Geophysical Journal International

Geophys. J. Int. (2019) **218,** 573–585
Advance Access publication 2019 April 08
GJI Seismology

doi: 10.1093/gji/ggz169

On the robustness of attenuation measurements on teleseismic *P* waves: insights from micro-array analysis of the 2017 North Korean nuclear test

M.J. Bezada [1], J. Byrnes[1] and Z. Eilon [2]

[1] *Department of Earth Sciences, University of Minnesota - Twin Cities, Minneapolis, MN 55104, USA. E-mail: mbezada@umn.edu*
[2] *Department of Earth Sciences, University of California - Santa Barbara, Santa Barbara, CA 93106, USA.*

Accepted 2019 April 6. Received 2019 April 5; in original form 2018 August 27

SUMMARY

Despite their importance as a fundamental constraint on Earth properties, regional-scale measurements of body-wave seismic attenuation are scarce. This is partially a result of the difficulty in producing robust estimates of attenuation. In this paper, we focus on measuring differential attenuation on records of teleseismic *P* waves. We examine a unique data set of five records of the North Korean nuclear test of 2017 measured at five broad-band seismic stations deployed within a few metres of each other but using different installation procedures. Given their extreme proximity, we expect zero differential intrinsic attenuation between the different records. However, we find that different attenuation measurement methods and implementation parameters in fact produce significant apparent differential attenuation ($\Delta t*$). Frequency-domain methods yield a wide range of $\Delta t*$ estimates between stations, depending on measurement bandwidth and nuances of signal processing. This measurement instability increases for longer time windows. Time domain methods are largely insensitive to the frequency band being considered but are sensitive to the time window that is chosen. We determine that signal-generated noise can affect measurements in both the frequency and time domain. In some cases, the range of results amounts to a significant fraction of the range of differential attenuation across the conterminous United States as determined by a recent study. We suggest some approaches to manage the inherent instability in these measurements and recommend best practices to confidently estimate body wave attenuation.

Key words: Fourier analysis; Time-series analysis; Body waves; Seismic attenuation; Site effects.

1 INTRODUCTION

Detailed geological characterization of the subsurface from measurements of its seismic properties is made difficult by the inherent ambiguity in interpreting seismic velocity (the most commonly determined seismic parameter). While crustal seismic velocities depend mostly on lithology, at subcrustal depths several different factors have comparable effects. There, wave speed can be affected not only by temperature (e.g. Goes *et al.* 2000), but also by melt fraction (e.g. Hammond & Humphreys 2000), and volatile content (e.g. Jacobsen *et al.* 2008). This ambiguity may be reduced, and additional insights may be gained, by placing robust constraints on seismic attenuation (e.g. Dalton *et al.* 2009; Eilon & Abers 2017). However, in contrast to the abundance of seismic velocity models at a variety of scales, models of seismic attenuation are relatively rare. This dearth of information on attenuation is partly a result of difficulties in making the relevant measurements. While measuring traveltime delays is straightforward, measurement of attenuation is complicated by a number of factors. The issues surrounding the measurement of attenuation on surface waves have been explored elsewhere (e.g. Bao *et al.* 2016). Here, we address the problem of measuring the attenuation of teleseismic *P* waves.

The most widely used method for estimating attenuation on teleseismic body waves is the spectral slope method of Teng (1968): Assuming frequency independence of the quality factor Q, the differential path-integrated attenuation between records at station i and j ($\Delta t*_{ij}$) is obtained by finding the best-fitting linear approximation to the natural logarithm of the ratio of two amplitude spectra (eq. 1).

$$ln\ \left(R(f)_{ij}\right) = A_0 - \pi f \Delta\, t^*_{ij}, \tag{1}$$

where $R(f)_{ij}$ is the ratio of the amplitude spectrum at station i to that at station j, A_0 is a term that accounts for the frequency-independent difference in amplitude between the two stations, and f is the frequency in Hz.

By considering only spectral ratios, this approach eliminates the need for characterizing the source spectrum as it is assumed to be common between stations and so cancels out.

When estimating the spectra and taking the ratio, a number of data processing choices need to be made that can affect the result (e.g. Adams & Humphreys 2010). For example, one must choose a frequency band over which to calculate the ratio and perform the linear fit. The low and high-end cut-offs of the frequency band are typically determined based on the frequency range that shows significantly larger amplitudes in a post-arrival window than in a pre-arrival window. Unfortunately, there is no objective way to define exactly what represents 'significantly larger' amplitudes. Perhaps more importantly, the pre-arrival trace is representative of 'random' (cultural and environmental noise) noise but not 'signal-generated' noise; the wavetrain caused by the interaction of the primary wavefield with very local structure (scatterers, surface and basement topography, natural or artificial structures, etc.) that is sometimes referred to as the 'site response'. Signal-generated noise has been found to impact attenuation measurements much more than random noise (Adams & Humphreys 2010) and its presence or absence in the amplitude spectra cannot be gleaned from the pre-arrival timeseries. Moreover, signal-generated noise can be consistent between measurements, and so will not average to zero as the number of measurements increases. Another source of measurement variability is the estimation and treatment of the spectra themselves. Researchers must choose whether or not to apply smoothing to the spectra before taking the ratios, whether to use the standard discrete Fourier transform, the multitaper method (Thomson 1982), or a different approach, and whether or not to apply tapers and zero-padding when windowing the timeseries. Finally, the time window selected for analysis must also be determined. Different studies use windows of different length, often varying for each individual event or even each individual record. All of these seemingly innocuous choices may affect the measurement and, we show here, can have profound effects on the overall results.

Recently, alternatives to the spectral ratio method have been developed. The comb method of Eilon & Abers (2017) extends the spectral ratio concept to include fitting of the phase spectra, thus utilizing the fact that attenuation requires dispersion to maintain causality. The comb method exploits this fact by trying to fit the frequency-dependent phase shifts that are predicted for any given $t*$. Specifics of this method are reviewed in a later section.

Finally, a waveform matching method based on fitting attenuated versions of the estimated source time function to the observed waveforms has been applied to data from Iberia and Morocco (Bezada 2017) as well as Australia (Bezada & Smale 2019), and the Central Appalachians (Byrnes *et al.* 2019) yielding results that are well-correlated to known geological features. This method is similar to the time-domain method of Adams & Humphreys (2010) and will also be described in more detail in a later section.

Here, we compare the effect of analysis choices on the results obtained with various methods on an unusual data set where stations were deployed in such close proximity that differential intrinsic attenuation is expected to be zero. The first key result is that differences in station installation parameters produce substantial differences in the post-arrival waveforms and spectra. Secondly, we find that choices in estimating the spectrum can produce changes in the measured differential attenuation, and that for some stations both the frequency and time-domain methods can find non-zero differential attenuation owing to the effect of the signal-generated noise on the spectra.

2 DATA

For the original purpose of evaluating different installation configurations, four broad-band direct burial digital seismographs (Nanometrics Meridian Compact 120 s) were installed near the legacy Transportable Array station SPMN in eastern Minnesota (Fig. 1). The goal of this micro-array deployment was to compare records from the direct burial instruments to those of the vault-installed SPMN station to optimize the direct-burial installation parameters for a future temporary array deployment.

The installation strategies for each station are shown in Fig. 1(a). Two of the direct-burial stations were deployed at a depth of 1 ft (32 cm) and two at a depth of 3 ft (96 cm). For each depth, one instrument was buried directly, in a strict sense, with nothing between the instrument and the surrounding soil. The second instrument was deployed within a PVC pipe, cut to be the same length as the hole, open on both ends, and with a diameter only slightly larger than the instrument itself. The PVC pipes were filled with clean sand and the hole covered; while the holes where the instruments were directly buried were filled back in with the original material. In the remainder of the text we will refer to the stations as 1DB, 1PVC, 3DB and 3PVC according to the depth of burial in feet ('1″ or '3″) and whether they were directly buried ('DB') or encased in a pipe ('PVC').

All stations were located in a ∼5 m x ∼3 m rectangular area (Fig. 1b), with the maximum interstation distance being <5 m. Because of the very close proximity of all the stations, we assume that there is no difference in the intrinsic attenuation experienced by teleseismic waves recorded in each of the stations, as they all sample identical Earth structure. Even if a low upper crustal *P*-wave velocity of 2 km s^{-1} is assumed, the width of the Fresnel zones would only approach the station spacing for frequencies in the kHz range.

The temporary stations were only deployed for 2 weeks. No natural earthquakes suitable for attenuation analysis occurred during the deployment, but the stations recorded the seismic event of 3 September 2017, in the Democratic People's Republic of Korea (DPRK, Fig. 1c) that was widely reported to be an underground nuclear test (e.g. CTBTO 2017; Gaebler *et al.* 2019). The waveform from the DPRK event is comprised of a main pulse ∼1.4 s in duration, without a significant coda in the SPMN record (Fig. 2). This, and all subsequent, seismograms are velocity waveforms. Interestingly, the record from this event is comparable to that of deep-focus earthquakes of similar magnitude (Fig. 2), which are especially useful for measurements of attenuation given the simplicity of their waveforms (Bezada 2017). Given the similarity between the waveforms of the DPRK event and a natural earthquake, we conclude that the DPRK event is a suitable stand-in for a natural earthquake, and that the insights derived from its analysis are transferable to the analysis of natural deep-focus earthquakes. The analysis presented here may not straightforwardly apply to earthquakes with high magnitudes and long durations, which have waveforms and spectral characteristics that are much different from the DPRK event. Nonetheless, given that deep-focus events are often preferred for teleseismic attenuation analysis (e.g. Hwang *et al.* 2009, 2011; Cafferky & Schmandt 2015; Bezada 2017) we judge this analysis to be an informative and useful exercise.

The data recorded by the direct-burial and TA stations are nearly indistinguishable, but the data recorded by the PVC-cased stations are markedly different (Figs 3a and b). Before the arrival, variations in the phase and amplitude of the long-period noise are apparent. Post-arrival, the main pulse is very similar in all the traces. However, ∼1.2 s after the first break, the records from 1PVC and 3PVC deviate

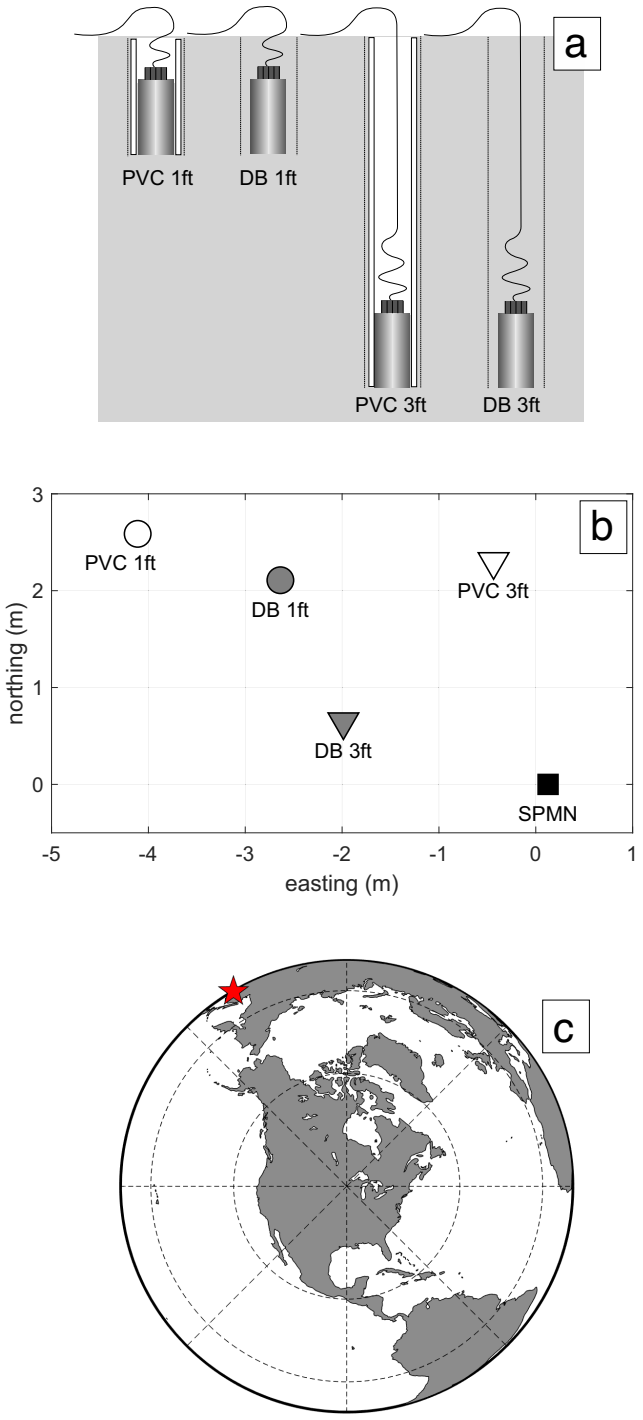

Figure 1. Installation sketch, Station locations and event location. (a) Sketch of installation design for temporary stations. Solid grey indicates local material, dotted pattern represents clean sand. Dashed lines show original post-hole and thick white lines represent the PVC pipe. Seismometers drawn approximately to scale. (b) Micro-array station distribution diagram showing the proximity of the stations. Inverted triangles represent stations buried at a depth of 3 ft (∼92 cm) and circles represent stations buried at a depth of 1 ft (∼31 cm). Grey symbols represent stations buried directly in the ground, while open symbol represent stations encased in a sand-filled PVC pipe. A black square represents the location of the legacy TA station SPMN. (c) Event location map centred on station SPMN. The location of the event is denoted by a red star.

substantially from the other traces. The character of this difference is distinct from the variations in the pre-arrival records (Fig. 3b), which leads us to conclude that it represents signal-generated noise. This part of the coda in 1PVC and 3PVC must arise in some way from the interaction of the wavefield with the installation hardware. Although an interesting problem in its own right, in this paper we will not seek to identify the mechanism responsible for this signal-generated noise, but rather evaluate how its presence affects measurements of attenuation. Since we know that zero differential attenuation is the correct result, this exercise gives us insight into the magnitude and nature of the errors we might expect with any one of the measurement techniques we test.

Comparing the pre- and post-arrival spectra (10 s windows) for SPMN, 1PVC and 3PVC (Fig. 3c) we find that the signal-to-noise ratio appears satisfactory at frequencies higher than 0.3 Hz, and lower than ∼1.7 Hz. The high end of the acceptable frequency range is somewhat ambiguous, and defensible choices range from ∼1.4 to ∼2 Hz. The effect of the choice of high-end cut-off frequency will be explored in a later section.

3 METHODS

We apply several techniques to estimate the differential attenuation between records of the DPRK nuclear explosion at each of the test stations and the SPMN permanent station, which we use as a reference owing to its high quality of installation. A suite of measurements is done using the spectral ratio method, varying both methodology and signal-processing parameters. These results are presented in Section 4. One of the choices is the technique used to estimate the spectra that go into producing the spectral ratios [a2AS$R(f)_{ij}$ in eq. 1). For one set of measurements, we estimate the spectra using the fast Fourier transform (hereafter FFT) without smoothing the spectra. A second set again relies on the FFT, but the resulting spectra are smoothed by applying a 5-point moving average (hereafter FFT-MAVG). A third set of measurements uses the multi-taper method (hereafter MTM) with seven Slepian tapers, which is expected to yield smoother spectra than the FFT (Thomson 1982).

For each of these different techniques, we explore a range of high-frequency cut-offs and degrees of tapering in the time domain, while always zero-padding with 200 samples on either end before estimating the spectra. The low-end cut-off of the relevant frequency range was uniformly fixed to 0.3 Hz. We varied the high-end cut-off from 1.4 to 2.0 Hz, with a 0.1 Hz interval. We initially use 10 s of data starting at the first arrival of seismic energy (defined by the first break); we explore the effect of choosing a different time window in Section 5. The data are tapered with a Tukey (cosine-tapered) window. For each technique and trial frequency range, we repeat the measurements using taper fractions of 0.01, 0.05, 0.10, 0.15 and 0.20. Combining the five windowing parameters and seven frequency ranges results in a total of 35 different measurements for each station and for each of the techniques being implemented.

We explore two further methods. The first is the hybrid time and frequency domain technique of Eilon & Abers (2017). Here, one filters seismograms with a 'comb' of narrow-band filters and then measures the relative amplitude and arrival times to produce the relative amplitude and phase spectra. This procedure differs from the standard spectral ratio measurement in two fundamental ways. First, the algorithm aims to find the $\Delta t*$ measurement that best explains the differences in both amplitude and phase, while the spectral ratio method only considers amplitude. In the case where

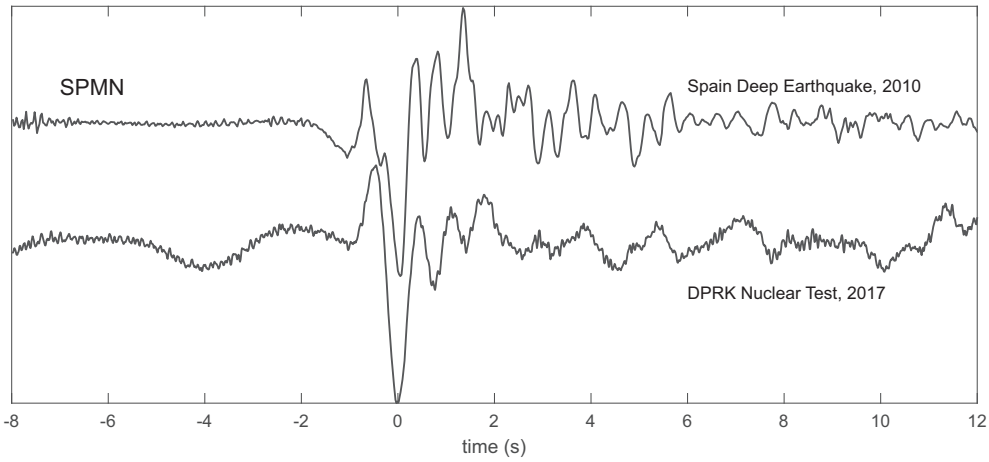

Figure 2. Velocity traces from the deep-focus Spanish event of 2010 (top) and DPRK nuclear test (bottom) as recorded by station SPMN. Note the similar character of the waveforms from the two events.

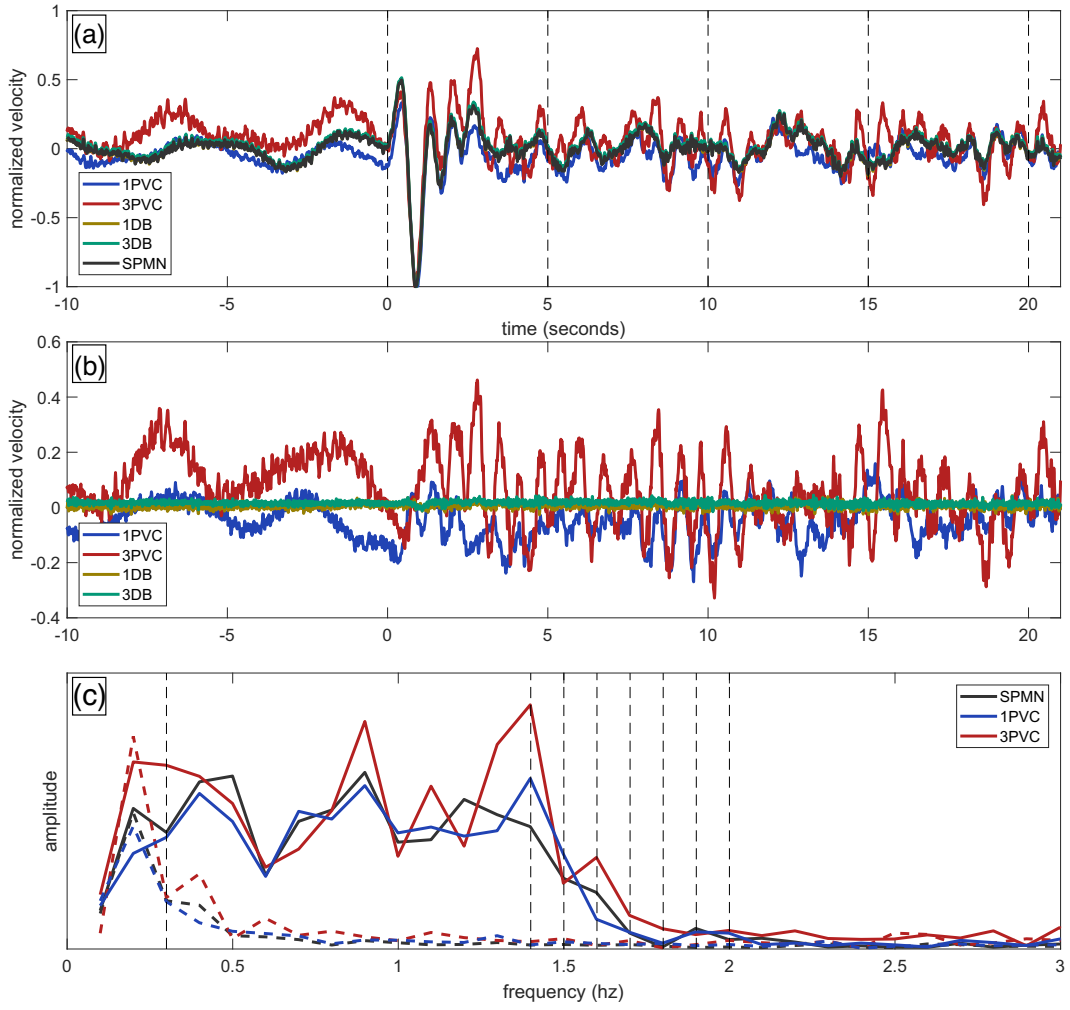

Figure 3. (a) Normalized velocity records of the DPRK nuclear test from all stations in our array. All traces are normalized by their maximum excursion and are colour-coded as indicated in the legend. (b) difference between each record and the reference record from SPMN. Note the change in vertical scale with respect to panel 'a'. Traces are colour-coded as in panel 'a'. (c) FFT-derived amplitude spectra from pre- (solid lines) and post-arrival (thick dashed lines) time windows recorded by the reference station and stations 1PVC and 3PVC (spectra for stations 1DB and 3DB are identical to that for SPMN). Thin vertical dashed lines in panel 'a' show limits of time windows used with frequency domain methods. Thin vertical dashed lines in panel 'c' indicate the low-end cut-off frequency and the different high-end cut-off frequencies that were considered in this study.

Q is considered to be frequency-independent in the frequency range of interest, the amplitude ratios are related to $\Delta t*$ by eq. (1), and the phase shifts ($\Delta\varphi$) are related to $\Delta t*$ by:

$$\Delta\varphi_{ij}(f) = \left[1 - \frac{ln(f)}{\pi}\right] \Delta t^*_{ij} + \delta t, \quad (2)$$

where δt is a constant offset that allows for frequency-independent traveltime variations; we ignore this parameter for this study. Secondly, this method produces $\Delta t*$ estimates by considering all the traces at the same time, rather than the pair-wise measurement we perform using the spectral ratio method. That is to say, the relative attenuation between all pairs of traces are determined simultaneously in a single inversion. For this technique, the same suite of 35 variations in the frequency range and the same Tukey window taper fraction was applied as for the spectral ratio measurements.

The second method is the time-domain approach of Bezada (2017). This method involves estimating the unattenuated source time function and variably attenuating it to match the observed waveforms using a grid-search approach. Numerical attenuation of the estimated source time function is done using the frequency-domain attenuation operator of Azimi *et al.* (1968)

$$F(f)_{att} = F(f)_{stf} \cdot exp\left[-\pi f \Delta t^* - \frac{i}{\pi} ln\left(\frac{f}{f_0}\right)\right], \quad (3)$$

where $F(f)_{stf}$ and $F(f)_{att}$ are the Fourier domain representations of the estimated source time function and attenuated trace, respectively. The attenuated trace is then converted back to the time domain to obtain the attenuated waveform.

Comparing eq. (3) with eqs (1) and (2), one can see that the comb and waveform matching techniques are equivalent in their formulation of the problem except for static time delays between the traces, which are solved for in the comb method (δt in eq. 2) and removed before calculating the misfit in the waveform-matching approach. The two methods are different in the way the best fit of the attenuation model is assessed. In the comb method, the misfit is determined in the frequency domain (weighted at each frequency by time-domain-similarity between narrow-bandpass-filtered waveforms), while in the waveform-matching approach the misfit is determined in the time domain.

For this study, we use the SPMN record as the source estimate since our aim is to measure any apparent differential attenuation relative to this station. An important distinction between the waveform-matching method and those performed in the frequency domain is the length of the time window that is being analysed. The time-domain approach seeks to minimize the misfit between the observed and synthetic (attenuated source time function) waveforms within a narrow time window that contains only the main pulse (Bezada 2017). Initially, we use a time window that is 0.75 s long (preferred window in Fig. S1) and present these results in Section 4. In order to emulate the variable frequency range used in the frequency-domain methods, we detrend and bandpass-filter all traces before performing the grid search. In each iteration, the low-frequency corner was held at 0.3 Hz, and the high-frequency corner was varied from 1.4 Hz to 2 Hz, mimicking the frequency ranges described above. The effect of using different time windows will be explored in detail in Section 5.

4 VARIATIONS TO PRE-PROCESSING AND FREQUENCY RANGE

Given the extreme proximity of all the stations involved in this study, we assert that differential intrinsic attenuation between these records, for frequencies below 2 Hz, should be zero. Indeed, for stations 1DB and 3DB, $\Delta t*$ estimates from both the frequency and time-domain methods are essentially zero in almost all cases tested (Figs 4–7). However, we find that $\Delta t*$ estimates for 1PVC and especially 3PVC can deviate from the expected value, indicating contamination of the measurement by signal-generated noise. Estimates of differential attenuation of 1PVC relative to SPMN from frequency-domain methods using a 10 s time window fall around –0.05 s, while for 3PVC most estimates are around –0.1 s, with some individual estimates having magnitudes exceeding 0.3 s (Figs 5 and 7). Negative $\Delta t*$ values would usually be interpreted as structure beneath stations 1PVC and 3PVC having substantially higher Q than beneath station SPMN. The estimates that most deviated from zero were found using the spectral ratio method on FFT-derived spectra without smoothing.

For our different implementations of the spectral ratio technique, we found the results to vary with the choices of high-end cut-off frequency (F_{high}) and Tukey window taper fraction (T_{frac}). The unsmoothed FFT measurements show the largest range of values (Fig. 7). For a T_{frac} value of 0.05, and variable F_{high}, the FFT method yielded $\Delta t*$ between 0 and –0.08 s at station 1PVC and –0.12 to –0.30 s at station 3PVC (Fig. 5). As T_{frac} is increased, this range is seen to increase. Using a technique that includes smoothing or stabilization of the spectra, whether it was moving-average smoothing (FFT-MAVG) or the multitaper method (MTM), was found to moderately reduce the sensitivity of $\Delta t*$ estimates to F_{high} at any given T_{frac} (Figs 5 and 7).

If instead we fix F_{high} at 1.5 Hz and vary the T_{frac} we see that this parameter also has a substantial control on measured $\Delta t*$ (Fig. 6). Using the FFT method, simply changing the taper fraction led to a range in $\Delta t*$ between 0.07 and 0.13 s for station 1PVC, while for station 3PVC it was 0.1–0.2 s. Again, the FFT-MAVG and MTM methods only moderately reduced these ranges (Figs 6 and 7).

The other frequency-domain method we tested, the comb method (Eilon & Abers 2017), produced $\Delta t*$ estimates comparable to those obtained by the spectral ratio method when using smoother spectra (either FFT-MAVG or MTM) and showed similar sensitivity to the choice of F_{high}. However, the comb method results were found to be largely insensitive to the choice of T_{frac} (Figs 4–6).

In contrast to the frequency-domain methods, the time-domain waveform matching approach is very stable with respect to frequency range. For all stations, variations on the $\Delta t*$ estimate with high-end cut-off frequency are negligible (Fig. 5). Importantly, using our preferred time window, the time-domain method always yields very small differential attenuation estimates for all stations ($\Delta t* \leq 0.02$) including 1PVC and 3PVC (Fig. 5). We did not explore the effect of Tukey windows on the time-domain estimates, as the windowing situation is considerably different in this case.

5 VARIATIONS TO THE LENGTH OF THE TIME WINDOW

In the previous section, we varied a number of parameters but maintained a constant time window for the analysis. The choice of time window is likely to impact both the time and frequency domain methods in distinct ways. For the time domain, the goal is to focus on the main excursion and how its width and shape change as the waveform is attenuated. However, there is leeway in defining where the main excursion begins and ends. In the frequency domain, a time window likewise needs to be chosen in order to produce the spectra. These time windows are typically longer than those used in

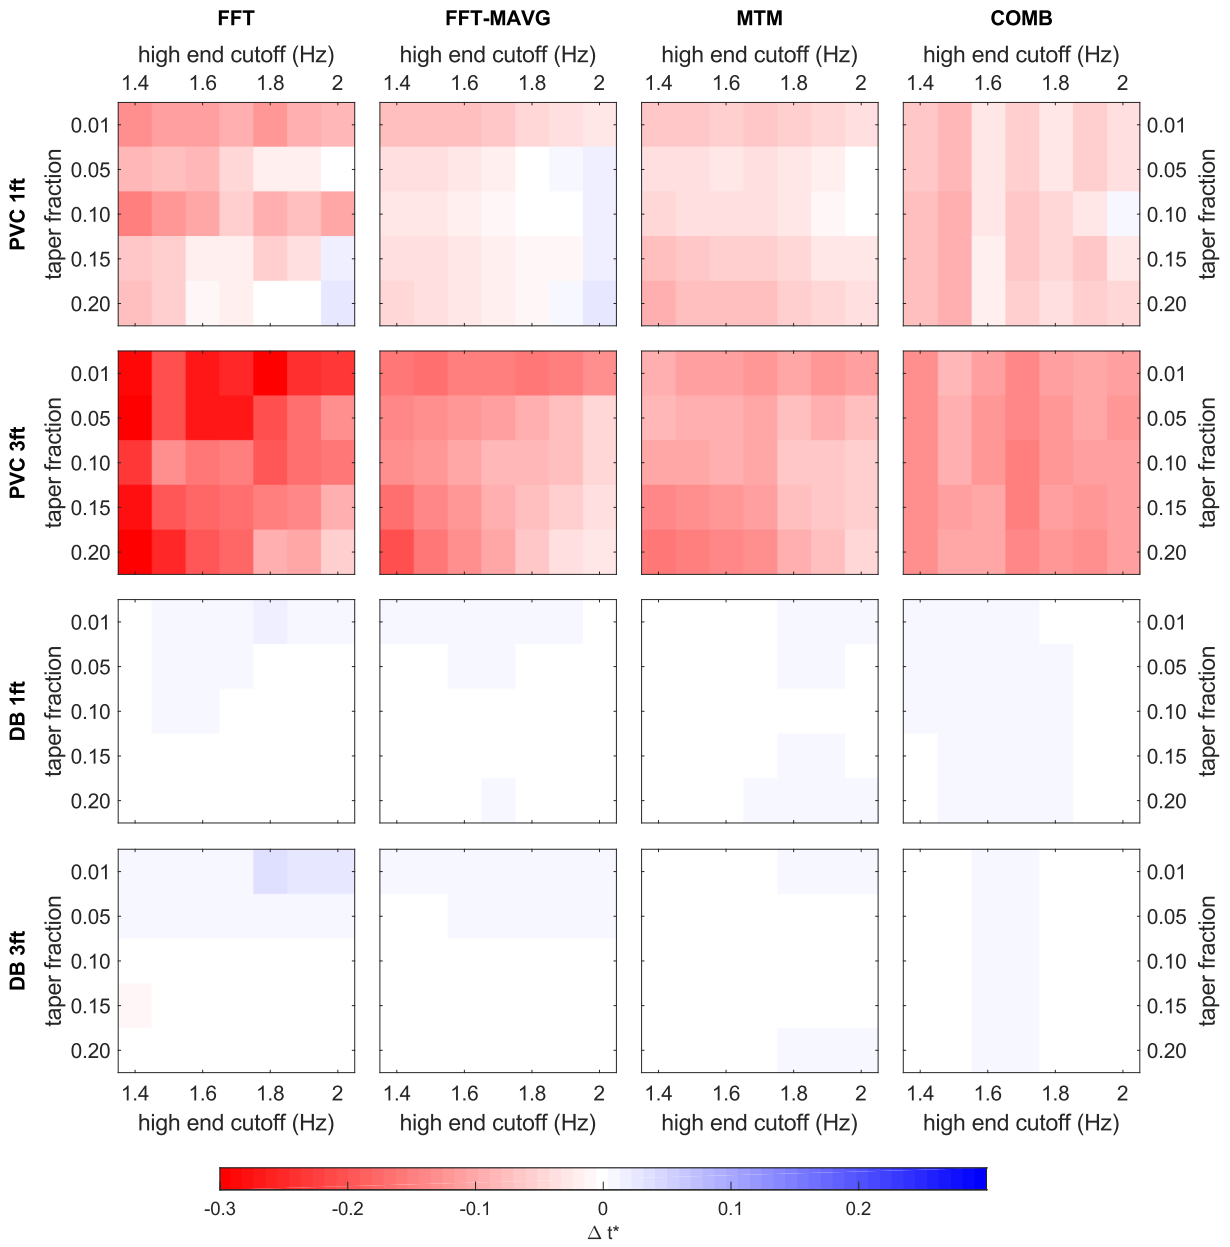

Figure 4. Comparison of $\Delta t*$ values obtained for four different stations (rows) using four different frequency domain methods (columns) and different parameters. Methods are labelled at the top of the figure. FFT—Spectral ratios using spectra derived from fft; FFT-MAVG—Spectral ratios using spectra derived from fft and smoothed using a 5-point moving average; MTM—Spectral ratios using spectra derived from the multitaper method (Thomson 1982); COMB—Simultaneous fitting of amplitude and phase information using the comb method (Eilon & Abers 2017). Stations are labelled on the left of the figure according to their depth of burial and whether they were encased in sand-filled PVC pipes (PVC) or not (DB). For each station-method pair, results are shown for different high-end cut-off frequencies and different Tukey window taper fractions as indicated on the plots.

the time domain since they should ideally capture multiple cycles of the lowest frequency of interest. We first explore the effect of varying the window over which the waveforms are fit in the time domain method, and then explore the effect of varying the time windows on the frequency domain methods.

We define nine different starting points and nine different ending points for the data time window (Fig. S1) and repeat the measurement for each of the stations and for all possible combinations of starting and ending points (Fig. 8), using a bandpass filter with corners at 0.3 and 1.7 Hz. Stations 1DB and 3DB show negligible differential attenuation, with $\Delta t*$ values of 0.00 or 0.01 s (Fig. 8). In contrast, stations 1PVC and 3PVC show substantial variability of $\Delta t*$, in many cases returning measurements significantly different from zero (Fig. 8). Using the time-domain method, the median $\Delta t*$ at station 3PVC for the various time windows is the correct value (zero), which eludes all the frequency-domain techniques. However, for some time windows the $\Delta t*$ estimates have magnitudes approaching or exceeding 0.16 s. For both 1PVC and 3PVC the range of $\Delta t*$ values returned by the time-domain method using the various time windows is comparable to the range seen for FFT-derived spectra using different frequency bands and a 10 s window (Fig. 7).

For the 3PVC records, different time windows result in $\Delta t*$ measurements that have different sign. To illustrate why this is the case,

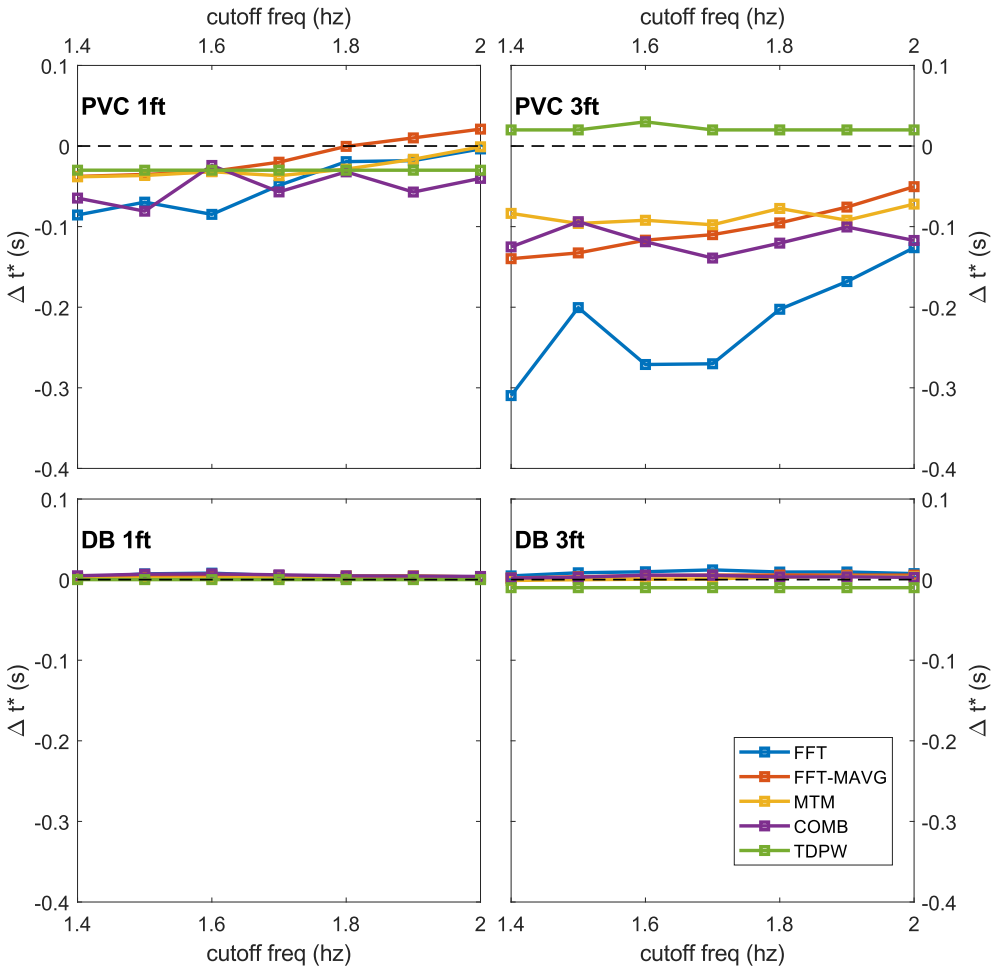

Figure 5. Comparison of the effect of changing the high-end cut-off frequency on $\Delta t*$ from different frequency domain methods for each of the stations as indicated in each panel. All results shown correspond to a Tukey window fraction of 0.05. Legend on bottom right panel applies to all panels. Frequency-domain method abbreviations are as in Fig. 4. TDPW—Time domain, preferred window.

we show the best-fitting numerically attenuated waveforms for the 3PVC record and time windows with negative and positive $\Delta t*$ results of 0.16 s (Fig. 9). The most important changes in the waveform are the slopes around the global minimum (main trough), and the relative amplitudes of the local maxima that precede and follow it. In the first case, where the waveform matching suggests the observed trace is less attenuated than the reference trace from station SPMN, the synthetic waveform minimizes the misfit by approximating the relative amplitude of the first local maximum. Doing so requires a negative $\Delta t*$ value, but this generates a trace with slightly steeper slopes and narrower width than in the reference trace. In the second case, the time window excludes both local maxima and the misfit is minimized by a positive $\Delta t*$ that slightly widens the main trough. This helps the synthetic waveform better approximate the slightly smaller slope on the earlier side of the trough. In the final case, the time window is such that the slopes on both sides of the main trough are included, but not the local maxima on either side. This seems to represent a good compromise, and returns a value of 0.02 s. However, it is clear that this time window could not be identified as ideal *a priori*.

These examples highlight that the waveform misfit is more sensitive to the relative amplitudes of earlier and later excursions than to the width and character of the main pulse. Unfortunately, excursions following the main pulse can—as in the case of 3PVC—be dominated by signal-generated noise (Fig. 3), but even the relative amplitudes of the first peak and trough can be affected by ambient noise with periods that are in the same range as those of the signal of interest, and thus cannot be removed by frequency filtering. On the other hand, working in the time domain or with the comb approach provides the advantage of being able to assess whether or not a synthetic waveform is accurately reproducing the desired characteristics of the observed waveform, beyond a simple metric that can be overwhelmed by parts of the waveform whose character is being determined by factors other than intrinsic attenuation. An analyst looking at Fig. 9(a), for example, would be able to classify that measurement as erroneous on account of the different width and slopes of the main excursion in the synthetic trace with respect to the observation. Similarly, in Fig. 9(b), the slope of the main trough immediately after the selected window is poorly matched by the synthetic, which could lead to the rejection of the measurement. We argue that this is a strength of the time domain and comb methods, a point on which we expand in the discussion section.

For assessing the effect of the time window on the frequency domain results, we repeated the analysis from the previous section (i.e. different combinations of F_{high} and T_{frac}) for each of the methods and for time windows starting at the first break and with lengths of 5, 15 and 20 s (in addition to the 10 s window results described above). For stations 1DB and 3DB, since the waveforms are nearly identical

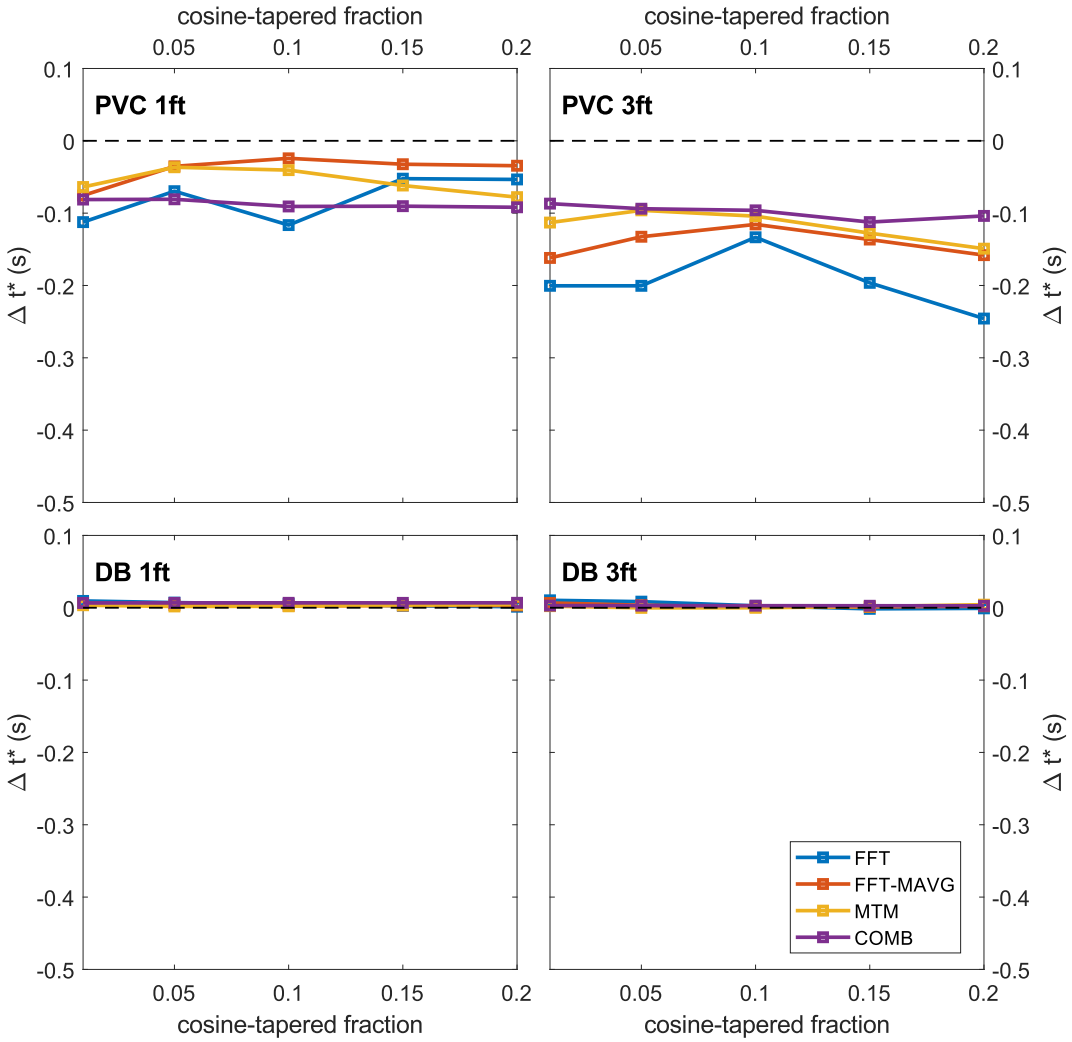

Figure 6. Comparison of the effect of changing the Tukey window taper fraction $\Delta t*$ from different frequency-domain methods for each of the stations as indicated in each panel. All results correspond to a high-end cut-off frequency of 1.5 Hz. Legend on bottom right-hand panel applies to all panels. Method abbreviations are as in Fig. 4.

to those of the reference station SPMN, the time window chosen had no impact on the results. For 1PVC and 3PVC however, we observed that the range of $\Delta t*$ values obtained for different implementation parameters increased for larger time windows (Fig. 10). In the most extreme case, the FFT method with 20 s of data produced a range of $\Delta t*$ of 0.6 s. The relative stability of the different frequency-domain methods is consistent with our findings in Section 4. Unsmoothed FFT-derived spectra resulted in larger scatter, smoothed FFT spectra and MTM spectra produced comparable scatter, and the scatter from the comb method was the smallest. The magnitude of the median value of the ensemble of observations for each method also increases with increasing length of the time window for the spectral ratio implementations. This is less true for the comb method, which appears to be more robust with respect to the length of the time window.

The spectral ratio method is successful in returning near-zero values for 1PVC and 3PVC when a short time window of 5 s is used (Fig. 10). We note however that typical time windows used in spectral ratio studies are considerably longer, owing to the fact that a 5 s window encompasses very few cycles of the relevant frequencies, making it difficult to robustly estimate the spectra. For all the time windows, the spectral ratio as a function of frequency has a 'v' shape, that is much more pronounced for longer windows (Fig. 11). In theory, each panel in Fig. 11 would display a horizontal line if there no signal-generated noise, and a straight line with non-zero slope if there was differential intrinsic attenuation between the records. We note that a straight line is a relatively poor approximation to the data in all cases, but unlike for the time domain, there is no basis to determine which parts or features of the observations are more likely to be caused by intrinsic attenuation.

6 DISCUSSION AND CONCLUSIONS

The scatter in $\Delta t*$ that we have observed between different stations and different measurement parameters is large relative to the range measured across different tectonic regions. For example, Cafferky & Schmandt (2015) found variations of $\Delta t*$ across the conterminous U.S. with a range of 0.4–0.8 s, depending on the frequency band being considered. It follows that the range of $\Delta t*$ (0.1–0.3 s) estimated for station 3PVC using a 10 s window and different implementations of the FFT spectral ratio is 40–75 per cent of the range across the conterminous U.S. Furthermore, the range of estimates

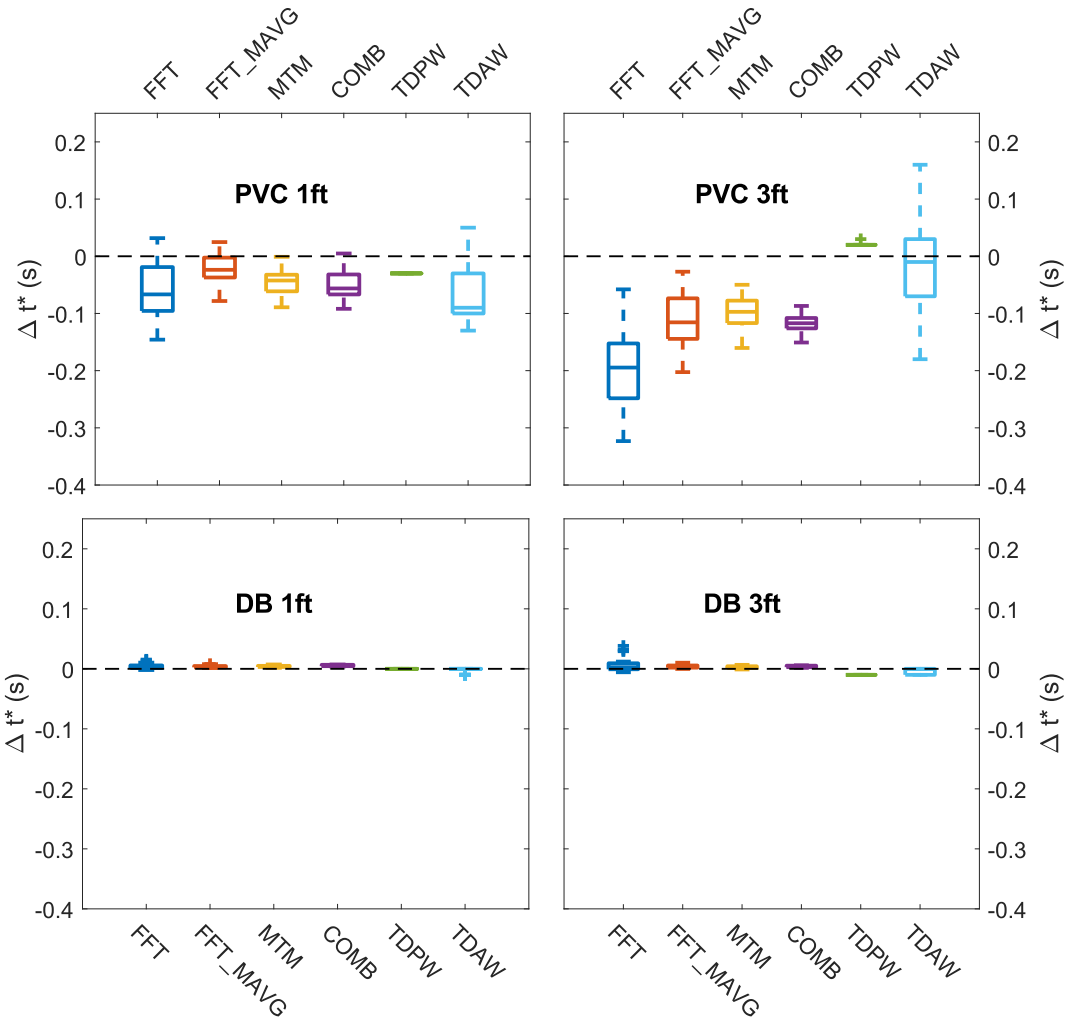

Figure 7. Summary of results of 35 combinations of processing parameter (presented in Fig. 4) plus results from the time domain method. Colour scheme corresponds to that used for Figs 5 and 6 for each method, methods are also indicated on the panels. Frequency-domain method abbreviations are as in Fig. 4. TDPW—Time domain, preferred window; TDAW—Time domain, all windows. Horizontal lines represent the median of each set of measurements, rectangles span the interquartile range and whiskers span the interdecile range. If present, outliers are represented by crosses.

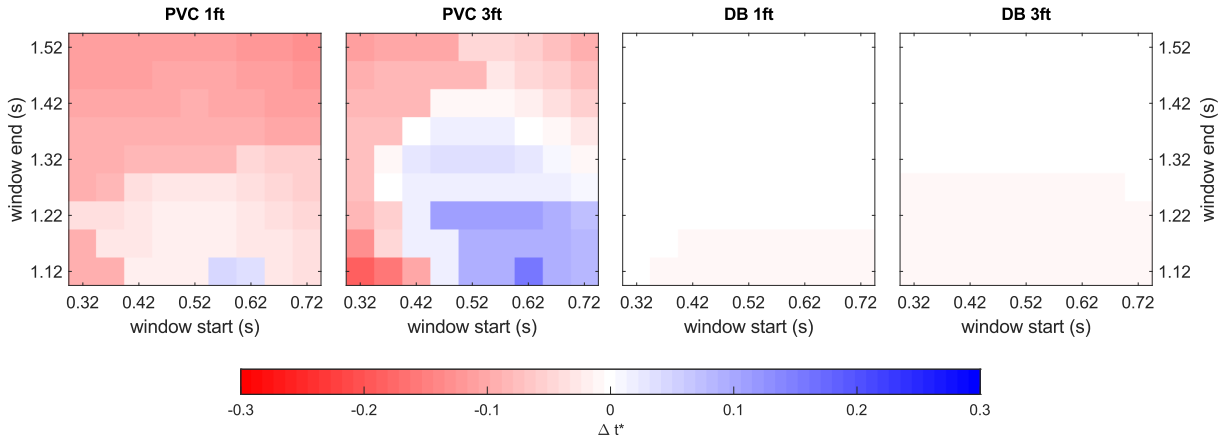

Figure 8. Comparison of $\Delta t*$ values obtained using the time domain method on different time windows and on each of the records as indicated above each plot.

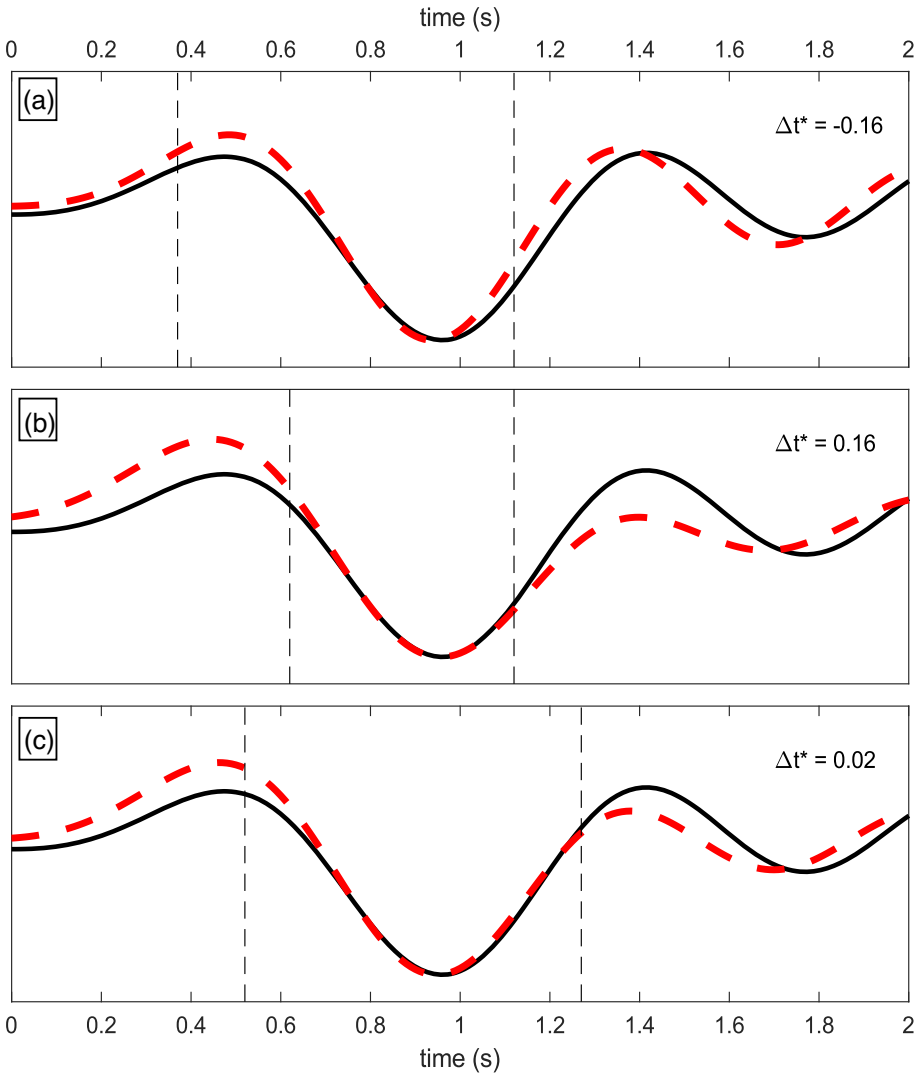

Figure 9. Observed traces (black) corresponding to station 3PVC and best-fit synthetic (dashed red) using three different time windows indicated by vertical dashed lines. In each case, $\Delta t*$ values used to produce the synthetic trace are shown in the upper right. Note that for the example with negative $\Delta t*$ the main pulse of the synthetic waveform is notably narrower than is observed.

using different windows in the time domain is similarly large. Even more worrisome is the fact that, when using 20 s windows, FFT-derived spectra produce estimates of $\Delta t*$ with a range equivalent to 75–150 per cent of the continental range. Notably, this large spread in measurements occurs in spite of the fact that we are measuring $\Delta t*$ using records of the same event, the same two stations, and that these two stations are separated by less than 3 m.

In the frequency domain, the instability of the measurement with respect to the parameters we tested is reduced when using smoother spectra. While Adams & Humphreys (2010) concluded that unsmoothed FFT spectra yielded better estimates of $\Delta t*$, we reach the opposite conclusion here. Our tests suggest that using the multitaper method or smoothing the FFT spectra with a moving-average filter at least provides a result that is more stable, though the impact of signal-generated noise cannot be easily corrected. Even for these more stable results, the ranges are still on the order of 0.1 s for both 1PVC and 3PVC (Fig. 7) using 10 s windows, and larger when using longer windows.

Using the comb method, the ranges of the measurements were comparable to those using the various implementations of the spectral ratio method, if somewhat smaller for the 3PVC case. Additionally, the comb ranges were less sensitive to the length of the time window. With respect to the median $\Delta t*$ estimate, values obtained with the comb method were similar to those obtained using the other frequency-domain approaches, indicating that consideration of the phase spectrum does not make this method totally immune to problems arising from signal-generated noise. Given the sensitivity of the frequency-domain methods to the specifics of their implementation, we would recommend that studies using these methods report the details of how their measurements were made. Although information such as the taper fraction of the Tukey window being used may seem like minutiae, we show here that this choice can affect the measurements and ultimately the conclusions of a study.

As a final estimate of $\Delta t*$ from each of the frequency-domain methods, we take the median of the 35 values obtained in each case (7 high-end cut-off frequencies and 5 Tukey window taper fractions, Figs 4 and 7). For 3PVC, these values deviate significantly from the expectation (zero differential attenuation). Once again referencing

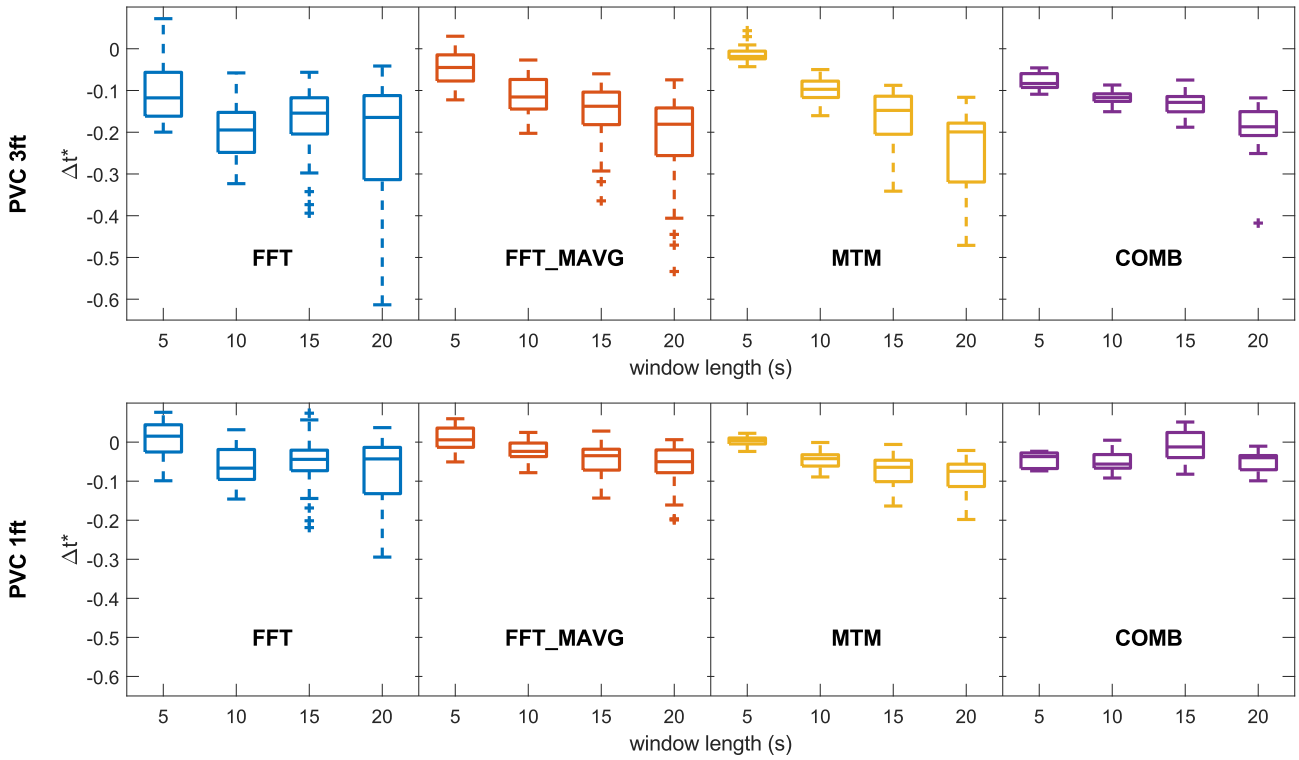

Figure 10. Summary of $\Delta t*$ results obtained using different frequency domain methods and time windows of different length. Colour scheme and method abbreviations are the same as in Fig. 4. Horizontal lines represent the median of each set of measurements, rectangles span the interquartile range and whiskers span the interdecile range. If present, outliers are represented by crosses.

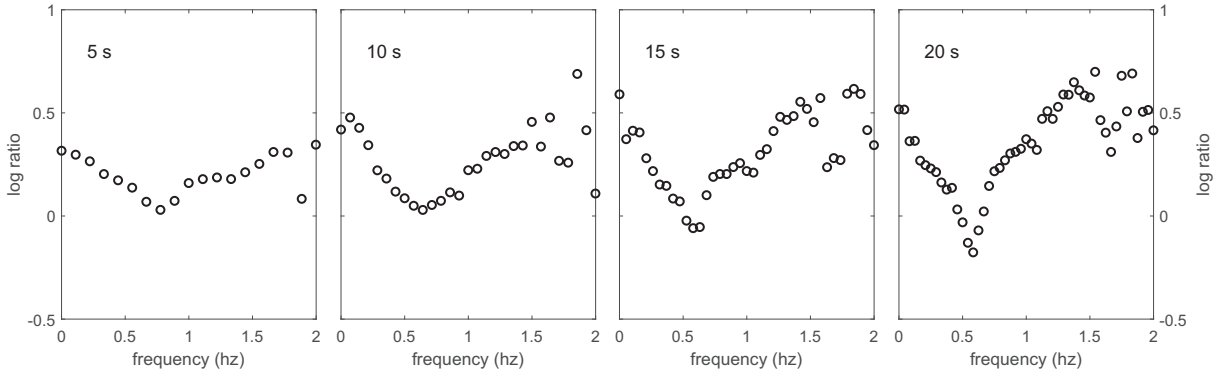

Figure 11. Log spectral ratio in the frequency range of interest using MTM-derived spectra and time windows of different length as indicated in each panel.

the Cafferky & Schmandt (2015) ranges described above, the $\Delta t*$ we find using FFT-derived spectra represents 25–50 per cent of the U.S. range, or 12–25 per cent of the range using smoother spectra, or the comb method. In either case, a deviation of this magnitude would affect whether the measurement represents a near-background value, a significant anomaly, or a highly prominent anomaly. For 1PVC, the deviation from zero is smaller and the ranges of measurements include $\Delta t*$ of zero.

Together, these observations highlight two different problems with measuring $\Delta t*$ on teleseismic *P* waves using frequency-domain methods: (1) the measurements are significantly dependent on the set of parameters and pre-processing choices made in the implementation of the technique and (2) the measurements may reflect the influence of signal-generated noise. As a result, studies relying on these methods may be difficult to reproduce and results may be misleading if there is not redundancy in the data. Furthermore, there are more parameters that can be varied, beyond what we have tested. In particular, the number of Slepian tapers used in the multitaper method, the number of samples used to pad the traces and whether or not to pre-filter before taking the spectra. Although we did not systematically test the effect of these choices, limited testing (not included) showed the number of tapers and the zero-padding had a noticeable effect, while the effect of pre-filtering was negligible. At the same time, arguably the best results were obtained with the spectral ratio method using MTM-derived spectra and short (5 s) time windows (Fig. 10). A problem with that approach is that the source time functions are only as long as a couple of cycles of the relevant frequencies, making it difficult to confidently estimate the spectra, and this result may not be generalizable. However, it does seem to be robust that when using the spectral ratio method, shorter time windows yield more stable results.

The time-domain method was more successful than the frequency-domain methods in returning near-zero values of differential attenuation for the records from stations encased in PVC (Fig. 5 and 7). However, this accuracy relies on identifying the exact time window that minimizes the effect of signal-generated noise. It would be difficult for an analyst to choose the 'correct' time window based on inspection of the waveforms. Modest changes to the time window used resulted in considerably different estimates of $\Delta t*$. However, a visual inspection can in some cases make clear whether or not the result is spurious. This represents an important advantage of the time domain approach since it is possible for a trained analyst to assess the quality of a given measurement by visually comparing the synthetic and observed waveforms. The drawback is that this can become a very time-intensive process, and care must be taken to avoid subjectivity in the quality control step. Hypothetically, a metric could be devised to automatically determine the optimal time window for each observed waveform and/or to automatically cull the results afterwards. Simple metrics were only partially successful in replicating the culling done by analysts in the study of Bezada (2017) and the particularities of individual waveforms have made it thus far impossible to find a universally applicable metric. On the other hand, once observations from many different events are combined, what seems like a very important problem on the scale of a single measurement decreases in significance. After a careful process of visual culling of results, Bezada (2017) obtained nearly identical maps of regional teleseismic attenuation regardless of whether all the measurements, or only those that had passed quality control, were used in the inversion.

One conclusion from this study is that direct burial of seismic instruments in the local medium, without casing, appears the preferable deployment method. Our results at first glance suggest that the $\Delta t*$ artefacts only occur for particular vault designs. This does not make our findings a mere curiosity of our micro-array for two reasons: First, the existence of modern data centres and open data practices have made seismic data from many different temporary deployments readily available to any researcher. As a result, seismologists often use data that they did not collect, and thus installation standards are unknown. Some of those installations may be inherently noisier, and our results show this can have an effect on the stability of attenuation measurements. Furthermore, we may jointly analyse data from multiple deployments which may have used different deployment standards or hardware. Correcting for instrument response is a common practice that accounts for the differences in the sensors, but the site response is much more difficult to characterize and correct for.

Secondly, and more importantly, vault design or installation parameters are only one factor that affects the site response. Other factors may include effects of very shallow subsurface structure, surface topography, natural and artificial structures. In the data set that we have analysed, installation was the only of these parameters that varied, and so we can attribute artefacts in the measurement solely to this factor. This does not, however, preclude the possibility that two stations may have different site responses despite having identical installations. An example that potentially illustrates this case is the study of Cafferky & Schmandt (2015). Very short wavelength variations in dt* were found in that study. In some instances, extremely low values in one station contrast with extremely high values in an adjacent one (see their fig. 4a). This occurs, in spite of the fact that all of the stations used in that study were part of EarthScope's Transportable Array; and had uniform, standardized, high-quality vault installations. We conclude that although this study is only demonstrating the effect of installation choices, its conclusions may be generalizable to variations in the site response, more broadly defined.

In this work, we were able to determine for both time and frequency-domain methods which parameters gave the results most consistent with the known value of zero differential attenuation. Clearly, this is not a luxury that we normally have. For this reason and given the range of values that can be obtained even in the best cases, it is likely that we must accept that the uncertainty associated with any measurement of teleseismic *P*-wave attenuation is large. Therefore, success depends both on having redundancy and on proper treatment of those errors.

Redundancy is important because it may help isolate the site effects from the intrinsic attenuation signal and can be achieved either by combining measurements from many events, by having a dense station distribution, or, ideally, both. Attenuation studies in subduction zones have the advantage of high redundancy in the form of up to hundreds of local events in the data set. Such data density enables a tomographic imaging approach where site effects can be included as a model parameter that is inverted for (e.g. Sarker & Abers 1998; Stachnik *et al.* 2004; Nakajima *et al.* 2013). Away from subduction zones, similar data density is very difficult to achieve, yet the problem of site effects can still be addressed. Cafferky & Schmandt (2015) use smoothing kernels to estimate regional intrinsic attenuation patters from noisy station-by-station estimates. Hwang *et al.* (2009) uses an inversion scheme that penalizes differences in $\Delta t*$ between stations closer than a reference distance, while Hwang *et al.* (2011) average results from stations within a certain radius. Bezada (2017) inverts for a smoothly varying intrinsic attenuation map and includes a 'station term' that absorbs very short wavelength structure and is akin to the site-effect model parameters in attenuation tomography studies. Either station or event terms were also considered in the Bayesian inversion approach adopted by Byrnes *et al.* (accepted). Given that in this study we analyse a single event, it is not possible to assess the relative merits of each of these approaches.

Our results show that site effects can go beyond a static shift in $\Delta t*$ and affect the stability of the measurement, and thus its uncertainty. To assess this uncertainty, measurements (whether in the time or the frequency domain) can be made multiple times using different parameters, and these errors should be taken into account when multiple measurements are combined. Adams & Humphreys (2010) tackle the problem of measurement instability by making many different $\Delta t*$ measurements (both in the time and frequency domains) for each event-station pair, treating them all as independent estimates and inverting them together, an approach also taken by Eilon & Abers (2017). Cafferky & Schmandt (2015), combine spectral ratio measurements from multiple events but analyse measurements from different frequency bands separately. Bezada (2017) and Bezada & Smale (2019) invert measurements from different sets of events and take the mean of the different inversion results as the preferred model, and the standard deviation at each model node as a measure of the uncertainty. Another promising option is using so-called hierarchical inversion approaches, where the uncertainties in the measurements are treated as additional unknowns. Byrnes *et al.* (accepted) implement this approach to infer Q structure, with encouraging results.

This study has focused on characterizing problems with the measurement of $\Delta t*$. In closing, we emphasize that better understanding of these problems helps to better manage them. With redundancy that helps isolate the site response, and with proper treatment of the measurement and related uncertainties, whether in the time or the frequency domain, one can produce robust estimates of differential

attenuation that can lead to more confident interpretations of Earth structure.

ACKNOWLEDGEMENTS

We are thankful to the property owner who gave us permission to install our stations in his land. Two anonymous reviewers provided comments that helped significantly improve the manuscript. This work was partially funded by the National Science Foundation (NSF) under grant EAR-1827277 to the University of Minnesota.

REFERENCES

Adams, D.C. & Humphreys, E.D., 2010. New constraints on the properties of the Yellowstone mantle plume from P and S wave attenuation tomography, *J. geophys. Res.: Solid Earth,* **115**(B12), doi:10.1029/2009JB006864.

Azimi, S.A., Kalinin, A.V., Kalinin, V.V. & Pivovarov, B.L., 1968. Impulse and transient characteristics of media with linear and quadratic absorption laws, *Izv. Phys. Solid Earth,* **2,** 88–93.

Bao, X., Dalton, C.A. & Ritsema, J., 2016. Effects of elastic focusing on global models of Rayleigh wave attenuation, *Geophys. J. Int.,* **207**(2), 1062–1079.

Bezada, M.J. & Smale, J., 2019. Mantle structure controls the location of intracontinental seismicity in Australia, in Press.

Bezada, M.J., 2017. Insights into the lithospheric architecture of Iberia and Morocco from teleseismic body-wave attenuation, *Earth planet. Sci. Lett.,* **478,** 14–26.

Byrnes, J., Bezada, M.J., Long, M.D. & Benoit, M., 2019. Thin lithosphere beneath the central Appalachian Mountains: constraints from seismic attenuation beneath the MAGIC array, *Earth planet. Sci. Lett,* in press.

Cafferky, S. & Schmandt, B., 2015. Teleseismic P wave spectra from USArray and implications for upper mantle attenuation and scattering, *Geochem. Geophys. Geosyst.,* **16**(10), 3343–3361.

CTBTO, 2017. CTBTO executive secretary lassina zerbo on the unusual seismic event detected in the Democratic People's Republic of Korea: CTBTO preparatory commission, Retrieved June 12, 2018, from https://www.ctbto.org/press-centre/press-releases/2017/ctbto-executive-secretary-lassina-zerbo-on-the-unusual-seismic-event-detected-in-the-democratic-peoples-republic-of-korea/.

Dalton, C.A., Ekström, G. & Dziewonski, A.M., 2009. Global seismological shear velocity and attenuation: a comparison with experimental observations, *Earth planet. Sci. Lett.,* **284**(1), 65–75.

Eilon, Z.C. & Abers, G.A., 2017. High seismic attenuation at a mid-ocean ridge reveals the distribution of deep melt, *Sci. Adv.,* **3**(5), e1602829, doi:10.1126/sciadv.1602829.

Gaebler, P. *et al.*, 2019. A multi-technology analysis of the 2017 North Korean nuclear test, *Solid Earth,* **10**(1), 59–78.

Goes, S., Govers, R. & Vacher, P., 2000. Shallow mantle temperatures under Europe from P and S wave tomography, *J. geophys. Res.,* **105**(B5), 11 153–11 169.

Hammond, W.C. & Humphreys, E.D., 2000. Upper mantle seismic wave velocity: effects of realistic partial melt geometries, *J. geophys. Res.: Solid Earth,* **105**(B5), 10 975–10 986.

Hwang, Y.K., Ritsema, J. & Goes, S., 2011. Global variation of body-wave attenuation in the upper mantle from teleseismic P wave and S wave spectra, *Geophys. Res. Lett.,* **38**(8), L06308, doi:10.1029/2011GL046812.

Hwang, Y.K., Ritsema, J. & Goes, S., 2009. Spatial variations of P wave attenuation in the mantle beneath North America, *J. geophys. Res.: Solid Earth,* **114**(B6), B06312, doi:10.1029/2008JB006091.

Jacobsen, S.D., Jiang, F., Mao, Z., Duffy, T.S., Smyth, J.R., Holl, C.M. & Frost, D.J., 2008. Effects of hydration on the elastic properties of olivine, *Geophys. Res. Lett.,* **35**(14), doi:10.1029/2008GL034398.

Nakajima, J., Hada, S., Hayami, E., Uchida, N., Hasegawa, A., Yoshioka, S., Matsuzawa, T. & Umino, N., 2013. Seismic attenuation beneath northeastern Japan: constraints on mantle dynamics and arc magmatism, *J. geophys. Res.: Solid Earth,* **118**(11), 5838–5855.

Sarker, G. & Abers, G.A., 1998. Deep structures along the boundary of a collisional belt: attenuation tomography of P and S waves in the Greater Caucasus, *Geophys. J. Int.,* **133**(2), 326–340.

Stachnik, J.C., Abers, G.A. & Christensen, D.H., 2004. Seismic attenuation and mantle wedge temperatures in the Alaska subduction zone, *J. geophys. Res.: Solid Earth,* **109**(B10), B10304, doi:10.1029/2004JB003018.

Teng, T.-L., 1968. Attenuation of body waves and the Q structure of the mantle, *J. geophys. Res.,* **73**(6), 2195–2208.

Thomson, D.J., 1982. Spectrum estimation and harmonic analysis, *Proc. IEEE,* **70**(9), 1055–1096.

SUPPORTING INFORMATION

Supplementary data are available at *GJI* online.

Figure S1. Normalized velocity records of the DPRK nuclear test from all stations in our array. All traces are normalized by their maximum excursion and are colour-coded as indicated in the legend. The origin on the time coordinate represents the first arrival time. Dashed lines on either side of the largest excursion represent start and end points for different time windows used in the time-domain estimation of $\Delta t*$. Results for all the time windows arising from the 81 possible combinations of starting and ending points are shown in Fig. 8 and the range of measurements for each record is shown in Fig. 7. Thicker dashed lines indicate the preferred time window used for the analysis described in Section 4.

Please note: Oxford University Press are 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.