- and *P*-velocity models are of minor importance.

Key words: Gravity anomalies and Earth structure; Inverse theory; Waveform inversion; Seismic tomography; Wave propagation; Wave scattering and diffraction.

1 INTRODUCTION

Density is one of the most important material properties influencing the dynamics of our planet's interior. Lateral density variations drive mass transport on all scales, from plate tectonics (e.g. Forsyth & Uyeda 1975; Chapple & Tullis 1977; Bunge *et al.* 2003; Liu & Gurnis 2008; Warners-Ruckstuhl *et al.* 2012) to whole-mantle convection (Turcotte & Schubert 2014). Knowledge of density is also required to distinguish between thermal and compositional heterogeneities (e.g. Trampert *et al.* 2004; Mosca *et al.* 2012). Despite its importance, however, variations of density inside the Earth remain poorly constrained, compared to variations of seismic wave speeds.

1.1 Why is density difficult to constrain?

The majority of tomographic studies traditionally rely on picked arrival times of specific seismic phases. While this is a tried and tested way to infer variations of seismic velocities inside the Earth, infinite-frequency ray theory predicts that arrival times of body waves are exactly insensitive to density (Cerveny 2001). At finite frequencies,

the scattering properties of density anomalies also indicate weak sensitivity of arrival times to density: while the scattered wave of a pure velocity anomaly travels forward with the seismic wave, the scattered signal of a pure density anomaly (velocities remaining constant) travels in the opposite direction, which means that the forward travelling wave remains unperturbed (e.g. Wu & Aki 1985; Tarantola 1986; Fichtner & Trampert 2011).

The amplitudes of reflected and transmitted waves depend directly on density contrasts (Aki & Richards 2002). There are, however, large trade-offs with velocity, because the impedance contrasts causing the partitioning of energy into reflected and transmitted signals are also a function of seismic velocity.

The frequency-dependent arrival times of Rayleigh waves are sensitive to density, but this sensitivity is oscillatory with depth (Takeuchi & Saito 1972). This means that a positive density anomaly may interact with both the positive and negative parts of the sensitivity kernel, thus strongly reducing the net effect. Normal modes display sensitivity to density due to the gravitational restoring force. However, this effect is strongest for the gravest normal modes and these are only sensitive to the longest wavelength structure

in the Earth (e.g. Dahlen & Tromp 1998; Woodhouse & Deuss 2007).

Gravity data are sensitive to mass distribution only, and as such are most directly related to density. However, because they are potential-field measurements, solutions to the gravitational inverse problem are inherently non-unique (Blakely 1995). This means that even in the case of full surface coverage and error-free data, an infinity of solutions remains—this in contrast to, for example, ray tomography, where the central-slice theorem dictates that a perfect reconstruction would be obtained in that case (Iyer & Hirahara 1993). Gravity on its own can thus provide only weak constraints on density structure, and as a result, strong prior information needs to be incorporated into the inverse problem (Blakely 1995; Tarantola 2005).

1.2 Previous work

1.2.1 Seismic data

Laterally averaged 1-D density models (e.g. Dziewoński *et al.* 1975; Dziewoński & Anderson 1981) are based on measurements of the Earth's total mass and moment of inertia and on frequencies of the gravest normal modes. Because of the persistent non-uniqueness of the resulting inverse problem, these models have to be supplemented with *a priori* assumptions about possible adiabaticity in different portions of the Earth, as well as about the density jumps at discontinuities. These assumptions may lead to biases, especially near discontinuities (Kennett 1998; de Wit & Trampert 2015). Information about the density contrast at such discontinuities can be obtained from the amplitudes of reflected and transmitted waves, but these suffer from trade-offs with velocity structure (Shearer & Flanagan 1999; Kato & Kawakatsu 2001).

The sensitivity of normal modes has also been used to study the long-wavelength 3-D density structure in the lower mantle (Ishii & Tromp 1999, 2001, 2004). However, the robustness of these results has been questioned by various authors (Resovsky & Ritzwoller 1999; Kuo & Romanowicz 2002; Resovsky & Trampert 2002). Koelemeijer *et al.* (2017) confirmed that earlier results from normal-mode data were not robust, but can be improved significantly with the addition of the most recent measurements.

For long-period surface waves (100–500 s), density may be resolvable in the upper 150–200 km (Tanimoto 1991). At these depths, an anticorrelation of density and *S* velocity is found in some regions, where continents appear light, but fast. The ratio of horizontal and vertical components (*H/V* ratios) of higher frequency surface waves can be used to map density as well (e.g. Lin *et al.* 2012). However, *H/V* sensitivity displays strong trade-offs with other parameters, and can only be used with significant regularization, or with the inclusion of *a priori* scaling relations.

Scaling relations are a commonly used method to incorporate density into seismic tomography. For instance, the magnitude of relative density anomalies is often scaled to that of relative *S*-wave velocity. Such scaling ratios $\text{dln}\rho/\text{dln}v_s$ typically range between 0.1 and 0.4 (e.g. Karato & Karki 2001; Simmons *et al.* 2009), but may be negative especially in (some parts of) the shallow and deep mantle. While in many regions, this results in reasonable models that may also agree well with gravity data, these scaling relations preclude the detection of those interesting regions where they do not hold. Such regions may be present in all parts of the Earth, from the crust to the deep mantle (e.g. Tanimoto 1991; Ishii & Tromp 1999; Simmons *et al.* 2009), and may point, for example, to chemical heterogeneity or the presence of melt.

On exploration scales, density is generally incorporated into studies by way of (contrasts in) impedance (e.g. Mora 1987) although some recent work has been done to study density separately (e.g. Jeong *et al.* 2012; Prieux *et al.* 2013).

1.2.2 Gravity data

Apart from seismic data, gravity measurements have been used to study the density distribution within the Earth—usually in the shallower crust and asthenosphere but sometimes on the global scale as well. Such gravity studies are often augmented with additional information to reduce non-uniqueness. Seismic models are used to provide a starting density model (e.g. Chaves & Ussami 2013) or to isolate the contribution of specific layers (e.g. Kaban *et al.* 2004; Herceg *et al.* 2016; Root *et al.* 2017). Both approaches may rely on a pre-defined density model or scaling between density and velocity, which precludes the discovery of those interesting cases where the scaling is different or may result in artefacts of the same size as the structure of interest (Herceg *et al.* 2016). Alternatively, the shapes of anomalies can be extracted from seismic velocity models, the density of which is then adjusted in order to fit the gravity data (e.g. Fadel *et al.* 2015). In this case, the assumption is made that the shapes of anomalies in seismic velocity coincide with the shapes of density anomalies. Finally, the contribution of modelled density anomalies can be analysed and compared to maps of the geoid, gravity field or gravity gradients (e.g. Panet *et al.* 2014).

1.2.3 Joint inversions

Another approach, aimed at exploiting the complementary information present in different data sets, consists of combining multiple types of observables into a joint inversion. These data can range from body and surface waves, to normal modes and geodynamic data, including estimates of past plate motions and the location of subducting slabs. Such joint inversions have been carried out with the aim of constraining velocity and density structure in the whole mantle (e.g. Ricard *et al.* 1993; Nataf & Ricard 1996; Tondi *et al.* 2009; Simmons *et al.* 2010; Moulik & Ekström 2016). Despite the additional data, however, in many cases such studies still rely on imposed scaling relations between velocity and density, or on subjective choices of damping in order to avoid artefacts and reduce trade-offs.

1.3 The effect of density on the seismic wavefield

Despite the issues outlined above, density anomalies do affect the seismic wavefield in a characteristic way, as illustrated in Fig. 1, which is in principle measurable (Płonka *et al.* 2016). For Fig. 1, wave propagation was calculated in 2-D for a single source in a medium with constant seismic velocities and a rectangular 10 per cent ($+260 \text{ kg m}^{-3}$) density anomaly. Such a setup is similar to scattering experiments conducted, for example, by Frankel & Clayton (1986) and, more recently, Prieux *et al.* (2013), and serves to visualize how structure affects the seismic wavefield. The radiation pattern of the source is such that mainly *S* waves travel in the horizontal direction. All boundaries are absorbing (Cerjan *et al.* 1985). As the *S* wave reaches the anomaly, it interacts with it (Figs 1a–c), causing perturbations in the wavefield travelling both in the forward and backward directions (Figs 1d–f). These wavefield perturbations are generated at the edges of the anomaly, where there is a contrast in density. With a density anomaly of 10 per cent, the differential

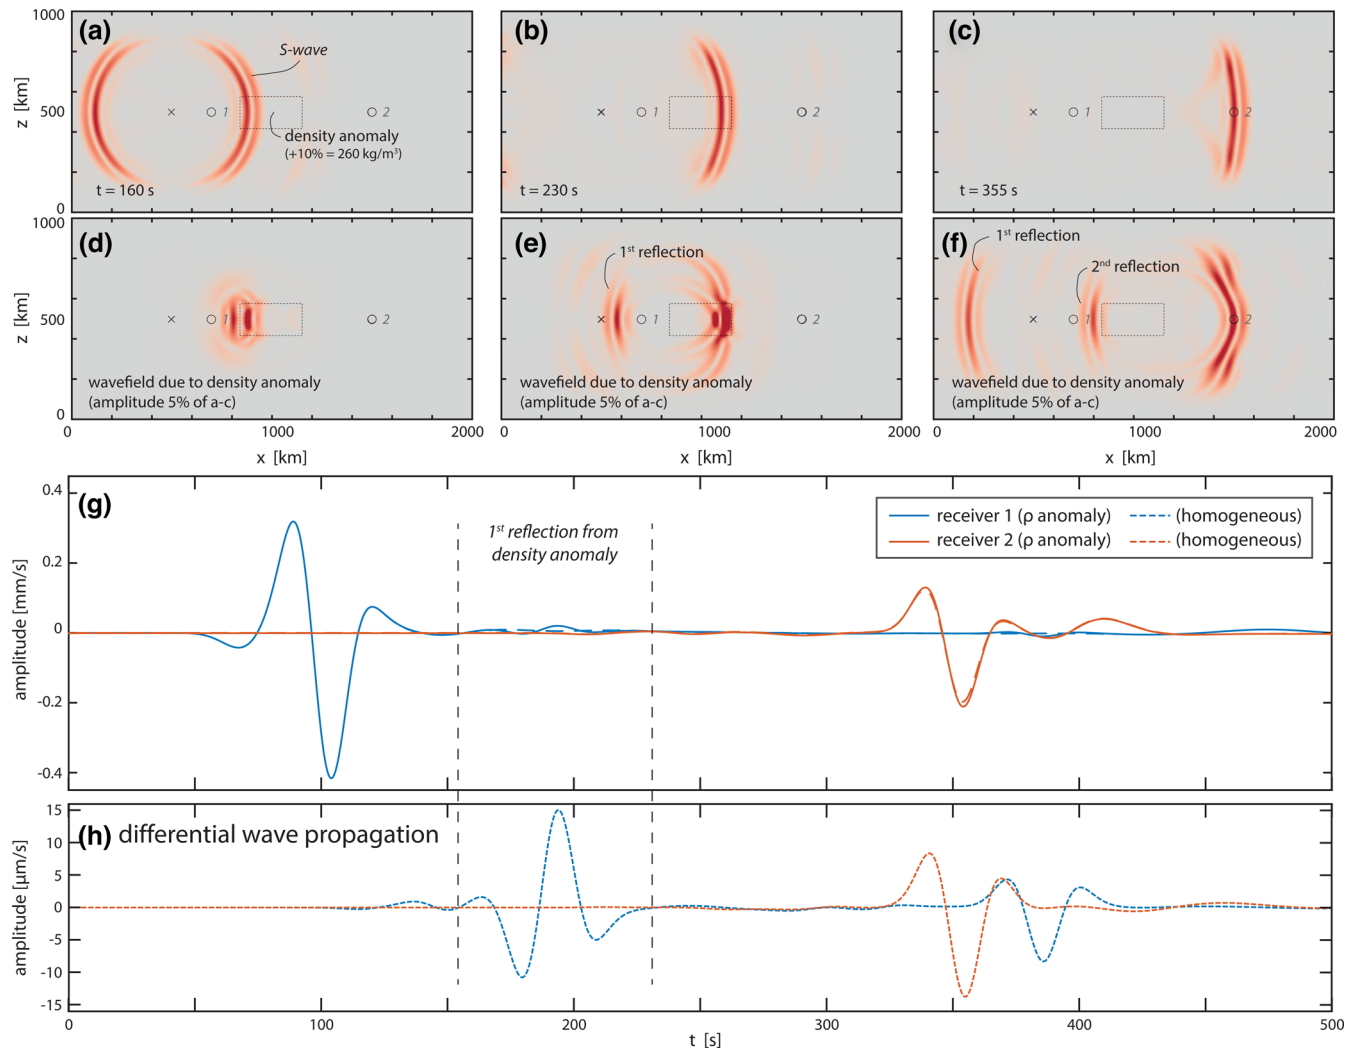


Figure 1. Snapshots of an S wave propagating past a rectangular 10 per cent ($+260 \text{ kg m}^{-3}$) density anomaly. The complete video of this wave propagation can be found in the online Supporting Information. (a)–(c) Velocity wavefield. (d)–(f) Differential wavefield caused by the density anomaly only—amplitudes here are 5 per cent of the amplitudes shown in (a)–(c). (g) Seismograms recorded at receivers 1 and 2 for both cases, with and without a density anomaly. At receiver 1, a clear separate arrival is visible caused by the reflection at the first density interface. (h) Differential seismograms from receivers 1 and 2, obtained by subtracting $v_{\text{diff}} = v_{\rho \text{ anomaly}} - v_{\text{homog}}$.

wavefield caused by the density anomaly has amplitudes of about 5 per cent of the original wavefield (see seismograms in Fig. 1g). This is not a very large signal, but it is measurable with current state-of-the-art broad-band sensors. A significant part of the energy is scattered in directions other than the propagation direction of the primary wave, most notably in the backwards direction, that is, as a reflection towards the source. As a result, a separate arrival is visible at receiver 1 after the direct wave has passed through (indicated by the dotted lines in Fig. 1g). This constitutes the largest amplitude differential signal (Fig. 1h).

The differential wavefield caused by the density anomaly travels with the same speed as the original wavefield, which means that for the wave travelling directly from source to receiver 2, it is mainly visible as an amplitude change. Nevertheless, the shape of the waveform is altered (Fig. 1g) and this expresses itself as a slight change in finite-frequency traveltimes of about 0.5 s, which can be measured, for instance, by cross-correlation (e.g. Luo & Schuster 1991; Dahlen *et al.* 2000; Płonka *et al.* 2016). In addition to these direct-wave effects, scattered waves emanating from the density heterogeneity can be measured as small-amplitude arrivals at receivers off the

direct wave path. As Fig. 1 illustrates, the strongest sensitivity of seismic waves to density lies in density contrasts or gradients, illustrating the close relationship of density with reflection and transmission coefficients on one hand, and with surface wave sensitivity on the other.

1.3.1 Our approach

The above observations, complemented by the recent study by Płonka *et al.* (2016) showing that Earth-like density variations on the crustal scale have measurable effects on seismograms, lead us to speculate that waveform inversion may be a viable method to study density in the Earth's interior at regional scales. Waveform inversion studies have existed for a long time (Woodhouse & Dziewoński 1984; Gauthier *et al.* 1986) and have the advantage that in principle, all information contained in the seismogram can be used (e.g. Fichtner & Igel 2008; Bozdağ *et al.* 2011; Rickers *et al.* 2012). In this paper, we will study to what extent waveform inversion can be used to recover density as an independent parameter.

Our philosophy is to introduce as few constraints on density as possible. To study the pure effect of density heterogeneity on wave propagation, we will not assume any scaling relationship between seismic velocity and density structure. Moreover, we will not assume any prior information on the distribution of density heterogeneity throughout our models. This is because we want to study to what extent it is possible to image density completely independently from other parameters.

In this study, we will investigate the conditions that must be met in order to reconstruct density on the regional to global scale, and the types of prior information and inversion strategies that can be exploited. For the sake of simplicity, we assume an isotropic, non-attenuating medium in which the source locations and signals are assumed to be known. This is in order to focus on the effect of density on seismic wave propagation. While some of these aspects are known to be important when studying the real Earth, with this study we simply want to explore the kind of quality level that must be reached in terms of data, *a priori* information, inversion scheme, etc.

First, we investigate to what extent it is possible to recover density from the seismic signal at all. We will also assess how the inversion is affected if density is ignored or (erroneously) scaled to S velocity. Next, we will focus on prior information in the form of starting models in S and P velocities, and the extent to which we can assume that these are correct, and thus invert for density only. We will study the effect of including gravity information into our inversion, and investigate the effect of noise on the recovery of density structure. Finally, we will check to what extent density can be recovered when no strong impedance contrasts are present, contrary to the scenario in Fig. 1 where they generated the most significant density signal. Additional tests on parametrization and anomaly amplitude can be found in the online Supporting Information.

2 METHODS

In the following sections, we will elaborate on the numerical experiments that we carry out in this study. To have a flexible workflow in which it is easy to test different types of inversion schemes and model setups, and still remain within reasonable bounds of computational time and effort, we use a 2-D Cartesian model setup in which we solve the elastic isotropic wave equation.

2.1 Model

Our 2-D model is designed to mimic the Earth's mantle. This choice is driven by our interest in density on the regional to global scales, but it should be noted that in principle, our experiments are scale-independent. The size of the domain is 2890 km vertically (approximately the depth of the Earth's mantle) by 6000 km horizontally. The 1-D background values for the seismic velocities and density are taken from PREM (Dziewoński & Anderson 1981). The top and bottom boundaries are reflecting, mimicking in a simplistic but effective manner the Earth's surface and the core–mantle boundary. The side boundaries are absorbing using the Gaussian taper method of Cerjan *et al.* (1985) to avoid artificial illumination by reflected waves that do not exist in the real Earth. Since waves propagating multiple times around the Earth are also not modelled, the illumination in this setup can be considered conservative. The model is divided into 430×207 blocks, resulting in a grid spacing of ~ 14 km.

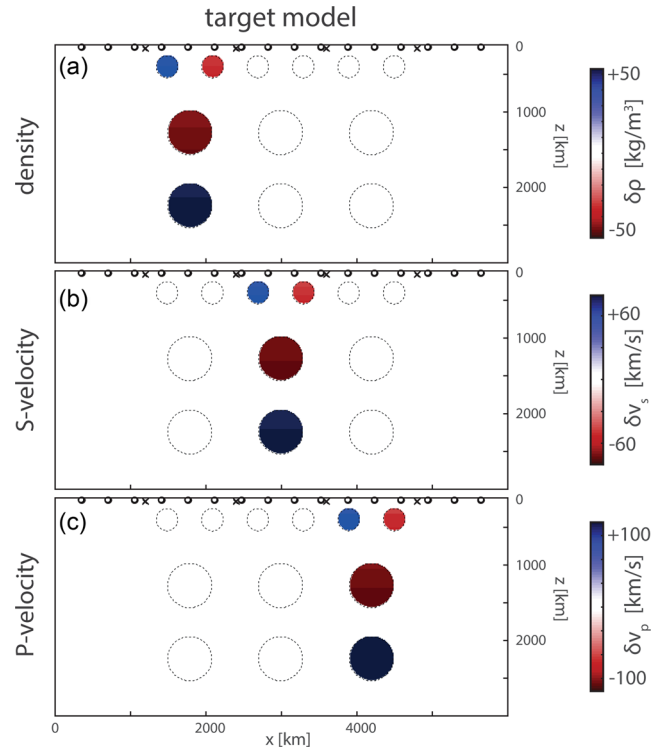


Figure 2. Target model. Three columns of positive and negative anomalies are present: density anomalies on the left, S -velocity anomalies in the middle and P -velocity anomalies on the right. All anomalies are 1 per cent of the background value (PREM, Dziewoński & Anderson 1981).

In the target model, separate perturbations of ± 1 per cent strength in density, and S and P velocities are superimposed onto the background model PREM (Fig. 2). These perturbations are completely uncorrelated to one another. This is done in order to make sure that, for instance, an apparently resolved density structure is not just an artefact of velocity structure leaking into density.

2.2 Sources and receivers

There are four seismic source locations at 56 km depth, at each of which two ‘events’ in the form of vector forces are excited. Their radiation patterns vary such that one event radiates predominantly P waves in the vertical direction, and the other predominantly S waves. Because of the absorbing boundaries, most of the energy propagating in the horizontal direction is lost, although it will serve to illuminate the upper mantle to some extent. This results in a total of eight events. Our choices concerning the number of sources and receivers are deliberately conservative, intended to prevent overly optimistic results.

Near the surface of the domain, 16 seismic receivers are located which are spaced ~ 350 km apart. This source–receiver configuration is computationally convenient and proves sufficient for our purposes, although in the real Earth, a much denser coverage, albeit more irregular can be expected. Because of the model grid spacing, the minimum source wavelength is limited to 100 km to ensure sufficiently accurate synthetics. This is achieved by band-pass filtering the source-time function at periods longer than 30 s. As a result, the signals have frequencies from 0.0067 to 0.032 Hz (periods 30–150 s), a frequency range for which broad-band seismometers generally supply good-quality data on the global scale. The seismic signal is recorded for 1200 s in order for both the

synthetic core reflections ‘PcP’ and ‘ScS’ to reach the receivers. At 20 km above the surface, gravity is measured at 51 separate gravity sensors. This is done in a point-mass approach: every gridpoint is assigned a mass, which contributes to the gravity field measured at each sensor. In order to remove artefacts that might result from this, the gravity sensors are not placed directly at the surface but 20 km above it.

2.3 Seismic waveform inversion

Within the model domain, we solve the wave equation

$$\rho(\mathbf{x}) \partial_t^2 \mathbf{u}(\mathbf{x}, t) - \nabla \cdot [\mathbf{C}(\mathbf{x}) : \nabla \mathbf{u}(\mathbf{x}, t)] = \mathbf{f}(\mathbf{x}, t), \quad (1)$$

where \mathbf{u} is the displacement field and \mathbf{f} is the forcing term. Furthermore, ρ denotes density and \mathbf{C} is the elastic tensor, which is in our case the isotropic elastic tensor $C_{ijkl} = \lambda(\delta_{ij}\delta_{kl}) + \mu(\delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk})$. In the following, we use \mathbf{m} to denote the ensemble of all model parameters in order to treat different parametrizations of the material properties in a unified way. We developed a 2-D finite-difference wave propagation code with a computationally efficient staggered grid (Virieux 1984, 1986) in which we define $\mathbf{x} = (x, z)$ where x is the horizontal and z the vertical direction. We only consider P – SV motion, that is, particle motion that is in the same plane as the one in which the model is defined.

To recover our target model, we compare the seismograms of test models to those of the target model. To this end, we construct an L_2 -norm misfit functional in which we compare the complete 1200 s of waveform data, summed over all components (c), receivers (r) and sources (s), and normalized by the initial misfit:

$$\min \quad J(\mathbf{m}) := \frac{\sum_{s,r,c} \int_t [\mathbf{u}(\mathbf{m}, t) - \mathbf{u}_{\text{target}}(t)]^2 dt}{\sum_{s,r,c} \int_t [\mathbf{u}(\mathbf{m}_0, t) - \mathbf{u}_{\text{target}}(t)]^2 dt}, \quad \text{subject to eq. (1).} \quad (2)$$

The advantage of an L_2 norm is that it automatically takes into account all the information, meaning that theoretically, every single oscillation can be exploited. The disadvantage is that because the pure waveform difference is the measurement, this metric is naturally dominated by the largest amplitude signals. It is thus expected that S velocity is imaged best, while the subtle density effects as shown in Fig. 1 have a much less pronounced effect. While the least-squares misfit is known to enhance nonlinearity (e.g. Gauthier *et al.* 1986; Luo & Schuster 1991) this is not a problem for the global scale considered here because heterogeneities are generally so small that the initial model is always sufficiently close to the optimum. We use a deterministic inversion approach in which consecutive updates are calculated starting from some initial model \mathbf{m}_0 . For a given model \mathbf{m}_i , an updated model \mathbf{m}_{i+1} that better explains the data are calculated from a descent direction \mathbf{s}_i , scaled by some step length $\gamma_i > 0$:

$$\mathbf{m}_{i+1} = \mathbf{m}_i + \gamma_i \mathbf{s}_i, \quad (3)$$

to ensure $J(\mathbf{m}_{i+1}) < J(\mathbf{m}_i)$. The descent direction \mathbf{s}_i is determined by the negative gradient of the misfit of the current model, preconditioned by an approximation of the inverse Hessian \mathbf{B}_i :

$$\mathbf{s}_i = -\mathbf{B}_i \cdot \nabla J(\mathbf{m}_i). \quad (4)$$

(Nocedal & Wright 2006). Using the adjoint approach (e.g. Taran-tola 1988; Liu & Tromp 2008; Fichtner *et al.* 2006), we calculate the gradients of the seismic data misfit $\nabla J(\mathbf{m}_i)$ using only two wave propagation simulations per event. The forward wavefield is stored at evenly spaced intervals (every 10 time steps, in our case).

Gradients, or sensitivity kernels, are then constructed on the fly during the time-reversed adjoint simulation by combining information from the forward and adjoint wavefields. Sensitivity kernels for density, S and P velocities are defined by

$$K_{\rho, \text{event}} = - \int_T \dot{\mathbf{u}}^\dagger \cdot \dot{\mathbf{u}} dt + (v_p^2 - 2v_s^2) \int_T (\nabla \cdot \mathbf{u})(\nabla \cdot \mathbf{u}^\dagger) dt + v_s^2 \int_T 2\epsilon^\dagger : \epsilon dt, \quad (5)$$

$$K_{v_s, \text{event}} = 2\rho v_s \int_T (\nabla \cdot \mathbf{u})(\nabla \cdot \mathbf{u}^\dagger) dt - 4\rho v_s \int_T 2\epsilon^\dagger : \epsilon dt, \quad (6)$$

$$K_{v_p, \text{event}} = 2\rho v_p^2 \int_T 2\epsilon^\dagger : \epsilon dt, \quad (7)$$

with $:$ denoting the elementwise scalar product of two tensors, ϵ the strain tensor corresponding to the wavefield \mathbf{u} , \mathbf{u}^\dagger the (time-reversed) adjoint wavefield and ϵ^\dagger the adjoint strain tensor corresponding to it (e.g. Fichtner 2010). For legibility, we omitted spatial and temporal dependences. The event kernels are summed to obtain the final kernels. Gradients are obtained from eqs (5)–(7) using the Riesz representation, that is, the projection of the kernel onto the finite-difference grid. Note that the definitions above are given for the physical parameters, but the actual inversion parameters are defined as relative perturbations with respect to a background model, that is, we use

$$\mathbf{m} = (m_1 \ m_2 \ m_3)^T \quad (8)$$

and define

$$\begin{aligned} \rho &= \rho_{\text{ref}} \cdot (1 + m_1), \\ v_s &= v_{s, \text{ref}} \cdot (1 + m_2), \\ v_p &= v_{p, \text{ref}} \cdot (1 + m_3). \end{aligned} \quad (9)$$

This non-dimensionalization serves to circumvent issues with differences between the magnitudes of the different parameters and improves the scaling of line search methods where all parameters are updated with the same step length. Partial derivatives with respect to \mathbf{m} can be computed easily using the sensitivity kernels defined in eqs (5)–(7) and applying the chain rule. On the kernels, this has the effect of scaling them by the reference model.

2.4 Descent method

For the iterative misfit reduction, we employ the L-BFGS scheme (Byrd *et al.* 1995; Nocedal & Wright 2006), which uses the models and gradients of previous iterations to construct an approximation of the inverse Hessian, that is, \mathbf{B}_i in eq. (4). This serves to include curvature information into the inversion. L-BFGS is matrix-free and requires only vector–vector operations. Thus, the additional costs of computing the search direction are negligible compared to conjugate-gradient methods. While this method is not guaranteed to converge towards the actual Hessian, experience shows that it strongly speeds up convergence. Another advantage is that the search direction is automatically scaled using curvature information which makes it possible to skip the computation of additional trial step lengths in most cases. Hence, for most updates, two simulations suffice: one forward and one adjoint. The L-BFGS minimization scheme (see Appendix) is implemented into our methodology via an interaction between our finite-difference wave propagation code and a separate custom-built optimization toolbox for seismic inversion.

Table 1. Overview of used frequency ranges. All inversions start with the lowest frequency band and after every 20 iterations, the scheme switches to the next band.

Frequency number	Frequency band (Hz)	Period range (s)
1	0.0067–0.0067	150–150
2	0.0067–0.0083	150–120
3	0.0067–0.0104	150–96
4	0.0067–0.0130	150–77
5	0.0067–0.0163	150–61
6	0.0067–0.0204	150–49
7	0.0067–0.0254	150–39
8	0.0067–0.0318	150–31

An L_2 -norm waveform misfit functional may suffer from cycle-skipping, for which reason we use a multiscale approach (Bunks *et al.* 1995), starting the inversion with the longest periods (150 s, or 0.0067 Hz), and adding higher frequency content in a stepwise manner across iterations, up to periods of ~ 30 s (0.032 Hz). The complete frequency progression is summarized in Table 1. As a mild form of regularization, we applied a ~ 30 km wide Gaussian smoothing of the sensitivity kernels to avoid the appearance of small-scale artefacts. We experimented with the number of iterations per frequency, and found that 20 is sufficient to explain the artificial data to such an extent that it prevents cycle skips when going to the next higher frequency band. This results in a total of 160 iterations (8 frequency bands \times 20 iterations).

3 NUMERICAL EXPERIMENTS

We carry out a number of tests assessing different inversion strategies. In this section, we will briefly introduce these tests, the motivation for doing them and some information on their construction. Their results will be discussed in Section 4, and for a number of additional tests in the on-line Supporting Information.

3.1 Reference case: inverting for density

As a very first experiment, we will assess to what extent density can be recovered as an independent parameter in addition to the seismic velocities v_S and v_P . In this case, we supply no *a priori* information other than the background and initial model PREM. All three parameters are free to vary. This experiment will provide a first indication of what part of density structure is recoverable.

3.2 Neglecting density

In seismic tomography, it is often the case that only S and/or P velocity are inverted for. While seismic traveltimes in the ray approximation are by construction insensitive to density, this is not the case for complete seismic waveforms, which include the phases and amplitudes of body, surface and scattered waves. In waveform inversion, therefore, this sensitivity to density, if ignored, may lead to biases in the recovered velocity model. In this test, we will compare the results of an inversion where only S and P velocities are variable, in contrast to our reference case where all three parameters are inverted for. One can keep density fixed by simply not updating it:

$$\begin{aligned}\rho &= \rho_{\text{ref}} \quad (\text{fixed}), \\ v_S &= v_{S,\text{ref}} \cdot (1 + m_2),\end{aligned}$$

$$v_P = v_{P,\text{ref}} \cdot (1 + m_3). \quad (10)$$

In this test, too, the starting model is PREM. Because density is kept fixed, there is only two-third of the number of free parameters compared to the reference case of Section 3.1.

3.3 Scaling density to S velocity

Alternatively, it is common to scale density to S velocity as

$$R_{\rho/S} = \text{dln}\rho / \text{dln}v_S. \quad (11)$$

Here, we will investigate what the effect is of constraining an inversion to a pre-defined scaling ratio between density and S velocity. Again, we will assess to which extent this constraint causes errors in the recovery of parameters. For the actual inversion parameters, this means that again, there are only two-thirds as many free parameters. Here, however, density is also a function of m_2 :

$$\begin{aligned}\rho &= \rho_{\text{ref}} \cdot (1 + R_{\rho/S} \cdot m_2), \\ v_S &= v_{S,\text{ref}} \cdot (1 + m_2), \\ v_P &= v_{P,\text{ref}} \cdot (1 + m_3)\end{aligned} \quad (12)$$

with an imposed scaling of $R_{\rho/S} = 0.2$. Note that the resulting sensitivity kernel for m_2 will now also include a term with K_ρ .

Our target model, however, has three differently scaled columns with $R_{\rho/S} = 0.4, 0.2$ and -0.2 . All of these lie within the assumed range of validity within the Earth (see e.g. Karato & Karki 2001; Simmons *et al.* 2009). This means that the imposed scaling will only be correct for one of the three columns, whereas it is incorrect for the other two.

3.4 Initial velocity models

While the previous tests assumed no initial information on any of the parameters, in most regions of the Earth some information is present on the velocity structure. We therefore also conduct a series of tests in which the target S - and P -velocity structures are included in the starting model to progressively better degrees. Our approach is to include a percentage of the magnitude of the target S - and P -velocity anomalies into the starting models. The rationale behind this is that tomographic inversions often recover the (long-wavelength) shape of anomalies reasonably well, but have problems constraining their amplitude. In our starting models, S - and P -velocity anomalies are 50, 75 and 100 per cent of their actual values, respectively. The more the starting model resembles the true model, the more of the signal is due to density only and so the better convergence should be.

3.5 Constraints on velocity structure

Because the effect of density heterogeneity on the seismic wavefield is small compared to the effect of velocity variations, the recovery of density will also be relatively slow. We therefore test here whether it is beneficial for the recovery of density to reduce the number of free parameters by inverting for density only. Fixing the velocity model may be reasonable if one is certain that it is accurate for a given region, but the question is what is ‘accurate’ enough. The aim of this test is therefore twofold: on one hand, to investigate whether the recovery of density is improved and/or sped up if it is the only parameter inverted for, and on the other hand, to assess the extent to which errors in the velocity starting model influence the recovery of density structure. Such errors may consist of artefacts, smearing and

a limited recovery in the amplitude of anomalies. Starting models are the same as in Section 3.4, but now, only density is allowed to vary. The percentages reflect the variation in recovered velocity anomaly strength that is often found across different tomographic studies and seems reasonable in the light of performed resolution analysis of waveform inversions (Fichtner & van Leeuwen 2015; Simutè *et al.* 2016). We here invert for density only, with v_S and v_P fixed to the initial values, only one parameter is inverted for, effectively reducing the model space:

$$\begin{aligned}\rho &= \rho_{\text{ref}} \cdot (1 + m_1), \\ v_S &= v_{S,\text{ref}}, \\ v_P &= v_{P,\text{ref}}.\end{aligned}\quad (13)$$

If the velocity model is correct, we expect the recovery of density to be both faster and better. However, because density has a much smaller effect on wave propagation than variations in seismic velocity (Fig. 1), we also expect that even small deviations from the correct velocity model will translate into significant errors in the density model.

3.6 Including gravity information

Because gravity has significant sensitivity to the distribution of mass that may be complementary to seismic sensitivity, we also test to what extent the inclusion of gravity data into the inversion aids the recovery of density structure. There are many ways in which gravity can be included into the inversion, depending on the chosen data types, constraints on the size and/or extent of anomalies, relationship of density to other parameters, artificial scaling of sensitivity, etc. Here, we take the simplest and most straightforward approach which does not make any explicit *a priori* assumptions on density distribution or its correlation with velocity structure. We therefore add a gravity term to the misfit functional in eq. (2):

$$J_{\text{seis+grav}} = J_{\text{seis}} + J_{\text{grav}} \quad \text{with} \quad J_{\text{grav}}(\mathbf{g}, \rho) := \frac{\sum_{r,c} [\mathbf{g}(\rho) - \mathbf{g}_{\text{target}}]^2}{\sum_{r,c} [\mathbf{g}(\rho_0) - \mathbf{g}_{\text{target}}]^2}. \quad (14)$$

This gravity misfit is constructed similarly to the seismic misfit that is also normalized by the initial model misfit. As a result, both are unitless and can be added up consistently. The full gravity vector \mathbf{g} is used in eq. (14), but we also perform tests using the gravity potential φ which is related to \mathbf{g} through $\mathbf{g} = -\nabla\varphi$ and is usually expressed in terms of the geoid height. For the relations of \mathbf{g} and φ to the distribution of density, we have

$$\mathbf{g}(\rho) = -G \int_V \frac{\rho(\mathbf{x})}{r(\mathbf{x})^3} \mathbf{r}(\mathbf{x}) \, d\mathbf{x}, \quad (15)$$

$$\varphi(\rho) = -G \int_V \frac{\rho(\mathbf{x})}{r(\mathbf{x})} \, d\mathbf{x}, \quad (16)$$

where \mathbf{r} is the vector between the receiver and \mathbf{x} , r is its length and G is the universal gravitational constant. Because gravity is unrelated to v_S and v_P , it only affects the density part of the gradients. The total density kernel corresponding to the joint misfit (14) is given by

$$K_{\text{total},\rho} = K_{\text{seis},\rho} + K_{\text{grav},\rho}. \quad (17)$$

The additional gradient $K_{\text{grav},\rho}$ now has to be calculated. The gradient for a single gravity sensor and component for measurements of \mathbf{g}

and φ with respect to density becomes:

$$K_{\mathbf{g},\rho} = -2G \cdot \frac{[\mathbf{g} - \mathbf{g}_{\text{target}}]}{\sum_{r,c} [\mathbf{g}(\rho_0) - \mathbf{g}_{\text{target}}]^2} \cdot \frac{\mathbf{r}}{r^3}, \quad (18)$$

$$K_{\varphi,\rho} = -2G \cdot \frac{[\varphi - \varphi_{\text{target}}]}{\sum_{r,c} [\varphi(\rho_0) - \varphi_{\text{target}}]^2} \cdot \frac{1}{r}. \quad (19)$$

These are summed over sensors and components in order to obtain the total gravity kernels. Due to the \mathbf{r}/r^3 and $1/r$ terms in eqs (18) and (19), gravity sensitivity to density is generally concentrated near the surface, closest to the receivers. The different dependencies on r imply that geoid kernels will have larger sensitivity to deep structure than gravity vector kernels evaluated at the same receiver location. This sensitivity is, however, also much smoother and thus less detailed.

In Fig. 3, we compare sensitivity to density for several measurements used in this study. Sensitivity kernels are calculated for a blank PREM starting model compared to data obtained from the standard target model (Fig. 2). This means that these kernels are effectively the contributions from each of these measurements to a first model update in a steepest descent inversion algorithm. In Figs 3(a)–(d), we compare the sensitivity to density for both the gravity vector \mathbf{g} and potential φ for measurements done at two different heights above the model. For the full gravity vector (Figs 3a and b, $1/r^2$ dependence), sensitivity is concentrated near the surface, and ‘sees’ the two upper-mantle anomalies: one positive and one negative. For the gravity potential (Figs 3c and d, $1/r$ dependence), sensitivity is less concentrated near the top surface, and hence ‘sees’ the lower mantle more strongly. Only the signal of the upper (negative) density anomaly in the lower mantle is visible in the kernels, as the lower (positive) one is deeper, and thus has a weaker signal which is completely hidden by the upper anomaly. A similar effect occurs when one applies upward continuation of the gravity data, which can be seen when comparing the kernels for measurements at different heights (compare Figs 3a and c to b and d).

Compared to gravity, the seismic kernel (Fig. 3e) shows much more detail. Sensitivity clearly focuses in the upper-mantle anomalies, and even a hint of the edges of the lower-mantle anomalies is visible. At the same time, however, there are more artefacts, a somewhat oscillatory pattern and some trade-off to S -velocity structure. Here, too, there is sensitivity especially to the upper-mantle anomalies.

3.7 The effect of noise

The outset of this study was to assess to what extent and under which conditions it is possible to recover density variations inside the Earth from geophysical data. Because the density signal is relatively small in amplitude (Fig. 1), it is important to assess what noise levels are allowable for the recovery of these anomalies. In the next series of tests, we will therefore study the effect of noise on the inversion results. We will investigate both uncorrelated noise (simulating instrumental noise) and correlated noise (simulating ambient noise).

3.7.1 Generating uncorrelated noise

Uncorrelated noise is generated as follows. For each source–receiver pair and each component, a noise trace is prepared by generating independent random noise traces for all recordings, and

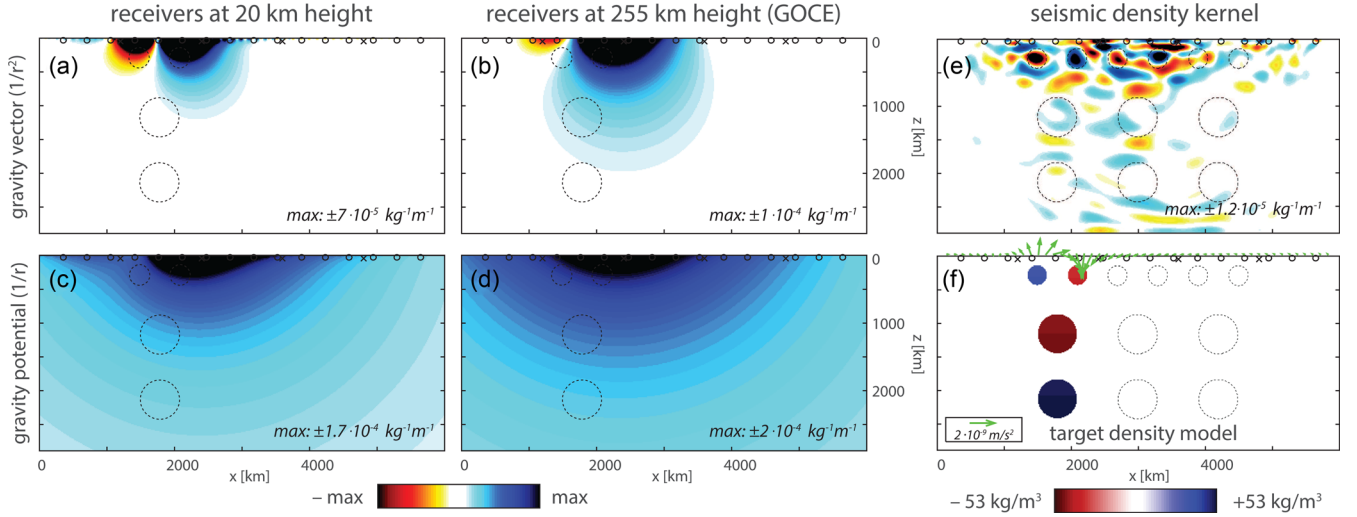


Figure 3. Density kernels for gravity and seismic data. The kernels depicted in panels (a)–(d) represent the negative descent direction for the gravity term in the misfit functional of eq. (14), meaning that in a steepest descent inversion strategy, this is what the gravity contribution to the first-iteration model update will be. Both using geoid instead of gravity vector data and evaluating the data at greater height have the effect of a low-pass filter. In contrast, sensitivity of seismic measurements (panel e) shows much more detail and clearly highlights the true upper-mantle anomalies. (a) Full gravity vector kernel from data measured at 20 km height above the top surface. (b) Same, at 255 km height. (c) Gravity potential kernel from data measured at 20 km height. (d) Same, but at 255 km height. (e) Seismic sensitivity kernel for data filtered at periods of 150 s. (f) Target density model. The difference between gravity vectors for the starting and target models $\mathbf{g} - \mathbf{g}_{\text{obs}}$ is plotted in green arrows.

filtering at the maximum frequency band of the data. The noise traces are scaled to a prescribed signal-to-noise ratio and then band-pass filtered again at each inversion frequency band individually (see Table 1). Because the low-pass filter does not have a sharp cut-off, the noise at lower frequencies will be relatively strong compared to that at higher frequencies. This is reminiscent of what is generally observed in real-data waveforms on the regional to global scale. Noise is then added to the traces.

3.7.2 Generating correlated noise

Correlated noise is generated by performing wave propagation through the target model from noise sources at the surface. For each event, random noise is emitted at each gridpoint along the top surface of the model domain, mimicking ambient noise generation. The noise emitted at each point is prepared using a random number generator, scaled to a random magnitude, and then bandpass filtered to remove frequencies which are too high for the numerical grid. The noise sources all radiate their signal simultaneously, for the whole duration of the wave propagation simulation (1200 s). This results in a separate noise event for each actual event. These are scaled and then bandpass filtered at each inversion frequency band individually (see Table 1). Similarly to the uncorrelated noise, this procedure results in a variable noise level for each frequency band.

3.8 The effect of impedance contrast

The effect of density on the seismic wavefield is most apparent at impedance contrasts (Fig. 1), with P and S impedance defined as

$$Z_P = v_P \cdot \rho, \quad Z_S = v_S \cdot \rho. \quad (20)$$

A lack of impedance contrast may thus be expected to cause a reduction in the recovery of density. We therefore design a test in which a density contrast is offset by a contrast in the seismic velocities, so that both P and S impedance variations are roughly zero. This means that if an increase in density across an interface

is offset by a reduction in velocity (and vice versa), the impedance remains constant. Our target model in this case only features one column of anomalies, in which density, v_S and v_P overlap such that the impedance effect of each density contrast is cancelled by a velocity contrast of opposite sign (this therefore also results in a negative $R_{p/s}$). The starting model is again PREM. Situations in which density and velocity change with opposite sign across an interface exist in nature; one such example is a transition from sandstone to salt, in which case seismic velocity increases while density decreases.

4 RESULTS

In this section, we will discuss the results of the experiments described above in Section 3.

4.1 Reference case

In our initial reference experiment, we assess to what extent it is possible to recover density alongside the seismic velocities using a starting model which contains no information on the target anomalies, and without including any additional information. Results are shown in Figs 4(d)–(f). All three parameters are indeed well recovered, especially in the upper mantle. While density is the least recovered of the three, the target anomalies are still clearly visible. The shapes of the anomalies are recovered almost perfectly and there is no mapping of any parameter into one of the others.

Interestingly, the recovered density model has notable ‘shadows’ of opposite sign around the recovered anomalies. This reflects the fact that it is *contrasts* in density that alter the waveforms most significantly (see Fig. 1). Similar ‘shadows’ are not present in the velocity models. The recovered S -velocity anomalies are somewhat sharper than the P anomalies—a result of the fact that P waves travel faster and (all other factors being equal) thus have longer wavelength signals at the same source frequency.

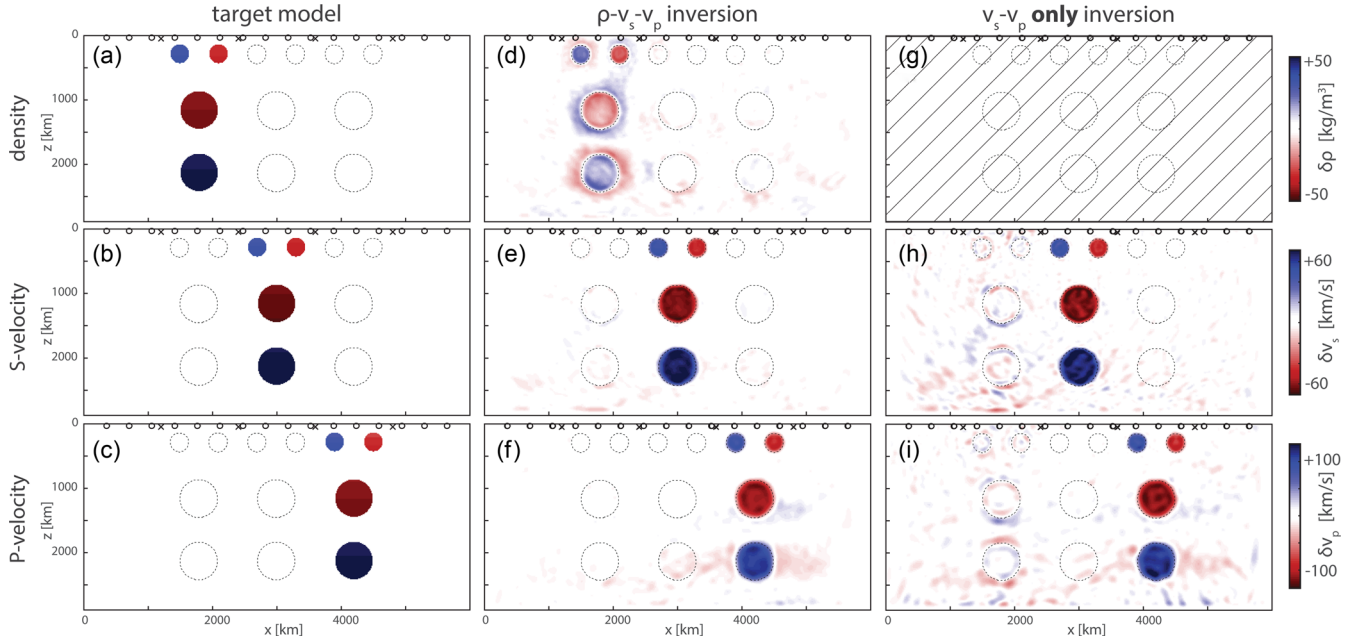


Figure 4. Results of tests 1 and 2, compared to the target model. (a)–(c) Target model. (d)–(f) Recovered models for the reference test (Section 4.1). All three parameters are recovered, although density is fainter than the others. (g)–(i) Recoverd models for the inversion where only S - and P -velocities are unconstrained (Section 4.2). Density remains fixed at PREM values, despite the fact that there is density structure in the target model. This density structure that is not accounted for results in small-scale artefacts in the other parameters.

4.2 Neglecting density

In the second test, only seismic velocity structure is updated, and potentially present density structure is ignored. Results for this test are shown in Figs 4(g)–(i). While the recovery of the seismic velocities is again very close to the target, there are now small-scale artefacts all throughout the model domain with a concentration in the lower mantle. In addition to that, circular artefacts appear at the locations of the edges of the actual density anomalies. These are especially apparent at the lowest frequencies and become narrower as the frequency content of the inversion is increased, but they do remain present (Figs 4h and i).

4.3 Scaling density to S velocity

In the third set of tests, we assess the effect of scaling density to S velocity in a fixed manner (eqs 11 and 12). Results are shown in Fig. 5. This is compared to a case where all parameters are allowed to update freely such that density is independent from S velocity. In this case, the three columns of differently scaled anomalies are all recovered to some extent. There are significant artefacts throughout the model domain, but the main shape, sign and relative strength of each of the three columns is visible. If on the other hand density is kept at a fixed scaling of $R_{\rho/S} = 0.2$, only the middle column is ‘correct’. The target model’s negative scaling in the rightmost column is completely overruled by the imposed $R_{\rho/S}$, and the scaling of the leftmost column is not strong enough compared to the target model. These violations of the actual model result in artefacts in P -velocity structure that are stronger than in the case where all parameters are free.

4.4 Initial velocity structure

The influence of prior information on P - and S -velocity structure on the recovery of density is shown in Fig. 6. The similarity between the

final density models (Figs 6a, d and g) shows that for the recovery of density, it matters little what information on S and P velocities is supplied. The most significant effect of including prior knowledge is that short-wavelength artefacts are reduced. Unsurprisingly, recovery of S and P velocities improves with the inclusion of prior information on these parameters, but in all cases the overall shapes are correct.

4.5 Constraints on velocity structure

The next set of tests describes the effect of fixing velocities to invert for density only. In contrast to the previous tests of Section 4.4, the velocity model has an enormous influence on the recovery of density if it is kept fixed. If the correct velocity model is taken, the recovery of density is indeed significantly better and faster than in the cases where all three parameters are allowed to vary (Fig. 7a, compare to Fig. 6g). This is a result of the fact that the seismic velocities, whose sensitivities are generally larger, are not allowed to vary.

This result, however, degrades quickly as soon as the velocity model does not have the exact target values. When the velocity model has anomalies of 75 per cent of their actual strength (Fig. 7b), numerous short-wavelength artefacts appear and missing S -velocity structure maps into the density model. At 50 per cent of the actual values, the recovered density structure is completely overshadowed by the artefacts and mapped S - and P -velocity structure.

4.6 Including gravity information

Fig. 8 shows the resulting final models for an inversion using only seismic information (a)–(c), using seismic and gravity vector data (d)–(f) and using seismic and geoid data (g)–(i). Surprisingly, the recovery of all three parameters decreases slightly when gravity data are added to the inversion. This effect is, however, clearest for




Figure 5. The effect of imposing a fixed scaling $R_{\rho/S}$ between density anomalies and S -velocity anomalies (see eq. 11 and Section 4.3). Models in density are shown in the large panels (a)–(c), while the corresponding S - and P -velocity models are shown in smaller panels to the right of these. (a) Target model in density. Each column has a different actual scaling $R_{\rho/S}$, that fall within the reported range of scalings between density and S -velocity (see, e.g. Karato & Karki 2001; Simmons *et al.* 2009). The target anomalies in S velocities are of the same size in each column (1 per cent just like in the reference model), while there are no target heterogeneities in P -velocity. (b) Recovered density model if all parameters are allowed to update freely. While numerous artefacts are present, the model is qualitatively similar to the target model. The recovered S -velocity model is very similar while there are some minor artefacts in the P -velocity model. (c) Recovered density model if it is scaled in a fixed manner to S -velocity with $R_{\rho/S} = 0.2$. Here, the target density structure of the left and right columns is completely overruled. Artefacts in P -velocity are now somewhat stronger but the recovered S -velocity model is similar to the previous test.

density. Despite the difference in kernel shape (Fig. 3), the density recovery of the full gravity vector case and the gravity potential case does not differ by much, although it is slightly worse using the smoother gravity potential information.

4.7 The effect of noise

As expected, the presence of data noise deteriorates the recovery of anomalies, as shown in Fig. 9. Noise levels are similar for correlated and uncorrelated noise, ranging from a maximum of ~ 10 per cent at the lowest frequencies to ~ 0.5 per cent at the highest frequencies (Figs 9j and k). These levels are intended to mimic real high-quality data commonly used in full-waveform inversion in those frequency ranges (e.g. Colli *et al.* 2013; Fichtner & Villaseñor 2015; Simutè *et al.* 2016).

There is not a large difference between the uncorrelated and correlated noise tests (Figs 9d and g). In both cases, the presence of data noise results in artefacts in the recovered models of all three parameters, but this effect is most pronounced for density. Here, the recovered anomalies are also most strongly affected. The shape of the artefacts is still relatively large because the noise level is more severe at the lower frequencies (Figs 9j and k). Recovery decreases

as the noise level increases. Because of the low receiver density (one per ~ 350 km), the effect of noise is relatively pronounced. When using more receivers, the results will improve, especially in the case with uncorrelated (random) noise. In this regard, our synthetic inversion results are somewhat conservative.

4.8 The effect of impedance

In our final test, a slightly different target model is used, in which density, S - and P -velocity anomalies lie on top of each other in such a manner that the impedance contrast is zero at the edges. Results of this are shown in Fig. 10. If there is no impedance contrast, the recovery of density is reduced (compare, e.g. to Figs 4d–f), but not removed. The recovery of the seismic velocities, too, is somewhat reduced.

5 DISCUSSION

In this study, we show that density heterogeneities can be recovered as an independent parameter in 2-D synthetic seismic waveform tomography. Our experiments are executed in a simple setup that

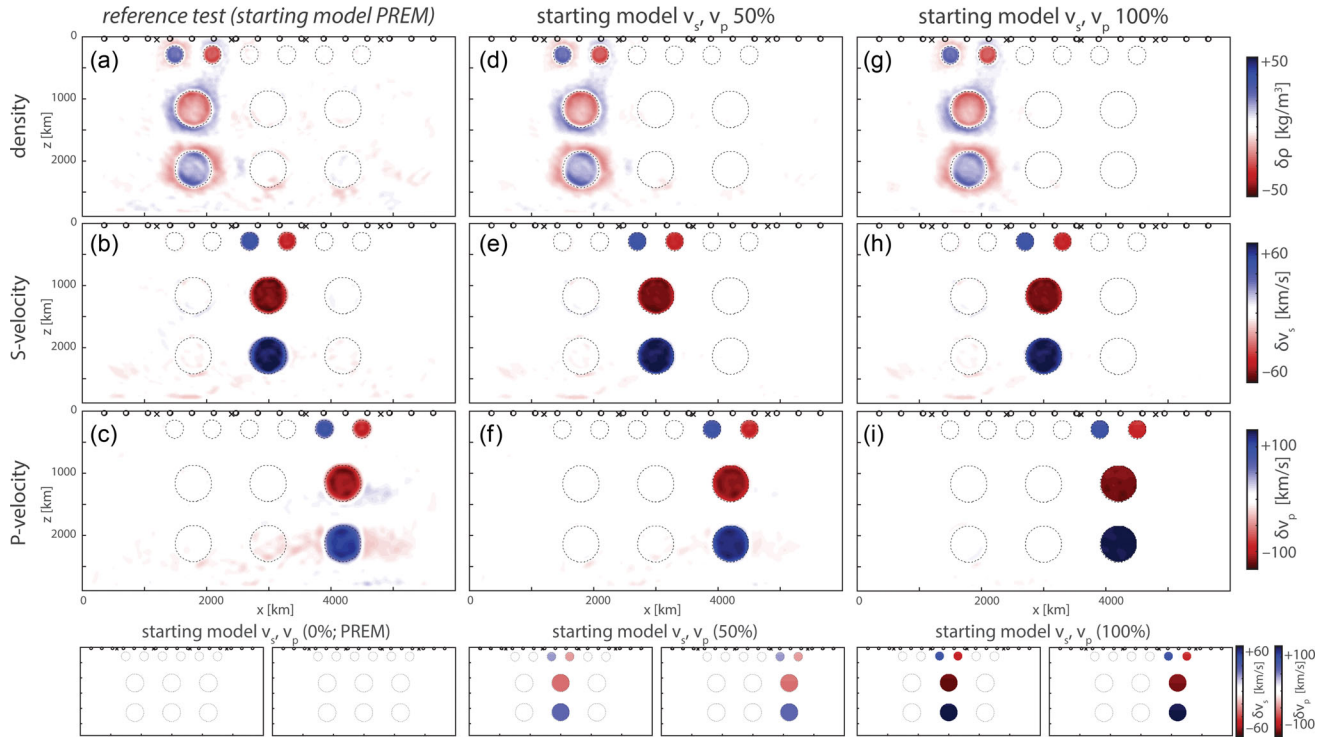


Figure 6. Recovered models for different starting models. (Section 4.4, shown for v_s and v_p in the small panels beneath each of the columns. The starting model for density is always PREM. The inclusion of extra *a priori* information on v_s and v_p does not have a great influence on the final model after 160 iterations.)

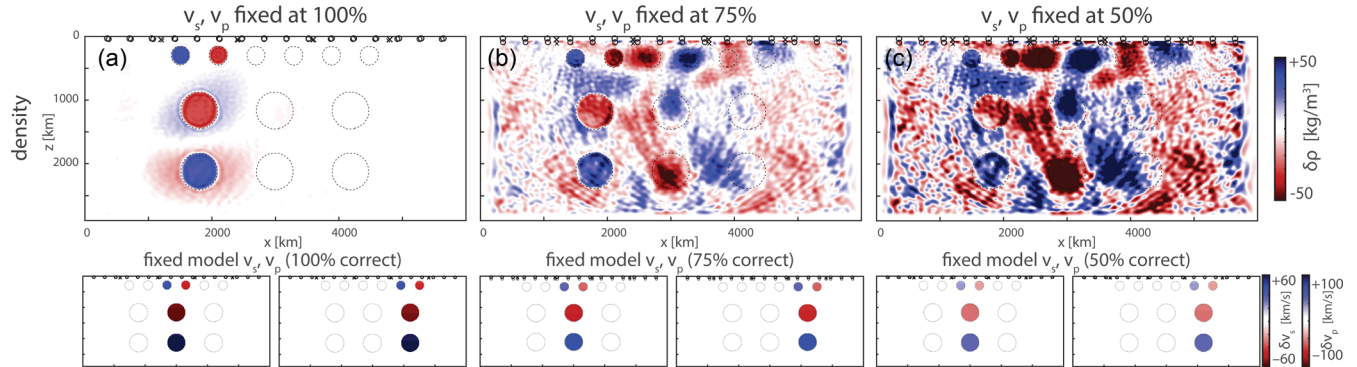


Figure 7. Recovered density models for different fixed S - and P -velocity models (Section 4.5). (a) If v_s and v_p are fixed at the actual values, the recovery of density is improved significantly compared to the reference case where all three parameters are free (Fig. 6g). (b) and (c) However, the more the S and P models deviate from the target model, the more artefacts are mapped into density, to the extent that these completely overpower the density signal.

does not take into account some factors that are known to influence the seismic wavefield (e.g. attenuation and uncertainties in sources and receivers). Nevertheless, we believe that our results provide good indications that the recovery of density may become possible in 3-D real-data inversions. This is relevant in order to more accurately quantify the forces driving plate tectonics and mantle convection and can help to distinguish between thermal and compositional heterogeneities. In the following discussion, we address key issues of density tomography using waveform inversion such as the effect of density anomalies on the wavefield, inversion strategies and data types used.

5.1 The effect of density on the seismic wavefield

Density contrasts create a scattered wavefield that travels both along with the direct wave and in other directions (Fig. 1). The influence

of density on the wavefield derives largely from impedance contrasts created by density contrasts. As a result, it is mainly the short-wavelength components of anomalies that are recovered in an inversion approach (Fig. 4). Inversion strategies aimed at recovering density thus rely crucially on the availability of low-frequency data. These should be used in a multiscale approach where the lowest frequencies are inverted for first (e.g. Bunks *et al.* 1995). This is in order to make sure that also the longer wavelength components of anomalies are recovered properly. While this is not a problem in regional- to global-scale seismology, this may be more difficult in the case of active-source experiments, where low frequencies are more difficult to obtain.

Frequencies should be chosen carefully so that no local minima are created when the inversion switches to a band containing higher frequencies. In the results presented here, this was achieved by increasing the maximum frequency by a factor of no more than 1.3 across frequency bands.

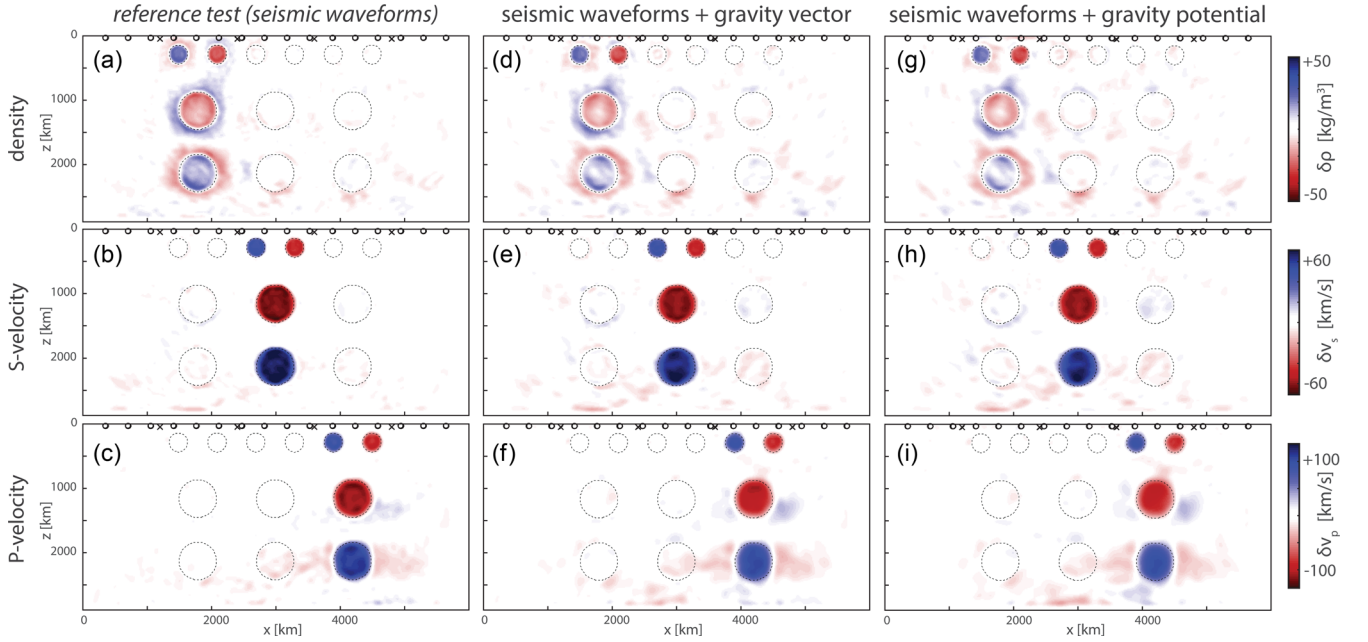


Figure 8. The effect of including gravity data into the inversion (Section 4.6). (a)–(c) Reference test. (d)–(f) Recovered model using seismic and gravity vector data. Artefacts seem slightly more pronounced than in the reference case that used only seismic data. (g)–(i) Recovered model using seismic and gravity potential data. Results are similar to those from the previous test with seismic and gravity vector data.

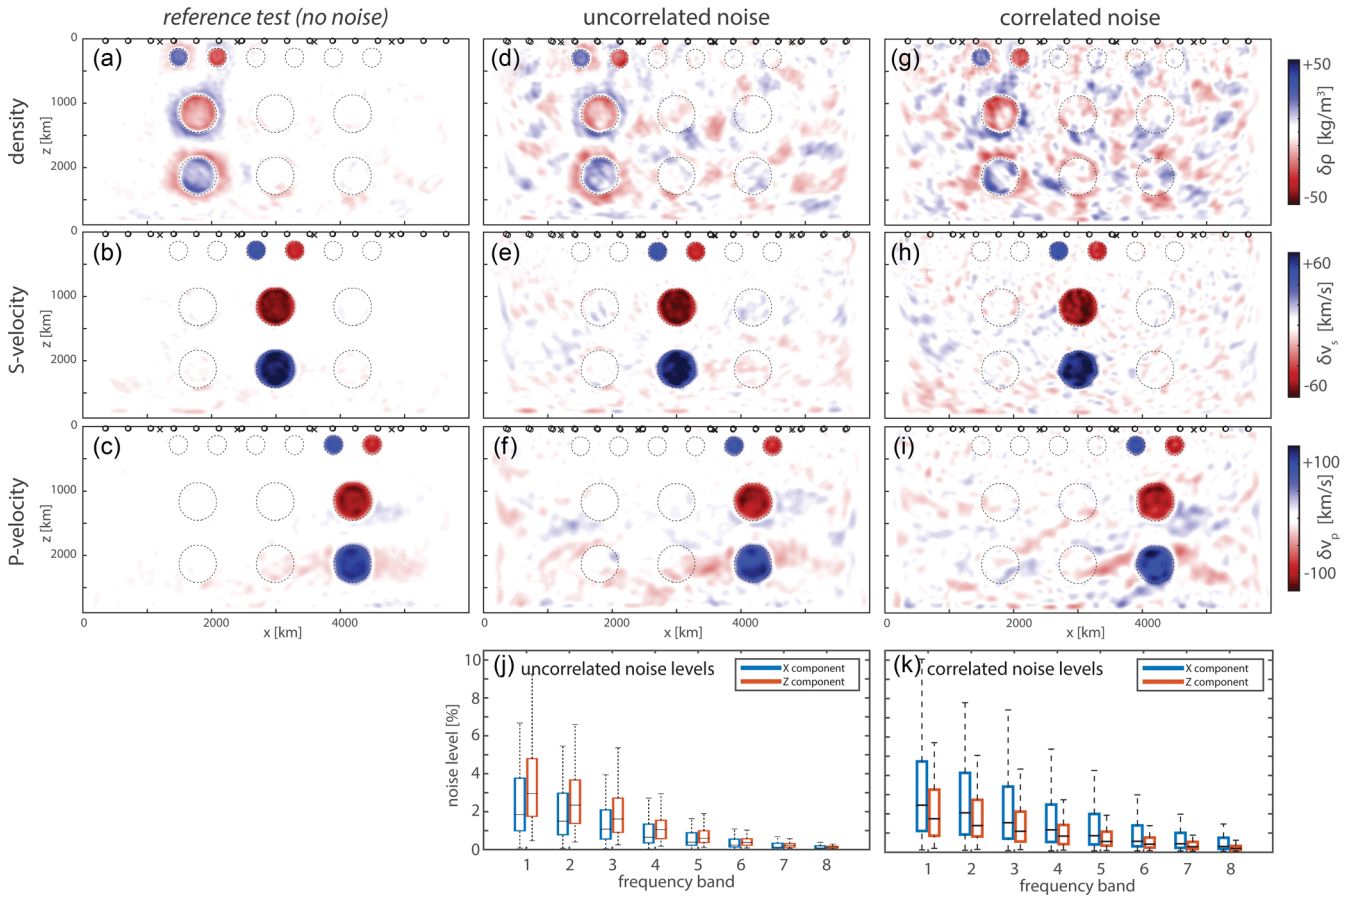


Figure 9. The effect of data noise on the recovery of model parameters (Section 4.7). (a)–(c) Reference test. (d)–(f) Recovered model for a test with *uncorrelated* noise with frequency-dependent amplitudes of about 0.5–10 per cent of the maximum data amplitudes (see panel j). The effect of noise is especially significant for density, the imprint of which on the wave propagation has low amplitude itself. (g)–(i) Recovered model for a test with *correlated* noise (see Section 3.7.2) with amplitudes of about 0.5–10 per cent of the maximum data amplitudes (see panel k). (j) Variation of random noise levels across frequencies. (k) Variation of correlated noise levels across frequencies.

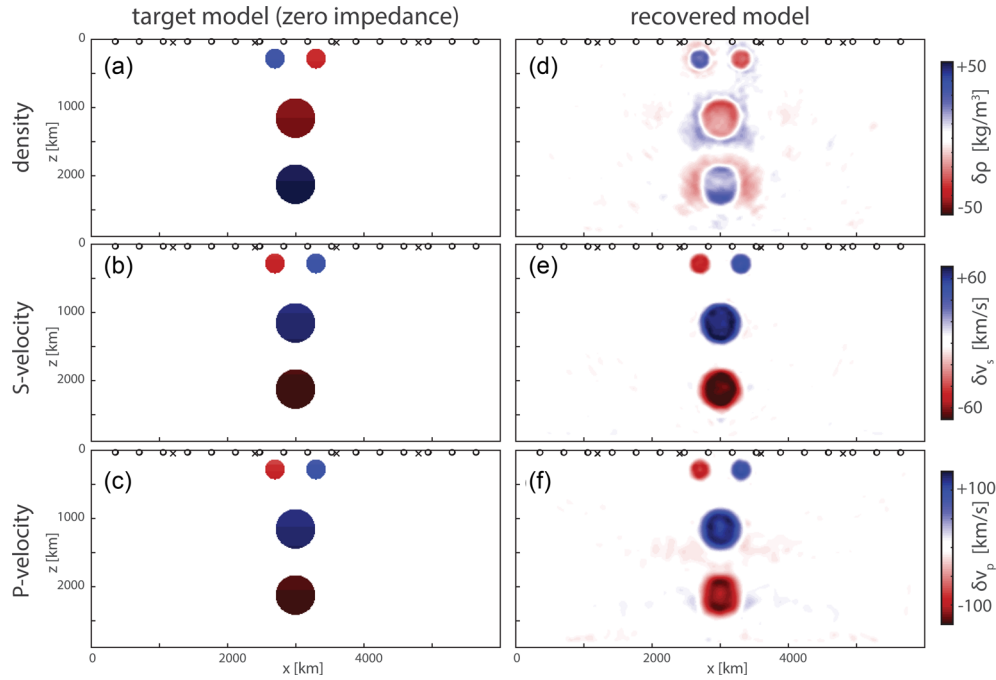


Figure 10. Results from Section 4.8, where the target model has a zero impedance contrast. (a)–(c) Target model. In contrast to all previous tests, the anomalies in density and S - and P -velocities lie in the same locations and are chosen such that the impedance contrasts across the anomaly edges are zero. (d)–(f) Recovered model. A lack of impedance contrast somewhat deteriorates the recovery of especially density (compare Figs 4d–f). The starting model was PREM.

While it is the effect on impedance contrast of density that has the most significant impact on the seismic wavefield, density can still be recovered even if the impedance contrast is zero (as a result of opposite contrasts in P and S velocities—see Fig. 10).

Common strategies in seismic tomography are to either ignore density altogether, or scale it to S velocity. In contrast to seismic traveltimes in the ray approximation, the full seismic waveforms do have a dependence on density. As a result, both ignoring and scaling density may lead to erroneous results. If ignored, the effect of density on the seismic wavefield may lead to artefacts in the recovery of S and P velocities (Fig. 4). The missing density structure, in particular the locations of density contrasts, is then mapped into S and P velocities, thus causing artefacts throughout the model domain. Similar effects occur if an incorrect scaling between density and S velocity is imposed (Fig. 5). Most problematically, however, potentially interesting density structure is completely overruled.

In the experiments presented here, we use an L_2 -norm waveform misfit functional, a metric that is most sensitive to the largest amplitude signals. As a result, it is not surprising that S -velocity structure is in general recovered best, while the density structure, whose effect on the wavefield is much more subtle, is recovered much more slowly.

Our measurements are sensitive to density partially because we include entire seismograms in the misfit functional. This means that also the scattered waves are included that are caused by the density anomaly (e.g. the reflected wave indicated in Fig. 1g). These are not normally picked up when one specifically studies the (traveltimes of) classical P , S and surface wave phases. Nevertheless, even if cross-correlation traveltimes differences are used to determine misfit of such ‘classical’ phases, measurable changes in phase occur as well (Plonka *et al.* 2016). This is because the density signal travelling with the direct wave alters the shape of the waveform. Similar results were obtained by Yuan *et al.* (2015).

The density effect on the wavefield is relatively small. Because of this, the effect of data noise is especially significant for the recovery of density (and stronger than on v_s and v_p —see Fig. 9). Nevertheless, even with a noise level of around 5 per cent at the lowest frequencies, the parameters are still recoverable to some extent. These are noise levels which are realistic in regional- to global-scale inversion setups, although it requires careful data selection and dense coverage. The noise results presented here can be viewed as conservative for these synthetic results, as we make use of only 16 receivers spread over a distance of 6000 km, resulting in an interstation distance of about 350 km. In realistic applications, the receiver coverage, although irregular, will be denser. Using more receivers will lead to more efficient cancelling of at the least the uncorrelated noise.

5.2 The influence of gravity data

When studying density in the Earth’s interior, gravity is a natural observable that is often employed (e.g. Kaban *et al.* 2004; Chaves & Ussami 2013; Panet *et al.* 2014; Fadel *et al.* 2015). In our results, however, the inclusion of gravity measurements does not improve the recovery of density compared to the seismic waveform-only inversions. This is despite the fact that gravity may have sensitivity to density that is additional and complementary to seismic sensitivity, which is mainly caused by density contrasts.

This is a consequence of the fact that solutions to inverse problems from potential fields such as gravity are physically non-unique and thus ill-posed: a given set of gravity measurements has an infinite variety of mass distributions that all explain it perfectly even when coverage is perfect and measurement errors are absent (Blakely 1995). In contrast, our inverse approach is deterministic, which means that a single direction of update is chosen. In principle, the inclusion of additional data should increase the curvature of the misfit functional near the optimum. Because of its non-uniqueness,

however, the addition of gravity data makes it more difficult to actually reach this minimum. The gravity component of the direction of update is one out of many, and thus unlikely to be in a direction that is favourable to convergence.

The non-uniqueness of potential-field inverse problems is physical, and would exist even if the data were perfect and continuously measured along the surface. It is also independent of the misfit functional that is chosen, and even in the perfect-data case, the solutions resulting from the different misfit functionals are simply different points in an infinitely large space of models that fit the gravity data exactly equally well.

The misfit functionals used in this study have a $1/r^2$ or $1/r$ dependence on distance from the receivers (eq. 14), which means that the gradients will also be strongest in the vicinity of the receivers, that is, at the surface (Fig. 3). However, the strength of the actual density anomalies we study does not have a $1/r^n$ dependence, so a misfit functional based on the gravity potential or gravity vector works to push the inversion into the wrong direction. The same is true for measurements of the gravity gradient, which would result in an even stronger concentration towards the surface. In general, there is no reason to assume that a $1/r^n$ density distribution is present in the actual Earth.

Of course, it is possible to pre-condition the gravity misfit such that the distance dependence is removed, but this does not remove the ill-posedness of the problem: it simply shifts the solution into a different direction in the infinite set of equally valid solutions. Including gravity without additional constraints thus slows down the inversion.

Several types of constraints can be included in order to reduce the non-uniqueness. The most generally taken approach when combining observables is to link density anomalies to seismic anomalies, for example, through some sort of scaling relation between the parameters (e.g. Maceira & Ammon 2009; Tondi *et al.* 2009; Lin *et al.* 2012; Chaves & Ussami 2013). This however excludes the interesting cases where velocities and density are not scaled—a decorrelation that might indicate the presence of chemical heterogeneity. In another variant, the shapes of the presumed density anomalies are chosen from *a priori* supplied (and often seismically obtained) models of the subsurface (e.g. Fadel *et al.* 2015). If this information is sufficiently reliable, this approach can work to reduce the non-uniqueness, but does not necessarily remove it completely. The problem with both these types of methods is that the use of a misfit based on any type of gravity data tends to overstress the contribution of near-surface heterogeneities. More problematically, these sorts of potentially incorrect prior knowledge may lead to incorrect results.

Other constraints include setting a maximum density contrast (which limits the extent to which the anomalies can get focused at the surface but does not remove the non-uniqueness) and a minimum-volume constraint. This latter constraint is, however, geologically not necessarily obvious: it is easy to think of geological situations in which masses are not expected to be concentrated into the smallest possible volume.

5.3 Other factors influencing the recovery of density

The effect of prior information on S and P velocities on the recovery of density is minimal. Even if these parameters are completely correct in the starting model, the recovery of density is not notably improved (see Fig. 6). This changes if the seismic velocities are then kept fixed and only density is allowed to update. The recovery

of density is very good if the velocities are fixed to the true model (Fig. 7)—significantly better than with v_S and v_P free. However, if the velocity model is not entirely correct, results deteriorate. This means that in practice, it will not be possible to keep velocities fixed, unless observables can be found that are uniquely sensitive to density. Finding such observables, for instance, using the optimal observable approach of Bernauer *et al.* (2014), will be the focus of future work.

Our method uses a simplified 2-D setup with only isotropic parameters and no attenuation, in order to study the first-order effects that density has on the seismic wavefield and on inversions. Especially attenuation should, however, be taken into account as well, as its effect on seismograms is in the same order of magnitude as that of density (Plonka *et al.* 2016). However, while attenuation only reduces the amplitudes of already-present seismic phases, density causes significant scattering in all directions (Fig. 1), thereby providing additional constraints. In our experiments, we furthermore assume that the sources are perfectly known, which will not be the case for real data in regional and global seismology.

Another aspect of our approach that must be kept in mind is that the amplitude information that we use in the L_2 -norm misfit may not be reliable for real data as it is susceptible to, for example, instrument and site effects. Nevertheless, as Fig. 1(g) shows, amplitude changes caused by a density anomaly may have a measurable effect on the phase and general waveform. Finally, additional arrivals caused by reflections can still be picked up as separate signals unconnected to the main phase.

6 CONCLUSIONS

In this study, we show that density can be recovered using waveform inversion on the global scale. Our study treats density as an independent parameter and does intentionally not introduce any prior constraints on its relation to seismic anomalies. This is relevant in order to be able to quantify the forces driving plate tectonics and mantle convection.

The effect of density on the seismic wavefield is small. This means that the presence of data noise affects the recovery of density more strongly than of S and P velocities. However, at noise levels to be expected on regional to global scales, density is still recoverable from the data. While the effect of density on the seismic wavefield lies mainly in impedance contrasts, density is still recoverable even if the impedance contrast is negligible.

We also show that the presence of density anomalies, if not accounted for, may negatively affect the recovery of other parameters. Furthermore, without strong additional constraints, the inclusion of gravity measurements into the misfit functional deteriorates the recovery of density and the other parameters. Finally, we show that prior information on the seismic velocities has little influence on the recovery of density, unless this prior model is fixed. However, the velocities can only be fixed if the velocity model is completely accurate, a situation is unlikely to occur in real-data cases. If the fixed velocity model is slightly wrong, errors therein severely map into density.

ACKNOWLEDGEMENTS

This research was funded by Netherlands Organisation for Scientific Research NWO grant under project no. 864.11.008. We want to thank the editor, Nathan Simmons and another anonymous reviewer, whose comments have really helped to greatly improve this

manuscript. We furthermore want to thank Agnieszka Plonka, Hanneke Paulssen, Jeannot Trampert, Andrew Valentine, Sanne Cottaar, Rein Baarsma, Theodoros Chatzivasileiadis, Jasper Leuven and Karin Sant for helpful discussions and suggestions. Both the finite-difference code used to simulate wave propagation and the optimization toolbox used for the inversions are available upon request from the authors.

REFERENCES

- Aki, K. & Richards, P., 2002. *Quantitative Seismology*, University Science Books.
- Bernauer, M., Fichtner, A. & Igel, H., 2014. Optimal observables for multi-parameter seismic tomography, *Geophys. J. Int.*, **198**, 1241–1254.
- Blakely, R.J., 1995. *Potential Theory in Gravity and Magnetic Applications*, Cambridge Univ. Press.
- Bozdağ, E., Trampert, J. & Tromp, J., 2011. Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements, *Geophys. J. Int.*, **185**, 845–870.
- Bunge, H.-P., Hagelberg, C.R. & Travis, B.J., 2003. Mantle circulation models with variational data assimilation: inferring past mantle flow and structure from plate motion histories and seismic tomography, *Geophys. J. Int.*, **152**, 280–301.
- Bunks, C., Saleck, F.M., Zaleski, S. & Chavent, G., 1995. Multiscale seismic waveform inversion, *Geophysics*, **60**, 1457–1473.
- Byrd, R.H., Lu, P., Nocedal, J. & Zhu, C., 1995. A limited memory algorithm for bound constrained optimization, *SIAM J. Sci. Comput.*, **16**(5), 1190–1208.
- Cerjan, C., Kosloff, D., Kosloff, R. & Reshef, M., 1985. A nonreflecting boundary condition for discrete acoustic and elastic wave equations, *Geophysics*, **50**, 705–708.
- Cerveny, V., 2001. *Seismic Ray Theory*, Cambridge Univ. Press.
- Chapple, W.M. & Tullis, T.E., 1977. Evaluation of the forces that drive the plates, *J. geophys. Res.*, **82**(14), 1967–1984.
- Chaves, C.A.M. & Ussami, N., 2013. Modeling 3-D density distribution in the mantle from inversion of geoid anomalies: application to the yellowstone province, *J. geophys. Res.*, **118**(12), 6328–6351.
- Colli, L., Fichtner, A. & Bunge, H.-P., 2013. Full waveform tomography of the upper mantle in the South Atlantic region: imaging westward fluxing shallow asthenosphere?, *Tectonophysics*, **604**, 26–40.
- Collis, S.S., Ghayour, K., Heinkenschloss, M., Ulbrich, M. & Ulbrich, S., 2001. Towards adjoint-based methods for aeroacoustic control, in *39th Aerospace Sciences Meeting and Exhibit*, January 8–11, 2001, Reno, Nevada, AIAA Paper, 2001–0821.
- Dahlen, F.A. & Tromp, J., 1998. *Theoretical Global Seismology*, Princeton Univ. Press.
- Dahlen, F., Hung, S.-H. & Nolet, G., 2000. Fréchet kernels for finite-frequency traveltimes—I. Theory, *Geophys. J. Int.*, **141**, 157–174.
- de Wit, R.W.L. & Trampert, J., 2015. Robust constraints on average radial lower mantle anisotropy and consequences for composition and texture, *Earth planet. Sci. Lett.*, **429**, 101–109.
- Dziewoński, A.M. & Anderson, D.L., 1981. Preliminary reference Earth model, *Phys. Earth planet. Inter.*, **25**, 297–356.
- Dziewoński, A.M., Hales, A.L. & Lapwood, E.R., 1975. Parametrically simple Earth models consistent with geophysical data, *Phys. Earth planet. Inter.*, **10**, 12–48.
- Fadel, I., van der Meijde, M., Kerle, N. & Lauritsen, N., 2015. 3D object-oriented image analysis in 3D geophysical modelling: analysing the central part of the East African Rift System, *Int. J. Appl. Earth Obs. Geoinf.*, **35**, 44–53.
- Fichtner, A., 2010. *Full Seismic Waveform Modelling and Inversion*, Springer.
- Fichtner, A. & Igel, H., 2008. Efficient numerical surface wave propagation through the optimization of discrete crustal models—a technique based on non-linear dispersion curve matching (DCM), *Geophys. J. Int.*, **173**, 519–533.
- Fichtner, A. & Trampert, J., 2011. Hessian kernels of seismic data functionals based upon adjoint techniques, *Geophys. J. Int.*, **185**, 775–798.
- Fichtner, A. & van Leeuwen, T., 2015. Resolution analysis by random probing, *J. geophys. Res.*, **120**, pp. 5549–5573.
- Fichtner, A. & Villaseñor, A., 2015. Crust and upper mantle of the western Mediterranean - Constraints from full-waveform inversion, *Earth planet. Sci. Lett.*, **428**, 52–62.
- Fichtner, A., Bunge, H.-P. & Igel, H., 2006. The adjoint method in seismology—I. Theory, *Phys. Earth planet. Inter.*, **157**, 86–104.
- Forsyth, D. & Uyeda, S., 1975. On the relative importance of the driving forces of plate motion, *Geophys. J. Int.*, **43**(1), 163–200.
- Frankel, A. & Clayton, R., 1986. Finite-difference simulations of seismic scattering: implications for the propagation of short period seismic waves in the crust and models of crustal heterogeneity, *J. geophys. Res.*, **91**, 6465–6489.
- Gauthier, O., Virieux, J. & Tarantola, A., 1986. Two-dimensional nonlinear inversion of seismic waveforms: numerical results, *Geophysics*, **51**, 1387–1403.
- Herceg, M., Artemieva, I. & Thybo, H., 2016. Sensitivity analysis of crustal correction for calculation of lithospheric mantle density from gravity data, *Geophys. J. Int.*, **204**(2), 687–696.
- 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.
- Ishii, M. & Tromp, J., 2001. Even-degree lateral variations in the Earth's mantle constrained by free oscillations and the free-air gravity anomaly, *Geophys. J. Int.*, **145**, 77–96.
- Ishii, M. & Tromp, J., 2004. Constraining large-scale mantle heterogeneity using mantle and inner-core sensitive normal modes, *Phys. Earth planet. Inter.*, **146**, 113–124.
- Iyer, H.M. & Hirahara, K., 1993. *Seismic Tomography, Theory and Practice*, Chapman & Hall.
- Jeong, W., Lee, H.-Y. & Min, D.-J., 2012. Full waveform inversion strategy for density in the frequency domain, *Geophys. J. Int.*, **188**(3), 1221–1242.
- Kaban, M., Schwintzer, P. & Reigber, C., 2004. A new isostatic model of the lithosphere and gravity field, *J. Geod.*, **78**(6), 368–385.
- Karato, S.-i. & Karki, B.B., 2001. Origin of lateral variation of seismic wave velocities and density in the deep mantle, *J. geophys. Res.*, **106**(B10), 21 771–21 783.
- Kato, M. & Kawakatsu, H., 2001. Seismological in situ estimation of density jumps across the transition zone discontinuities beneath Japan, *Geophys. Res. Lett.*, **28**, 2541–2544.
- Kennett, B.L.N., 1998. On the density distribution within the Earth, *Geophys. J. Int.*, **132**, 374–382.
- Koelemeijer, P.J., Deuss, A. & Ritsema, J., 2017. Density structure of Earth's lowermost mantle from Stoneley mode splitting observations, *Nat. Comm.*, **8**, 15241, doi:10.1038/ncomms15241.
- Kuo, C. & Romanowicz, B., 2002. On the resolution of density anomalies in the Earth's mantle using spectral fitting of normal mode data, *Geophys. J. Int.*, **150**, 162–179.
- Lin, F.-C., Schmandt, B. & Tsai, V.C., 2012. Joint inversion of Rayleigh wave phase velocity and ellipticity using USArray: constraining velocity and density structure in the upper crust, *Geophys. Res. Lett.*, **39**, doi:10.1029/2012GL052196.
- Liu, L. & Gurnis, M., 2008. Simultaneous inversion of mantle properties and initial conditions using an adjoint of mantle convection, *J. geophys. Res.*, **113**, B08405, doi:10.1029/2007JB005594.
- Liu, Q. & Tromp, J., 2008. Finite-frequency sensitivity kernels for global seismic wave propagation based upon adjoint methods, *Geophys. J. Int.*, **174**, 265–286.
- Luo, Y. & Schuster, G.T., 1991. Wave-equation traveltime inversion, *Geophysics*, **56**, 645–653.
- Maceira, M. & Ammon, C.J., 2009. Joint inversion of surface wave velocity and gravity observations and its application to central asian basins shear velocity structure, *J. geophys. Res.*, **114**(B2), doi:10.1029/2007JB005157.
- Mora, P., 1987. Nonlinear two-dimensional elastic inversion of multioffset seismic data, *Geophysics*, **52**, 1211–1228.

- Mosca, I., Cobden, L., Deuss, A., Ritsema, J. & Trampert, J., 2012. Seismic and mineralogical structures of the lower mantle from probabilistic tomography, *J. geophys. Res.*, **117**, doi:10.1029/2011JB008851.
- Moulik, P. & Ekström, G., 2016. The relationships between large-scale variations in shear velocity, density, and compressional velocity in the earth's mantle, *J. geophys. Res.*, **121**(4), 2737–2771.
- Nataf, H.C. & Ricard, Y., 1996. 3SMAC: an a priori tomographic model of the upper mantle based on geophysical modelling, *Phys. Earth. planet. Inter.*, **95**, 101–122.
- Nocedal, J. & Wright, S., 2006. *Numerical Optimization*, Springer Science & Business Media.
- Panet, I., Pajot-Métivier, G., Greff-Lefftz, M., Métivier, L., Diamant, M. & Mandea, M., 2014. Mapping the mass distribution of Earth's mantle using satellite-derived gravity gradients, *Nat. Geosci.*, **7**(2), 131–135.
- Plonka, A., Blom, N. & Fichtner, A., 2016. The imprint of crustal density heterogeneities on regional seismic wave propagation, *Solid Earth*, **7**(6), 1591–1608.
- Prieux, V., Brossier, R., Operto, S. & Virieux, J., 2013. Multiparameter full waveform inversion of multicomponent ocean-bottom-cable data from the Valhall field. Part 1: Imaging compressional wave speed, density and attenuation, *Geophys. J. Int.*, **194**, 1640–1664.
- Resovsky, J. & Ritzwoller, M., 1999. Regularisation uncertainty in density models estimated from normal mode data, *Geophys. Res. Lett.*, **26**, 2319–2322.
- Resovsky, J. & Trampert, J., 2002. Reliable mantle density error bars: an application of the Neighbourhood Algorithm to normal-mode and surface wave data, *Geophys. J. Int.*, **150**(3), 665–672.
- Ricard, Y., Richards, M., Lithgow-Bertelloni, C. & Le Stunff, Y., 1993. A geodynamic model of mantle density heterogeneity, *J. geophys. Res.*, **98**(B12), 21 895–21 909.
- Rickers, F., Fichtner, A. & Trampert, J., 2012. Imaging mantle plumes with instantaneous phase measurements of diffracted waves, *Geophys. J. Int.*, **190**, 650–664.
- Root, B.C., Ebbing, J., van der Wal, W., England, R.W. & Vermeersen, L.L.A., 2017. Comparing gravity-based to seismic-derived lithosphere densities: a case study of the British Isles and surrounding areas, *Geophys. J. Int.*, **208**(3), 1796–1810.
- Shearer, P.M. & Flanagan, M.P., 1999. Seismic velocity and density jumps across the 410- and 660-kilometer discontinuities, *Science*, **285**(5433), 1545–1548.
- Simmons, N.A., Forte, A.M. & Grand, S.P., 2009. Joint seismic, geodynamic and mineral physical constraints on three-dimensional mantle heterogeneity: implications for the relative importance of thermal versus compositional heterogeneity, *Geophys. J. Int.*, **177**(3), 1284–1304.
- Simmons, N.A., Forte, A.M., Boschi, L. & Grand, S.P., 2010. GyPSuM: a joint tomography model of mantle density and seismic wave speeds, *J. geophys. Res.*, **115**, B12310, doi:10.1029/2010JB007631.
- Simutè, S., Steptoe, H., Cobden, L., Gokhberg, A. & Fichtner, A., 2016. Full-waveform inversion of the Japanese Islands region, *J. geophys. Res.*, **121**(5), 3722–3741.
- Takeuchi, H. & Saito, M., 1972. Seismic surface waves, *Methods Comput. Phys.*, **11**, 217–295.
- Tanimoto, T., 1991. Waveform inversion for three-dimensional density and S wave structure, *J. geophys. Res.*, **96**, 8167–8189.
- Tarantola, A., 1986. A strategy for nonlinear elastic inversion of seismic reflection data, *Geophysics*, **51**, 1893–1903.
- Tarantola, A., 1988. Theoretical background for the inversion of seismic waveforms, including elasticity and attenuation, *Pure appl. Geophys.*, **128**, 365–399.
- Tarantola, A., 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation*, 2nd edn, SIAM, Philadelphia.
- Tondi, R., Achauer, U., Landes, M., Davi, R. & Besutiu, L., 2009. Unveiling seismic and density structure beneath the Vrancea seismogenic zone, Romania, *J. geophys. Res.*, **141**, doi:10.1029/2008JB005992.
- Trampert, J., Deschamps, F., Resovsky, J. & Yuen, D., 2004. Probabilistic tomography maps chemical heterogeneities throughout the lower mantle, *Science*, **306**, 853–856.
- Turcotte, D.L. & Schubert, G., 2014. *Geodynamics*, Cambridge Univ. Press.
- Virieux, J., 1984. SH wave propagation in heterogeneous media: velocity-stress finite difference method, *Geophysics*, **49**, 1933–1942.
- Virieux, J., 1986. P-SV wave propagation in heterogeneous media: velocity-stress finite difference method, *Geophysics*, **51**, 889–901.
- Warners-Ruckstuhl, K.N., Govers, R. & Wortel, R., 2012. Lithosphere–mantle coupling and the dynamics of the Eurasian Plate, *Geophys. J. Int.*, **189**(3), 1253–1276.
- Woodhouse, J.H. & Dziewoński, A.M., 1984. Mapping the upper mantle: three-dimensional modeling of Earth structure by inversion of seismic waveforms, *J. geophys. Res.*, **89**, 5953–5986.
- Woodhouse, J.H. & Deuss, A.F., 2007. Theory and observations Earth's free oscillations, in *Treatise on Geophysics*, Elsevier, doi:10.1016/B978-044452748-6.00002-X.
- Wu, R. & Aki, K., 1985. Scattering characteristics of elastic waves by an elastic heterogeneity, *Geophysics*, **50**, 582–595.
- Yuan, Y.O., Simons, F.J. & Bozdağ, E., 2015. Multiscale adjoint waveform tomography for surface and body waves, *Geophysics*, **80**(5), R281–R302.

SUPPORTING INFORMATION

Supplementary data are available at [GJIRAS](https://doi.org/10.1016/j.gjiras.2017.05.001) online.

Figure 1. The effect of inversion parametrization on the recovery of density alongside seismic velocities for the ‘reference’ target model (given in Fig. 2).

Figure 2. The effect of inversion parametrization for a target model in which density is scaled to S velocity with different scalings. This is the same model as used in the ‘Scaling relations’ experiment of our paper (Fig. 5).

Figure 3. The effect of increased anomaly strength on the recovery of density and seismic velocities from waveform inversion using an L2-norm. As the anomalies are ten times the amplitude of those in the reference test, the colour scale is similarly adapted.

Figure 4. Reproduction of Fig. 4 in the paper.

Supplementary data in the form of extra experiments on anomaly strength and inversion parametrisation, as well as a video of wave propagation displayed in Figure 1 are available at GJIRAS online. All synthetic data used in this study as well as the code used to generate them can be found at doi.org/10.5281/zenodo.268929 (Blom 2017).

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the paper.

APPENDIX: L-BFGS

Here follows a brief description of the L-BFGS algorithm as used in our inversions.

In order to compute the L-BFGS update \mathbf{s}_i in eq. (4), we require the models and gradients from the last k iterations to construct

the approximation of the inverse Hessian. Note that the matrix \mathbf{B}_i is never constructed explicitly, but we only compute its action on the negative gradient by using a two-loop recursion outlined in Algorithm 1. The implementation mostly follows algorithm 7.4 in Nocedal & Wright (2006). However, there are two important modifications to the classical textbook implementation. First, the optimization variables represent discretized continuous fields (e.g. density) evaluated at the gridpoints of the numerical mesh. It is important to use an inner product in a gradient-based descent algorithm that takes these spatial dependencies into account, because the choice of inner product can largely influence the convergence behaviour (e.g. Collis *et al.* 2001). In particular, using the standard Euclidean inner product often results in badly scaled gradient directions. Hence, we use a diagonal mass matrix \mathbf{D} that incorporates the spatial grid sizes dx and dz of the finite-difference discretization for all vector–vector multiplications. Second, any smoothing or depth scaling should not be applied to the gradient itself, because this would result in inconsistent search directions. Instead, we can either incorporate those operations into the inner product as well or define a scaling and smoothing matrix \mathbf{S} that is used as L-BFGS initialization matrix, which is done in line 13 of Algorithm 1.

Algorithm 1 L-BFGS search direction

```

1: Input:      models  $\mathbf{m}_i, \dots, \mathbf{m}_{i-k}$       and      gradients
                $\nabla J(\mathbf{m}_i), \dots, \nabla J(\mathbf{m}_{i-k})$ 
2: Output: search direction  $\mathbf{s}$  to update the model
3: for  $j = i - k$  to  $i - 1$  do
4:    $\mathbf{s}_j \leftarrow \mathbf{m}_{j+1} - \mathbf{m}_j$ 
5:    $\mathbf{y}_j \leftarrow \nabla J(\mathbf{m}_{j+1}) - \nabla J(\mathbf{m}_j)$ 
6:    $\rho_j \leftarrow (\mathbf{y}_j^T \mathbf{D} \mathbf{s}_j)^{-1}$ 
7: end for
8:  $\mathbf{q} \leftarrow -\nabla J(\mathbf{m}_i)$ 
9: for  $j = i - 1$  to  $i - k$  do
10:   $\alpha_j \leftarrow \rho_j \mathbf{s}_j^T \mathbf{D} \mathbf{q}$ 
11:   $\mathbf{q} \leftarrow \mathbf{q} - \alpha_j \mathbf{y}_j$ 
12: end for
13:  $\mathbf{s} \leftarrow (\mathbf{s}_{j-1}^T \mathbf{D} \mathbf{y}_{j-1})(\mathbf{y}_{j-1}^T \mathbf{D} \mathbf{y}_{j-1})^{-1} \mathbf{S} \mathbf{q}$ 
14: for  $j = i - k$  to  $i - 1$  do
15:   $\beta \leftarrow \rho_j \mathbf{y}_j^T \mathbf{D} \mathbf{s}$ 
16:   $\mathbf{s} \leftarrow \mathbf{s} + (\alpha_j - \beta) \mathbf{s}_j$ 
17: end for
18: return  $\mathbf{s}$ 

```
