Estimation of Spectral Power Laws in Time-Uncertain Series of Data with Application to the GISP2 \(δ^{18}\)O Record

Estimation of Spectral Power Laws in Time-Uncertain Series of Data with Application to the GISP2 δ^{18}O Record

A. Rhines,[1] P. Huybers,[1]

P. Huybers, Department of Earth and Planetary Science, Harvard University, 20 Oxford St., Cambridge, MA 02138, USA.

A. Rhines, Department of Earth and Planetary Science, Harvard University, 20 Oxford St., Cambridge, MA 02138, USA. (arhines@fas.harvard.edu)

[1]Department of Earth and Planetary Science, Harvard University, 20 Oxford St., Cambridge, MA 02138, USA.

Abstract. Errors in the timing assigned to observations degrade estimates of the power spectrum in a complicated and non-local fashion. It is clear that timing errors will smear concentrations of spectral energy across a wide band of frequencies, leading to uncertainties in the analysis of spectral peaks. Less understood is the influence of timing errors upon the background continuum. We find that power-law distributions of spectral energy are largely insensitive to errors in timing at frequencies much smaller than the Nyquist frequency, though timing errors do increase the uncertainty associated with estimates of power-law scaling exponents. These results are illustrated analytically and through Monte Carlo simulation, and are applied in the context of evaluating the power-law behavior of oxygen isotopes obtained from Greenland ice cores. Age-errors in layer counted ice cores are modeled as a discrete and monotonic random walk that includes the possibility of biases toward under- or over-counting. The $\delta^{18}O_{ice}$ record from the Greenland Ice Sheet Project 2 is found to follow a power-law of 1.40 ± 0.19 for periods between 0.7 and 50 ky, and equivalent results are also obtained for other Greenland ice cores.

1. Introduction

Power-law behavior, i.e. when spectral power scales proportionately with frequency raised to an exponent, has proven a useful description for climate over a wide range of timescales (e.g. Wunsch 1972; Vyushin and Kushner 2009; Shackleton and Imbrie 1990). In order to span a wider range of timescales, some studies have combined multiple spectral estimates from low-resolution, long-record proxy data and high-resolution, modern instrumental data. Harrison [2002] produced a patchwork spectrum from many sea-level records that generally followed a power law with an exponent of minus two extending over periods from $\sim$ 1 yr to $\sim$ 600 Myr. Notable, however, is that sea level variability scaled more nearly with a power law of -1.4 at periods shorter than 100 years. Using a similar patchwork approach, Huybers and Curry [2006] compared many records reflecting sea- and land-surface temperature from the instrumental era and paleorecord and found that temperature variability followed power laws ranging from -0.6 (tropical) to -0.4 (high-latitudes) at decadal to centennial timescales, whereas steeper power laws from -1.6 (tropical) to -1.3 (high-latitudes) existed at longer periods.

In both Harrison [2002] and Huybers and Curry [2006], the lower frequency, more steeply scaling variability is from paleoclimate data, while the higher-frequency and more shallow scaling variability is generally from instrumental data. The question arises whether the steepening of the power law at centennial timescales might be an artifact of the errors present in certain proxy timeseries.

There are many potential sources of error in any proxy timeseries. Among other complications, the data are sparse, representative of quantities integrated over poorly defined

geographical areas, generally encoded as a function of multiple physical and possibly biological variables, and uncertain in measurement magnitude (e.g. Bradley 1999). Proxies are also subject to pervasive uncertainty in timing. Here we focus on the influence of timing errors upon spectral estimates of the background continuum because such errors are common but have received relatively little attention.

Studies of error propagation in spectral analysis have primarily addressed the influence of measurement noise. Indeed, most of the standard methods were developed for engineering applications where the assumption of perfect timing is normally adequate. However, timing errors are generally non-negligible in paleoclimate data. For example, even the meticulously layer-counted GISP2 record has time-uncertainty equal to about 2% of the estimated age (Alley et al. 1997). The case of jitter (white timescale noise) was explored by Moore and Thomson [1991], who showed that even small timing errors can result in large changes in the power spectral estimate of an oceanographic dataset. Extensions by Thomson and Robinson [1996] suggested that more realistic correlated errors have greater consequences for spectral estimation, although their approach was not tractable outside the assumption of nearly uniform sampling. Mudelsee et al. [2009] developed statistical tests to estimate the frequency and significance of time-uncertain spectral peaks using Monte-Carlo methods with the Lomb-Scargle periodogram, applying bootstrap to correct the estimator bias. This small literature represents an important step forward in grappling with the ubiquitous issue of time uncertainty in all but the most recent instrumental climate records. However, the effect of age model errors such as those encountered in paleoclimate timeseries on the estimation of power-law climate spectra has not yet been explored.

2. Time-induced Changes in the Power Spectrum

The power spectrum, $P(f)$, of a continuous signal, $x(t)$, can be estimated using the periodogram (Bracewell 1986),

$$P(f) = |F(f)|^2 \equiv \left| \int_{-\infty}^{\infty} x(t) e^{-2\pi i f t} dt \right|^2 . \tag{1}$$

The expectation of the periodogram, $E[P(f)]$, is said to exhibit power-law scaling if,

$$E[P(f)] = a f^{\beta}. \tag{2}$$

To the extent that the power spectrum of a climate timeseries exhibits power-law scaling, the logarithm behaves linearly, $\log(P) = \beta \log(f) + \log(a)$. Below we explore the implications of replacing the signal, $x(t)$, with a time-uncertain version, $x(t')$. Here, x is not a function, but rather a representation of a series of measurements placed on a timescale, t'. We define this uncertain estimate of the timescale as, $t' \equiv t + \epsilon(t)$, where $\epsilon(t)$ is the time error.

Errors in t' distort the integral in Eq. 1 because changes in the timescale alter the frequency and phase of the Fourier components of the signal. We wish to determine the ways in which these timing errors alter the inferred spectrum, $P'(f)$, of a time-uncertain power-law signal, beginning with an illustrative example. Although real age errors will typically take the form of a random walk, we first consider a simpler case where time error grows linearly between the initial time, t_i, and the switch time, t_s, and then shrinks linearly between t_s and the final time, t_f,

$$\epsilon(t) = \begin{cases} \gamma_1 t & \text{if } t_i \le t \le t_s, \\ \gamma_1 t_s + \gamma_2 (t - t_s) & \text{if } t_s < t \le t_f. \end{cases} \tag{3}$$

The error rate, γ, is equal everywhere to $d\epsilon/dt$, and γ_2 is here defined as $-\gamma_1 t_s/(t_f - t_s)$, such that the total length of the timeseries is unchanged. This leads to a distorted representation of the signal — the first half is stretched, while the second half is compressed. See Fig. 1a,c for an illustration of this timing error applied to a red-noise signal. How will such timing errors influence the spectral estimate of narrow and broadband features present in $x(t)$?

Our approach is to examine the two segments of the record characterized by different temporal distortions independently, and then combine their spectra to estimate the spectrum of the full signal. That is, the signal can be decomposed into two segments by applying rectangular windows,

$$x(t) = x(t)\Pi\left(t, t_i, t_s\right) + x(t)\Pi\left(t, t_s, t_f\right),$$

where the windowing function, Π, is defined as,

$$\Pi(t, t_1, t_2) = \begin{cases} 1 & \text{if } t_1 \leq t < t_2, \\ 0 & \text{otherwise.} \end{cases}$$

Such windowing introduces sidebands due to the Gibbs phenomenon (e.g., Priestley 1994). Furthermore, the sum of the spectral estimates of the individual segments will differ from the spectral estimate obtained from the entire segment owing to differences in frequency resolution and interactions of the phase across the two segments, but in the synthetic experiments described later, we show that the average influence of these effects is negligible. Note that segmenting timeseries, computing their spectral estimates, and then averaging is a common procedure for estimating the spectrum of a noisy timeseries (Bartlett 1950).

If $x(t)$ contains a periodic component with frequency, $f_\circ$, the time errors (Eq. 3) will shift the variability to lower and then higher frequencies, f_1 and f_2, defined by,

$$f_\circ = (1+\gamma_1)f_1 = (1+\gamma_2)f_2, \tag{4}$$

and the resulting spectral estimate will split the original peak in two,

$$P' \approx P'_1 + P'_2 = a_1\delta(f-f_1) + a_2\delta(f-f_2), \tag{5}$$

where $\delta(f)$ is the Dirac delta function. Here a_1 and a_2 are positive constants whose magnitude will depend upon the length of the record segments and the normalization conventions that are used in reporting spectral power. In practice, the samples are taken over finite window lengths, so that the peaks at the inferred frequencies are sinc functions whose resolution will depend on the scope of time errors and the length of the record. If the difference between the two frequencies is small, the two peaks may not be resolved and the effect would be to simply blur the original peak.

Interestingly, while time errors significantly distort estimates of the power spectrum in the vicinity of spectral peaks, power-law scaling estimates obtained from stretched and squeezed time series appear largely intact (Fig. 1b,d). This insensitivity of power-law scaling estimates to time errors can be understood from the self-similarity of power-law signals. If P_1 and P_2 are power-law spectra as in Eq. 2, their inferred spectra are simply scaled and frequency shifted in proportion with the rate of change of the time error (Eq. 4),

$$P'_1 = a_1\,(1+\gamma_1)^\beta\, f^\beta, \tag{6}$$

as follows from the similarity theorem (e.g., Bracewell 1986, pp.101—103), and likewise for P_2'. The logarithm of the resulting spectral estimate is then,

$$\begin{aligned} \log(P') &\approx \log(P_1' + P_2') \\ &= \beta \log(f) + \log\left(a_1(1+\gamma_1)^\beta + a_2(1+\gamma_2)^\beta\right), \end{aligned} \tag{7}$$

where the identity that $\log(a+b) = \log(a) + \log(1+b/a)$ is used. The constant value in Eq. 7 is complicated, but the logarithmic scaling of $P'(f)$ with frequency according to β is unaffected when compared with Eq. 2. Although a simple example, Eq. 7 illustrates how power-law scaling can remain invariant in the presence of timing errors. A linear rescaling of the timescale of a signal does not affect a spectral power law. If the power law is an approximate description of a noisy discrete spectrum (as is typically the case), the estimate of that power law is also unaffected by a linear rescaling of the timescale.

This line of reasoning can be extended to a more general case, in which the rate of time error changes numerous times over the course of a record. As with the two-segment case, we view a timeseries which has been variously stretched and squeezed by N changes in γ as a composite of N shorter segments $x_n(t)$. Using a similar segmenting approach, the power spectrum of the individual segments will follow the same frequency scaling as Eq. 7, and give an expected power spectral estimate of $x(t')$ that remains proportional to f^β.

Segments of a signal following a spectral power law still display that same power-law after being differentially compressed or stretched, at least over the resolved frequencies and for the simple piece-wise manner in which the spectrum is estimated. The suggestion is that time errors do not distort the expectation of estimates of β. In the next section we examine more general timing errors and more general estimates of the power spectrum — and find similar behavior.

3. Synthetic Experiments

We now wish to determine whether the simple result from the previous section holds in practice, and to examine the influence of more realistic time-uncertainty upon more complex spectral structures. We adopt a Monte Carlo approach of generating random signals with a known spectral structure, distorting them in time, and then examining the resulting spectral estimate. Records are initially generated at very high resolution, in order to better approximate continuous signals and avoid sampling and edge effects. We model the time error as a finite-length random walk arising from cumulative counting errors. Though the counting error distribution is not Gaussian, its variance is finite and the expected cumulative error approaches a normal distribution after tens of counted layers. Details and physical motivation for this model are provided in the Appendix.

Though we apply an error model suitable for discretely layer-counted records, other tests using continuous error models suitable for chronologies based on accumulation rates (Huybers and Wunsch 2004) or using piece-wise linear errors as discussed in the foregoing section, all yield consistent results. Timing errors with a periodic or quasi-periodic component, or errors correlated with the value of the signal also provide equivalent results, despite their large effect on narrowband variability (Herbert 1994).

There are several possible ways to estimate power-law scaling and the value of β, whose results are not necessarily equivalent, particularly in the case of noisy and sparse data (Clauset et al. 2007). Our approach is to use an ordinary least-squares estimate of the spectral slope of $\log(P)$ versus $\log(f)$, where the mean of $\log(P)$ and $\log(f)$ is first subtracted so that the y-intercept is zero — the covariance that otherwise arises between the y-intercept and β makes it more difficult to interpret the results. P is estimated

using a standard periodogram. Another popular method is detrended fluctuation analysis (e.g. Vyushin and Kushner 2009) but which can be shown to be equivalent to the more common Fourier transform methods used here (Heneghan and McDarby 2000) up to differences in how the detrended fluctuations are weighted in estimating the slope.

Once a new timescale is generated, the timeseries must be resampled on a regular grid. Many methods are available for this interpolation, including mean, linear interpolation, random, or bootstrap infilling (Wilson et al. 2003; Mudelsee et al. 2009). In these tests, linear interpolation is used for the sake of simplicity, and because its distortion is easily identified and contained. Interpolation reduces the variance of a signal, but these effects are confined to the highest frequencies, i.e. near the Nyquist frequency, $f_{\mathrm{Ny}} \equiv 1/2\Delta t_{max}$. Thus, biases in power law fits of the continuum background can be minimized by using the appropriate frequency cutoff. Based on our experience with power-law signals, we find that a safe rule of thumb is to use a cutoff of $f_{\mathrm{Ny}}/2$, though the details associated with the signal structure and time-error could yield cases where other cutoffs are more appropriate. More generally, computing statistics using a range of cutoffs and determining the sensitivity of the result appears prudent when substantial time error is suspected.

First, an ensemble of 1,000 randomly generated $\beta = -2$ power law signals are sampled on timescales t' produced using the counting errors described in the Appendix (Fig. 2). The underlying timeseries have ten times the resolution of the signals used in the analysis, in order to avoid the high-frequency sampling bias discussed above. The average fit of the power-law across these randomly generated signals is unaffected by the errors in timing, remaining at -2 to within the precision of the fit. We do note, however, that the distribution of realized power laws is 8% wider when subject to timing errors, t',

of 5% of the length of the record than when compared against the ensemble of power laws not subject to timing errors. For the sake of comparing the spectra from different realizations, the total length of the signal is then constrained to the original length by subtracting the linear trend in time error between the first and last data point, making the discrete frequency axis identical for each realization. As shown in Section 2, such scaling in the time domain does not influence the power law in the frequency domain. The error structure then takes the form of a Brownian Bridge, discussed in more detail by Huybers and Wunsch [2004].

Next we examine a mixed timeseries, having periodic and power-law variability. The imposition of timing errors results in spectral distortion in the vicinity of the peak, while the remainder of the spectral estimate maintains the original power-law scaling (Fig. 2). Effects similar to those of narrowband distortion are observed when multiple background scaling regimes are present. For example, in a spectral break between two power-law scaling exponents, the distribution of power about the knee of the spectrum is smoothed out while the power-law regions are unchanged (Fig. 2). If discontinuities in the spectrum are rapid or numerous, much of the narrowband detail can be obscured by this sort of smoothing.

4. Application to GISP2

Insofar as the spectrum of the climate record scales as a power-law (or several power-law regimes), Sections 2 and 3 suggest that time-uncertainty will not affect estimates of β away from the Nyquist frequency of the largest timestep, at least in the expectation. Narrowband variations will be distorted by time errors, but the example of Section 2 suggests that their influence will tend to be localized in frequency. It is therefore useful

to investigate the uncertainty in the estimation of β for a real climate record due to age-model errors: is it best characterized as a power law — which is relatively insensitive to time errors — or to a noisy collection of narrowband processes, which can be distorted significantly by modest time errors? This question is explored by applying realistic time errors discussed in the Appendix to the GISP2 δ^{18}O record (Fig. 3) and examining the scaling of the resulting power spectra.

We evaluate the power spectral estimate of the GISP2 δ^{18}O record, using the counting error described in the Appendix to perturb the standard age-model record (Fig. 3a). The record is limited to 50 ky ago through the present, due to the larger and more poorly understood timing errors in deeper sections of the core. We note that there is no significant concentration of climatic precession energy. This could stem from a lack of sensitivity to precession forcing, nonlinearities, or the relative shortness of the record making it difficult to resolve bands with 21 ky periods. A fit is obtained for β in each realization, with spread evident under different age-models (Fig. 4a). The residuals of the ordinary least-squares fits are used to estimate a normal probability distribution of β for each realization, and these distributions are combined to produce an estimate of the uncertainty in β (Fig. 4c). For the most recent 50 ky of GISP2, the original timescale produces an estimate of $\beta_0 = -1.41 \pm 0.17$. When time-uncertainty is considered, the distribution shifts and broadens slightly such that $\beta_{\text{est}} = -1.40 \pm 0.19$. This is consistent with the slightly greater spread in realizations of β obtained when time errors were introduced into the synthetic records. Similar results are obtained when the timescale error is correlated with the δ^{18}O magnitude or, e.g., with orbital eccentricity or other climate forcing signals — such complications do not appear to influence the result in any significant way, nor do

they appreciably modify the power-law spectra obtained in Section 3. A similar analysis performed on the North Greenland Ice Core Project (NGRIP) core (Svensson et al. 2006) yields results equivalent to those of GISP2 when the same base time period and sampling rate are used for both records. Along the same lines, an analysis of the Greenland Ice Core Project (GRIP) record also produces results which agree with those of Ditlevsen et al. [1996] — namely, a spectral slope of -1.6 for periods greater than 200 yr — when the same time intervals and cutoff frequencies are used in analyzing both records. For both NGRIP and GRIP, inclusion of higher-frequency data made available by the higher sampling rate than GISP2 allows the break in the spectrum at centennial timescales to be resolved. This leads to much shallower power law estimates, apparently not as a consequence of distortion of the power spectrum, but because a linear fit is being improperly attempted over two distinct scaling regimes.

We find that the scaling exponent is approximately invariant under the expected time-uncertainty. Resampling the record over 1,000 realizations for a range of prescribed expected fractional error $E[|(t-t')/t|]$ at the oldest point, we estimate β for each timeseries (Fig. 5). When f_{max} equals $f_{\mathrm{Ny}}/2$, the fit remains within 5% of the unperturbed age-model fit until the age error is 6%, exceeding the estimated counting error by a factor of three, indicating that the scaling is robust under the expected time-uncertainty. Under more extreme age-model errors of 10% or more, there is greater spread in the estimates of β with the standard deviation growing from 0.17 to 0.2 and bias appears that can exceed 5%. In practice, we then expect relatively large time-uncertainty of 10% or more to increase the likelihood that scaling of the power spectral estimate will be incorrectly estimated due to interpolation biases if our rule of thumb is used. In contrast, interpo-

lation errors are important for much smaller expected cumulative timing error when the spectrum is estimated out to the highest possible frequencies.

5. Discussion and Conclusion

Estimates of power-law scaling exponents are insensitive to time-uncertainty in the expectation, and this invariance was demonstrated upon synthetic records (Section 3) and for the GISP2 δ^{18}O record (Section 4). This invariance can be understood from the power-law being preserved under shifts, stretches, and squeezes of a timescale (Section 2). Although time uncertainty is inevitable in paleoclimate records, magnitudes comparable to that in the GISP2 ice core do not appreciably affect estimates of power-law scaling. In particular, examination of the GISP2 power-law behavior under many plausible age model realizations yielded results virtually identical with those obtained using published age models. If errors exceed 10%, the distribution widens by more than 15% and the expectation begins to be affected through a bias introduced by interpolation. Furthermore, individual, realistic age-model realizations can result in power spectra that diverge significantly from the expectation, so that examination of power laws under a wide range of plausible timescales is prudent, especially if narrow-band concentrations of energy may be present.

A practical issue which will be encountered when resampling any record is that interpolating sample values at intermediate points reduces high-frequency variance, and this region of the spectrum should be avoided in subsequent analysis of power-laws. Limiting the analysis to frequencies below half the Nyquist frequency seems to be a useful rule of thumb, at least for the random walk age distortion explored here. This is important for

paleoclimate timeseries, which are often difficult to obtain at a high temporal resolution and are generally sampled non-uniformly in time.

For paleoclimate proxy data, the appropriate choice of a time error model differs according to the type of proxy and the manner in which its age was estimated. The error model presented in the Appendix should be broadly applicable for counted timescales, such as those associated with varved sediments, tree rings, and annually banded ice core records, all of which are expected to fundamentally follow a random walk pattern. The insensitivity of power law estimates to timing error holds for this counting-error model, as well as for continuous random walk error models and piece-wise error models. We have found no form of time errors, other than those with very large magnitudes, that give rise to significant changes in either the expected value or spread of power law estimates. It thus appears that timing error is not responsible for the steeper power-law scaling identified in paleoclimate records, relative to the scaling at higher frequencies that can be examined using instrumental records (Harrison 2002; Huybers and Curry 2006), though it remains to be seen whether the steeper power-law scaling can be attributed directly to dynamical processes.

Appendix

In order to generate appropriate timing errors, we require a description of the process by which age models are created. Paleoclimate signals are generally recorded in some accumulating medium, e.g., ocean sediments, lake varves, glacier ice, corals, speleothems, or tree trunks. For purposes of specificity, we develop a time-error model that is relevant to layer counted ice cores, and the Greenland Ice Sheet Project 2 (GISP2) core in particular.

See Huybers and Wunsch [2004] for a development in the context of a marine sediment core.

The Meese/Sowers depth-age scale for GISP2 was derived by counting annual layers with several independent optical, chemical, and electrical techniques (Meese et al. 1994). The GISP2 core is exceptionally well-dated because the high accumulation rate makes discontinuities in stratigraphy relatively unlikely, and the multiparameter continuous count method reduces the probability of missing or over-counting years (Meese et al. 1997). Errors were estimated by intercomparison with volcanic ash layers and independently published age-models (Alley et al. 1993). Estimates place the error in the upper 2500m ($\sim$0—58 ky) at an absolute maximum of 10%, while the errors are in fact believed to be smaller than 2% (Alley et al. 1997). This error increases through 2500—2800 m depth ($\sim$58—110 ky), where discontinuities in the core lead to a layer undercount of up to 20% (Meese et al. 1997). Thus, in order to limit the analysis to perturbations of a well-dated record, we focus our attention to the most recent 50 ky of the core, in which the expected age error is less than 2%. The limiting case of 10% error is also considered, but only as a worst-case scenario.

Annual layers were counted to discern the flow of time with depth in the GISP2 core (Alley et al. 1997). Seasonal alternations in optical properties of ice occur because of changes in the concentration of dust, aerosols, and other impurities over the course of the seasonal cycle as well as changes in bubble density associated with the seasonal cycle in accumulation, temperature, and solar insolation. Lighter bands in Greenland ice tend to be associated with summer hoar complexes, while darker and more transparent layers are associated with uninterrupted winter accumulation (Gow et al. 1997; Alley et al.

299 1997). In some portions of the core, springtime dust layers are also clearly visible. These
300 optical markers, in conjunction with electrical conductivity measurements, permit for a
301 multiparameter layer count. Note, however, that bubbles no longer exist in a gaseous
302 phase at depths greater than 1400 m, instead forming clathrates and eliminating one of
303 the key visual markers. Coupled with dynamic flow thinning, this makes it increasingly
304 difficult to count annual layers in deeper sections of the core.

The errors associated with counting annual layers are cumulative and, therefore, naturally modeled as a random walk. Starting from the top and counting layers downward, counted time accrues at a rate of one layer per year, $t'_{n+1} = t'_n + \tau_n$, where τ_n represents the possibility that the annual band was correctly counted once, $\tau_n = 1$, a layer was missed, $\tau_n = 0$, or that more than one year was counted, $\tau_n = 2, 3, 4....$ Counts are confined to integer numbers, so that the error structure is described by a random walk on a lattice. We define P_1 as the probability of correctly counting a given true annual layer, $\tau_n = 1$, α_u as the probability of not counting it, $\tau_n = 0$, and α_o as the probability of counting an extra layer within the true annual band, $\tau_n = 2$, conditional on one layer already having been counted. Assuming that the conditional probability of counting an additional layer is constant, the probability of counting $m - 1$ extra layers is then $\alpha_o^{m-1} P_1$. For the moment assume that the mean of the distribution is one, so that the number of years missed, on average, balances the number of extra years counted. These assumptions, along with normalization, lead to the coefficient values,

$$\alpha_u = \alpha_o = 1 - \sqrt{P_1},$$

305 and thus to the probability distribution,

$$\Pr(\tau) = \begin{cases} \alpha_u & \text{if } \tau = 0 \\ P_1 \alpha_o^{\tau-1} & \text{if } \tau \geq 1 \\ 0 & \text{if } \tau \leq -1. \end{cases} \tag{A1}$$

Eq. A1 is a mixed distribution that is geometric for $\tau \geq 1$. The variance of the distribution is finite, and the random walk age error which is generated by accumulation of these counting errors, $\epsilon(t_n)$, grows proportionately to $\sqrt{t_n}$ (Fig. 6). Thus, in this symmetric scenario, the expected fractional error between true and estimated time, $\langle t_n - t'_n \rangle / t_n$, will in fact *shrink* as $1/\sqrt{t_n}$. This would imply that the time error grows at a slower-than-linear rate, in contradiction to previously reported error estimates (Alley et al. 1997). In order to obtain errors upwards of 2% at 50 ky, one must set the parameter P_1 to be 0.015, which is a much lower probability of correctly counting a layer than seems plausible (e.g., Gow et al. 1997).

Interestingly, Eq. A1 is consistent with the expected error for atomic clocks, where much of the error arises from biases toward under- or over-counting. Introduction of a bias parameter allows for a more general representation of cumulative timing error and makes it straight-forward to account for the error estimates from the literature. Bias is represented by setting the mean rate of counting to differ from one. This bias, b, can be constant, stationary, or nonstationary, depending on the physical situation. For a long ice core record the bias can be expected to drift with depth as the condition of the ice changes and, importantly, as the Holocene calibration loses accuracy.

Similar to the symmetric case, normalization and the requirement that the expected value of the distribution is equal to $1 + b$ leads the determination of the coefficients, which

now depend on the bias parameter b in addition to P_1,

$$\begin{aligned} \alpha_u &= 1 - \sqrt{P_1(1+b)}, \\ \alpha_o &= 1 - \sqrt{\tfrac{P_1}{1+b}}. \end{aligned}$$

Over many steps, the expected cumulative error $\epsilon(t)$ approaches a normal distribution centered on b, as follows from the central limit theorem. By computing many realizations, the variance of the distribution can then be used to numerically determine P_1 such that the desired 2% expected error of 1 ky is achieved at 50 ky. We model the bias as an autoregressive order one process, with an autoregressive coefficient of 0.999 (corresponding to a decorrelation time of 2 ky) and noise parameter of 7.5×10^{-3}. This produces an error structure close to that described by Alley et al. [1997] when P_1 is set to 0.73, a value which is near the estimated 'worst case' ability to identify annual layers (Rasmussen et al. 2006). The bias parameter is given upper and lower limits, $P_1 - 1 \leq b \leq 1 - P_1$, in order to maintain consistency with the prescription of P_1.

Note that Eq. A1 assumes that the probability of under- or over-counting layers is independent of previous counting errors, which provides for simplicity, but fails to account for the expectation of a relatively constant accumulation rate that tends to curtail the likelihood of long strings of under- or over-counts. The high probabilities of miscounting an individual annual layer and miscounting strings of annual layers may make this error model something of a worst-case scenario, but which would then underscore the finding that power-law estimates are insensitive to timing error.

Acknowledgments. The authors would like to acknowledge helpful comments from Martin Tingley, Carl Wunsch, Jake Gebbie, Nathan Arnold, and Ethan Butler. Support was provided by NSF P2C2 grant (130710).

References

Alley, R., et al., 1993: Abrupt increase in Greenland snow accumulation at the end of the Younger Dryas event. *Nature*, **362**, 527–527.

Alley, R. B., et al., 1997: Visual-stratigraphic dating of the GISP2 ice core: Basis, reproducibility, and application. *J. Geophys. Res.*, **102 (C12)**, 26 367–26 38.

Bartlett, M., 1950: Periodogram analysis and continuous spectra. *Biometrika*, **37 (1)**, 1–16.

Bender, M., T. Sowers, M. Dickson, J. Orchardo, P. Grootes, P. Mayewski, and D. Meese, 1994: Climate correlations between Greenland and Antarctica during the past 100,000 years. *Nature*, **372 (6507)**, 663–666.

Bracewell, R. N., 1986: *The Fourier Transform and Its Applications*. McGraw-Hill.

Bradley, R. S., 1999: *Paleoclimatology*. Academic Press.

Clauset, A., C. Rohilla Shalizi, and M. E. J. Newman, 2007: Power-law distributions in empirical data. *arXiv:0706.1062v2*

Ditlevsen, P., et al., 1996: Contrasting atmospheric and climate dynamics of the last-glacial and Holocene periods. *Nature*, **379**, 810–812.

Gow, A., D. Meese, R. Alley, J. Fitzpatrick, S. Anandakrishnan, G. Woods, and B. Elder, 1997: Physical and structural properties of the Greenland Ice Sheet Project 2 ice core: A review. *J. Geophys. Res.*, **102**, 26.

Harrison, C. G. A., 2002: Power spectrum of sea level change over fifteen decades of frequency. *Geochem., Geophys., Geosyst.*, **3 (8)**, 1–17.

Heneghan, C. and G. McDarby, 2000: Establishing the relation between detrended fluctuation analysis and power spectral density analysis for stochastic processes. *Phys. Rev.*

E, **62 (5)**, 6103–6110.

Herbert, T., 1994: Readings orbital signals distorted by sedimentation: models and examples. *Orbital Forcing and Cyclic Sequences*, P. de Boer and D. Smith, Eds., Blackwell Scientific Publications, 483–507.

Huybers, P. and W. Curry, 2006: Links between annual, milankovitch and continuum temperature variability. *Nature*, **441 (7091)**, 329–332.

Huybers, P. and C. Wunsch, 2004: A depth-derived pleistocene age model: Uncertainty estimates, sedimentation variability, and nonlinear climate change. *Paleoceanography*, **19**.

Meese, D., et al., 1994: The accumulation record from the GISP2 core as an indicator of climate change throughout the Holocene. *Science*, **266 (5191)**, 1680.

Meese, D., et al., 1997: The Greenland Ice Sheet Project 2 depth-age scale: Methods and results. *J. Geophys.Res.*, **102**, 26 411–26 424.

Moore, M. I. and P. J. Thomson, 1991: Impact of jittered sampling on conventional spectral estimates. *J. Geophys. Res.*, **96**, 18 519–18 526.

Mudelsee, M., D. Scholz, R. Röthlisberger, D. Fleitmann, A. Mangini, and E. W. Wolff, 2009: Climate spectrum estimation in the presence of timescale errors. *Nonlin. Proc. Geophys.*, **16 (1)**, 43–56.

Priestley, M., 1994: *Spectral analysis and time series (Probability and mathematical statistics)*. Academic Press Limited, London.

Rasmussen, S., et al., 2006: A new Greenland ice core chronology for the last glacial termination. *J. Geophys. Res*, **111**, D06 102.

Shackleton, N. and J. Imbrie, 1990: The δ 18O spectrum of oceanic deep water over a five-decade band. *Climatic Change*, **16 (2)**, 217–230.

Svensson, A., et al., 2006: The Greenland ice core chronology 2005, 15-42 ka. Part 2: Comparison to other records. *Quaternary Sci. Rev.*, **25**, I23/24, 3258–3267.

Thomson, P. J. and P. M. Robinson, 1996: Estimation of second-order properties from jittered time series. *Ann. Inst. Statist. Math.*, **48 (1)**, 29–48.

Vyushin, D. I. and P. J. Kushner, 2009: Power-law and long-memory characteristics of the atmospheric general circulation. *J. Clim*, **22 (11)**, 2890–2904.

Wilson, P. S., A. C. Tomsett, and R. Toumi, 2003: Long-memory analysis of time series with missing values. *Phys. Rev. E*, **68 (1)**, 017 103.

Wunsch, C., 1972: Bermuda sea level in relation to tides, weather, and baroclinic fluctuations. *Rev. Geophys.*, **10 (1)**, 1—-49.

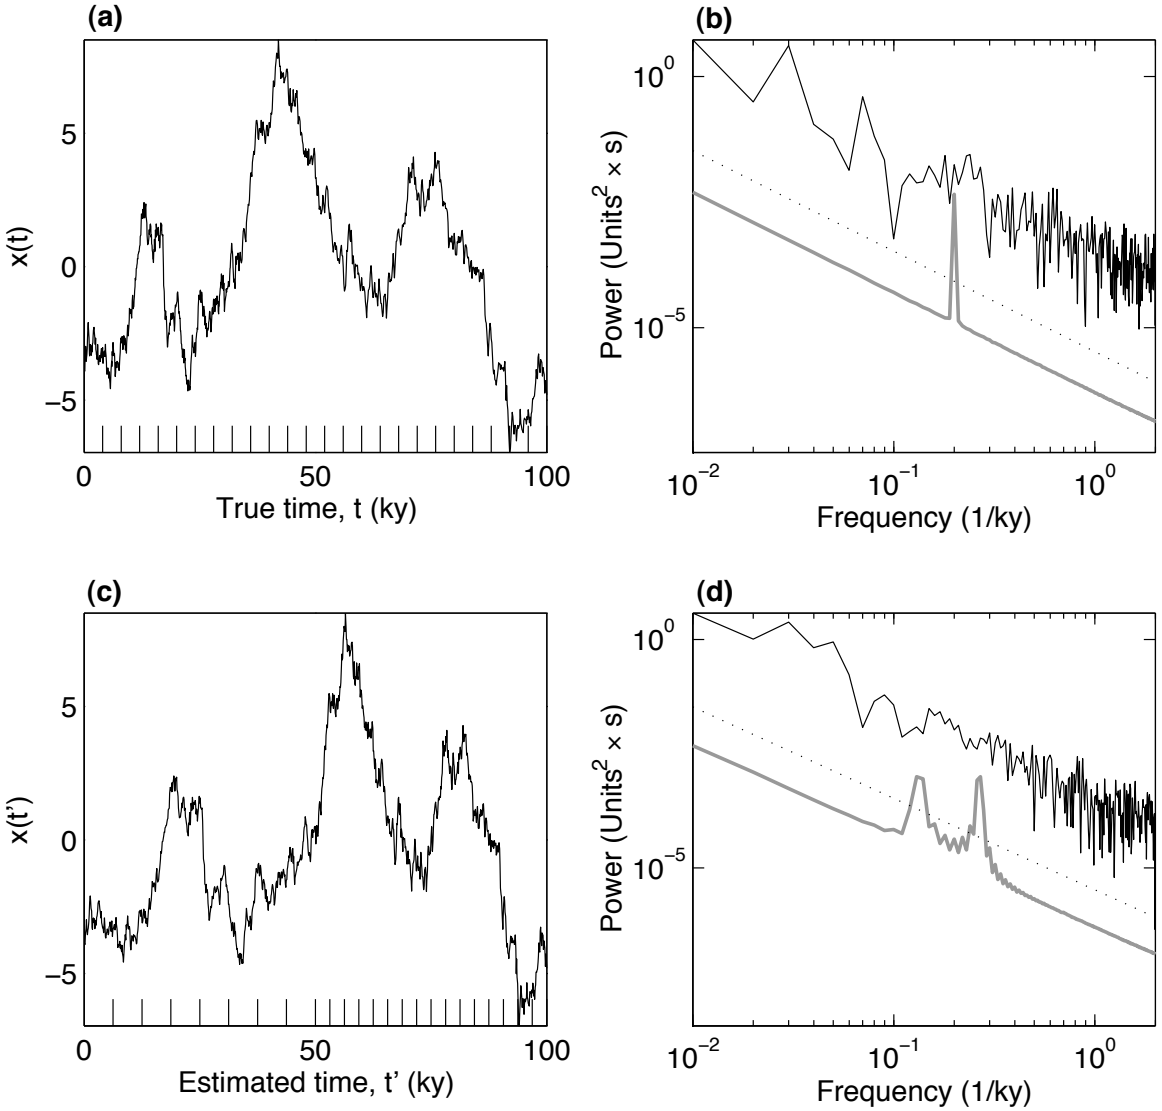

Figure 1. Example of the effect of time-errors on spectral estimates. (a) Measurements from a core section nominally spanning 100 ky and containing a power-law signal with a 0.2 ky^{-1} narrowband component. (b) The power spectral estimate of one realization on the correct timescale (black), with the mean over 1,000 realizations (gray, shifted downward by three decades for visual clarity), and a minus-two power law for reference (dotted). (c) The measurements on an incorrect timescale where time error grows at 1/3 yr/yr between 0 and 50 ky of estimated time and then at $-1/3$ yr/yr between 50 and 100 ky, leading to non-uniform sampling in actual time. Ticks correspond to the same sequence of points in (a). (d) The power spectral estimate of the measurements on the incorrect timescale for one realization (black) and the mean over 1,000 realizations of random signals composed of a power law plus narrowband variability and subject to the same time error (gray), with a minus-two power law for reference (dotted). The narrowband component is split into two broadened peaks, while the power-law background is only affected near frequencies having narrowband energy. The majority of the background remains a -2 power-law in the expectation.

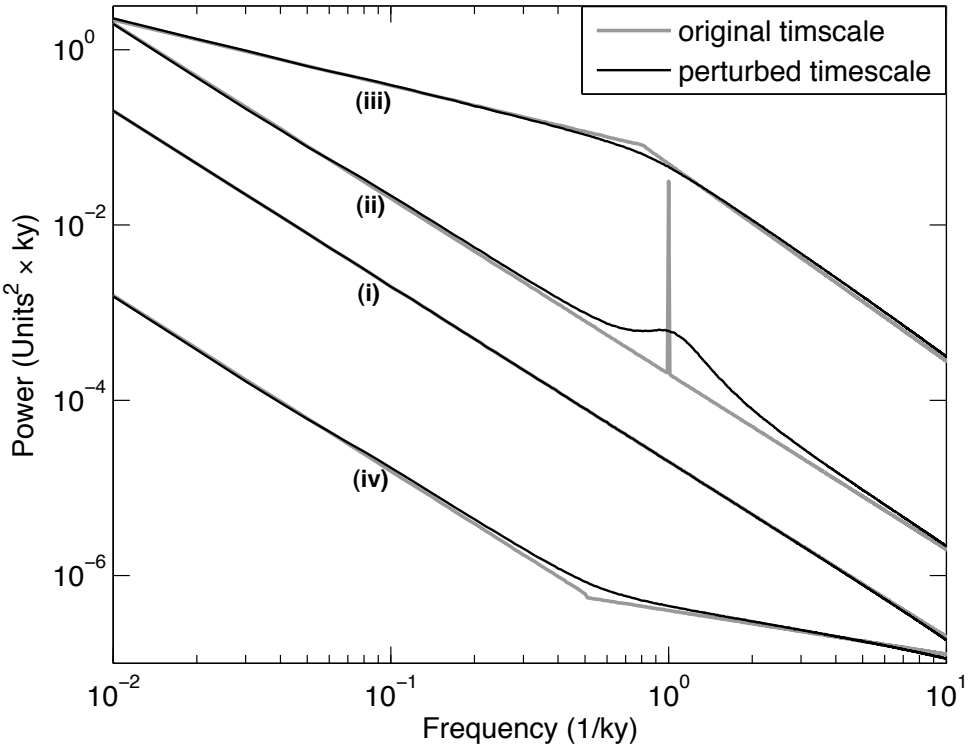

Figure 2. Illustration of the sensitivity of power spectral estimates to time-errors. Signals are nominally 100 ky long, and the average estimate of the power spectrum of each over 1,000 realizations is plotted on the correct (gray) and perturbed (black) timescales. The perturbed timescales have an expected error equal to 5% of the timeseries length (see the Appendix). The timescale error is then detrended so that all spectra can be plotted on a common set of axes (see text). (i) An ensemble of $\beta = -2$ power-law signals are perturbed. The resulting expectation of the spectrum is unchanged. (ii) Narrowband energy of 1/ky is embedded in an ensemble of $\beta = -2$ power-law signals, and the same timescale errors are applied. The spectral estimate in the vicinity of the peak is distorted as the power in the peak is scattered over nearby frequencies. (iii,iv) Similarly, discontinuities in scaling exponents are smoothed by errors in timing.

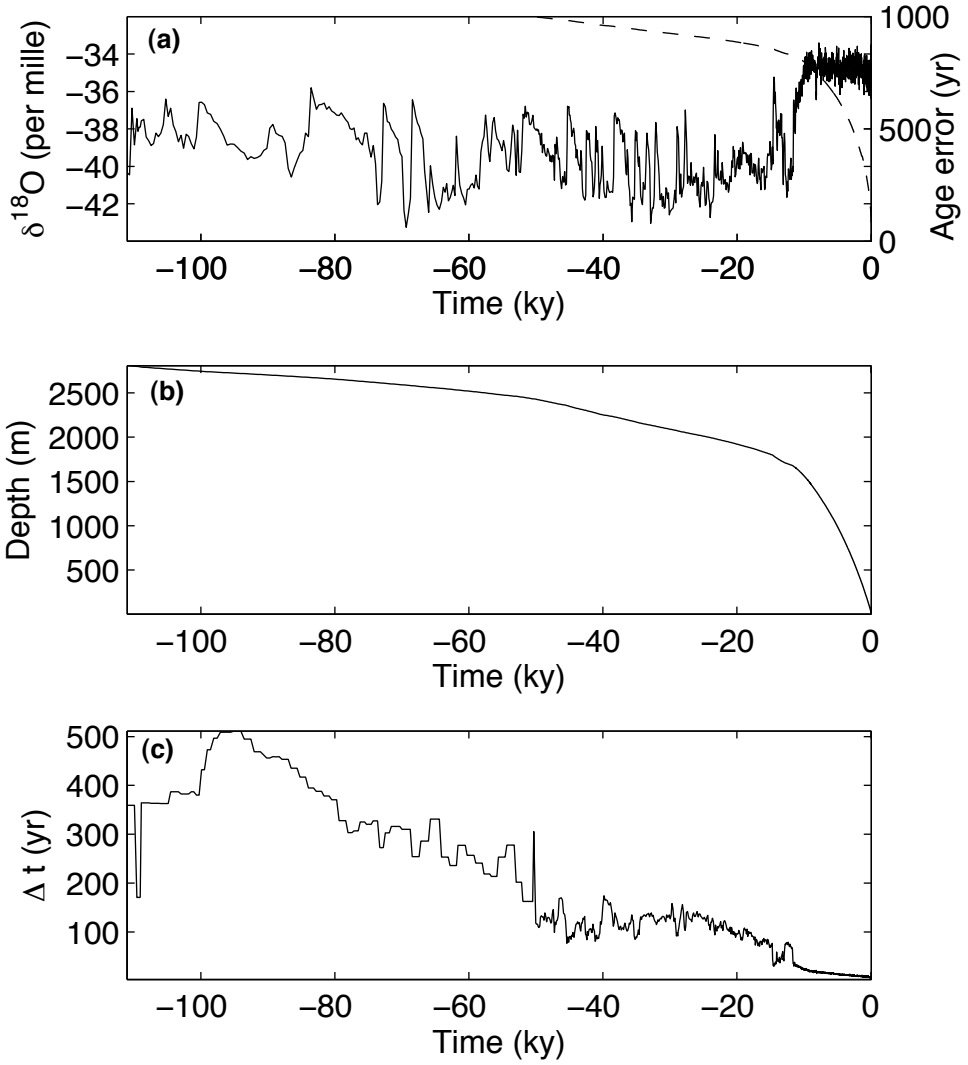

Figure 3. (a) The Greenland Ice Sheet Project 2 δ^{18}O record (solid). Modeled cumulative counting errors in the most recent 50 ky lead to an expected age-error curve which grows with the square root of age (dashed). (b) The Meese/Sowers depth-age scale (Meese et al. 1994). (c) Because of compaction in the core, the sampling interval increases with age, limiting the frequency resolution in older sections of the record.

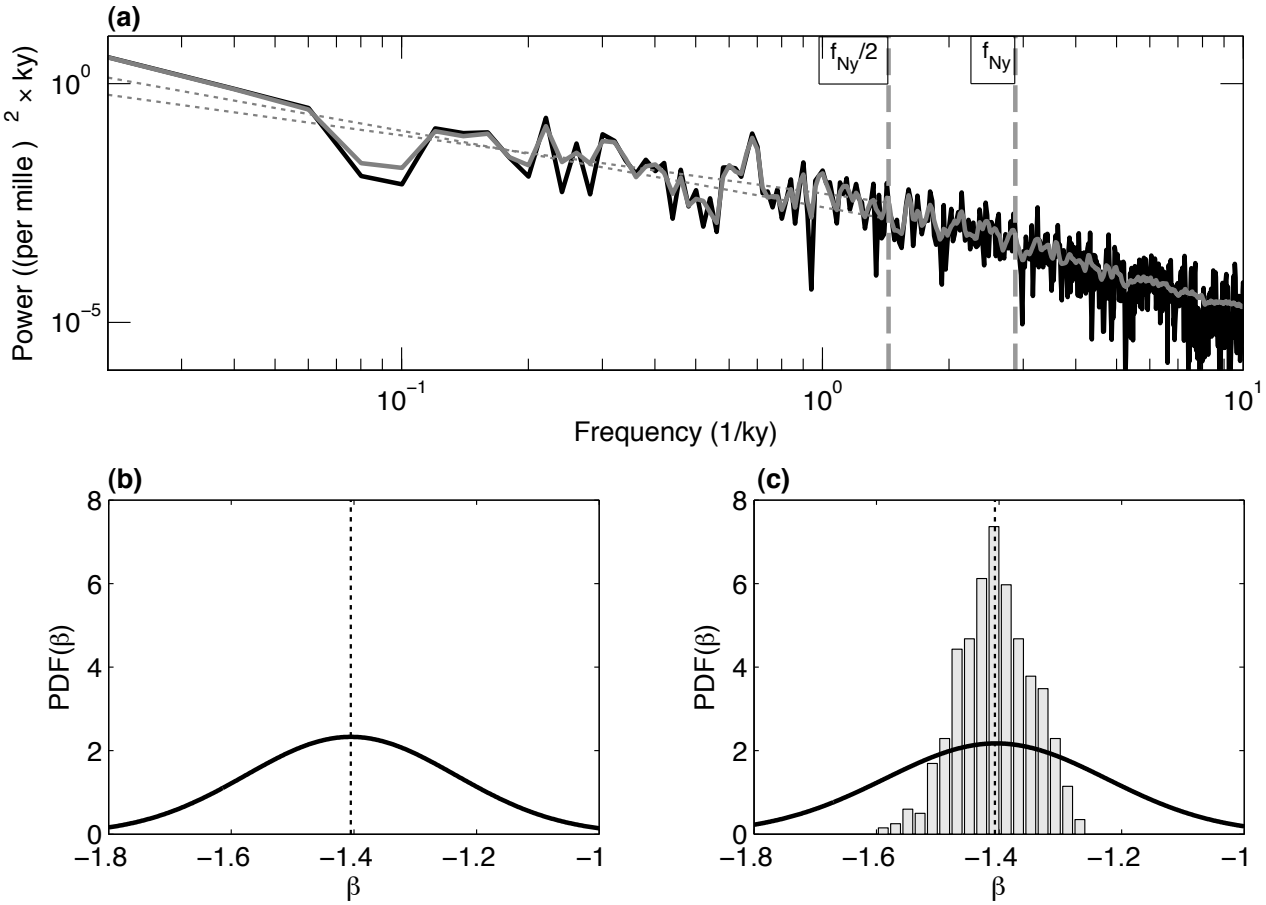

Figure 4. The effect of age-model errors on the power spectral estimate of the last 50 ky of the GISP2 δ^{18}O record. 10,000 age-model realizations are drawn from the cumulative counting error model (as discussed in the Appendix), the error of which grows with age to an expected relative error at the oldest point of 2%. (a) The power spectral estimate of the standard age-model (black), along with the mean power in each frequency band over the different age-model realizations (gray). The 95% confidence intervals of the β estimates are computed at frequencies below $f_{\mathrm{Ny}}/2$ (dotted lines). (b) The least-squares maximum likelihood estimate of β for δ^{18}O using the original age model (dotted line) and its distribution (solid), which is assumed to be normal. (c) Normalized histogram of the age-uncertain maximum likelihood β estimates from each time-error realization (bars), and the combined uncertainty now accounting for the distribution associated with each maximum likelihood estimate (solid). The original timescale gives $\beta_0 = -1.41 \pm 0.17$, whereas the ensemble of perturbed timescales gives $\beta_{\mathrm{est}} = -1.40 \pm 0.19$. The majority of the uncertainty comes from the estimation procedure, not time errors.

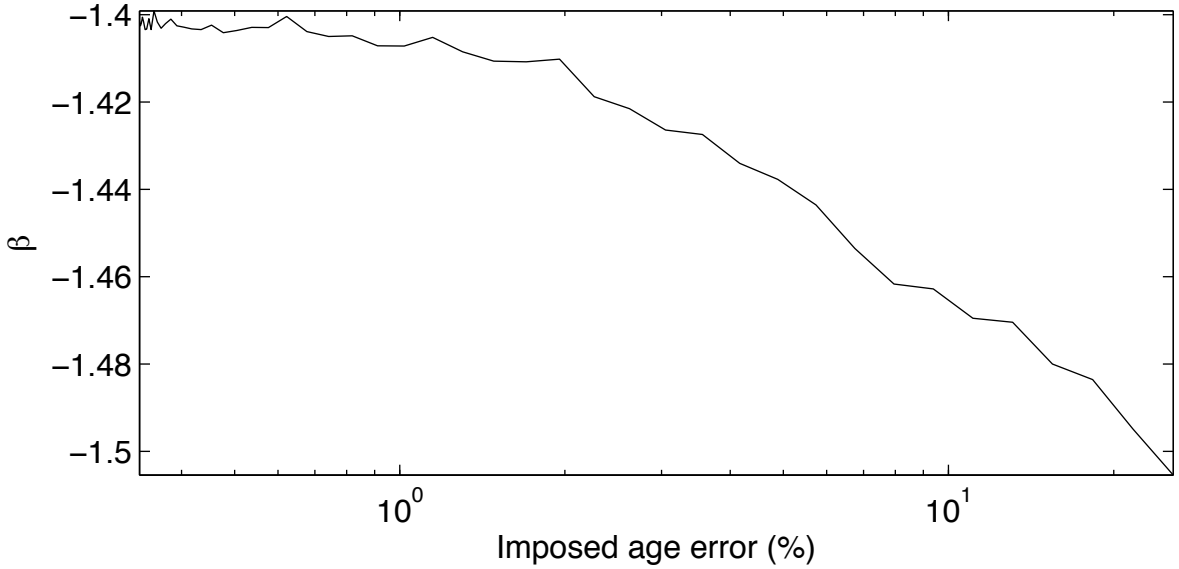

Figure 5. The sensitivity of the estimate of β to different levels of time-uncertainty. The mean estimate of β is plotted as a function of the expected relative age-model error of the oldest data point. The means of the time-uncertain β estimates for $f_{max} = f_{\mathrm{Ny}}/2$ are shown with expected errors reaching extreme levels of 25%. The fit does not deviate significantly from that of the original signal until the expected age-model error is in the vicinity of 10 %.

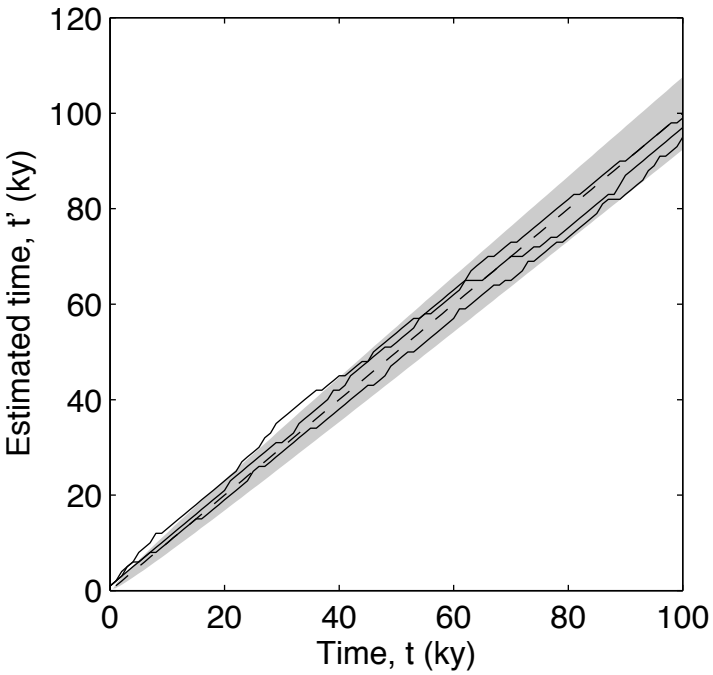

Figure 6. Three realizations of random walk timescales subject to counting errors (solid lines), are compared with the true timescale (dashed line). The shaded region indicates the region within which 95% of age-model points are expected to fall.