JGR Solid Earth

RESEARCH ARTICLE
10.1029/2023JB026353

A Bayesian Formulation for Estimating the Composition of Earth's Crust

Gailin Pease[1], **Anne Gelb**[2], **Yoonsang Lee**[2], **and Brenhin Keller**[1]

[1]Department of Earth Sciences, Dartmouth College, Hanover, NH, USA, [2]Department of Mathematics, Dartmouth College, Hanover, NH, USA

Key Points:

- We use a Bayesian formulation to estimate the composition of Earth's crust with error bounds
- The estimated composition is very sensitive to the prior distribution of compositions
- Our upper crust estimate is more mafic than those found by previous work

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

G. Pease and B. Keller,
gailin.l.pease@dartmouth.edu;
brenhin.keller@dartmouth.edu

Citation:

Pease, G., Gelb, A., Lee, Y., & Keller, B. (2023). A Bayesian formulation for estimating the composition of Earth's crust. *Journal of Geophysical Research: Solid Earth*, *128*, e2023JB026353. https://doi.org/10.1029/2023JB026353

Received 1 JAN 2023
Accepted 17 JUL 2023

Author Contributions:

Conceptualization: Gailin Pease, Brenhin Keller
Data curation: Brenhin Keller
Formal analysis: Gailin Pease, Anne Gelb, Yoonsang Lee
Investigation: Gailin Pease
Methodology: Gailin Pease
Software: Gailin Pease, Brenhin Keller
Supervision: Brenhin Keller
Validation: Gailin Pease
Visualization: Gailin Pease
Writing – original draft: Gailin Pease
Writing – review & editing: Anne Gelb, Yoonsang Lee, Brenhin Keller

Abstract

Due to the inaccessibility of Earth's deep interior, geologists have long attempted to estimate the composition of the continental crust from its seismic properties. Despite numerous sources of error including nonuniqueness in the mapping between composition and seismic properties, the corresponding uncertainties have typically been estimated qualitatively at best. We propose a Bayesian approach that uses mineralogical modeling to combine prior knowledge about the composition of the crust with seismic data to give a posterior distribution of the predicted composition at any location, combined with a Monte Carlo simulation to estimate the average composition of the Earth's crust. Our approach yields an estimated composition of 59.5% silica in the upper crust (90% credible interval 58.9 %–60.1%), 57.9% in the middle crust (90% credible interval 57.2%–58.6%), and 53.6% in the lower crust (90% credible interval 53.0%–54.2%). Our estimate exhibits less compositional stratification over depth and a more intermediate composition in the upper and middle crust than previous estimates. Testing our approach on a simulated crust reveals the importance of prior assumptions in estimating the composition of the crust from its seismic properties, and suggests that future work should focus on quantifying those assumptions.

Plain Language Summary

The composition of the continental crust is important for understanding Earth's evolution on global and regional scales. We build on previous work estimating Earth's composition with a new approach using Bayesian statistics, where prior information from rocks sampled at Earth's surface is combined with information from seismic properties to estimate Earth's crustal composition at depth. To connect seismic properties and compositions, we model the seismic properties of rock compositions. Our approach provides rigorous uncertainty estimates for crust composition at depth. By quantifying the sources of this uncertainty, we are able to propose a path forward for future work to further improve compositional estimates for Earth's crust at depth.

1. Introduction

Earth's continents are important for understanding many Earth processes. They contribute to the evolution of the biosphere and climate via the negative feedback between atmospheric CO_2 concentration and terrestrial silicate weathering which acts to stabilize Earth's climate on geologic timescales (J. C. Walker et al., 1981). The composition of the continental crust is key to understanding crust formation and evolution, as it informs the mass-balance constraints on the relative roles of processes such as arc and rift magmatism, fractional crystallization, relamination, delamination, and erosion (Alonso-Perez et al., 2009; Grove et al., 2003; Hacker et al., 2011; Kay & Kay, 1993; Keller et al., 2015; Liu et al., 2001). In addition, crustal composition directly informs our understanding of the partitioning of Earth's overall budget of major and trace elements, including inorganic nutrients such as phosphorous which are critical for the surface biosphere (T. Walker & Syers, 1976), and heat-producing elements like K, Th, and U which influence the thermal stability of the crust (Mareschal & Jaupart, 2013).

Many previous studies have attempted to estimate the composition of Earth's crust. Perhaps the most widely cited review, and the estimate of the composition of the continental crust that is most widely used today, is that of Rudnick and Gao (2014), shown beside other estimates in Figure 1. Estimates of Earth's upper crustal composition can use surveys of surface compositions, but estimates in the lower crust must rely on indirect data (Rudnick & Gao, 2014). The seismic properties of the Earth's crust and present-day surface heat flow are the physical observations most often used to inform estimates of the composition of the deep crust.

Such estimates are possible because a rock's composition influences its seismic properties (Behn & Kelemen, 2003; Christensen, 1996). However, this relationship is not unique. Very different compositions can have similar seismic

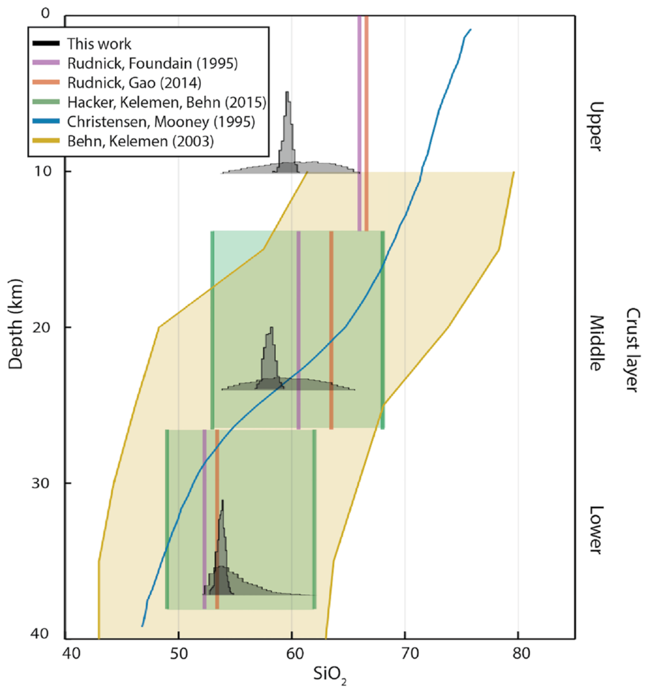

Figure 1. Previous estimates of weight % SiO_2 in the continental crust to 40 km depth. Upper, middle, and lower crust boundaries are the average boundary depths from `Crust1.0` Laske et al. (2013). Upper, middle, and lower crust estimates are shown for Rudnick and Fountain (1995) (purple) and Rudnick and Gao (2014) (orange). The most felsic and most mafic models for the middle and lower crust found by Hacker et al. (2015) are in green; they do not model the upper crust. The depth-composition relationship of Christensen and Mooney (1995) is shown in blue. The high and low weight % SiO_2 bounds of Behn and Kelemen (2003) are shown in yellow; they do not estimate surface composition. The results of this work (gray histograms) are shown for the upper, middle, and lower crust (specific depth of layers depends on location). Our results are shown using a surface distribution (narrower distribution, reflective of seismic information and prior knowledge about compositions exposed at the surface) and an uninformative prior (wider distribution, reflective of seismic information alone).

properties, and a single composition can have different seismic properties depending on formation conditions, alteration or cracking, and the current temperature and pressure conditions of the rock.

Estimates using seismic properties use one of two approaches to manage this inherent nonuniqueness: either the implicit incorporation of prior knowledge (Christensen & Mooney, 1995; Rudnick & Fountain, 1995), resulting in a single composition estimate without uncertainty, or exclusion of prior knowledge (Behn & Kelemen, 2003; Hacker et al., 2015), resulting in a wide range of possible compositions with no estimate of their relative likelihoods. Implementation of both approaches usually lacks a rigorous statistical framework, instead using a variety of relatively ad-hoc algorithms to match seismic properties to compositions.

Hacker et al. (2015) is one approach which is agnostic to prior information. They identify the most mafic and felsic compositions that could possibly form the middle and lower crust, producing wide boundaries of possible crust composition but no guidance for identifying what within those bounds is most likely. At the other extreme, Rudnick and Fountain (1995)identify overlapping bands of seismic properties for each composition they consider for the lower crust. They implicitly incorporate prior information when choosing from among the multiple compositions which match the seismic data, where they use qualitative ideas about which compositions are more likely to occur at depth. For a full review of existing approaches, see Huang et al. (2013).

In many previous approaches, prior knowledge is incorporated into estimates of the crust's composition at depth via observations of exposed rocks and assumptions about which exposures are representative of compositions at depth. Although these assumptions add valuable information to the estimates they produce, their qualitative nature makes it difficult to determine the extent to which they influence the resulting estimates. Because results using prior knowledge are able to estimate narrow ranges of compositions compared to the wide range of compositions allowed by the seismic properties (Hacker et al., 2015), we infer that the prior influences the accuracy and precision of resulting estimates. Our method allows us to interrogate the extent to which prior assumptions determine resulting composition estimates.

Some attempts have been made toward a probabilistic approach to this problem. Sammon et al. (2022). use a distribution of surface exposed granulite terrains and xenoliths in the Southwestern United States as the range of possible compositions at depth in the region, and use a probabilistic framework to compare the modeled seismic properties of those granulites to the seismic properties of the crust. However, the distribution of surface granulites was not treated as a true prior distribution, nor was the relationship between this prior and the resulting estimate formally expressed.

Bayesian statistics provides a formal approach for combining uncertain knowledge with uncertain observations, exactly the problem we face in estimating the composition of the deep crust. We begin by assuming that compositions at Earth's surface provide a prior for the compositions we might find at depth, and combine that prior with seismic data describing the deep crust to form estimates of crust composition. Our approach addresses the limitations of previous work by providing an estimate of the composition of the crust with meaningful uncertainty, and by providing a way to explicitly include prior information or assumptions about the composition of the crust.

In Section 2 we provide a general overview of our approach, including how we formulate crust composition estimates as a Bayesian problem. We then provide details on each key step of our approach in Section 3, including tests we did to compare the robustness of our approach to key methodological choices. Finally, in Section 4 we present our results along with their geologic implications.

10.1029/2023JB026353

2. Problem Formulation: A Bayesian Approach

Using the Bayes rule, the posterior distribution of compositions $\mathbb{P}(\vec{c}_l|\vec{s}_l)$ at a location l with seismic properties $\vec{s}_l$ is:

$$\mathbb{P}(\vec{c}_l|\vec{s}_l) = \frac{\mathbb{P}(\vec{s}_l|\vec{c}_l)\mathbb{P}(\vec{c}_l)}{\mathbb{P}(\vec{s}_l)} \tag{1}$$

This formulation illuminates the centrality of the two sources of information that will drive the discussion in much of this paper: first, the relationship of a composition to its seismic properties, and second, the prior distribution $\mathbb{P}(\vec{c}_l)$.

The physical properties $\vec{s}_l$ could include any property of the crust measurable from the surface. Since they are readily modeled from surface measurements of wave arrival times (Pasyanos et al., 2014; N. A. Simmons et al., 2021), here we consider seismic properties, specifically the pressure wave velocity (V_p), the ratio of pressure to shear wave velocity (V_p/V_s), and the density (ρ). We note that V_p is the seismic property most commonly used to predict composition (Rudnick & Gao, 2014), while V_p/V_s can differentiate between rock types with similar V_p (Christensen, 1996), and has been used along with V_p to predict composition with crust seismic properties (Hacker et al., 2015; Holbrook et al., 1992).

This framework allows the use of any distribution, either empirical or theoretical, of compositions as a prior. Here, we use the distribution of compositions observed at the surface, collected in the `Earthchem` database and compiled by (Keller & Schoene, 2012) as the prior distribution of compositions in the upper, middle, and lower crust. Since the same processes of magmatic differentiation are responsible for the compositional diversity of crustal rocks both at the surface and at depth, features of the surface distribution can inform us about features of the distribution of compositions at depth (Alonso-Perez et al., 2009; Grove et al., 2003; Keller et al., 2015). For example, the relative scarcity of intermediate igneous rocks compared to mafic and felsic rocks, a widely observed phenomenon sometimes called the Daly Gap (Chayes, 1963; Daly, 1925; Dufek & Bachmann, 2010; Keller et al., 2015; Keller & Harrison, 2020), is present in our prior (Figure S1 in Supporting Information S1).

Some previous investigations use databases of rocks limited to those thought to originate in the lower crust to estimate lower crustal composition (Huang et al., 2013; Rudnick & Fountain, 1995). The exclusion of many surface compositions is supported by arguing based on geological evidence that some rocks found at the surface are unlikely to make up most of the crust at depth. To test the effect of this assumption, we compare our results using the `Earthchem` prior to results using a prior from Huang et al. (2013) (Figure S1 in Supporting Information S1).

Figure S2 in Supporting Information S1 shows the major data products and key computational steps of our proposed approach. The process begins with the calculation of the empirical prior distribution from a large database of igneous and metaigneous rocks (Keller & Schoene, 2012), which we resample to correct for sampling biases. We then model the seismic properties of each composition $\vec{c}$ in the prior (Figure S3 in Supporting Information S1) using the thermodynamic software `Perple_X` at temperature and pressure conditions θ, described below, and porosity and exhumation conditions Φ, described in Section 3.3:

$$f_{\vec{s}}(\vec{c}; \theta, \Phi) \tag{2}$$

The parameters θ and ϕ are important because of the dependence of seismic properties on a rock's current temperature and pressure and on physical alteration since the rock's formation, including porosity and exhumation. For temperature and pressure, included in θ, we use estimates based on geotherms, discussed further in Section 3.3. Unlike θ, Φ includes parameters which affect seismic properties but which are not well constrained across the crust, including porosity and exhumation. Our choices for these parameters, including a sensitivity analysis, are discussed further in Section 3.3. For simplicity, we drop the parameters θ and Φ when writing $f_{\vec{s}}(\vec{c})$ in the remainder of this paper.

We formulate the likelihood $\mathbb{P}(\vec{s}_l|\vec{c}_l)$ through a noise model between the observed properties $\vec{s}_l$ and the modeled properties $f_{\vec{s}}(\vec{c}_l)$, where increasing differences between $\vec{s}_l$ and $f_{\vec{s}}(\vec{c}_l)$ result in decreasing likelihood described by $f_{\vec{0},\Sigma}(f_{\vec{s}}(\vec{c}_l) - \vec{s}_l)$, derived in Section 3.4. The properties of the crust, $\vec{s}_l$, include its seismic velocities and density and are provided by the `SPiRaL` and `Litho1.0` models (Artemieva, 2006; Laske et al., 2013; N. A. Simmons et al., 2021).

Journal of Geophysical Research: Solid Earth 10.1029/2023JB026353

Implicit in the comparison of modeled seismic properties to the properties of the crust is the assumption that averaging of seismic properties and compositions is commutative. That is, we assume that the average seismic properties of a volume of crust, which is the aggregate of many compositions, is the same as the seismic properties of the average composition of that volume of crust. A reasonable assumption is that this holds for randomly or uniformly mixed compositions, where mineral modes are mixed through a heterogeneous section of crust similarly to the way they are spread through a homogeneous composition which itself is formed of mixed minerals. However, there are two common cases where it may not hold. First, at some pressures, temperatures, and compositions, different mineral modes may be present and stable in two separate compositions than if those compositions were mixed. Second, some patterns of heterogeneity might themselves induce different seismic behavior; for example, the repeated layering of compositions with different wave speeds. As in previous investigations, our work assumes that composition averaging is commutative, but future work could explore the limits of this assumption.

Since there is no analytical solution for the posterior $\mathbb{P}\left(\vec{c}_l \,|\, \vec{s}_l\right)$ at each location, we use rejection sampling to sample from the posterior distribution (Smith & Gelfand, 1992) (Section 3.5). Using samples from location-specific posterior distributions, we use a Monte Carlo simulation [gp]for to calculate our global mean composition estimate (Section 3.6). We first randomly select N uncorrelated locations on Earth's continental crust. We sample a composition from the posterior at each location and compute the average of these samples. Repeating this process provides the empirical distribution of estimates of the global average crustal composition for each crustal layer.

Because our approach is analytical, we are able to test it by estimating the composition of a synthetic model of Earth's crust with a known true composition, as well as testing with a variety of priors. We find that the accuracy of the composition estimate of the simulated crust depends on the prior distribution and the amount of seismic variables used. The latter of these observations illustrates the impact of the Bayesian likelihood, while the former illustrates the importance of the Bayesian prior. Improving either improves the quality of our estimates.

3. Methods

3.1. Defining a Prior Distribution

We use a dataset compiled by Keller and Schoene (2012) from `Earthchem`, an online geochemistry database (http://portal.Earthchem.org/), as the surface composition distribution. The dataset contains 68,696 samples with major element compositions, lithology, and sampling location. We reduce the spatial and temporal biases of the original database using weighted resampling, where each sample is weighted in inverse proportion to its spatial and temporal distance from other samples. We sample according to these weights and add Gaussian noise to each resampled composition. Details are provided in Text S1 in Supporting Information S1, and spatio-temporal distributions before and after resampling are shown in Figure S9 in Supporting Information S1. Resampling decreases biases, but is not effective in places with very few samples, for example, Antarctica and much of western Africa. Despite these biases, the dataset provides a measure of the relative probability of all major rock types observed on the surface.

The major element composition $\vec{c}$ includes the 10 most common oxides in the crust, namely, SiO_2, TiO_2, Al_2O_3, FeO, CaO, Na_2O, K_2O, H_2O, and CO_2. We limit $\vec{c}$ to these major elements because they control the mineral composition and so the seismic properties of a rock.

We use the distribution of `Earthchem` samples because we believe that the distribution of rocks at the surface are informative of the rocks found at depth. However, a reasonable alternative might be to assume that a dataset of rocks exhumed from deep within the crust are better representative of the middle and lower crust. To test the impact of this assumption, we use the Huang et al. (2013) distribution as an alternative prior. Another reasonable alternative, which we do not test here, could be to assume the mean crust compositions of upper, middle, and lower crust from previous experiments, and to use prior distributions adjusted to center at those average compositions. Using different priors in the upper, middle, and lower crust could also be used to produce crust composition estimates that align with theories of density-driven compositional stratification.

A final option, for those interested in understanding what the seismic data alone can tell us about the crust, is to use an uninformative prior, which is a common tool in Bayesian analysis. It provides a result dependent only on the likelihood. Because of the Monte Carlo step in the method presented here, a flat composition distribution is not an uninformative prior, since the global average composition result will be biased toward the center of even a uniform prior. Instead, we use a Gaussian mixture model to produce a distribution of random priors, where the distribution of prior means is uniform, as an uninformative prior within the bounds of the distribution of prior means, here 51.27%–68.69% SiO_2. These bounds likely limit our lower crust estimates, since we see probability densities clustered around the lower bound in our average compositon estimates (Figure 5) (The reverse may be true in the upper crust in our synthetic test, see Figure 6). This limitation could be resolved by using a more flexible model to create the uninformative prior. Even given the limitation, our results on the uninformative prior provide a meaningful estimate of what seismic information alone can tell us about crust composition.

3.2. Crust Datasets

To apply the methodology used here to the Earth's crust, we need a set of global datasets describing its seismic properties. In selecting these datasets, we consider several features. Wherever possible, we wanted to use models of crust properties derived from seismic data, not models of crust type, to avoid relying on existing models of crust type and composition. This rules out the widely-used `Crust1.0` (Laske et al., 2013; Mooney et al., 1998), where seismic properties are assigned according to crust age and tectonic setting where seismic constraints are unavailable. The seismic properties should also not be assigned according to assumptions about composition at depth, because this is in fact what we are trying to learn. This prevents us from using the V_p data from `Litho1.0`, which is derived from `Litho1.0`'s modeled V_s using constant V_p/V_s ratios from `Crust1.0` (Pasyanos et al., 2014). Despite its being informed by `Crust1.0` crust type assumptions, we use density from `Litho1.0`, since crust density estimates are required for the `Perple_X` seismic modeling employed in our investigation. Future work could explore alternative modeling approaches to avoid this dependency on a density dataset. Our interest here is in global crust properties, including global averages, so we do not use available high-quality, high-resolution regional seismic models, although our methodology could be applied to those datasets to estimate regional compositions. We use the `SPiRaL` seismic model for V_p and V_s (N. A. Simmons et al., 2021) (Figure S5 in Supporting Information S1), which fits the above criteria.

The seismic datasets we use divide the crust into three layers. The layer depths in `SPiRaL` are inherited from `Crust1.0`. These divisions may obfuscate more complex compositional stratification within the crust, which some higher resolution regional datasets could reveal. Our approach could be extended to consider more crust layers, but would need to be adapted to efficiently handle many layers.

`Litho1.0` reports a measure of model uncertainty, while `SPiRaL` does not. The uncertainty provided by `Litho1.0` is powerful because it can be propagated through estimates of other properties, including the composition estimation proposed here. However, because only one of the datasets we use provides uncertainty estimates, we do not propagate seismic dataset uncertainty here, although it could easily be added to future work.

We use the thermal TC1 model of Artemieva (2006) as our model of crust geotherms. The TC1 model is derived from borehole heat flow measurements, where they are available, combined with thermal models extending to areas where direct measurements are not available. The TC1 model is the most spatially limited of the three data models used here, lacking coverage of the continental shelves and limiting the extent of our composition estimates.

3.3. Modeling the Seismic Properties of Compositions

To estimate the seismic properties of a given composition, we use `Perple_X`, a thermodynamic software package that calculates the mineral assemblage and physical properties of any major element composition $\vec{c}$ at any temperature and pressure θ (Connolly, 2005). `Perple_X` has been used in previous estimates of crust composition (Behn & Kelemen, 2003; Sammon et al., 2022). To more accurately apply `Perple_X` throughout the crust, we further model the impact of metastability and cracks on `Perple_X`-calculated seismic properties. We denote as Φ all of the parameters defining a rock's porosity and exhumation, so that the modeled properties are given by $f_{\vec{s}}(\vec{c};\theta,\Phi)$ (Equation 2).

It is computationally intractable to model the seismic properties of each composition in the prior at every set of conditions θ_l. We therefore divide crust geotherms into 10 evenly spaced bins (Figure S4 in Supporting Information S1, Section 3.2). We model the seismic properties of each composition $\vec{c}$ in the prior at 10 crust configurations θ. Each θ describes the geotherm (Δtemperature/Δdepth) at the center of each bin and the median crust depth of all locations with that geotherm. Assuming an average crustal density of 2900 km/m^2, we can then calculate crustal temperature and pressure in each bin.

Unlike θ, which includes parameters known for the crust at each location, Φ includes parameters which affect seismic properties but which are not well constrained across the crust, including porosity and exhumation.

Metastability is the propensity of rocks to be exhumed to conditions far from those at which they reached equilibrium. Metastability is an important feature of the crust, with implications for mountain building and crust strength (Jackson et al., 2004). Because `Perple_X` calculates the thermodynamic equilibrium of a given composition, it is incapable of modeling exhumed rocks. To model these rocks, we use the `Perple_X`-calculated temperature and pressure derivatives of the density ρ and bulk modulus κ of a rock at formation conditions A, defined by θ and the exhumation parameter in Φ, to estimate the properties of that rock at a different, shallower location in the crust with conditions described by θ (Guerri et al., 2015). We then use the temperature and pressure-dependent shear moduli of Holland and Powell (1998) of each of the `Perple_X`-calculated minerals at A to calculate the shear modulus of the rock at θ.

In the upper crust, where rocks are above the brittle-ductile transition, porosity introduces an additional divergence between an idealized `Perple_X` model and observed seismic properties (Rudnick & Jackson, 1995). Porosity has been directly observed to depths of 12 km in the Kola Superdeep Borehole, where it affects regional seismic studies of the upper crust (Ganchin et al., 1998). Although most locations on Earth lack superdeep boreholes, it is reasonable to assume that rocks in the upper crust, particularly above the brittle-ductile transition, generally feature cracking and pore space (Vitovtova et al., 2014). We model pore space, both dry and fluid-filled, using the approximations derived by (David & Zimmerman, 2011) and porosities derived from (Vitovtova et al., 2014), described further in Text S1 in Supporting Information S1. Cracking, pore space, and alteration on average decrease the density and p-wave velocity of each sample.

In our central result, we use values at the center of possible ranges for all parameters in Φ: 5 km exhumation, a formation temperature of 550°C. In the upper crust, we assume there is total porosity of 0.7% (Chen et al., 2020; Vitovtova et al., 2014), with 5% of that being low-aspect ratio pore space (cracks) and alteration of 1%. We assume upper crust porosity is dry.

We perform a sensitivity analysis to better understand the dependence of our result on the poorly constrained parameter (Figure S12 in Supporting Information S1). We find that our result is somewhat sensitive to both exhumation depth and formation temperature, with results most sensitive to formation temperature in the lower crust. In the upper crust, where we assume there is porosity, we find that our result is somewhat sensitive to the amount of porosity, but not sensitive to other porosity parameters (Figure S7 in Supporting Information S1). The use of single values for the uncertain properties Φ is a limitation of our approach. A better approach might be to consider a distribution of possible Φ, then marginalize the resulting distribution $\mathbb{P}(c, \Phi | s, \theta)$ over Φ. This would require adaptations to our rejection sampling approach.

3.4. The Likelihood

A key part of the Bayesian framework proposed here is the likelihood $\mathbb{P}\left(\vec{s}_l | \vec{c}_l\right)$, the probability of observing seismic properties $\vec{s}_l$ if the crust at that location is composed of some composition $\vec{c}_l$.

We estimate the likelihood by assuming that our modeled seismic properties for a composition are not exactly the true properties of a rock with that composition in the crust, but that we can estimate the range of errors we expect between modeled and actual properties by looking at the range of errors between modeled and laboratory samples. Specifically, we model the relationship between the true seismic properties of a sample, , and the modeled seismic properties $f_{\vec{s}}\left(\vec{c}\right)$ as

$$\vec{s} = f_{\vec{s}}\left(\vec{c}\right) + \epsilon \tag{3}$$

where $\epsilon \sim \mathcal{N}(0, \Sigma)$, a multivariate Gaussian distribution with parameters estimated by the difference statistics between laboratory and modeled seismic properties (V_p, V_p/V_s, and ρ) for the samples described in (Kern et al., 1999). Σ is given in Table S1 in Supporting Information S1.

We can then formalize the likelihood $\mathbb{P}\left(\vec{s}_l | \vec{c}_l\right)$ as the probability of observing an error of a certain size, specifically

$$\mathbb{P}\left(\vec{s}_l|\vec{c}_l\right) = \mathbb{P}\left(\epsilon = f_{\vec{s}}\left(\vec{c}_l\right) - \vec{s}_l\right) = f_{0,\Sigma}\left(f_{\vec{s}}\left(\vec{c}_l\right) - \vec{s}_l\right) \tag{4}$$

where $f_{0,\Sigma}$ is the normal pdf of the error ϵ.

To parameterize the likelihood distribution, we compare our calculated seismic properties to laboratory measurements for 30 samples from (Kern et al., 1999) (Figure S10 in Supporting Information S1). We find a systematic bias in the errors, with the V_p errors between calculated and laboratory measurements distributed around $V_p = 0.12$, and V_p/V_s errors distributed around $V_p/V_s = -0.019$ (these means exclude granulite samples, for which the formation temperature and pressure used for the comparison are inappropriately low). This systematic bias has been previously described in modeling deep crustal samples (Hacker et al., 2015; Rudnick & Jackson, 1995). The differences are usually attributed to irreversible alteration in laboratory samples during exhumation.

To explore the source of our observed systematic error, we use the mineral modes listed for each (Kern et al., 1999) and the seismic property toolbox of Hacker and Abers (2004) to find a modeled estimate of the seismic properties of each sample without any error from mineralogical modeling (Figure S11 in Supporting Information S1). The resulting errors are similar to our systematic error using `Perple_X`, supporting our interpretation that the systematic error between `Perple_X` and laboratory samples is not due to incorrect mineral assemblages or `Perple_X`-specific errors but instead primarily due to alteration and cracking of the laboratory samples not present at depth. We therefore assume that the true distribution of errors between `Perple_X`-calculated properties and the true properties of a rock at depth is centered around zero. We assume that the variance of errors between `Perple_X` and laboratory properties (Figure S10) is representative of the variance of errors between `Perple_X`-modeled properties and the true properties of the same composition at depth, so we use the variance Σ of a multidimensional normal distribution fit to the errors between `Perple_X` and laboratory samples to parameterize $f_{0,\Sigma}$.

3.5. Rejection Sampling the Posterior

We use rejection sampling to sample from the posterior distribution. The general algorithm for rejection sampling from a distribution with pdf $\mathbb{P}(x) = f(x)$ is to sample from a proposal distribution with $\mathbb{P}(x) = g(x)$. Each sample x from the proposal distribution is accepted with a probability proportional to $f(x)/g(x)$, the ratio of the probabilities of observing that sample in the proposal and target distributions. The accepted samples are equivalent to samples drawn from the target distribution $f(x)$.

In this case, we are interested in sampling from $\mathbb{P}\left(\vec{c}|\vec{s}\right)$. We use the prior distribution $\mathbb{P}\left(\vec{c}\right)$ as the proposal distribution, so the acceptance probability for a sample is proportional to

$$\frac{\mathbb{P}\left(\vec{c}|\vec{s}\right)}{\mathbb{P}\left(\vec{c}\right)} = \frac{\mathbb{P}\left(\vec{s}|\vec{c}\right)\mathbb{P}\left(\vec{c}\right)}{\mathbb{P}\left(\vec{c}\right)\mathbb{P}\left(\vec{s}\right)} = \frac{\mathbb{P}\left(\vec{s}|\vec{c}\right)}{\mathbb{P}\left(\vec{s}\right)} \tag{5}$$

where equation 5 is simplified using equation 1. Equation 5 can be further simplified using the fact that $\mathbb{P}\left(\vec{s}\right)$ is constant for a given , so we can consider $\mathbb{P}\left(\vec{s}\right)$ as a normalization constant. The conditional probability $\mathbb{P}\left(\vec{s}|\vec{c}\right)$ is the likelihood described above. We can then calculate equation 5 for each step of the rejection sampling algorithm, allowing us to sample from $\mathbb{P}\left(\vec{c}|\vec{s}\right)$ at any location on the crust. Algorithm pseudocode is supplied in Supporting Information S1.

3.6. A Monte Carlo Algorithm for Estimating Global Average Composition

The rejection sampling algorithm described above samples points from the posterior density $\mathbb{P}\left(\vec{c}|\vec{s}\right)$ at one location. However, we are also interested in the global average composition of the entire crust, both because it allows comparison with prior work that attempts to estimate global averages and because of its importance for biogeochemical, petrological, and geophysical questions concerning Earth's major and trace element balance. We use a Monte Carlo approach, where we sample M times from each of N locations on the crust and find the distribution of average compositions across locations. Critically, sampling from all available 22,530 locations would far underestimate the uncertainty of the global average composition, because the compositions of those 22,530 locations are not statistically independent.

There are spatial correlations between the compositions at different locations on Earth, with closer locations more likely to be similar than distant locations. To estimate this correlation, we use the spatial covariance of

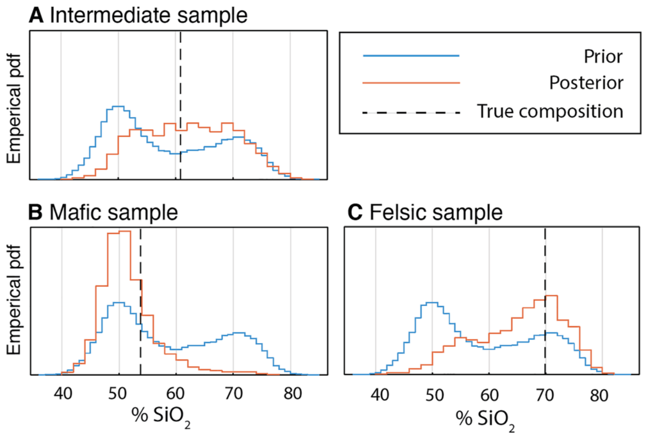

Figure 2. Posterior distributions of predicted composition at three locations on a test crust, where the true composition is intermediate (a), mafic (b), or felsic (c), using the Bayesian model proposed here. The true composition of each location is shown as a dotted line. The modeled seismic properties of the three locations are, for the mafic location, $\rho = 2{,}976$ kg/m^3, $V_p = 5.658$ km/s, $V_p/V_s = 1.585$); for the intermediate location, $\rho = 2{,}732$ kg/m^3, $V_p = 5.634$, $V_p/V_s = 1.532$); and for the felsic location $\rho = 2672$ kg/m^3, $V_p = 5.73$ km/s, $V_p/V_s = 1.572$). The posterior distribution is 10-dimensional, since all major elements are included, but only silica is shown here. The prior distribution is shown in blue and is the same in all three panels. The uncertainty in the posterior distributions comes from two factors: (1) very different compositions can have similar seismic properties; and (2) the seismic properties of a single composition can vary significantly due to differences in formation conditions.

surface samples (Figure S13 in Supporting Information S1). Although there is no way to measure the spatial correlation of compositions at depth, it is reasonable to assume that correlation also exists at depth, and we assume that correlation at depth occurs up to the same scale as correlation at the surface. Because we use surface correlation to calculate regions, we may over- or under-estimate spatial correlation in the middle and lower crust, so we are less confident in our uncertainty estimates in the middle and lower crust.

In our Monte Carlo approach, instead of sampling from all available locations, we take the average of samples from the posterior distributions at each of N uncorrelated locations on the continental crust, where N is an approximation of the number of uncorrelated regions in the continental crust. To calculate N, we divide the land area of the continents, approximately 149 million km^2, by the area of the minimum uncorrelated scale (Figure S13 in Supporting Information S1), about 769.5 km^2. This calculation gives $N = 252$ regions.

Each Monte Carlo run then results in one global average composition estimate $\vec{C}_m$:

$$\vec{C}_m = \frac{1}{N} \sum_{l=1}^{N} \vec{c}_l \tag{6}$$

where $\vec{c}_l \sim \mathbb{P}\left(\vec{c}|\vec{s}_l\right)$ and l iterates through the N uncorrelated locations on the crust. The distribution of these means provides a distribution of the global average crust composition. We do not currently weight the average by the volume of each location, which varies as crust thickness varies. Further work could address this shortcoming.

Larger N, or less spatial correlation, would narrow the uncertainty of the mean, while smaller N would have the opposite effect (Figure S14 in Supporting Information S1). Although N affects the uncertainty of the estimate, it does not change the center of the distribution of estimated crust compositions.

Because our distributions are of compositions, they are limited both theoretically and practically. The components of each composition vector must sum to one, and each composition must appear in a sample dataset (either `Earthchem` or Huang et al. (2013)). This means that the silica compositions of our samples are generally limited to between 30 and 90 (Figure S1 in Supporting Information S1). Because of this, the posterior distribution for a felsic sample will have a tail toward mafic samples, but not vise versa, since there are relatively fewer felsic samples to make up a tail on the other side of the distribution's mode (Figure 2). The opposite is true for mafic samples. When sampling from many such distributions, as in the Monte Carlo approach here, the sampled results will be biased in the direction of the tail, which is toward intermediate samples. This effect partially explains the bias toward intermediate average compositions seen in our result across priors (Figure 5) and in a test on a synthetic crust (Figure 6). Although this effect limits the most extreme mafic or felsic composition estimates, we believe this accurately reflects the fact that it is unlikely that all of Earth's crust is composed of the most mafic or most felsic possible compositions.

3.7. Testing Our Approach on a Synthetic Crust of Known Composition

We test our method on a simulated crust to explore the accuracy and limitations of our approach. We first calculate the seismic properties of a simulated planetary crust (Guerri et al., 2015), then use our Bayesian approach to estimate the composition of this simulated crust. We can then compare our estimates to the true composition of the simulated crust. Our insights from this test allow us to propose promising avenues to reducing uncertainty and increasing both the accuracy and precision of crust composition estimates.

We use the estimates of Rudnick and Gao (2014) as the global average compositions for our simulated crust, namely, 66.6% SiO_2 in the upper crust, 63.5% in the middle crust, and 53.4% in the lower crust. This choice of test composition allows us to ascertain whether our approach could accurately estimate the composition of a crust with significant compositional stratification with depth. An additional benefit of this choice is that it allows us to test whether we could recover the composition of Earth's crust if the most commonly used estimate in the existing literature is accurate.

For each layer in the simulated crust, we randomly select N compositions from the `Earthchem` prior whose global average composition is within 0.01% of the target for that layer. Each location on Earth's continents is assigned a composition, with an associated geotherm from TC1 (Artemieva, 2006). Seismic properties of each sample are calculated as described in Section 3.3.

When using our Bayesian methodology on the synthetic crust, we use the same priors as on real data, namely, the `Earthchem` distribution of surface samples, the Huang et al. (2013) prior, and an uninformative prior. We do not re-estimate the distribution of surface samples, even though the upper and middle simulated crust are significantly more felsic than the mean of the `Earthchem` prior and the lower crust is more mafic.

Because of the non-unique nature of the seismic data used here, our compositionally stratified synthetic test cannot be recovered using the same prior in all layers. This suggests that we would not be able to accurately estimate the true crust composition using the same prior in all layers if it is as highly stratified as the Rudnick and Gao (2014) estimate. This test demonstrates that the prior is an important control on our results. In Section 4.1, we discuss the implications of this sensitivity for interpreting the crust composition estimates of this and other work.

4. Results and Discussion

We present a new estimate of Earth's crustal composition using the `Earthchem` prior and seismic data from the `SPiRaL` (N. A. Simmons et al., 2021) and `Litho1.0` (Pasyanos et al., 2014) models (Figure S5 in Supporting Information S1, Section 3.2). The major element breakdown of our estimate is given in Table 1. The prior has a large impact on our estimate (Figure 5), with the more mafic Huang et al. (2013) prior resulting in more mafic composition throughout the crust. An uninformative prior (in contrast to the results from both the `Earthchem` and Huang et al. (2013)) yields an imprecise estimate. This demonstrated importance of the prior means that identifying and using a good prior is essential to an accurate estimate of Earth's crust composition. Here, we focus on our estimate using the `Earthchem` prior, because, as discussed above, we believe it to be a contain information about the rocks present throughout the crust. That is, while future work should absolutely consider alternative priors, we believe igneous and metaigneous rocks to be a reasonable starting point considering that (a) given that they must form at the surface, sedimentary and metasedimentary rocks are likely volumetrically secondary at depth and (b) may require different treatment during equilibrium thermodynamic and seismic property modeling. For applications where an estimate informed by seismic information alone is appropriate, we provide the major element breakdown of our result using an uninformative prior in Table 2.

Our approach allows us to break down results spatially and temporally, as well as to consider the global average composition. When we consider the composition of crust of different ages (Figure S6 in Supporting Information S1), we find that there is a trend toward more felsic younger crust in the upper and middle crust, while lower crust composition is more stable over time. However, these trends are small relative to the uncertainty in crust composition estimate, with differences over time bins generally smaller than the width of the 90% credible interval. Future work could explore which of these trends are significant and which formation processes they might be related to.

We show the composition of the crust over space in Figure 4. Large batholiths in the Sierra Nevada, the Andes, and the Himalayas are visible as felsic regions (Figure 4, panels a–c). The felsic tops (Figure 4, panel a) and mafic roots (panel c) of subduction arcs including Japan, New Zealand, and Indonesia are also visible. The spatial distribution of standard deviations (Figure 4, panels d–f), with less variance in more mafic regions, reveals that our uncertainty is correlated with predicted composition. This may be because the seismic properties of felsic rocks are less unique (which we speculate may be attributed to the relatively similar seismic properties of many felsic minerals and/or greater diversity in felsic magma genesis), leading to broader posterior distributions in those locations.

To compare our results with previous work presenting results by crust region (Christensen & Mooney, 1995; Rudnick & Fountain, 1995; Sammon et al., 2021), we present the breakdown of our results by category in Table

S4 in Supporting Information S1. We use the categories of Sammon et al. (2021), facilitating comparison to an estimate using a probabilistic but non-Bayesian approach. Our result is more mafic throughout all middle and lower crust regions, which is likely due to methodological differences. Based on the central importance of the prior, as demonstrated in this work, the difference in results may be due to the lack of a formal prior in Sammon et al. (2021).

When we consider our average crust composition estimates, we find that the middle crustal composition is slightly more mafic, on average, than the upper crust, while the lower crust is significantly more mafic (Figure 5, Table 1). This trend agrees with previous work modeling the Earth as increasingly mafic with depth; however, we find less variation in composition between the upper and middle crust than previous work. This less differentiated crust is consistent with some work arguing that seismic boundaries in the crust are the result of phase transitions, not compositional differences (Guerri et al., 2015).

When the same prior distribution is used for all layers of the crust, as in our work, we find intermediate average crust compositions in the upper, middle, and lower crust. This demonstrates that the seismic properties of the crust alone do not determine the composition estimates of Rudnick and Gao (2014) and other similarly stratified models; instead, a combination of prior assumptions about the crust and the seismic properties of the Earth are responsible for current estimates of crust composition. Although previous estimates do not explicitly use a "prior", they still build on theories about compositional stratification to narrow down composition estimates from non-unique seismic data. In our Bayesian framework, evidence for composition stratification in the crust could be used to formulate a prior that reflects this stratification (i.e., a more mafic lower crust prior and a more felsic upper crust prior). This could be achieved either by using a prior in the lower crust of exclusively xenoliths and exhumed terraines, or by resampling a prior of surface compositions in accordance with theoretical compositional stratification.

Our upper crust estimate is the most different of the layers from that of Rudnick and Gao (2014), perhaps due to a difference in methodology. The upper crust estimate of Rudnick and Gao (2014) is derived not from seismic properties but from sediment averages, which may be biased by mechanical or chemical weathering processes, and surface exposure data, which may be biased toward cratons. Our approach, which uses in situ data about upper crust seismic properties to estimate composition, avoids these biases. However, using upper crustal seismic properties introduces uncertainties around pore space not present in the middle and lower crust. Although the upper crust has some porosity, the amount and shape of pore space throughout the upper crust is poorly constrained (Vitovtova et al., 2014). We test our approach with a range of pore space parameterizations and find that increasing porosity results in more mafic estimates of upper crustal composition (Figure S7 in Supporting Information S1). This sensitivity to poorly constrained parameters means that we are less confident in our estimate of upper crustal composition than middle and lower crustal composition. However, even an estimate with no porosity results in an upper crust significantly more mafic than that estimated by Rudnick and Gao (2014), indicating that the difference in our estimate is not only due to uncertainty stemming from porosity. Most of the difference appears to be due to discrepancies in the prior assumptions, as described above.

Although the major elements affect the mineral composition and seismic properties of the crust, they are not the only elements of interest to geoscientists. The trace element makeup of the crust is important for understanding its heat production and evolution (Rudnick & Gao, 2014) To that end, previous estimates of crust composition use correlations between major and trace elements to estimate trace element composition in the crust (Sammon et al., 2021). We take a similar approach, propagating trace elements defined for samples in our prior distribution through to the posterior distributions at each location on the crust. We can then use our Monte Carlo algorithm to estimate the global trace element composition of Earth's crust (Table S5 in Supporting Information S1).

4.1. Sensitivity of Results to the Prior

We expect the composition estimate to be influenced by our prior, and we observe this in our results on the simulated crust, where the estimated composition is biased toward the center of the prior distribution. In the upper and middle layers, where the simulated Earth's composition is more felsic than the center of the `Earthchem` prior distribution, the distribution of predicted averages is biased toward mafic compositions (Figure 6). The opposite is true in the lower crust, where the simulated Earth's composition is more mafic than the center of the prior. This illustrates a limitation of using the same prior distribution in the upper, middle, and lower crust. If there are

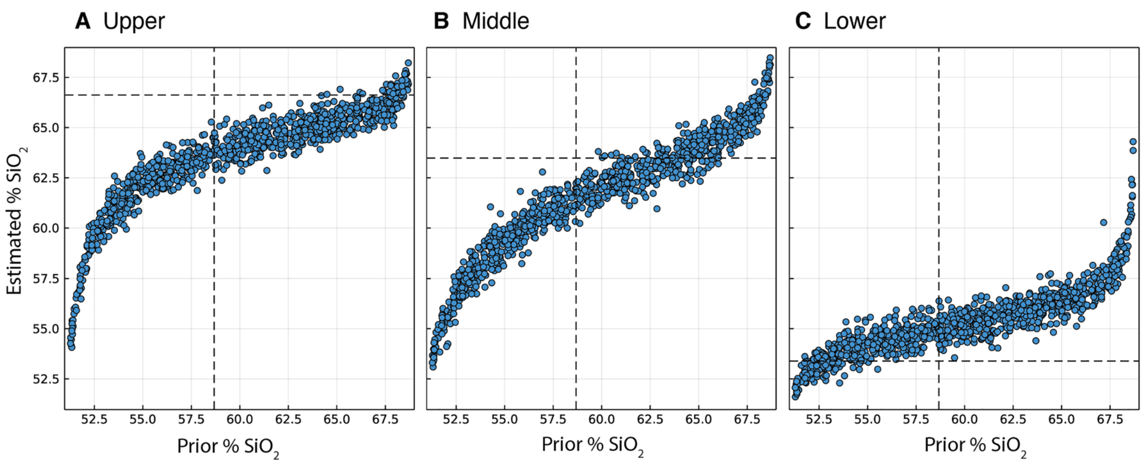

Figure 3. Estimated compositions of upper (a), middle (b), and lower (c) crust of simulated crust using different randomly generated priors. Each composition estimate was calculated using a different prior. The mean of each prior is plotted on the *x*-axis, while the estimated mean crustal composition is plotted on the *y*-axis. Each panel also shows the true simulated crust composition by a horizontal line, and the mean composition of the `Earthchem` prior as a vertical line. The histograms shown in Figure S6 in Supporting Information S1 can be thought of as vertical slices through the scatter plots shown here.

significant differences between the layers, as proposed in Rudnick and Gao (2014) and modeled in our simulated crust, different priors might be appropriate in the upper, middle, and lower crust.

When we use the more mafic Huang et al. (2013) prior, we estimate all three layers of the simulated crust to be more mafic relative to their estimates using the `Earthchem` prior, as seen in Figure 6. These results agree with prior work finding that seismic observations alone cannot constrain a precise estimate of crust composition (Hacker et al., 2015). The prior is therefore not incidental to the problem but a central influence on the accuracy of the result.

We extend our analysis of the impact of the prior by testing our simulated earth with a wide range of artificial priors whose means are uniformly distributed between 51.27% and 68.69% SiO_2. During each Monte Carlo run, we use as our prior a random distribution from this family of priors. The mean of the prior affects the estimated crust composition, with priors with more felsic average composition yielding more felsic average crust estimates. At each layer in the crust, a range of priors with means around 5% of the true simulated crust composition all result in composition estimates close to the true composition of simulated crust (Figure 3). This insight can illuminate the source of the error observed in our estimates of the composition of the simulated crust. The intersection of the mean of the `Earthchem` prior and the true composition in the upper and lower crust is far from

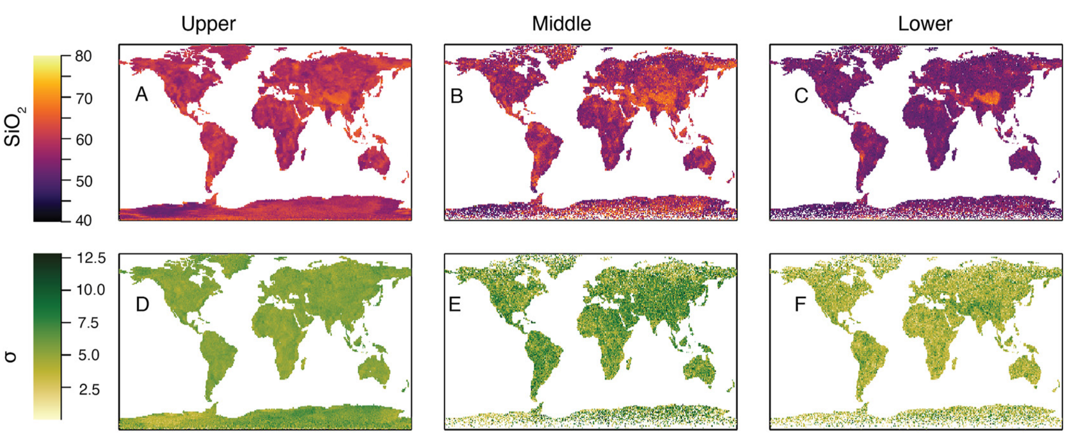

Figure 4. The results of the algorithm described here shown over space. The top panels show the % SiO_2, while the bottom show the standard deviation of the results at each latitude and longitude. 2,975,000 samples from the posterior distributions at each of 2,975,000 randomly chosen locations are binned into 1° × 1° latitude and longitude bins, and the mean and standard deviation of silica composition in each latitude/longitude bin in the upper, middle, and lower crust are visualized.

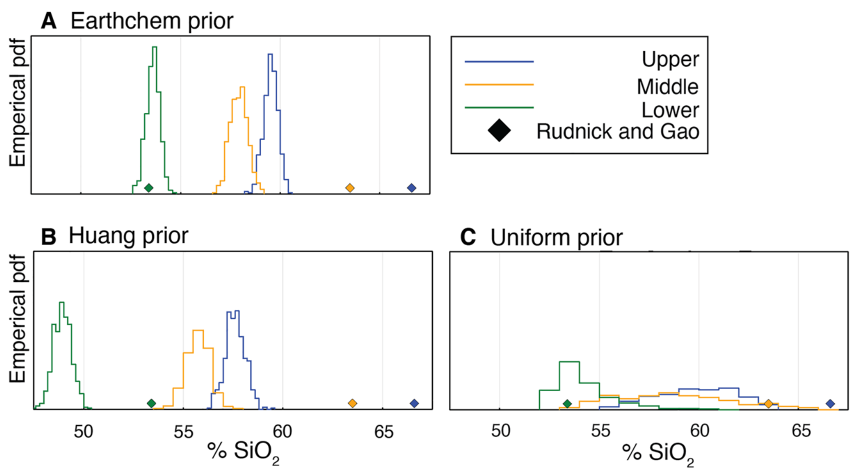

Figure 5. The composition of the crust using different priors. Panel (a) shows the result using the `Earthchem` prior (prior mean silica is 58.7%), panel (b) shows the result using the (Huang et al., 2013) (prior mean silica 52.3%), and panel (c) shows the result using an uninformative prior (prior means uniformly distributed from 51.3% silica to 68.7% silica). The lower crust estimate from the uninformative prior is artificially narrower because it is bounded below by the range of the uninformative prior, which is a limitation of how we designed the uninformative prior. Each empirical pdf is normalized such that its area sums to one. Different panels use different *y*-axis scales.

the band of crust estimates (Figure 3), indicating that the `Earthchem` prior is too far from the true composition of the simulated crust to yield accurate results in the upper and middle crust. In the real crust, we must rely on geologic understanding to ensure that our prior is "good enough". We interpret this dependence on the prior not as a weakness of our approach, but instead as a necessary effect of the non-uniqueness of seismic data. Given this, any method must rely on prior assumptions to narrow down the possible options. Our approach quantifies this effect.

We find that the uninformative prior leads to a much wider range of compositions, and a range of compositions which, in all layers of the crust, includes the true composition of the simulated crust (Figure 6). This is an important check on our method, illustrating that the use of an uninformative prior provides results that include the true composition, even if the result is much less precise than a result using an informative prior. Even with less information from the prior, the Bayesian approach was still able to identify differences between the simulated crust layers, with more mafic predictions in the lower crust than the upper and middle crust, and somewhat more felsic predictions in the upper crust than the middle crust.

While seismic properties alone, with an uninformative prior, can provide wide bounds on composition estimates, narrow composition estimates are the result of combining information from the prior with the seismic properties of the crust. We find that the prior distribution has a large effect on both accuracy and precision. Average crust composition estimates are biased toward the center of the prior, and precise composition estimates require the use of an informative prior. The centrality of prior assumptions in estimates of crust composition suggests that improving crust composition estimates relies not only on the gathering of additional physical data on the crust but also improved priors.

4.2. Sensitivity of the Results to Seismic Variables

The information in a Bayesian approach comes from two sources, the prior and the likelihood. The likelihood, $\mathbb{P}(\vec{s}|\vec{c})$, describes the information the model gets from the observation . Using more information in the likelihood would make for a narrower posterior distribution $\mathbb{P}(\vec{c}|\vec{s})$, and ultimately a more precise Monte Carlo estimate of the global average composition. There are two ways to decrease uncertainty through the likelihood.

The simplest way to add information in the likelihood is to expand the seismic properties . We explore the impact of this by estimating the composition of our simulated crust with only V_p, then with both V_p and V_p/V_s, and comparing the quality of those results to the result of using all three of V_p, V_p/V_s, and ρ. As each additional seismic variable is added, the composition estimates become closer to the true composition of the simulated crust

10.1029/2023JB026353

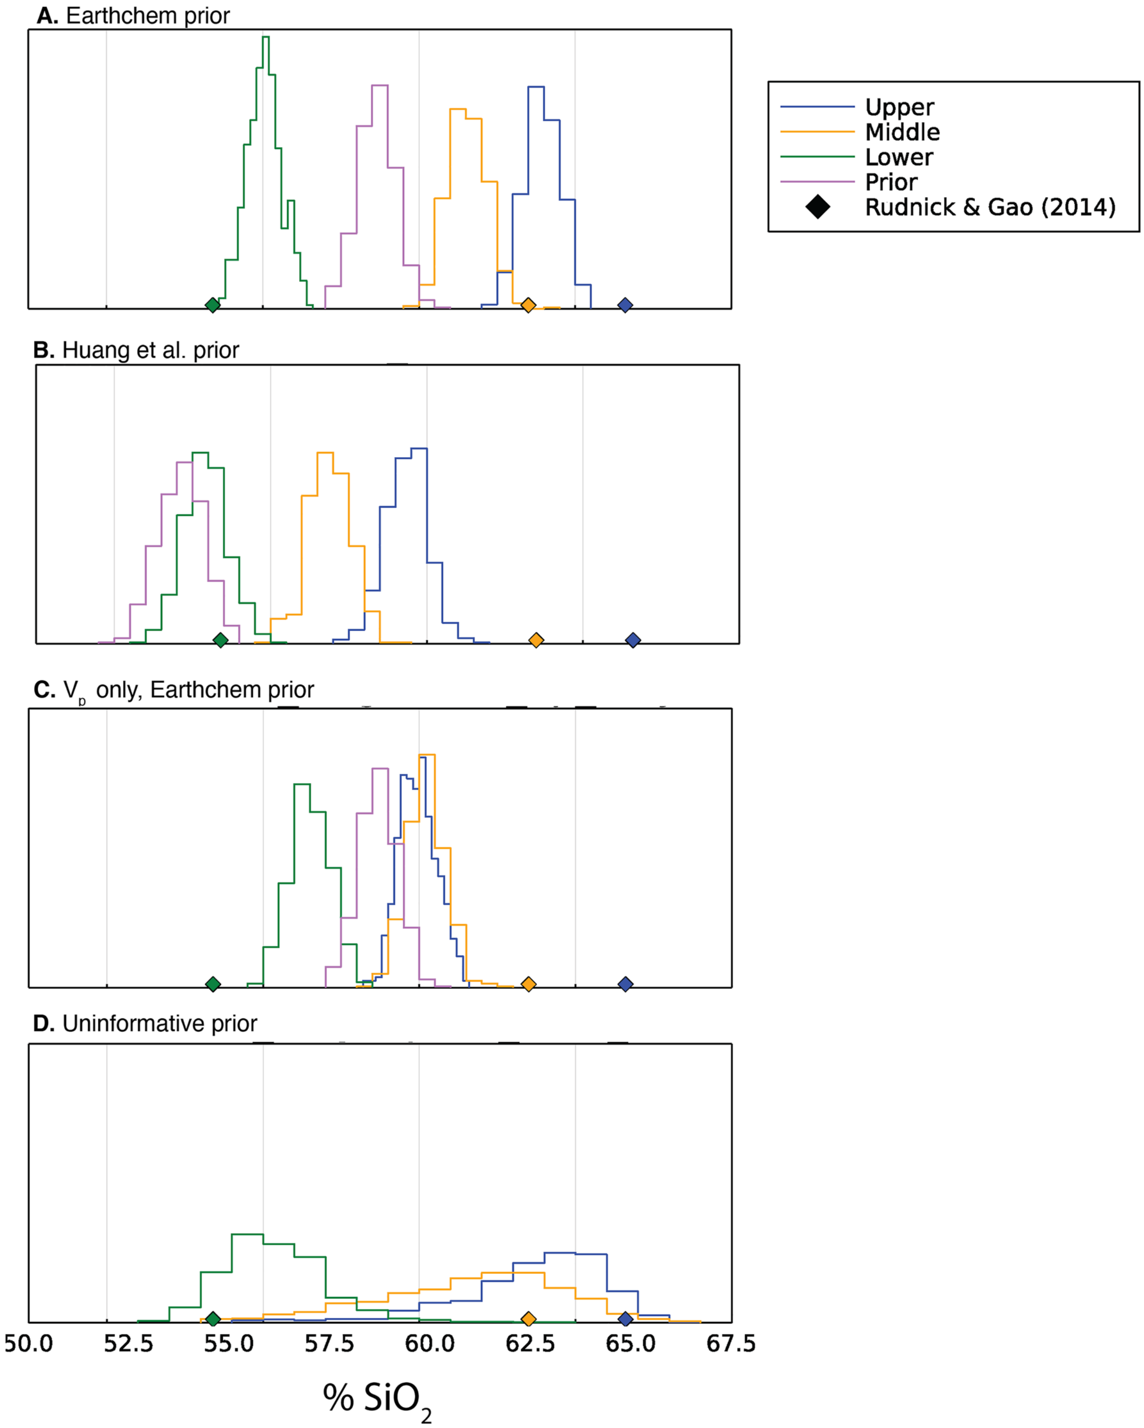

Figure 6. Estimated composition of simulated crust using three variations on the Bayesian method proposed here. Simulated crust has the same average silica composition as (Rudnick & Gao, 2014), shown as diamonds. Histograms show compositions from 500 Monte Carlo runs. (a) shows estimate using V_p, V_p/V_s, and ρ of simulated crust and using the `Earthchem` prior composition distribution; this is the basic method proposed here. (b) shows the result of using the more mafic prior composition distribution from (Huang et al., 2013), which causes composition estimates to be shifted toward more mafic compositions. (c) shows the result of using only V_p of the simulated crust, using the same `Earthchem` prior as (a), which performs worse at estimating the composition of the simulated crust. The purple histogram in each figure shows the distribution of crust composition estimates using the prior alone (`Earthchem` in (a) and (c), Huang in (b)) without any seismic data. In panel (d), each histogram shows the distribution of 1,000 estimates of the mean composition of a simulated crust with the true mean composition of Rudnick and Gao (2014), shown as diamonds. Each estimate was produced using a randomly generated prior distribution. The resulting distributions are wider, reflecting the uncertainty introduced by an uncertain prior.

Table 1
Estimated Mean, Fifth Percentile of Estimates, and 95th Percentile of Estimates for Each Major Element in Each Layer of the Crust Using the `Earthchem` *Prior*

Layer		SiO_2	TiO_2	Al_2O_3	FeO	MgO	CaO	Na_2O	K_2O	H_2O	CO_2
Upper	Mean	59.5	0.901	17.3	6.53	2.88	5.39	4.14	3.24	0.0673	0.0615
	5th	58.9	0.843	17.0	6.23	2.66	5.08	3.98	3.0	0.0442	0.0461
	95th	60.1	0.965	17.6	6.86	3.1	5.73	4.3	3.47	0.0942	0.0806
Middle		57.9	0.975	16.1	7.16	4.46	6.56	3.6	2.13	0.719	0.379
		57.2	0.906	15.8	6.82	4.13	6.23	3.46	1.93	0.633	0.316
		58.6	1.05	16.4	7.49	4.81	6.88	3.74	2.34	0.813	0.447
Lower		53.6	1.07	15.6	8.26	7.11	8.21	3.15	1.21	1.1	0.68
		53.0	0.996	15.2	7.94	6.59	7.88	3.01	1.04	0.971	0.582
		54.2	1.16	16.0	8.58	7.59	8.51	3.3	1.39	1.21	0.779

(Figure 6). This is a general property of a Bayesian posterior, which is always made more certain by additional informative observations.

The alternative to using additional seismic variables could be narrowing the noise model of modeled seismic properties. Functionally, this increases the penalty for compositions whose modeled seismic properties are further from the crust's seismic properties. However, when we explored the potential effects of this by using an error distribution with half the standard deviation of the actual error found in Section 3.3, we find only a minimal impact on the results (Figure S8 in Supporting Information S1). We conclude that more data would be a more productive direction for further work than narrowing the error distribution.

Two potential sources of information about the crust not used here are anisotropy and heat flow. Heterogeneity, aligned cracking, and platey minerals all contribute to anisotropy throughout the crystalline crust (Maupin et al., 2007) and may contain information about crust composition. `SPiRaL`, unlike previous global-scale models, explicitly models anisotropy by allowing the vertical and horizontal velocities of the crust to vary separately (N. Simmons & Myers, 2018). We do not make use of this data here, but it is an opportunity for future work. Heat flow, like anisotropy, is related to crust composition and has available global models (Artemieva, 2006), so its inclusion could also improve the accuracy of crust composition estimates.

Table 2
Estimated Mean, Fifth Percentile of Estimates, and 95th Percentile of Estimates for Each Major Element in Each Layer of the Crust Using an Uninformative Prior

Layer		SiO_2	TiO_2	Al_2O_3	FeO	MgO	CaO	Na_2O	K_2O	H_2O	CO_2
Upper	5th	55.7	0.586	15.4	4.86	1.71	3.22	3.42	1.89	1.08	0.519
	Median	60.8	0.835	15.6	6.25	3.19	4.75	3.67	2.8	1.42	0.642
	95th	65.1	1.14	15.9	7.9	4.97	6.64	3.91	3.48	1.84	0.798
Middle		55.4	0.624	15.5	4.97	1.99	3.63	3.55	1.62	0.862	0.535
		59.9	0.874	15.7	6.28	3.69	5.3	3.77	2.47	1.34	0.666
		64.4	1.13	16.0	7.59	5.36	6.93	4.03	3.25	1.83	0.818
Lower		53.0	0.821	15.2	6.74	3.99	5.68	3.1	1.09	0.867	0.519
		54.8	1.02	15.7	7.78	6.26	7.55	3.34	1.54	1.35	0.651
		58.8	1.15	16.2	8.37	7.35	8.36	3.82	2.56	1.62	0.768

Note. The prior is limited to mean compositions between 50.27% and 68.69% SiO_2, but falls off quickly outside of those bounds. In the lower crust, relatively large probability density at the mafic end of the distribution (see Figure 5) indicates that these limits may artificially constrain the more mafic end of the composition estimate.

5. Conclusions

We formulate a novel Bayesian approach for estimating the composition of Earth's combination from a prior distribution of compositions and measured seismic data. This approach formalizes the conceptual framework often used in previous work, where assumptions of the relative likelihood of various compositions at depth informed results but lacked a mathematical framework.

We use the `Earthchem` distribution of surface compositions as our prior and seismic data from the `SPiRaL` and `Litho1.0` models to generate a posterior distribution. We suggest that using the full range of observed surface compositions as a prior may provide a good first-order solution as it involves fewer assumptions than most other more opinionated priors. We find a more mafic upper and middle crust than previous work, likely driven by the intermediate average composition of our prior. Our lower crust composition estimate is similar to previous work. For all layers, we provide uncertainty distributions.

Using our approach, we are able to quantify the effect of using different prior distributions. We find that an uninformative prior results in a wide range of composition estimates, spanning the estimates in previous work, while a more mafic distribution of compositions from Huang et al. (2013) results in more mafic estimates. Future work could explore using opinionated priors that reflect theories of crustal differentiation or incorporate additional sources of evidence.

Data Availability Statement

The crust-composition code repository used for calculating the composition results presented here is preserved at https://doi.org/10.5281/zenodo.7407276, available via MIT license, and developed openly at https://github.com/gailin-p/crust-composition-release (Pease & Keller, 2022).

Acknowledgments
This work is partially supported by AFOSR F9550-18-1-0316 (Gelb), NSF-DMS 1912685 (Gelb), ONR N00014-20-1-2595 (Gelb and Lee), NSF-DMS 1912999 (Lee). GP and CBK acknowledge funding from Dartmouth College. Thanks to Akshay Mehra, Leslie Sonder, and Laura Sammon for valuable discussion. Thanks to Laura Sammon and an anonymous reviewer for constructive reviews that substantially improved the manuscript.

References

Alonso-Perez, R., Müntener, O., & Ulmer, P. (2009). Igneous garnet and amphibole fractionation in the roots of island arcs: Experimental constraints on andesitic liquids. *Contributions to Mineralogy and Petrology*, *157*(4), 541–558. https://doi.org/10.1007/s00410-008-0351-8

Artemieva, I. M. (2006). Global 1× 1 thermal model tc1 for the continental lithosphere: Implications for lithosphere secular evolution. *Tectonophysics*, *416*(1–4), 245–277. https://doi.org/10.1016/j.tecto.2005.11.022

Behn, M. D., & Kelemen, P. B. (2003). Relationship between seismic p-wave velocity and the composition of anhydrous igneous and meta-igneous rocks. *Geochemistry, Geophysics, Geosystems*, *4*(5). https://doi.org/10.1029/2002gc000393

Chayes, F. (1963). Relative abundance of intermediate members of the oceanic basalt-trachyte association. *Journal of Geophysical Research*, *68*(5), 1519–1534. https://doi.org/10.1029/JZ068i005p01519

Chen, J., Kuang, X., & Zheng, C. (2020). An empirical porosity–depth model for Earth's crust. *Hydrogeology Journal*, *28*(7), 1–2339. https://doi.org/10.1007/s10040-020-02214-x

Christensen, N. I. (1996). Poisson's ratio and crustal seismology. *Journal of Geophysical Research*, *101*(B2), 3139–3156. https://doi.org/10.1029/95jb03446

Christensen, N. I., & Mooney, W. D. (1995). Seismic velocity structure and composition of the continental crust: A global view. *Journal of Geophysical Research*, *100*(B6), 9761–9788. https://doi.org/10.1029/95jb00259

Connolly, J. A. (2005). Computation of phase equilibria by linear programming: A tool for geodynamic modeling and its application to subduction zone decarbonation. *Earth and Planetary Science Letters*, *236*(1–2), 524–541. https://doi.org/10.1016/j.epsl.2005.04.033

Daly, R. A. (1925). The geology of Ascension island. *Proceedings of the American Academy of Arts and Sciences*, *60*(1), 3–80. https://doi.org/10.2307/25130043

David, E., & Zimmerman, R. W. (2011). Elastic moduli of solids containing spheroidal pores. *International Journal of Engineering Science*, *49*(7), 544–560. https://doi.org/10.1016/j.ijengsci.2011.02.001

Dufek, J., & Bachmann, O. (2010). Quantum magmatism: Magmatic compositional gaps generated by melt-crystal dynamics. *Geology*, *38*(8), 687–690. https://doi.org/10.1130/G30831.1

Ganchin, Y., Smithson, S., Morozov, I., Smythe, D., Garipov, V., Karaev, N., & Kristofferson, Y. (1998). Seismic studies around the kola superdeep borehole, Russia. *Tectonophysics*, *288*(1–4), 1–16. https://doi.org/10.1016/s0040-1951(97)00280-1

Grove, T. L., Elkins-Tanton, L. T., Parman, S. W., Chatterjee, N., Müntener, O., & Gaetani, G. A. (2003). Fractional crystallization and mantle-melting controls on calc-alkaline differentiation trends. *Contributions to Mineralogy and Petrology*, *145*(5), 515–533. https://doi.org/10.1007/s00410-003-0448-z

Guerri, M., Cammarano, F., & Connolly, J. A. (2015). Effects of chemical composition, water and temperature on physical properties of continental crust. *Geochemistry, Geophysics, Geosystems*, *16*(7), 2431–2449. https://doi.org/10.1002/2015gc005819

Hacker, B. R., & Abers, G. A. (2004). Subduction factory 3: An excel worksheet and macro for calculating the densities, seismic wave speeds, and H_2O contents of minerals and rocks at pressure and temperature. *Geochemistry, Geophysics, Geosystems*, *5*(1). https://doi.org/10.1029/2003gc000614

Hacker, B. R., Kelemen, P. B., & Behn, M. D. (2011). Differentiation of the continental crust by relamination. *Earth and Planetary Science Letters*, *307*(3–4), 501–516. https://doi.org/10.1016/j.epsl.2011.05.024

Hacker, B. R., Kelemen, P. B., & Behn, M. D. (2015). Continental lower crust. *Annual Review of Earth and Planetary Sciences*, *43*(1), 167–205. https://doi.org/10.1146/annurev-earth-050212-124117

Holbrook, W. S., Mooney, W. D., & Christensen, N. I. (1992). The seismic velocity structure of the deep continental crust. *Continental lower crust*, *23*, 1–43.

Holland, T., & Powell, R. (1998). An internally consistent thermodynamic data set for phases of petrological interest. *Journal of Metamorphic Geology*, *16*(3), 309–343. https://doi.org/10.1111/j.1525-1314.1998.00140.x

Huang, Y., Chubakov, V., Mantovani, F., Rudnick, R. L., & McDonough, W. F. (2013). A reference Earth model for the heat-producing elements and associated geoneutrino flux. *Geochemistry, Geophysics, Geosystems*, *14*(6), 2003–2029. https://doi.org/10.1002/ggge.20129

Jackson, J., Austrheim, H., McKenzie, D., & Priestley, K. (2004). 07) Metastability, mechanical strength, and the support of mountain belts. *Geology*, *32*(7), 625–628. https://doi.org/10.1130/G20397.1

Kay, R. W., & Kay, S. M. (1993). Delamination and delamination magmatism. *Tectonophysics*, *219*(1–3), 177–189. https://doi.org/10.1016/0040-1951(93)90295-u

Keller, C. B., & Harrison, T. M. (2020). Constraining crustal silica on ancient Earth. *Proceedings of the National Academy of Sciences*, *117*(35), 21101–21107. https://doi.org/10.1073/pnas.2009431117

Keller, C. B., & Schoene, B. (2012). Statistical geochemistry reveals disruption in secular lithospheric evolution about 2.5 gyr ago. *Nature*, *485*(7399), 490–493. https://doi.org/10.1038/nature11024

Keller, C. B., Schoene, B., Barboni, M., Samperton, K. M., & Husson, J. M. (2015). Volcanic–plutonic parity and the differentiation of the continental crust. *Nature*, *523*(7560), 301–307. https://doi.org/10.1038/nature14584

Kern, H., Gao, S., Jin, Z., Popp, T., & Jin, S. (1999). Petrophysical studies on rocks from the dabie ultrahigh-pressure (uhp) metamorphic belt, central China: Implications for the composition and delamination of the lower crust. *Tectonophysics*, *301*(3–4), 191–215. https://doi.org/10.1016/s0040-1951(98)00268-6

Laske, G., Masters, G., Ma, Z., & Pasyanos, M. (2013). Update on crust1. 0—A 1-degree global model of Earth's crust. *Geophysical Research Abstracts*, *15*, 2658.

Liu, Y.-S., Gao, S., Jin, S.-Y., Hu, S.-H., Sun, M., Zhao, Z.-B., & Feng, J.-L. (2001). Geochemistry of lower crustal xenoliths from neogene hannuoba basalt, north China craton: Implications for petrogenesis and lower crustal composition. *Geochimica et Cosmochimica Acta*, *65*(15), 2589–2604. https://doi.org/10.1016/s0016-7037(01)00609-3

Mareschal, J.-C., & Jaupart, C. (2013). Radiogenic heat production, thermal regime and evolution of continental crust. *Tectonophysics*, *609*, 524–534. https://doi.org/10.1016/j.tecto.2012.12.001

Maupin, V., Park, J., Romanowicz, B., & Dziewonski, A. (2007). Theory and observations—Wave propagation in anisotropic media. Seismology and the structure of the Earth. *Treatise on Geophysics*, *1*, 289–321. https://doi.org/10.1016/b978-044452748-6/00007-9

Mooney, W. D., Laske, G., & Masters, T. G. (1998). Crust 5.1: A global crustal model at 5×5. *Journal of Geophysical Research*, *103*(B1), 727–747. https://doi.org/10.1029/97jb02122

Pasyanos, M. E., Masters, T. G., Laske, G., & Ma, Z. (2014). Litho1. 0: An updated crust and lithospheric model of the Earth. *Journal of Geophysical Research: Solid Earth*, *119*(3), 2153–2173. https://doi.org/10.1002/2013jb010626

Pease, G., & Keller, B. (2022). Crust composition code. Version 1. Software. https://doi.org/10.5281/zenodo.7407276

Rudnick, R. L., & Fountain, D. M. (1995). Nature and composition of the continental crust: A lower crustal perspective. *Reviews of Geophysics*, *33*(3), 267–309. https://doi.org/10.1029/95rg01302

Rudnick, R. L., & Gao, S. (2014). Composition of the continental crust. *Treatise on Geochemistry (Second Edition)*, 1–51. https://doi.org/10.1016/b978-0-08-095975-7.00301-6.4

Rudnick, R. L., & Jackson, I. (1995). Measured and calculated elastic wave speeds in partially equilibrated mafic granulite xenoliths: Implications for the properties of an underplated lower continental crust. *Journal of Geophysical Research*, *100*(B6), 10211–10218. https://doi.org/10.1029/94jb03261

Sammon, L. G., Gao, C., & McDonough, W. (2022). Lower crustal composition in the southwestern United States. *Journal of Geophysical Research: Solid Earth*, *125*(3), e2019JB019011. https://doi.org/10.1029/2019jb019011

Sammon, L. G., McDonough, W., & Mooney, W. D. (2021). Estimating compositions of the deep continental crust. In *Egu general assembly*.

Simmons, N., & Myers, S. (2018). Spiral 1.0: Global tomography model of travel times and surface waves with transversely isotropic crust and mantle. In *Egu general assembly conference abstracts*.10642.

Simmons, N. A., Myers, S. C., Morency, C., Chiang, A., & Knapp, D. R. (2021). SPiRaL: A multiresolution global tomography model of seismic wave speeds and radial anisotropy variations in the crust and mantle. *Geophysical Journal International*, *227*(2), 1366–1391. https://doi.org/10.1093/gji/ggab277

Smith, A. F., & Gelfand, A. E. (1992). Bayesian statistics without tears: A sampling–resampling perspective. *The American Statistician*, *46*(2), 84–88. https://doi.org/10.1080/00031305.1992.10475856

Vitovtova, V., Shmonov, V., & Zharikov, A. (2014). The porosity trend and pore sizes of the rocks in the continental crust of the Earth: Evidence from experimental data on permeability. *Izvestiya - Physics of the Solid Earth*, *50*(5), 593–602. https://doi.org/10.1134/s1069351314040181

Walker, J. C., Hays, P., & Kasting, J. F. (1981). A negative feedback mechanism for the long-term stabilization of Earth's surface temperature. *Journal of Geophysical Research*, *86*(C10), 9776–9782. https://doi.org/10.1029/jc086ic10p09776

Walker, T., & Syers, J. K. (1976). The fate of phosphorus during pedogenesis. *Geoderma*, *15*(1), 1–19. https://doi.org/10.1016/0016-7061(76)90066-5