Aliasing in $1/f^{\alpha}$ noise spectra: Origins, consequences, and remedies

James W. Kirchner*
Department of Earth and Planetary Science, University of California, Berkeley, California 94720-4767, USA
(Received 24 November 2003; revised manuscript received 22 November 2004; published 14 June 2005; corrected 22 June 2005)

The scaling exponent of a $1/f^{\alpha}$ noise time series is commonly estimated from the power-law slope of its Fourier power spectrum. Here I show that because $1/f^{\alpha}$ noises typically have significant power above the Nyquist frequency, measurements of their power spectra will often be severely distorted by aliasing, not only near the Nyquist frequency, but also far below it. I show that spectral aliasing typically leads to large systematic biases in the scaling exponents, and thus the fractal dimensions, that are estimated from the power-law slopes of $1/f^{\alpha}$ noise spectra. I describe a simple spectral filtering method that corrects the distortions introduced by spectral aliasing, and recovers the broadband spectrum of $1/f^{\alpha}$ noises. Like a Wiener filter, this filtering method does not require that the correct spectrum is known in advance. I illustrate this filtering technique using two environmental noise spectra that are distorted by aliasing.

DOI: 10.1103/PhysRevE.71.066110 PACS number(s): 89.75.Da, 95.75.Pq, 05.10.−a

I. INTRODUCTION

The internal dynamics of many complex natural systems are often investigated through the fluctuation scaling of their time-series behavior. Noisy time series' fluctuation scaling is often characterized using spectral analysis. Of particular interest are noise series whose power spectra are inverse power functions of frequency, the so-called $1/f^{\alpha}$ noises. $1/f^{\alpha}$ noises are found in many natural systems, and the mathematics of fractals provides a natural framework for their interpretation and analysis. Examples are found in natural phenomena as diverse as heartbeats [1], walking gaits [2], animal population dynamics [3], granular creep [4], geomagnetic intensity [5], atmospheric temperature [6], hydrological time series [7], and paleoceanographic isotopic records [8].

Correctly interpreting $1/f^{\alpha}$ noises requires accurately determining their scaling exponents α, which are typically estimated from the log-log slopes of their Fourier power spectra. Here I show that spectral aliasing can significantly distort the apparent scaling exponents of $1/f^{\alpha}$ noises. $1/f^{\alpha}$ noises are particularly vulnerable to aliasing because their spectral power decreases relatively slowly with increasing frequency. Accurately measuring and interpreting $1/f^{\alpha}$ noise spectra requires detecting and correcting the aliasing artifacts embedded in them.

The goals of this paper are (1) to quantify how aliasing distorts measurements of $1/f^{\alpha}$ noise spectra, and (2) to demonstrate a spectral filtering technique for correcting those distortions. I briefly review the mathematics of spectral aliasing in undersampled time series. Theoretical calculations predict, and numerical experiments confirm, that spectral aliasing in $1/f^{\alpha}$ noises can substantially inflate measurements of spectral power, leading to significant underestimation of the power-law slope α, even for modest degrees of undersampling. I describe how these distortions can be corrected by a spectral filtering method analogous to Wiener filtering. This filtering method involves (1) modeling the $1/f^{\alpha}$ noise spectrum, (2) calculating the aliases that would be created by sampling such a $1/f^{\alpha}$ noise process, and (3) multiplying the measured real-world spectrum by a spectral filter, formed by the ratio of the modeled $1/f^{\alpha}$ spectrum with and without aliases. Like a Wiener filter, this spectral filter is effective even if the modeled spectrum is inaccurate; thus it does not require that the correct spectrum is known in advance. I demonstrate this filtering technique by applying it to two environmental noise spectra that are distorted by undersampling.

II. SPECTRAL ALIASING IN UNDERSAMPLED TIME SERIES

Any continuous function of time $x(t)$ has a corresponding Fourier transform $X(f)$, defined by its convolution with sine and cosine waves of frequency f,

$$X(f) \equiv \int_{-\infty}^{\infty} x(t) e^{-i2\pi ft}\, dt = \int_{-\infty}^{\infty} x(t)\cos(2\pi ft)dt + i\int_{-\infty}^{\infty} x(t)\sin(2\pi ft)dt. \quad (1)$$

For many natural phenomena, the phase information contained in the Fourier transform is not important; instead, the feature of interest is the magnitude of fluctuations, as measured by the spectral power $S_X(f)$,

$$S_X(f) \equiv |X(f)|^2 = \left(\int_{-\infty}^{\infty} x(t)\cos(2\pi ft)dt\right)^2 + \left(\int_{-\infty}^{\infty} x(t)\sin(2\pi ft)dt\right)^2. \quad (2)$$

A power spectrum—a plot of $S_X(f)$ versus frequency, usually on log-log axes–provides a graphical summary of the fluctuation scaling of a noise series. The power spectrum of a $1/f^{\alpha}$ noise series [for which $S_X(f) \sim f^{-\alpha}$] will plot as a straight line on log-log axes, and its slope provides a straightforward estimate for the scaling exponent α.

*Electronic address: kirchner@eps.berkeley.edu

The continuous function $x(t)$ is usually not directly observable; instead, the available data consist of measurements of $x(t)$ at an evenly spaced set of sampling times. To emphasize the important distinction between the continuous function $x(t)$ and its discretely sampled counterpart, I denote the latter as $y(t_j)=x(j/f_s)$, where f_s is the sampling frequency and j is an integer-valued index variable. The set of discrete samples $y(t_j)$ necessarily contains less information than the continuous function $x(t)$—infinitely less information, in the general case. An exception to this general rule arises if $x(t)$ has no fluctuations at any frequency above half the sampling frequency [that is, if $S_X(f)=0$ for all $f>0.5f_s$]. In this case the Nyquist sampling theorem demonstrates that the set of discrete samples $y(t_j)$ will contain equivalent information to the continuous function $x(t)$, in the formal sense that both $x(t)$ and its Fourier transform $X(f)$ could be exactly determined from the $y(t_j)$ alone [9,10]. This important special case has given rise to special terminology: the critical frequency $0.5f_s$ is termed the Nyquist frequency f_N, and $x(t)$ is said to be band-limited to frequencies below f_N. If the sampling rate is less than 2 times the highest frequency component in $x(t)$, or equivalently if $x(t)$ is not band-limited to frequencies below f_N, the Nyquist sampling theorem is not satisfied and the time series $y(t_j)$ is said to be *undersampled*.

If the time series $y(t_j)$ is undersampled, the loss of information relative to the continuous function $x(t)$ will be manifested as a critical ambiguity: a sinusoidal wave of frequency f_0 will yield exactly the same set of evenly spaced measurements $y(t_j)$ as a sinusoidal wave of any *other* frequency $kf_s \pm f_0$, for any integer k. Thus the apparent spectral power measured at f_0 will reflect both the true spectral power of the continuous function $x(t)$ at f_0, and the spectral power of all its aliases at frequencies $kf_s \pm f_0$.

Figure 1 illustrates the essence of this problem. The sampled points shown in Fig. 1(a) perfectly describe a sinusoid of frequency $f_0=0.3f_s$. However, the same points also exactly describe sinusoids of frequencies $f_s-f_0=0.7f_s$ and $f_s+f_0=1.3f_s$ [Figs. 1(b) and 1(c), respectively], as well as any other sinusoid of frequency $kf_s \pm f_0$. The discretely sampled waves at all frequencies $kf_s \pm f_0$ [black curves in Figs. 1(b) and 1(c)] will appear as aliases at frequency f_0 [gray curves in Figs. 1(b) and 1(c)]. A complex waveform that combines the three frequency components of Figs. 1(a)–1(c) [black curve, Fig. 1(d)], when discretely sampled, will appear to describe a wave at frequency f_0 [gray curve, Fig. 1(d)], but with greater amplitude than the true signal component at f_0 [dotted curve, Fig. 1(d)]. The measured power at f_0 will reflect both the true signal component at f_0 [Fig. 1(a)] and its aliases from the higher-frequency components f_s-f_0 and f_s+f_0 [Figs. 1(b) and 1(c)].

One can straightforwardly quantify how aliasing affects the power spectrum of a known function $x(t)$ subjected to discrete sampling [11,12]. The Fourier transform of the discretely sampled time series $y(t_j)$ is

$$Y(f) = \int_{-\infty}^{\infty} x(t)\mathrm{III}(t)e^{-i2\pi ft}dt, \tag{3}$$

where $\mathrm{III}(t)$ is the comb distribution [11], composed of an evenly spaced array of delta functions,

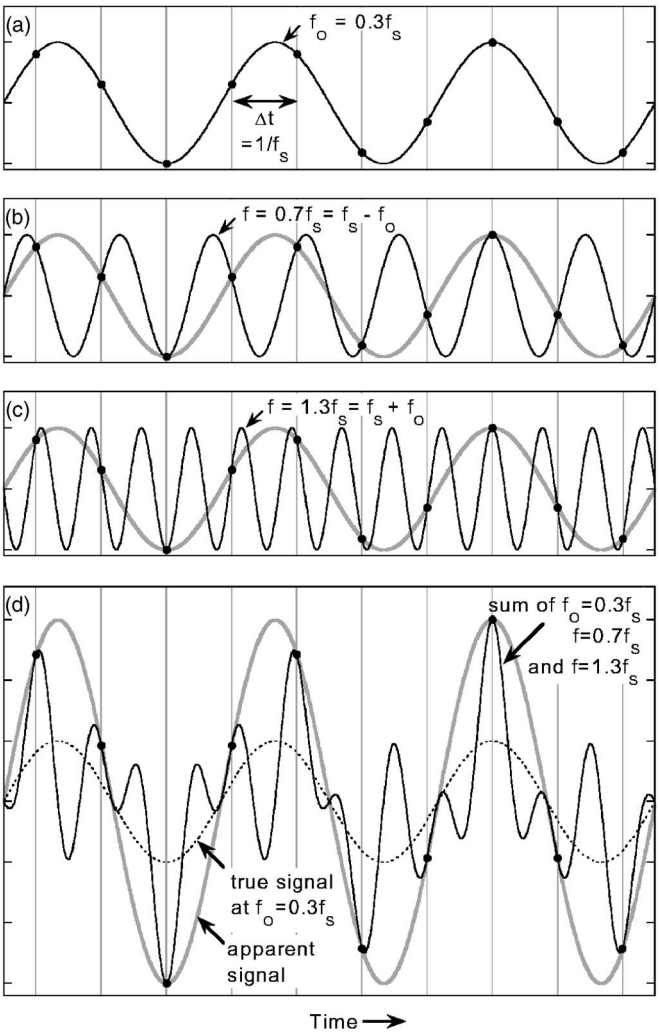

FIG. 1. Aliasing illustrated in the time domain. Vertical bars indicate sampling times; the sampling interval Δt is determined by the sampling frequency f_s. (a) The sampled data points (dots) are consistent with a sine wave of frequency $f_0=0.3f_s$, below the Nyquist frequency $f_N=0.5f_s$. (b) A sine wave of frequency $f=0.7f_s=f_s-f_0$, when sampled at the same times, will yield the same data points as in (a). For any frequency f_0 below the Nyquist frequency, there are an infinite number of frequencies $f=nf_s \pm f_0$ above the Nyquist frequency that yield identical data when sampled at the sampling frequency f_s; panel (c) shows another such alias, at a frequency of $f=1.3f_s=f_s+f_0$. (d) The solid black curve shows a waveform composed of the sine wave from (a), shown as a dotted line, summed with its two aliases from panels (b) and (c). When sampled at the sampling frequency f_s, this waveform yields data points that are consistent with a sine wave of frequency $f_0=0.3f_s$ (gray curve), but at an amplitude much greater than that frequency's actual contribution to the waveform (dotted curve).

$$\mathrm{III}(t) = \sum_{j=-\infty}^{\infty} \delta(f_s t - j), \tag{4}$$

where δ is the Dirac delta function. [For dimensional consistency between $Y(f)$, $y(t_j)$, and $x(t)$, the sampling distribution $\mathrm{III}(t)$ must be dimensionless. Because the delta function's

dimension is the reciprocal of its argument's dimension, the argument of δ must be specified in the nondimensional form used in (4). Superficially similar forms with dimensional arguments, such as $\delta(t-j\Delta t)$, appear widely in the literature but lead to dimensional inconsistencies.] The Fourier transform $Y(f)$ will reflect both the spectral characteristics of $x(t)$ and the distortions introduced by the sampling function $\mathrm{III}(t)$. One can evaluate those distortions using the following line of analysis. Since $\mathrm{III}(t)$ is a periodic function, repeating every $\Delta t=1/f_s$, it can be reexpressed as a Fourier series,

$$\mathrm{III}(t)=\sum_{k=-\infty}^{\infty} c_k e^{i2\pi k f_s t}, \tag{5}$$

where, using the identity $\delta(f_s t)=\delta(t)/f_s$, one can show that the Fourier coefficients c_k are all 1,

$$c_k=\frac{1}{\Delta t}\int_{-\Delta t/2}^{\Delta t/2}\delta(f_s t)e^{-i2\pi k f_s t}dt=\frac{1}{\Delta t}\frac{1}{f_s}\int_{-\Delta t/2}^{\Delta t/2}\delta(t)e^{-i2\pi k f_s t}dt$$
$$=\frac{1}{\Delta t}\frac{1}{f_s}=1. \tag{6}$$

Combining (3)–(6), one can recast the Fourier transform of $y(t_j)$ as

$$Y(f)=\int_{-\infty}^{\infty}\sum_{k=-\infty}^{\infty} e^{i2\pi k f_s t}x(t)e^{-i2\pi f t}dt$$
$$=\int_{-\infty}^{\infty}\sum_{k=-\infty}^{\infty} x(t)e^{-i2\pi(f-kf_s)t}dt. \tag{7}$$

Because the summation in (7) is taken over all k, one can replace the term $-kf_s$ with kf_s in the exponential without loss of generality. Interchanging the order of summation and integration in (7) then yields,

$$Y(f)=\sum_{k=-\infty}^{\infty}\int_{-\infty}^{\infty} x(t)e^{-i2\pi(f+kf_s)t}dt=\sum_{k=-\infty}^{\infty} X(f+kf_s). \tag{8}$$

That is, $Y(f)$, the Fourier transform of the sampled time series $y(t_j)$, equals the Fourier transform $X(f)$ of the continuous function $x(t)$ at the frequency of interest f, summed with the Fourier transforms of all of the aliases $X(f+kf_s)$,

$$\underbrace{Y(f)}_{\substack{\text{sampled}\\ \text{function}}}=\underbrace{X(f)}_{\substack{\text{continuous}\\ \text{function}}}+\underbrace{\sum_{k\neq 0} X(f+kf_s)}_{\text{aliases}} \tag{9}$$

Each of the aliases at $f+kf_s$ may be either in phase or out of phase with the true signal at f. Thus each alias may either increase or decrease the apparent amplitude at any frequency f depending on the phase angle ϕ between the signal and the alias,

$$|X(f)+X(f+kf_s)|^2=|X(f)|^2+|X(f+kf_s)|^2$$
$$+2|X(f)||X(f+kf_s)|\cos\phi. \tag{10}$$

In noise spectra, the phase angle ϕ between any two frequency components will be random; this is the essence of what noise means. As a result, the last term of (10) will randomly take on both positive and negative values for each particular frequency f, but its expected average value will be zero. Thus,

$$E[|X(f)+X(f+kf_s)|^2]=|X(f)|^2+|X(f+kf_s)|^2, \tag{11}$$

from which it directly follows that the expected value of the power spectrum of the sampled signal $y(t_j)$ is the power spectrum of the continuous function $x(t)$ plus the sum of all of its aliases:

$$\underbrace{E(S_Y(f))}_{\substack{\text{sampled}\\ \text{function}}}=\underbrace{S_X(f)}_{\substack{\text{continuous}\\ \text{function}}}+\underbrace{\sum_{k\neq 0} S_X(f+kf_s)}_{\text{aliases}} \tag{12}$$

Because $x(t)$ is a real function, the complex Fourier transform $X(f)$ is Hermitian, so $S_X(f)=S_X(-f)$. As a result (12) can be rewritten for positive frequencies only, in the form

$$\underbrace{E(S_Y(f))}_{\substack{\text{sampled}\\ \text{function}}}=\underbrace{S_X(f)}_{\substack{\text{continuous}\\ \text{function}}}+\underbrace{\sum_{k=1}^{\infty} S_X(kf_s-f)+S_X(kf_s+f)}_{\text{aliases}} \tag{13}$$

Figure 2 illustrates how the distortions introduced into the measured spectrum $S_Y(f)$ by aliasing will depend on the high-frequency tail of the $S_X(f)$. From Fig. 2 and Eq. (13), several key observations can be made. First, if the continuous function $x(t)$ is band-limited to frequencies below the Nyquist frequency $f_N=0.5f_s$, aliasing will not occur; $S_Y(f)$ will exactly equal $S_X(f)$ because there will be no spectral power at any frequencies $kf_s\pm f$ that could be aliased [Fig. 2(a)]. Second, if this criterion is not met [Fig. 2(b)], aliasing will at least double the measured spectral power $S_Y(f_N)$ at the Nyquist frequency because the aliased spectral power from $S_X(f_s-f_N=f_N)$ will equal the true signal power $S_X(f_N)$. Third, the measured spectrum $S_Y(f)$ will be symmetrical around $f=f_N$, which in turn implies that the high-frequency tail of $S_Y(f)$ will resemble white noise as f approaches f_N. Although white-noise tailing is often ascribed to quantization error or measurement noise, it is also an expected by-product of aliasing under quite general conditions, as demonstrated by Fig. 2 and the analysis presented above. Fourth, the effects of aliasing are not necessarily confined to the high-frequency tail of the power spectrum. As Fig. 2(c) shows, if the signal $x(t)$ contains frequency components above the sampling fre-

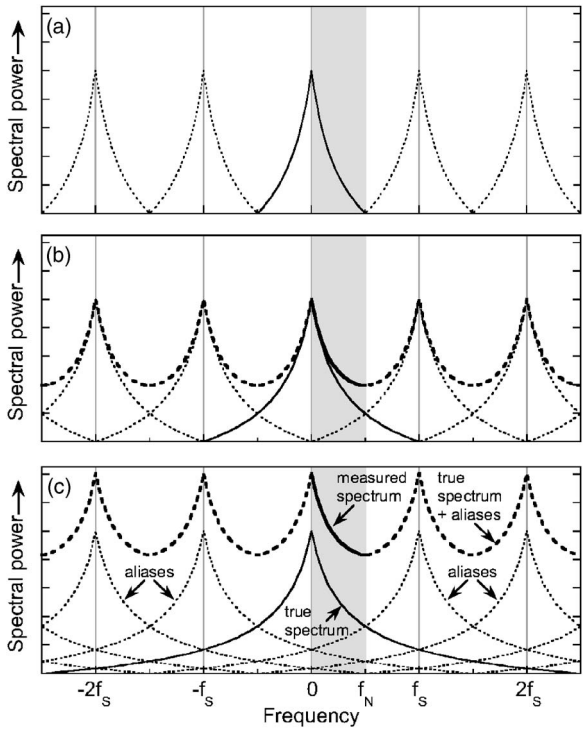

FIG. 2. Aliasing illustrated in the frequency domain. The true power spectrum for a continuous function is shown by the thin black line. Sampling this continuous function at a sampling frequency f_s, creates aliases of the true spectrum, shown by dotted lines, centered at $\pm f_s$, $\pm 2f_s$, $\pm 3f_s$, etc. Spectral power is typically calculated within the frequency range $0<f<f_N$ (where $f_N=0.5f_s$ is the Nyquist frequency), indicated by the broad shaded band. (a) If the true spectrum has no significant power at frequencies above f_N, measures of power at frequencies less than f_N will reflect the true power spectrum. (b) If the true spectrum has significant power at frequencies up to f_s, frequencies in the range $f_N<f<f_s$ will be aliased into the range $0<f<f_N$. For noise spectra (in which the aliases and the true signal at any frequency will be randomly phase-shifted relative to one another), the measured spectral power at any frequency f will be the sum of the true signal and the aliases at that frequency. This sum is shown by the heavy dashed curve and is highlighted by the thick black line over the range that the spectrum would normally be measured. (c) If the true spectrum has significant power at frequencies above f_s, multiple aliases will be superimposed on all frequencies $0<f<f_N$. The measured spectrum will be affected by aliasing over its entire frequency range, and aliased power will exceed the true power at the Nyquist frequency.

quency f_s, the measured spectrum $S_Y(f)$ will be distorted by multiple overlapping aliases across its entire frequency range.

III. SPECTRAL ALIASING IN POWER-LAW NOISES

The effects of spectral aliasing are particularly severe in $1/f^\alpha$ noises, because a power-law spectrum $[S_X(f)=S_0f^{-\alpha}$, where S_0 is an arbitrary constant that sets the scale of the spectral power] will have relatively large power in its high-frequency tail, compared to (for example) an exponential $[S_X(f)=S_0e^{-\alpha f}]$ or Gaussian $[S_X(f)=S_0e^{-\alpha f^2}]$ spectrum. In theory, if Eq. (13) were applied to a power-law noise spectrum,

$$S_Y(f) = S_0f^{-\alpha} + \sum_{k=1}^{\infty} S_0(kf_s+f)^{-\alpha} + \sum_{k=1}^{\infty} S_0(kf_s-f)^{-\alpha}, \tag{14}$$

the infinite summation terms would fail to converge for any $\alpha \leqslant 1$. However, in the real world a $1/f^\alpha$ spectrum with $\alpha \leqslant 1$ cannot extend to infinitely high frequency; its total energy, and thus the variance of its time series, would be infinite. Instead, real-world $1/f^\alpha$ spectra must eventually fall off faster than $1/f^\alpha$ above some threshold frequency; either they roll over to a steeper power-law slope, or they transition to a non-power-law upper tail. Often little can be known about the high-frequency truncation of such spectra, because it often occurs well above the sampling frequency. True $1/f^\alpha$ noises can exist for $\alpha>1$, and for such noises the aliasing terms in Eq. (14) will converge. As Eq. (14) suggests, the severity of aliasing in $1/f^\alpha$ noise spectra will depend on the value of α and the transition frequency (if any) at which the spectrum steepens beyond $\alpha=1$.

To illustrate the effects of aliasing on power-law noise spectra (Fig. 3), I generated artificial $1/f^\alpha$ noises using the straightforward procedure [13],

$$y(t_j) = \sum_{k=1}^{m} (kf_0)^{-\alpha/2} \sin(2\pi k f_0 t_j + \theta_k), \quad t_j = j/f_s, \quad j = 1,\ldots,n, \tag{15}$$

where $\alpha=0.5$ for Fig. 3, f_s is the sampling frequency, f_0 is the fundamental frequency (set equal to the reciprocal of the length of the series to be simulated), and θ_k is a random phase that takes on a different random value $(0 \leqslant \theta_k < 2\pi)$ for each of the m frequency components. Equation (15), like most methods for generating colored noise, effectively truncates the spectrum at the highest frequency component $f_{\max} = mf_0$, as it generates no fluctuations at higher frequencies. Thus the choice of m determines the highest frequency component in the synthetic noise.

If the cutoff frequency $f_{\max}$ is half the sampling frequency f_s or less, the Nyquist sampling theorem is satisfied and spectral analysis yields exactly the $1/f^\alpha$ spectrum used to generate the synthetic noise in the first place. By contrast, if the cutoff frequency equals the sampling frequency, every frequency component $0 \leqslant f \leqslant f_N$, with spectral power $S_0f^{-\alpha}$, will be accompanied by one alias from the counterpart frequency f_s-f, with spectral power $S_0(f_s-f)^{-\alpha}$. The spectral power of the time series at each frequency f will be, from Eq. (10),

$$\begin{aligned} S_Y(f) &= S_X(f) + S_X(f_s-f) + 2\sqrt{S_X(f)S_X(f_s-f)}\cos\phi \\ &= S_0[f^{-\alpha} + (f_s-f)^{-\alpha} + 2f^{-\alpha/2}(f_s-f)^{-\alpha/2}\cos\phi], \end{aligned} \tag{16}$$

where ϕ is the (random) phase angle between each frequency component and its alias. As ϕ varies randomly, the spectral power of the time series fluctuates within the envelope

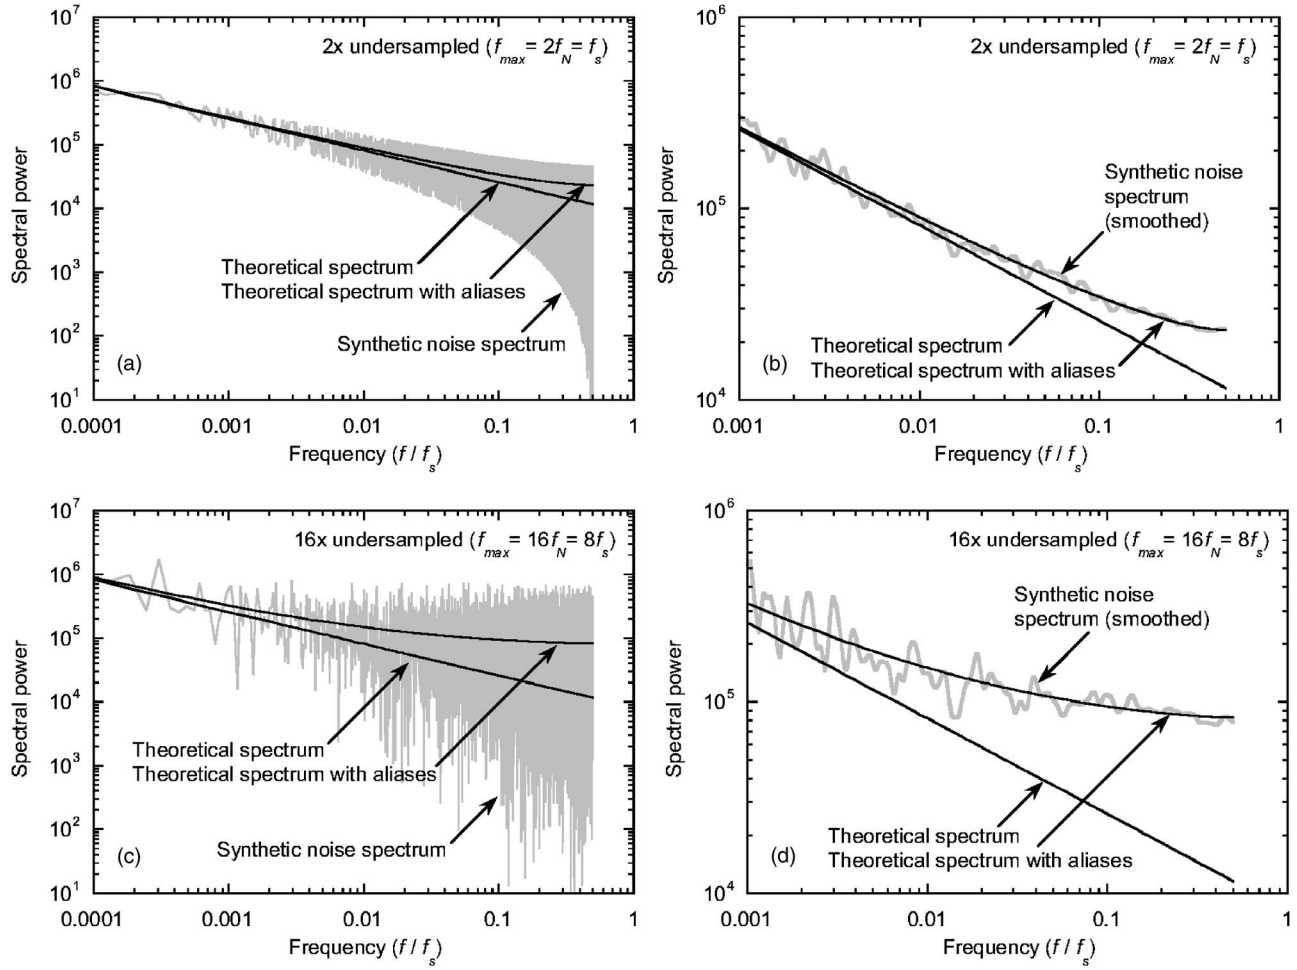

FIG. 3. Effects of undersampling on power spectra of synthetic power-law noises. Spectra of synthetic power-law noises (gray lines) deviate systematically from the theoretical spectra from which they were generated [straight black lines, $S_X(f) \sim 1/f^{0.5}$] when time series are undersampled, i.e., the highest frequency in the sampled process ($f_{\max}$) is greater than the Nyquist frequency of the sampling ($f_N = 0.5 f_s$). Left-hand panels (a) and (c) show unsmoothed spectra of synthetic power-law noises composed of $2^{14} = 16\,384$ points. Right-hand panels (b) and (d) show spectra after smoothing by taking weighted averages of spectral power over a Gaussian smoothing window with a scale factor of 0.05 log units. Scaling the smoothing window logarithmically means that at higher frequencies, more points are averaged together and thus the average spectrum is more precisely defined. If the noise process is undersampled by only twofold (top panels), each frequency f has only one alias, at frequency $f_s - f$. In this case [see panel (a)] the spectral power forms an envelope between an upper limit of $|X(f)|^2 + |X(f_s - f)|^2 + 2|X(f)||X(f_s - f)|$ and a lower limit of $|X(f)|^2 + |X(f_s - f)|^2 - 2|X(f)||X(f_s - f)|$, depending on whether the alias from frequency $f_s - f$ is in phase or out of phase with the true signal at frequency f. Aliasing inflates the spectrum by an average amount $S_X(f_s - f)$ [see panel (b)], in agreement with Eq. (17). Under more severe undersampling (bottom panels), the random variation in the power spectrum is larger and more widespread [panel (c)]. The average spectral power is elevated above the true $1/f^{0.5}$ spectrum across a wide range of frequencies [panel (d)], in agreement with Eq. (13).

shown in Fig. 3(a). The average spectral power, averaged over the fluctuations in ϕ, is described by

$$E(S_Y(f)) = S_X(f) + S_X(f_s - f) = S_0[f^{-\alpha} + (f_s - f)^{-\alpha}]. \tag{17}$$

As Eq. (17) predicts, and Fig. 3(b) demonstrates, the average spectral power deviates systematically from the unaliased spectrum. At the Nyquist frequency $f_N = 0.5 f_s$, aliasing doubles the average spectral power, $E(S_Y(f_N)) = S_X(f_N) + S_X(f_s - f_N) = 2S_X(f_N)$. Aliasing systematically elevates the average spectral power above the unaliased spectrum, and the size of the deviation increases with frequency.

If the cutoff frequency $f_{\max}$ is a large multiple of the sampling frequency, many overlapping aliases are superimposed on the true signal, resulting in large scatter in the power spectrum [Fig. 3(c)]. In such cases, aliasing will distort the spectral power over a wide range of frequencies. In Fig. 3(d), for example, $f_{\max} = 8 f_s$ and aliasing significantly inflates the spectral power at frequencies over two orders of magnitude below the Nyquist frequency. Equation (13) correctly predicts how aliasing distorts the average power spectrum of the discretely sampled synthetic noise series, as Fig. 3(d) illustrates.

As Fig. 3 shows, aliasing inflates the spectral power most at the highest frequencies. As a result, log-log regressions

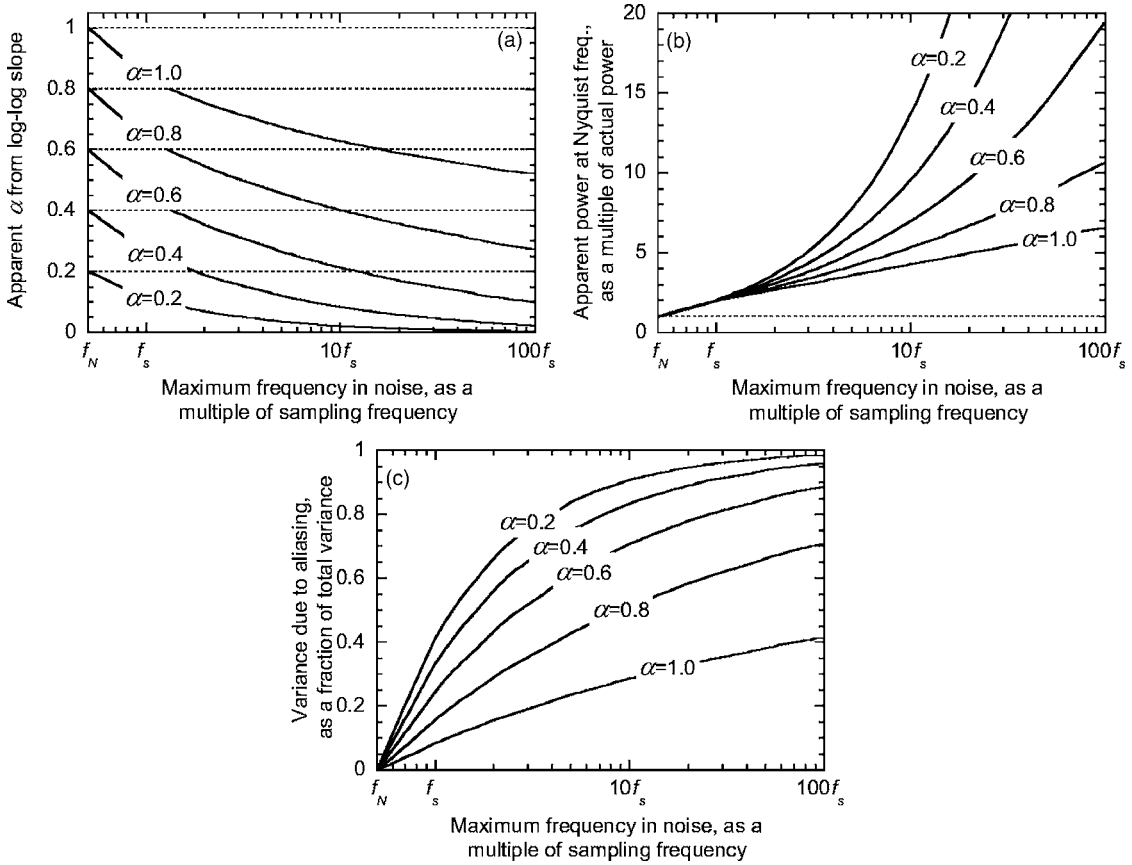

FIG. 4. Effects of aliasing on power spectra of $1/f^{\alpha}$ noises. (a) Apparent α, as indicated by the log-log slopes of power spectra measured over a frequency range of three orders of magnitude ($0.001f_N$ to f_N). (b) Inflation of measured power at the Nyquist frequency f_N. (c) Variance due to aliases, as a fraction of the total variance in the measured time series (estimated as the integral of spectral power from $0.001f_N$ to f_N). Curves are shown for $1/f^{\alpha}$ noises ranging from $\alpha=0.2$ to $\alpha=1.0$. Dotted lines show log-log slopes [panel (a)] and power levels [panel (b)] that would be measured in the absence of aliasing. At left edge of the plots, the maximum frequency in the noise equals the Nyquist frequency and aliasing does not occur. As the maximum frequency in the $1/f^{\alpha}$ noise becomes a large multiple of the sampling frequency, the power levels (a) and log-log slopes (b) of the measured spectra become highly distorted, and aliases comprise a significant fraction of the total variance (c) in the sampled time series.

from spectral plots like Figs. 3(b) and 3(d) will produce artifactually low estimates of the power-law slope α. The effect of aliasing on the regression slope is particularly strong because on log-log axes, most of the data points are clustered at the high-frequency end of the spectrum, where the artificial flattening of the spectrum by aliasing is most severe. As examples of the distortions that aliasing can produce, the log-log regression slopes of the smoothed spectra in Figs. 3(b) and 3(d) are 0.347 ± 0.0016 and 0.166 ± 0.003, respectively. In both cases the regression slopes deviate from the true value of $\alpha=0.5$ by about 100 times their standard errors, owing to the large bias introduced by aliasing.

Using Eq. (13) to estimate the expected aliased power spectrum for a range of scaling exponents α and cutoff frequencies $f_{\max}$, I calculated the apparent scaling exponent that would be obtained from least-squares regression in spectral plots. As Fig. 4(a) shows, aliasing can severely distort estimates of the scaling exponent α, even for time series that are only modestly undersampled (that is, even when the highest frequency in the noise is only a small multiple of the Nyquist frequency). Even modest undersampling can lead to measurements of spectral power at the Nyquist frequency that exceed the true spectral power by large multiples [Fig. 4(b)]. In undersampled power-law noises, a significant fraction of the total variance [Fig. 4(c)]—and thus a significant fraction of the total spectral power and the information content—will consist of aliases rather than real signals.

IV. CORRECTING FOR ALIASING BY FILTERING

Because aliasing introduces severe distortions in the power spectra of power-law noises, it would be desirable to correct for the effects of aliasing and retrieve, as nearly as possible, the underlying unaliased spectrum. In practice, one will only have the sampled values $y(t_j)$ and their power spectrum $S_Y(f)$ to work with; neither the continuous function $x(t)$ nor its power spectrum $S_X(f)$ will be directly observable. Equation (13) above shows that if the true power spectrum $S_X(f)$ were known, the aliased spectrum $S_Y(f)$ could be calculated directly. But $S_X(f)$ is not known, and thus one faces an ill-posed inverse problem: given the measured power spectrum $S_Y(f)$, which includes some degree of aliasing, how can one recover the best possible estimate of $S_X(f)$?

Here I propose a filtering approach to this problem. The approach is conceptually similar to Wiener filtering [14–16], but with several key differences as detailed below. A computer code that implements this filtering procedure is presented in Ref. [17].

The Fourier transform $Y(f)$ of the discrete measurements $y(t_j)$ can be considered as the sum of the signal $X(f)$, which one would like to approximate as closely as possible, and the noise due to aliasing [see Eq. (9)], which one would like to filter out. The task is to find a real-valued frequency-domain filter $\Phi(f)$ which, when multiplied by the alias-corrupted Fourier transform $Y(f)$, yields the best possible estimate $\hat{X}(f)$ for the Fourier transform of the (unknown) continuous function $x(t)$,

$$\hat{X}(f) = \Phi(f) Y(f). \tag{18}$$

Because $\hat{X}(f)$ will primarily be used for power-spectrum scaling analyses, the expected value of its spectral power $S_{\hat{X}}(f)$ should conform as closely as possible to the true (and unknown) $S_X(f)$. Combining Eqs. (18) and (13), one can directly obtain

$$\begin{aligned} E(S_{\hat{X}}(f)) &= (\Phi(f))^2 E(S_Y(f)) \\ &= (\Phi(f))^2 \left(S_X(f) + \sum_{k=1}^{\infty} S_X(kf_s - f) + S_X(kf_s + f) \right) \\ &\approx S_X(f). \end{aligned} \tag{19}$$

Here one sees the first difference with Wiener filtering approaches, for which the objective is to minimize the deviation between $\hat{x}(t)$ and $x(t)$, rather than between $S_{\hat{X}}(f)$ and $S_X(f)$. Equation (19) directly yields an estimate for the filter $\Phi(f)$,

$$\Phi(f) = \sqrt{\frac{S_X(f)}{S_X(f) + \sum_{k=1}^{\infty} S_X(kf_s - f) + S_X(kf_s + f)}}, \tag{20}$$

which can be seen to be the square root of the Wiener filter (the ratio of signal power to signal-plus-noise power).

At first glance, it does not appear that Eq. (20) is particularly helpful; it proposes that in order to obtain an estimate $S_{\hat{X}}(f)$ for the unknown $S_X(f)$, one needs to know $S_X(f)$ itself, including all its aliases. However, one can estimate $\Phi(f)$ in a similar fashion to a Wiener filter, by modeling the signal and its aliases, and estimating $\Phi(f)$ from the ratio of the *modeled* signal power to the *modeled* signal-plus-alias power. Whereas a Wiener filter requires specifying the signal and noise spectra individually, in the present case the signal and noise (i.e., aliases) arise from the same spectrum; thus implementing Eq. (20) requires only a single spectral model for $S_X(f)$. As with a Wiener filter, the filter estimate provided by Eq. (20) can be useful even if the underlying spectral model is inexact. The approach used here is to estimate a spectral model for $S_X(f)$, such that this model plus its aliases,

$$\begin{aligned} S_{Y_{\text{model}}}(f) = S_{X_{\text{model}}}(f) &+ \sum_{k=1}^{\infty} S_{X_{\text{model}}}(kf_s - f) \\ &+ S_{X_{\text{model}}}(kf_s + f), \end{aligned} \tag{21}$$

is a reasonable approximation to the measured spectrum $S_Y(f)$. One then estimates the alias filter $\Phi(f)$ as

$$\Phi(f) = \sqrt{\frac{S_{X_{\text{model}}}(f)}{S_{Y_{\text{model}}}(f)}}, \tag{22}$$

and estimates the alias-filtered spectrum as

$$S_{\hat{X}}(f) = \Phi(f)^2 S_Y(f) = \frac{S_{X_{\text{model}}}(f)}{S_{Y_{\text{model}}}(f)} S_Y(f). \tag{23}$$

Here I illustrate this approach with environmental monitoring data from a long-term watershed study in Wales [18]. This study has monitored many chemical constituents in rainfall and streamflow, including chloride, a natural hydrological tracer. Watersheds have recently been shown to act as fractal filters, converting white-noise rainfall chloride spectra into fractal $1/f^{\alpha}$-noise streamflow chloride spectra [19]. This fractal filtering behavior has important implications for contaminant transport [19,20], and several physical mechanisms have been proposed to explain it [21–24].

Detecting this fractal filtering phenomenon requires accurately measuring the power spectra of tracer concentrations in streamflow, but streamflow chemistry is only measured weekly in typical monitoring programs, raising the possibility of significant aliasing of higher-frequency fluctuations. The Welsh watershed study is unique worldwide, in that it includes 3 years of daily measurements of chloride concentrations, and thus allows a direct assessment of the aliasing that can result from weekly sampling. Figures 5(a) and 5(b) show three-year time series of chloride concentrations in one of the Welsh study streams, sampled at weekly and daily intervals, respectively. The weekly measurements of chloride concentrations appear to exhibit $1/f^{0.5}$ scaling [Fig. 5(c)], whereas daily measurements over the same time span exhibit clear $1/f^{1.0}$ scaling [Fig. 5(d)]. The daily measurements clearly demonstrate that the chloride fluctuations have significant spectral power above the Nyquist frequency for the weekly sampling, implying that the $1/f^{0.5}$ scaling observed in the weekly data may be an artifact of spectral aliasing. Thus the problem at hand is whether, if one only had the weekly data, one could filter out the spectral effects of aliasing and correctly infer the $1/f^{1.0}$ scaling exhibited by this stream.

Estimating the alias filter requires a model for the spectrum, so that the ratio of the signal power and the signal-plus-alias power can be estimated. I used a simple spectral model,

$$S_{X_{\text{model}}}(f) = \frac{S_0 f^{-\alpha}}{1 + (f/f_c)^2} \approx \begin{cases} S_0 f^{-\alpha}, & f \ll f_c, \\ S_0 f_c^2 f^{-(\alpha+2)}, & f \gg f_c, \end{cases} \tag{24}$$

as input to Eqs. (21)–(23). This model spectrum scales as $1/f^{\alpha}$ below some specified corner frequency f_c, then rolls over to a steeper spectral slope of $1/f^{\alpha+2}$ at higher frequen-

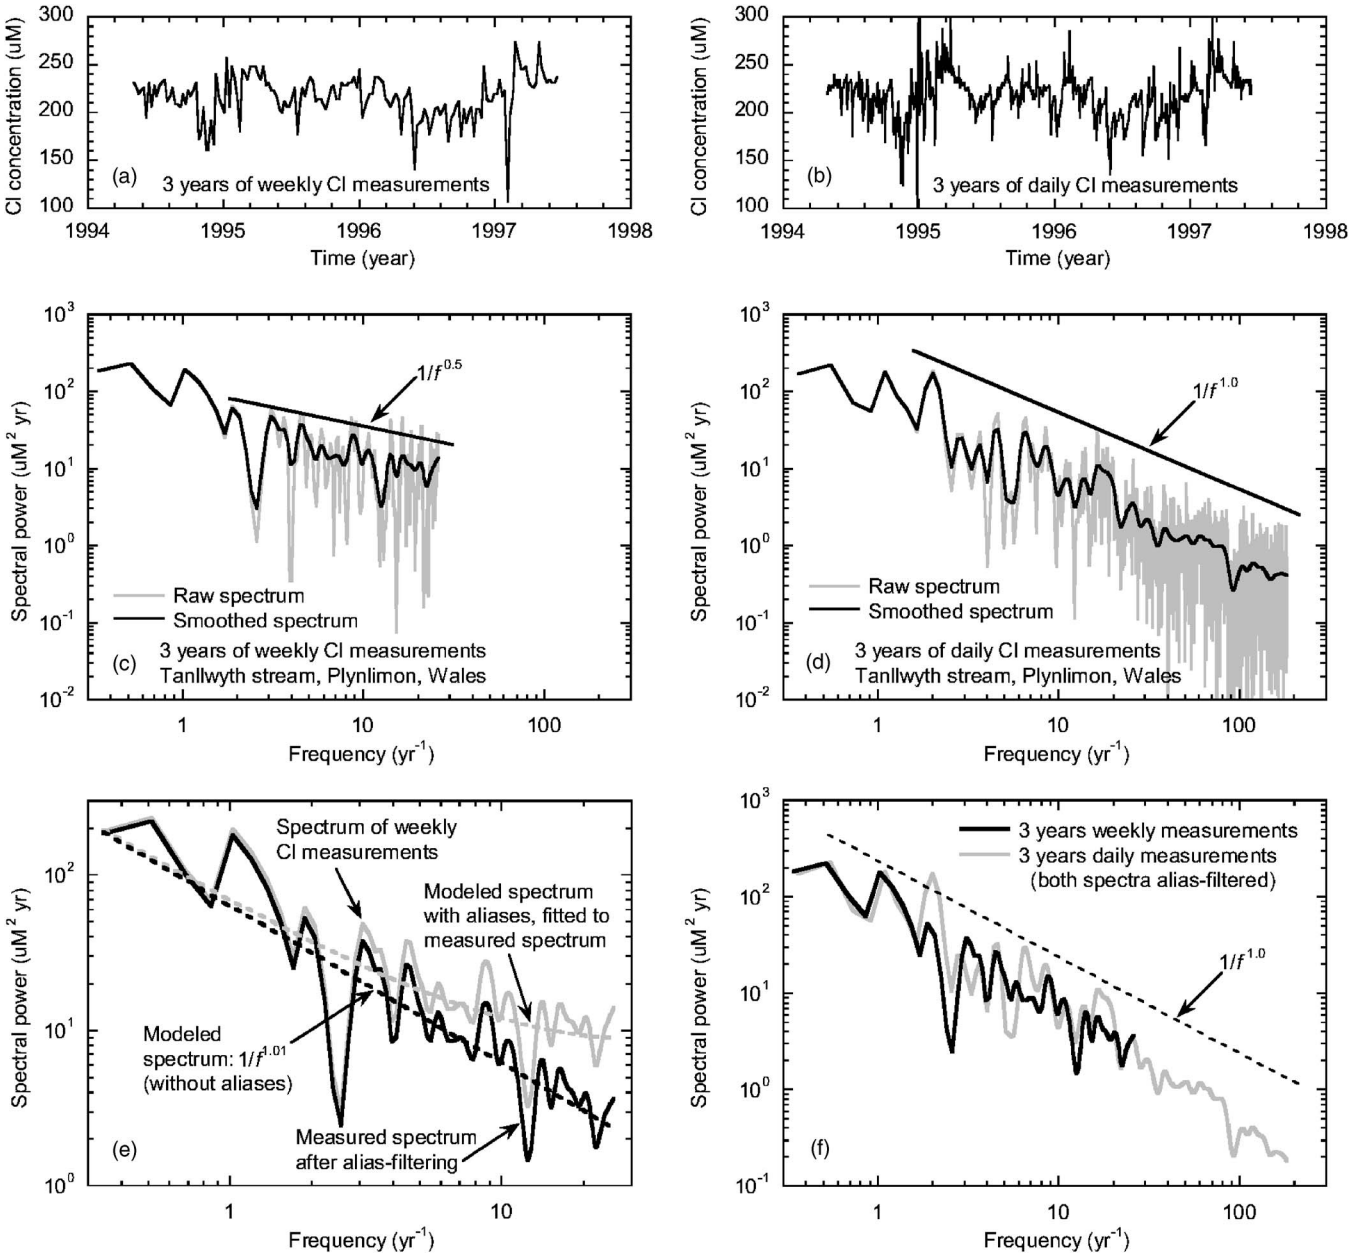

FIG. 5. Alias filtering applied to concentration fluctuations of chloride in streamflow at Plynlimon, Wales. Weekly measurements of chloride concentrations (a) appear to exhibit $1/f^{0.5}$ scaling (c), whereas daily measurements (b) over the same time span exhibit clear $1/f^{1.0}$ scaling (d). (e) Illustration of alias filtering (axes are expanded to show detail). A power-law model [Eq. (21)], including its aliases (dotted gray curve) is fitted to the smoothed spectrum of the weekly measurements (solid gray curve). The alias-filtered spectrum (solid black curve) is calculated from the measured spectrum by multiplying by the ratio of the power-law model with and without aliases (gray and black dashed lines). (f) Alias-filtered spectra for the weekly and daily data show consistent spectral scaling. In comparing panels (a) and (b), note that the data sets do not exactly coincide because the weekly samples are taken at a different time of day than the daily measurements; thus they are not an exact subset of the daily data.

cies. This steeper high-frequency tail roughly approximates the high-frequency damping that must occur as a result of dispersion; mathematically, it guarantees that the summation in Eq. (21) will converge for all $\alpha > -1$, including white noise ($\alpha=0$) and all of the conventional $1/f^{\alpha}$ noises ($\alpha > 0$). In practice, without high-frequency sampling, little can be known about the high-frequency tail of the spectrum. Fortunately, the form of the high-frequency tail has little effect on the alias-filtered spectrum, as long as the transition to that high-frequency tail occurs well above the sampling frequency f_s.

This spectral model has three parameters, α, f_c, and S_0. The scale factor S_0 divides by itself in the alias filter [Eq. (22)], and thus has no effect on the alias-filtered spectrum. The corner frequency f_c must typically be specified *a priori*, since unless f_c is well below the Nyquist frequency, the spectral roll-over in Eq. (21) will not be directly observable. Typically it will be possible to estimate f_c from simple physi-

cal arguments; in this particular case, one can argue that the spectrum should fall off steeply at some time scale ranging from the duration of a typical storm (on the order of 1 day) and the hydrologic response time of the watershed (on the order of several hours); in Fig. 5, f_c is assumed to be 365/year (i.e., 1/day). In any case, as long as the corner frequency f_c is substantially above the sampling frequency f_s, the exact value of f_c will have little effect on the alias-filtered spectrum.

The scaling exponent α is estimated by a nonlinear fitting procedure, which finds the value of α for which the signal-plus-alias spectrum $S_{Y\,\mathrm{model}}(f)$ from Eq. (21) matches the measured spectrum $S_Y(f)$ as closely as possible; that is, it minimizes the mean-square deviation of $\log[S_{Y\,\mathrm{model}}(f)]$ from $\log[S_Y(f)]$. Figure 5(e) demonstrates the fitting procedure applied to the weekly chloride concentrations. In this case, the best-fit value of α is 1.01, for which the modeled signal-plus-alias spectrum (the gray dashed line) matches the measured spectrum (the gray solid line) as closely as possible. One can see from the form of Eq. (23) that on a logarithmic scale, the filtering procedure is equivalent to shifting the measured spectrum (the gray solid line) downward by the difference between the modeled spectrum and the modeled signal-plus-alias spectrum (the gray and black dashed lines). The alias-filtered spectrum (the black solid line) exhibits $1/f^{1.0}$ scaling, consistent with the scaling behavior expected from the daily measurements [Fig. 5(d)]. The spectrum of the daily measurements can also be alias-filtered by the same procedure; the alias-filtered daily and weekly spectra exhibit the same scaling behavior, and coincide almost exactly [Fig. 5(f)].

As with a Wiener filter, the alias filter [Eq. (22)] can perform well even if the spectral model [Eq. (24)] is imprecise. For example, changing the corner frequency f_c from 365/year (i.e., 1/day) to 8760/year (i.e., 1/hour) would change the best-fit value of α in Fig. 5(e) by only 7%, from 1.01 to 1.08, with almost no discernable effect on the alias-filtered spectrum. As an extreme example of mis-specification of the spectral model, consider the black dashed line in Fig. 6(a), which shows the signal-plus-alias spectrum for the spectral model of Eq. (24) with α fixed at $\alpha=0.5$ (rather than the best-fit value of $\alpha=1.01$, which would yield the gray dashed line). The signal-plus-alias spectrum for this mis-specified spectral model does not correspond closely to the measured spectrum (the gray solid line). Nevertheless, the alias-filtered spectrum that is derived from this mis-specified model [black solid line, Fig. 6(b)] exhibits almost exactly the same scaling as the alias-filtered spectrum derived from the best-fit value of α [gray solid line, Fig. 6(b)]. Power-law fits to the two alias-filtered spectra yield almost the same log-log slopes (0.98 versus 1.01). Because the spectral filter $\Phi(f)$ is based on the ratio of the model spectrum to the model signal-plus-alias spectrum, it removes the (proportional) effect of the modeled aliasing, rather than forcing the alias-filtered spectrum to conform to the model spectrum. Thus mis-specification of the model spectrum has only a small effect on the results of the alias filtering procedure.

As a second illustration of this alias filtering procedure, consider the streamflow spectra shown in Fig. 7. Figures 7(a) and 7(b) show excerpts from a 28-year streamflow time series, sampled at daily and hourly intervals, for the same Welsh study stream shown in Fig. 5. The 28-year record of daily instantaneous streamflows for this site appears to exhibit $1/f^{0.40}$ scaling over the frequency range 3/year–182/year [Fig. 7(c)], whereas streamflows sampled hourly over the same period exhibit clear $1/f^{0.67}$ scaling over the same range of frequencies [Fig. 7(d)]. The hourly spectrum also reveals an upper bound to the power-law scaling regime, rolling off at frequencies above roughly 200/year.

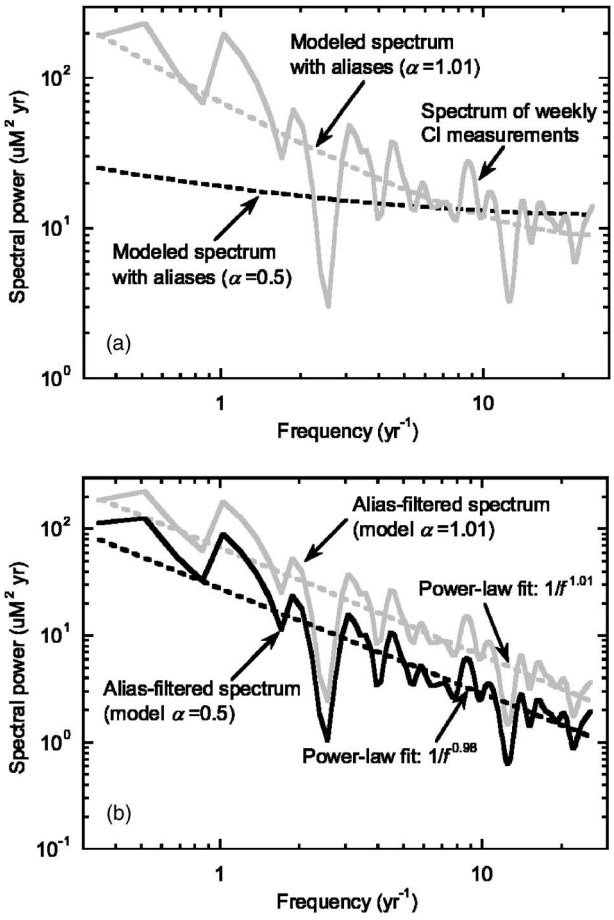

FIG. 6. Alias filtering with a mis-specified model spectrum, illustrated with the weekly chloride spectrum of Fig. 5. (a) Specifying a model spectrum [Eq. (24)] with a scaling exponent of $\alpha=0.5$ rather than the best-fit value of $\alpha=1.01$ yields a modeled signal-plus-alias spectrum (dashed black line) that clearly deviates from the spectrum of the weekly chloride measurements (solid gray line); the best-fit signal-plus-alias model (dashed gray line) is shown for comparison. (b) Nevertheless, the mis-specified model yields an alias-filtered spectrum (solid black line) whose scaling behavior is almost exactly the same as that derived from the best-fit model (solid gray line).

Can the alias-filtering technique outlined above correct for the effects of aliasing in the spectrum of the daily flow data, even though the power-law scaling extends only across a limited range, as one can see from the hourly data? Assuming, as before, a corner frequency f_c of 365/year, and minimizing the mean-square deviation of $\log[S_{Y\,\mathrm{model}}(f)]$ from $\log[S_Y(f)]$ over the frequency range 3/year–182/year, yields

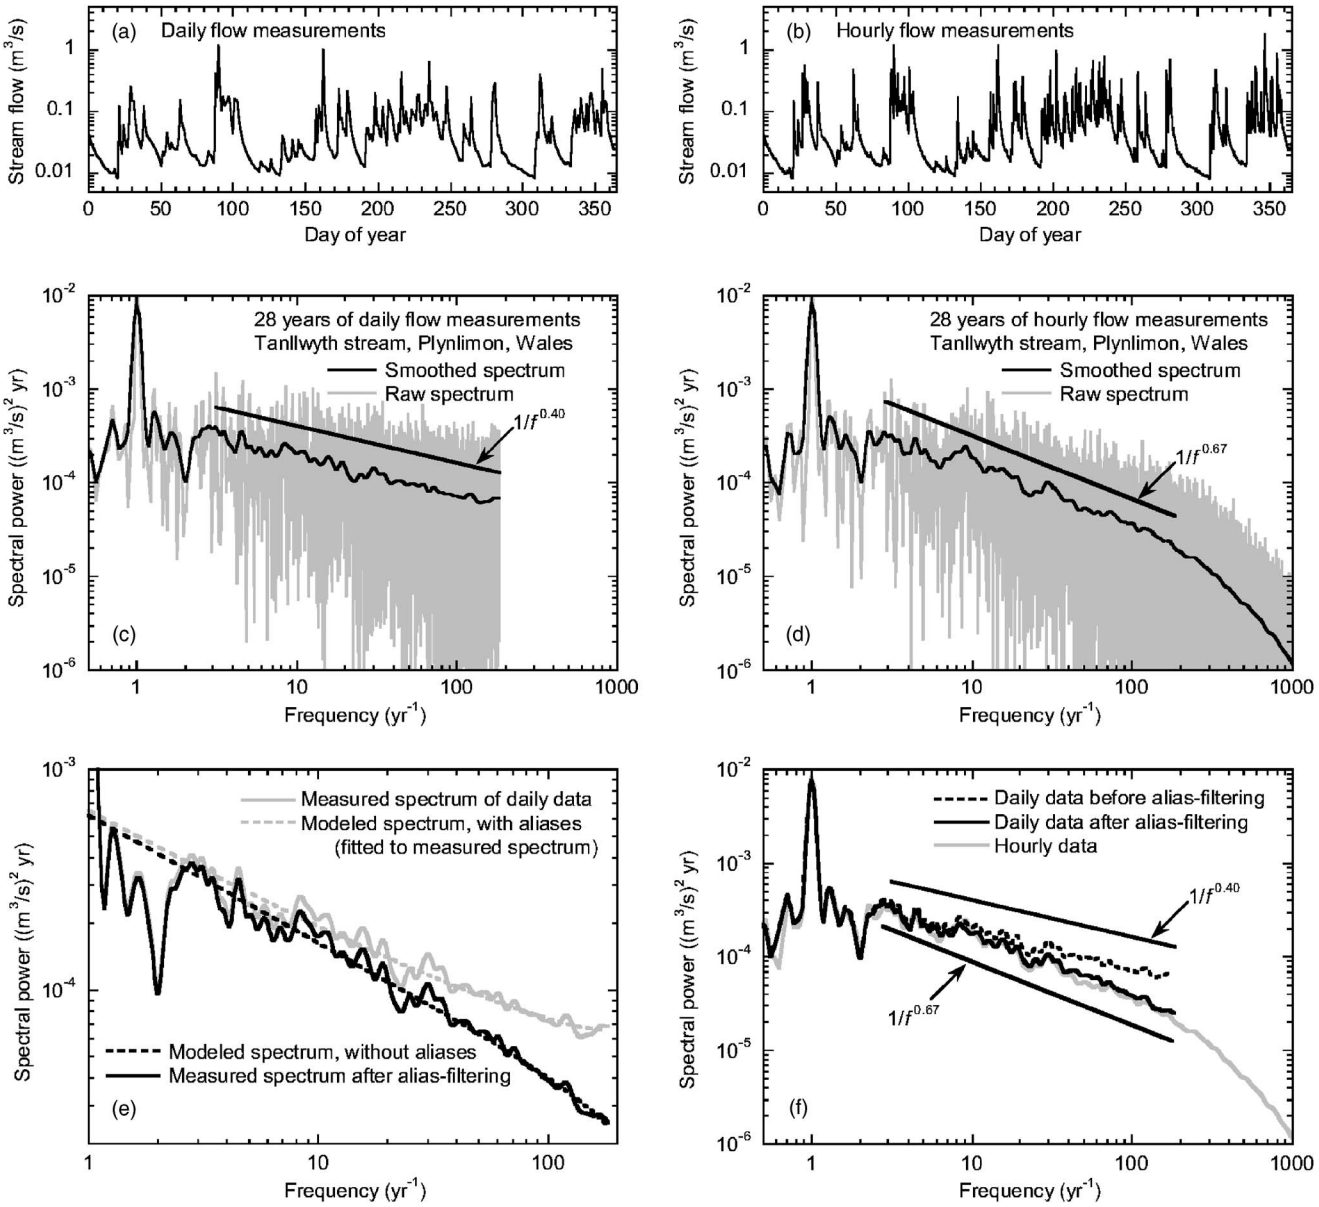

FIG. 7. Alias filtering applied to stream discharge records at Plynlimon, Wales. Daily (a) and hourly (b) streamflow records are similar, with the hourly records showing more fine detail (records for 1985 are shown as an example). Daily instantaneous flows appear to exhibit $1/f^{0.40}$ scaling (c) between frequencies of 3/year and 183/year(=0.5/day), whereas hourly measurements exhibit clear $1/f^{0.67}$ scaling (d) over the same frequency range. Alias-filtering (e) is achieved by fitting a power-law model [Eq. (21)], including its aliases (dashed gray curve), to the smoothed spectrum of weekly measurements (solid gray curve). The alias-filtered spectrum (solid black curve) is derived by multiplying the measured spectral power (solid gray curve) by the alias filter, formed from the ratio of the modeled spectra with and without aliases (dashed gray line and dashed black line). After alias-filtering, the daily and hourly spectra show consistent spectral scaling (f), whereas the unfiltered daily spectrum diverges from the hourly spectrum.

a best-fit value of $\alpha=0.59$ [gray dashed line in Fig. 7(e)]. The alias-filtered spectrum [black solid line in Fig. 7(e)] exhibits $1/f^{0.66}$ scaling, consistent with the $1/f^{0.67}$ scaling observed over this frequency range in the hourly data. The alias-filtered spectrum of the daily data is almost visually indistinguishable from the spectrum of the hourly data [solid black and gray lines in Fig. 7(f)], whereas the unfiltered daily spectrum [dotted line in Fig. 7(f)] deviates significantly from the spectrum of the hourly data.

As in the previous example, the alias filter [Eq. (22)] can perform well even if the spectral model [Eq. (24)] is inaccurately specified. For example, Fig. 8(a) shows two different spectral models for the daily streamflow data. The best-fit value of $\alpha=0.59$ yields a modeled signal-plus-alias spectrum (the dashed gray line) that corresponds closely to the measured spectrum of the daily data over the frequency range 3/year–182/year. Fixing α at $\alpha=1$, however, yields a signal-plus-alias spectrum (the dashed black line) that does not resemble the spectrum of the daily data. Nevertheless, when these two spectral models are used in the alias filter [Eq. (23)], they yield almost identical alias-filtered spectra

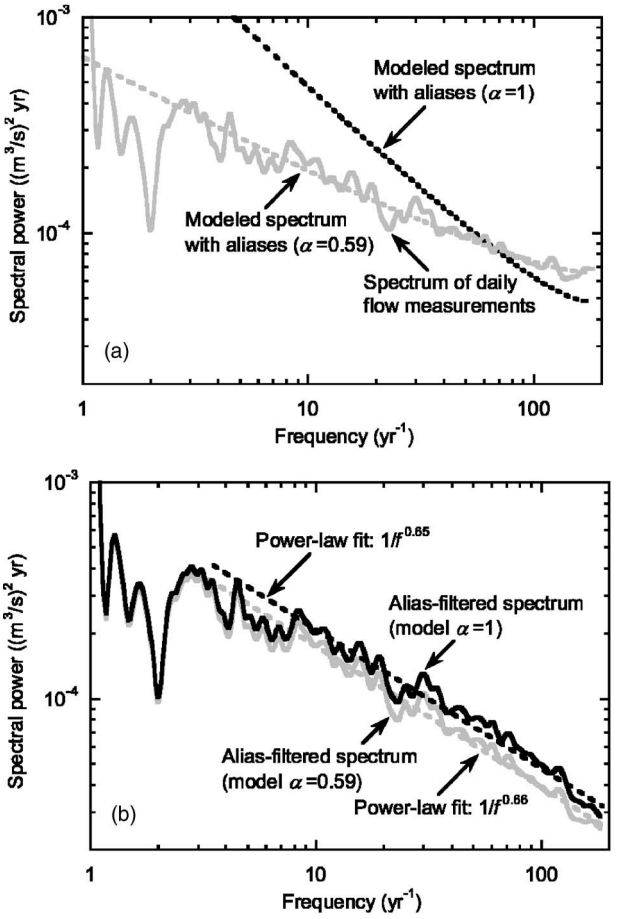

FIG. 8. Alias filtering with a mis-specified model spectrum, illustrated with the daily streamflow spectrum of Fig. 7. (a) Specifying a model spectrum [Eq. (24)] with a scaling exponent of $\alpha=1$ rather than the best-fit value of $\alpha=0.59$ yields a modeled signal-plus-alias spectrum (dashed black line) that clearly deviates from the spectrum of the daily streamflow measurements (solid gray line); the best-fit signal-plus-alias model (dashed gray line) is shown for comparison. (b) Nevertheless, the mis-specified model yields an alias-filtered spectrum (solid black line) whose scaling behavior is almost exactly the same as that derived from the best-fit model (solid gray line).

[Fig. 8(b)] with almost identical log-log slopes of 0.65 and 0.66.

Unlike the previous example, however, in this case the value of the corner frequency f_c has a significant effect on the alias-filtered spectrum. For example, changing the corner frequency f_c from 365/year (i.e., 1/day) to 8760/year (i.e., 1/hour) changes the best-fit value of α in Fig. 7(e) from $\alpha=0.59$ to $\alpha=0.88$, and changes the average log-log slope of the alias-filtered spectrum from 0.66 to 0.89. In this case, the corner frequency f_c affects the alias-filtered spectrum because the hourly streamflow data exhibit a spectral roll-over near the sampling frequency $f_s=365$/year of the daily measurements; this roll-over must be modeled correctly, in order to estimate the right ratio of alias power to signal power near the Nyquist frequency. In the previous example, the spectral roll-over occurred well above the sampling frequency, and thus had little effect on the spectral power that was aliased into the Nyquist interval.

V. EFFECTS OF TIME-AVERAGED SAMPLING ON ALIASING

Another strategy for avoiding spectral aliasing is to analog-filter the signal before sampling, such that the filtered signal has insignificant power above the Nyquist frequency. A form of analog filtering is automatically performed whenever one makes time-averaged measurements. Some measurements of natural systems are intrinsically time averaged. For example, measurements of rainfall rates are often derived from the total volume of rain captured during a finite sampling interval, and thus are inherently time-averaged over that sampling interval. Time-averaged sampling damps fluctuations at frequencies that are near, or above, the sampling frequency (Fig. 9), and thus tends to suppress aliasing of high-frequency components of the signal.

Time-averaged sampling can be viewed as discrete sampling of the continuous running mean of the signal (dotted lines in Fig. 9), where that running mean is averaged over a time interval equal to the reciprocal of the sampling frequency,

$$x_{\mathrm{avg}}(t) = f_s \int_{t-1/f_s}^{t} x(t')dt'. \tag{25}$$

Since the running mean, $x_{\mathrm{avg}}(t)$, is formally the convolution of the continuous signal $x(t)$ with a boxcar function of height f_s and width $1/f_s$, the power spectrum of $x_{\mathrm{avg}}(t)$ can be found straightforwardly by multiplying the power spectra of $x(t)$ and the boxcar function [11],

$$S_{X_{\mathrm{avg}}}(f) = S_X(f)\left(\frac{\sin(\pi f/f_s)}{\pi f/f_s}\right)^2 = S_X(f)\mathrm{sinc}^2(f/f_s). \tag{26}$$

For frequencies $f \ll f_s$, $\mathrm{sinc}^2(f/f_s) \approx 1$, so time averaging has little effect on the power spectrum. At the Nyquist frequency, $\mathrm{sinc}^2(f/f_s)=4/\pi^2$, and time-averaging reduces the spectral power by slightly more than half. At frequencies well above the Nyquist frequency, $\mathrm{sinc}^2(f/f_s)$ oscillates between $f_s^2/(\pi^2 f^2)$ and zero, with an average value of $f_s^2/(2\pi^2 f^2)$. Thus time-averaged sampling strongly depresses the spectral power at frequencies above the Nyquist frequency, and therefore greatly reduces aliasing of undersampled signals.

In practice one will not know the continuous running mean $x_{\mathrm{avg}}(t)$; one will only know its value at discrete, evenly spaced points. As before, the power spectrum of this discretely sampled time series $y_{\mathrm{avg}}(t_j)$ will depend on the convolution of the continuous function $x_{\mathrm{avg}}(t)$ with the comb distribution $\mathrm{III}(t)$,

$$Y_{\mathrm{avg}}(f) = \int_{-\infty}^{\infty} x_{\mathrm{avg}}(t)\mathrm{III}(t)e^{-i2\pi ft}dt. \tag{27}$$

Following the outline of Eqs. (3)–(13), one can show that the expected spectral power of the discrete sample of time-averaged measurements is

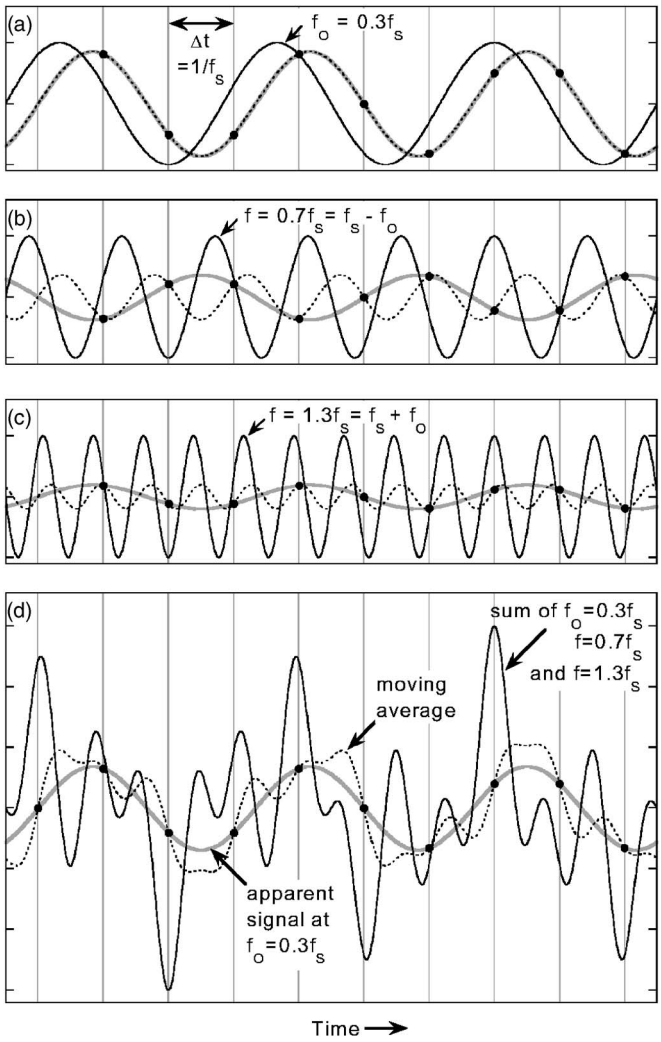

FIG. 9. Reduction of aliasing by time-averaging before sampling. Thin lines indicate unfiltered signals. Dotted lines indicate moving averages over an interval of Δt; these moving averages lag the signal that they are averaging. Dots indicate averages sampled at the sampling interval Δt, and gray lines indicate the signal that would be inferred from the sampled averages. (a) The signal inferred from the sampled averages is slightly damped and phase-lagged relative to the true signal, at frequencies below the Nyquist frequency $f_N=0.5f_s$. (b), (c) Frequency components above the Nyquist frequency are more strongly damped by averaging, and thus their aliases, shown by the gray curves, are much smaller than they would be under point sampling (compare Fig. 1). (d) Thus averages of a complex waveform will primarily reflect the lower-frequency components, and will reduce aliasing of the frequency components that lie above the Nyquist frequency. The averages in panel (d) yield an inferred signal that resembles the true signal at that frequency [panel (a)], whereas point sampling [see Fig. 1(d)] does not suppress the aliases and leads to an inflated estimate of signal strength.

$$\underbrace{E(S_{Y_{\text{avg}}}(f))}_{\substack{\text{sampled}\\ \text{averaged}\\ \text{function}}} = \underbrace{S_{X_{\text{avg}}}(f)}_{\substack{\text{averaged}\\ \text{function}}} + \underbrace{\sum_{k=1}^{\infty} S_{X_{\text{avg}}}(kf_s - f) + S_{X_{\text{avg}}}(kf_s + f)}_{\text{aliases}} = \underbrace{S_X(f)}_{\substack{\text{continuous}\\ \text{function}}} \underbrace{\text{sinc}^2(f/f_s)}_{\substack{\text{filtering by}\\ \text{averaging}}} + \underbrace{\sum_{k=1}^{\infty} S_X(kf_s - f)\text{sinc}^2[(kf_s - f)/f_s]}_{\text{aliases}}$$
$$+ \underbrace{\sum_{k=1}^{\infty} S_X(kf_s + f)\text{sinc}^2[(kf_s + f)/f_s]}_{\text{aliases}}. \tag{28}$$

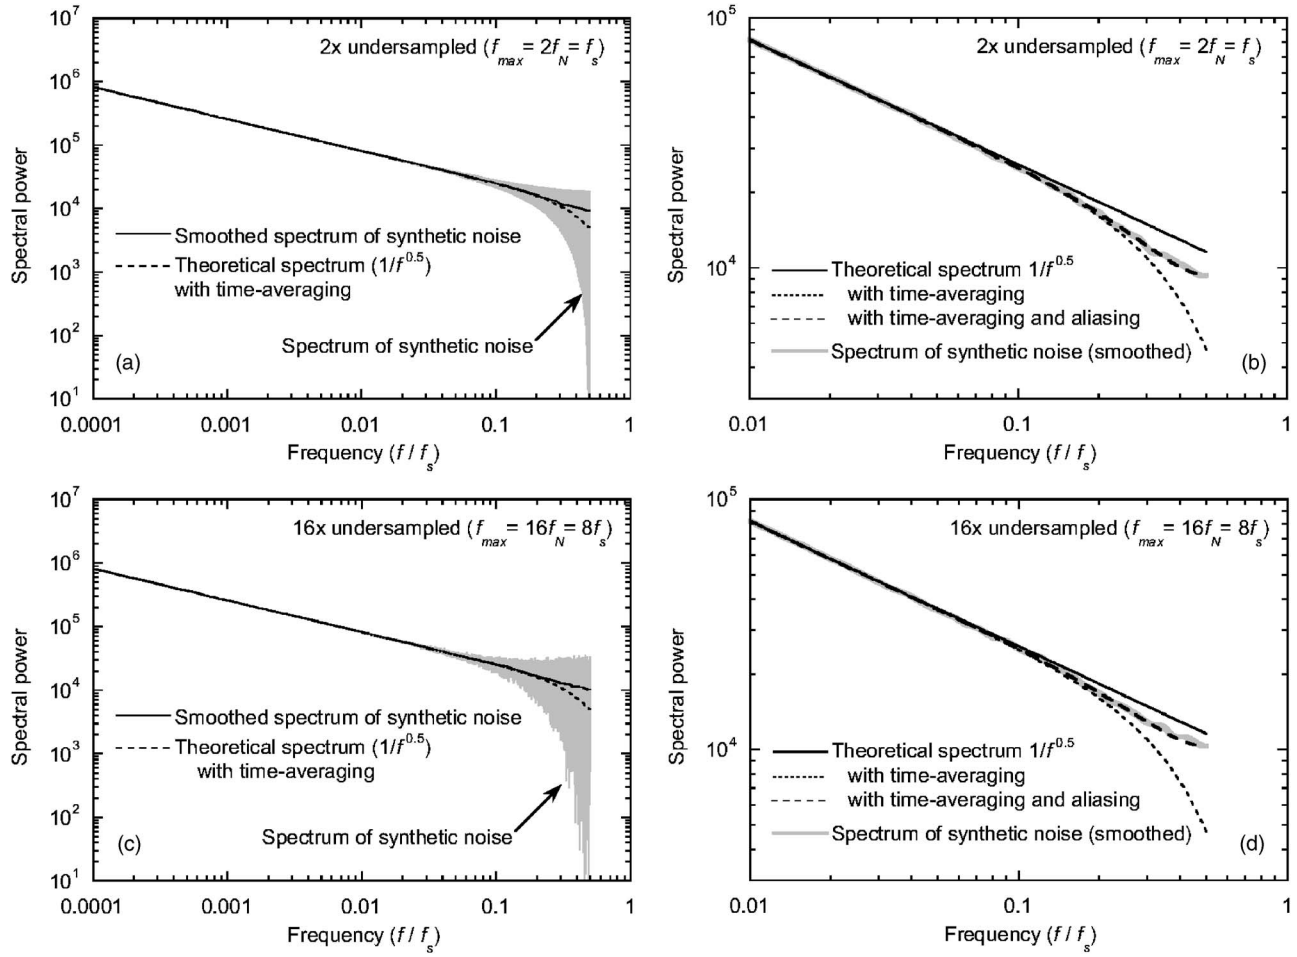

FIG. 10. Effects of time-averaged sampling on power spectra of synthetic power-law noises. The gray line is the spectrum of synthetic noise generated, as in Fig. 3, from an ideal $1/f^{0.5}$ spectrum over a frequency range that extends to the sampling frequency $f_{\max}=f_s$, top panels) or 8 times the sampling frequency ($f_{\max}=8f_s$, bottom panels). The gray lines in the left-hand panels are unsmoothed spectra; those in the right-hand panels (with axes expanded to show detail) are smoothed by averaging over 10 adjacent frequencies. The dotted black lines show the ideal $1/f^{0.5}$ spectrum with the damping of high-frequency fluctuations that would be expected from time-averaged sampling [Eq. (26)], in the absence of aliasing. The dashed black lines in the right-hand panels show the theoretical spectra that would result from the combined effects of time-averaged sampling and aliasing [Eq. (28)]; these correspond closely to the average spectral power of the synthetic noise (gray line). Note that, in comparison to Fig. 3, time-averaged sampling greatly reduces the scatter in the measured power spectrum (left-hand panels) that arises from aliasing of high-frequency noise components. As a result, aliasing has very little effect on the spectrum below roughly $f=0.1f_s$, even when the signal is severely undersampled [panel (c)]. Time averaging leads to damping of the expected spectrum (dashed black line) near the Nyquist frequency $f_N=0.5f_s$; this is partly offset by aliasing of higher-frequency components of the signal. As a result, the average spectral power of the synthetic noise (solid black line in left-hand panels, gray line in right-hand panels) closely approximates the ideal $1/f^{0.5}$ spectrum.

As Figs. 10(b) and 10(d) show, Eq. (28) accurately predicts the average power spectrum of synthetic power-law noise under the combined effects of aliasing and time-averaged sampling. Time-averaged sampling depresses the spectral power near and above the Nyquist frequency. This reduction in spectral power is partly offset by aliasing of the time-averaged signal. The net result is that spectra estimated from time-averaged sampling can closely approximate the true power-law noise spectrum, even if the signal is severely undersampled.

Figure 11 shows that the combined effects of aliasing and time-averaged sampling lead to a small but systematic overestimate of the log-log slopes of measured spectra of$1/f^{\alpha}$ noises. This distortion in the measured spectrum becomes worse in the absence of aliasing because, as described above, time-averaging reduces the measured spectral power near and above the Nyquist frequency, thus steepening the spectral slope; aliasing partially offsets this loss of spectral power. A comparison of Fig. 11 and Fig. 4 shows that time-averaged sampling sharply reduces the spectral artifacts that are generated by aliasing of undersampled signals [note the 100-fold difference in the scale of Fig. 11(b) compared to Fig. 4(b)].

The methods outlined in the preceding section can be used to filter a spectrum that has been distorted by aliasing and time-averaged sampling, thus recovering the best possible estimate of the undistorted spectrum. One simply uses Eq. (28) in place of Eq. (21) to estimate the effects of alias-

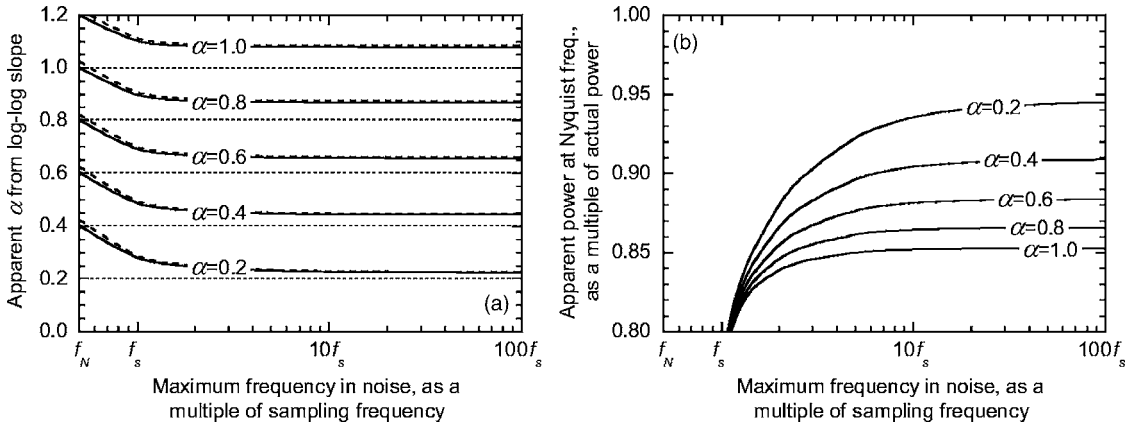

FIG. 11. Effects of aliasing and time-averaging on power spectra of $1/f^{\alpha}$ noises that are time averaged (over the sampling interval) before sampling (compare with Fig. 4). (a) Apparent α, as indicated by the log-log slopes of power spectra measured over a frequency range of three orders of magnitude ($0.001f_N$ to f_N). (b) Reduction in measured power at the Nyquist frequency f_N; note that vertical axis is expanded 100-fold compared to Fig. 4(b). Curves are shown for $1/f^{\alpha}$ noises ranging from $\alpha=0.2$ to $\alpha=1.0$; labels on curves indicate true α values [also shown as dotted lines in panel (a)]. At left edge of the plots, the maximum frequency in the noise equals the Nyquist frequency and aliasing does not occur, but time-averaging reduces the high-frequency spectral power (to 40.5% of true power at f_N), and consequently steepens the measured log-log slope. As the maximum frequency in the $1/f^{\alpha}$ noise becomes a large multiple of the sampling frequency, aliasing inflates the high-frequency power and thus partly offsets the effects of time averaging, and thus the log-log slopes of the measured spectra come closer to the true α values.

ing and time averaging on the model spectrum, then uses Eq. (23) to estimate the required spectral filter. The computer code in Ref. [17] includes an implementation of this procedure.

VI. DISCUSSION

It is important to recognize that the spectral distortions introduced by aliasing are systematic biases rather than random uncertainties. Thus it is no help to simply collect more data by sampling over longer spans of time. Only sampling at higher frequencies—specifically, at frequencies higher than twice the highest frequency component of the noise process—will eliminate spectral aliasing once and for all. Practical guides to time-series analysis (e.g., Ref. [25]) often recommend a "real-world sampling rate" five times higher than the sampling frequency specified by the Nyquist sampling theorem.

That kind of advice is often difficult to implement in practice. In many cases, the cutoff frequency of the noise process of interest is simply unknown, or is higher than any feasible sampling frequency. In other cases, such as long-term environmental time series, one has no choice but to work with whatever data have already been collected, at whatever sampling frequency has been used. These practical constraints imply that in many cases, it will not be possible to eliminate spectral aliasing through high-frequency sampling. In these cases, it will be necessary to filter the distortions introduced by aliasing, in order to recover the underlying noise spectrum as accurately as possible.

The degree to which aliasing can distort power-law noise spectra has remained largely unrecognized. Aliasing has most often been recognized as a problem when spectral peaks above the Nyquist frequency create spurious spectral peaks below the Nyquist frequency (e.g., Ref. [26]). By contrast, the potential for aliasing to distort $1/f^{\alpha}$ noises and other broad-band spectra has mostly gone unnoticed (but for an exception see Ref. [8]), even though many of the $1/f^{\alpha}$ spectra reported in the literature become flatter at their high-frequency ends, just as one would expect from aliasing (see Fig. 3).

It has long been recognized that aliasing can, in principle, be suppressed by sampling the signal at unevenly spaced points in time (e.g, Refs. [27,28]). However, uneven sampling does not eliminate aliased spectral power, but instead simply distributes its effects over a broader range of frequencies than if the signal were evenly sampled [29]. Therefore, while uneven sampling can be effective in suppressing the aliasing of strong spectral peaks from periodic signals, it appears unlikely to be effective in suppressing the aliasing of broad-band noise spectra like those considered here. It is also worth noting that many real-world "unevenly sampled" time series are just evenly sampled time series with gaps where one or more samples have been missed. These gapped time series will be vulnerable to aliasing for the same reasons that evenly sampled time series are.

Nor is aliasing likely to be suppressed by wavelet-based spectral estimation methods. As Fig. 1 illustrates, aliasing arises because of the loss of information in an undersampled time series, relative to the continuous-time behavior that it tries to represent. This loss of information is inherent in the sampled data; it is not an algorithmic problem and does not have an algorithmic solution.

The results reported here show that aliasing can result in large systematic biases in scaling exponents estimated from $1/f^{\alpha}$ spectra (Fig. 4). These results contradict published studies that imply that spectral methods accurately measure the scaling exponents α of synthetic $1/f^{\alpha}$ noises [30,31]. The synthetic $1/f^{\alpha}$ noises used in those studies exhibited $1/f^{\alpha}$ scaling only up to the Nyquist frequency, where their spectral

power dropped discontinuously to zero, as a consequence of the way they were synthesized. Aliasing was thus prevented because the synthetic $1/f^{\alpha}$ noises were band limited to exactly the Nyquist frequency. This lucky coincidence is unlikely to arise in real-world data. The numerical experiments reported here show that synthetic $1/f^{\alpha}$ noises that extend above the Nyquist frequency—as real-world noises often do—exhibit exactly the expected degree of aliasing. In real-world time series, $1/f^{\alpha}$ scaling below the Nyquist frequency must be assumed to imply, unless there is evidence to the contrary, that $1/f^{\alpha}$ scaling continues above the Nyquist frequency as well. If a measured power spectrum does not fall steeply away toward zero as it approaches the Nyquist frequency, this fact should be taken as *prima facie* evidence that aliasing is present. Unless the signal can be analog-filtered in the time domain before sampling, frequency-domain filtering methods like those outlined here provide the only hope for correcting the spectral distortions caused by aliasing in broad-band noise signals.

It bears emphasizing that the alias filtering procedure described here does not force the filtered spectrum to conform to the spectral model $S_{X_{\text{model}}}(f)$ [e.g., Eq. (24)] that is used to construct the filter. From Eq. (23) one can see that on logarithmic axes, the filtered spectrum $S_{\hat{X}}(f)$ will deviate exactly as much from the signal model $S_{X_{\text{model}}}(f)$ as the measured spectrum $S_Y(f)$ does from the signal-plus-aliases model $S_{Y_{\text{model}}}(f)$. The filtered spectrum will conform to the spectral model only if the signal-plus-aliases model conforms closely to the measured spectrum (which inherently contains both the signal and its aliases). Because the spectral filter proposed here does not force the filtered spectrum to conform to the spectral model, mis-specification of the spectral model has relatively little effect on the results, as Figs. 6 and 8 illustrate.

It is important to remember that the strength of aliasing in a broad-band noise spectrum is not an independent parameter that can be tuned; instead, it is inherent in the form of the spectrum and in the sampling frequency. Thus, for example, while one might be tempted to say that the measured spectrum of the weekly chloride data shown in Fig. 5(c) is consistent with a $1/f^{0.5}$ spectrum with no aliasing, this interpretation has no real-world relevance. One never measures the spectrum of the signal alone; instead one always measures the spectrum of the signal plus its aliases. For example, the measured $1/f^{0.5}$ spectrum shown in Fig. 5(c) could only represent a real-world $1/f^{0.5}$ noise process if it were somehow band-limited at exactly the Nyquist frequency. Instead, the real-world system that generated this particular time series exhibits significant fluctuations on time scales down to several hours. All spectra measured on such undersampled time series will be distorted by aliasing. For example, as Figs. 5(e) and 5(f) show, the process behind the weekly chloride data scales roughly as $1/f^{1.0}$ rather than $1/f^{0.5}$. Thus aliasing has reduced the log-log slope of the spectrum in Fig. 5(c) by roughly a factor of 2.

The spectral consequences of aliasing in $1/f^{\alpha}$ noise spectra are similar to those of additive white noise, such as would be expected from quantization error or other random measurement noise. The spectral power of measurement noise in Fig. 5(c) is of order 0.1 μM^2 yr, and thus is insignificant compared to the aliased power, which is of order 10 μM^2 year. Where additive white noise is a more significant component of the measured spectrum, it can be incorporated in the filtering procedure by adding a constant term to the signal-plus-aliases model $S_{Y_{\text{model}}}(f)$.

The severity of aliasing in $1/f^{\alpha}$ noise spectra will vary with the scaling exponent α. The larger the scaling exponent α, the faster the spectral power will fall off with increasing frequency, and the less high-frequency power will be available to be aliased. For example, for moderate degrees of undersampling ($f_{\max} \approx 20 f_N$), aliasing will flatten the log-log slope of a $1/f^{2.0}$ spectrum only about half as much as it will flatten the log-log slope of a $1/f^{1.0}$ spectrum. And of course the $1/f^{2.0}$ spectrum is twice as steep to begin with, so as a percentage of the slope, the aliasing distortion is smaller still, by another factor of 2. However, many $1/f^{\alpha}$ noises of interest in the real world have scaling exponents $\alpha < 2$ and even $\alpha < 1$, for which aliasing is expected to significantly distort the measured spectral slope (Fig. 4). For white or nearly white noises ($\alpha \ll 1$), the measured power below the Nyquist frequency can be utterly dominated by aliased noise.

The clear implication of this analysis is that some of the $1/f^{\alpha}$ noise spectra reported in the literature have been severely distorted by aliasing, and that their apparent scaling exponents α, calculated from their log-log spectral slopes, underestimate their true α values, often by large factors. In such cases, time-averaged sampling (or some other time-domain low-pass analog filtering procedure) provides the best defense against spectral aliasing. Where this is not possible, the frequency-domain alias filter presented here provides a method for correcting the distortions introduced by spectral aliasing, and recovering unbiased estimates of the broad-band spectra of $1/f^{\alpha}$ noises.

ACKNOWLEDGMENTS

I thank David Brillinger for his helpful advice, Xiahong Feng for her encouragement, and Sarah Godsey and Taylor Perron for their comments on the paper. This work was supported by NSF Grant No. EAR-0125550 and by the Miller Institute for Basic Research. I thank the Plynlimon field staff for collecting the data presented here; sample collection and analysis were supported by the Natural Environment Research Council (U.K.), the Environment Agency of England and Wales, and the Forestry Commission (UK).

[1] L. A. N. Amaral *et al.*, Phys. Rev. Lett. **81**, 2388 (1998).
[2] J. M. Hausdorff *et al.*, J. Appl. Physiol. **82**, 262 (1997).
[3] P. Inchausti and J. Halley, Evol. Ecol. Res. **4**, 1033 (2002).
[4] J. J. Roering *et al.*, Geology **29**, 143 (2001).
[5] J. D. Pelletier, Proc. Natl. Acad. Sci. U.S.A. **99**, 2546 (2002).
[6] J. D. Pelletier, Earth Planet. Sci. Lett. **158**, 157 (1998).
[7] J. Pelletier and D. Turcotte, J. Hydrol. **203**, 198 (1997).
[8] C. Wunsch, Clim. Dyn. **20**, 353 (2003).
[9] H. Nyquist, Trans. Am. Inst. Electr. Eng. **47**, 617 (1928).
[10] C. E. Shannon, Proc. IRE **37**, 10 (1949).
[11] R. N. Bracewell, *The Fourier Transform and its Applications*, 3rd ed. (McGraw-Hill, Boston, 2000).
[12] J. T. Broch, *Principles of Experimental Frequency Analysis* (Elsevier Applied Science, London, 1990).
[13] K. Y. R. Billah and M. Shinozuka, Phys. Rev. A **42**, 7492 (1990).
[14] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series, with Engineering Applications* (MIT Press, Cambridge, MA, 1949).
[15] E. A. Robinson and M. T. Silva, *Digital Foundations of Time Series Analysis, Vol. 2: Wave-Equation Space-Time Processing* (Holden-Day, San Francisco, 1981).
[16] W. H. Press *et al.*, *Numerical Recipes: the Art of Scientific Computing* (Cambridge University Press, New York, 1986).
[17] See EPAPS Document No. E-PLEEE8-71-182505 for a MATLAB computer code containing an implementation of the alias-filtering procedure. This document can be reached via a direct link in the online article's HTML reference section or via the EPAPS homepage (http://www.aip.org/pubservs/epaps.html).
[18] C. Neal *et al.*, Hydrology Earth Syst. Sci. **1**, 583 (1997).
[19] J. W. Kirchner, X. Feng, and C. Neal, Nature (London) **403**, 524 (2000).
[20] C. P. Stark and M. Stieglitz, Nature (London) **403**, 493 (2000).
[21] J. W. Kirchner, X. Feng, and C. Neal, J. Hydrol. **254**, 81 (2001).
[22] R. Haggerty, S. M. Wondzell, and M. A. Johnson, Geophys. Res. Lett. **29**, 1640 (2002).
[23] B. Berkowitz *et al.*, Water Resour. Res. **38**, 1191 (2002).
[24] H. Scher *et al.*, Geophys. Res. Lett. **29**, 1061 (2002).
[25] H. J. Weaver, *Applications of Discrete and Continuous Fourier Analysis* (Krieger, Malabar, FL, 1992).
[26] C. Wunsch, Paleoceanography **15**, 417 (2000).
[27] D. R. Brillinger, *Time Series: Data Analysis and Theory* (Holden-Day, San Francisco, 1981).
[28] J. D. Scargle, Astrophys. J. **263**, 835 (1982).
[29] D. L. Wu, P. B. Hays, and W. R. Skinner, J. Atmos. Sci. **52**, 3501 (1995).
[30] B. D. Malamud and D. L. Turcotte, Adv. Geophys. **40**, 1 (1999).
[31] B. D. Malamud and D. L. Turcotte, J. Stat. Plan. Infer. **80**, 173 (1999).

```
function [alpha, filtered_pwr, model_pwr, aliased_pwr]=alias_filter(freq, pwr, fs, fc, f_limit, avgs);
%function alias_filter: Performs alias filtering of power spectrum, using power-law model
%Receives the following as arguments:
%           freq -- vector of frequencies in power spectrum
%           pwr -- vector of spectral power corresponding to frequencies "freq"
%           fs -- sampling frequency
%           fc -- corner frequency for 1/f^2 steepening of power spectrum
%   f_limit -- lower frequency limit for estimating misfit of model-plus-alias spectrum vs. measured power
%           avgs -- flag for whether spectrum is derived from instantaneous point measurements (avgs<>1)
%                                   or from measurements averaged over each sampling interval (avgs==1)
%Returns the following:
%           alpha -- best-fit exponent of power-law model
%           filtered_pwr -- vector of alias-filtered spectral power
%           model_pwr -- vector of modeled spectral power
%           aliased_pwr -- vector of modeled spectral power, plus aliases

log_pwr = log(pwr); %log of power
freq_mask = freq > f_limit; %creates frequency mask vector with 1's if freq is above lower bound for fitting

options = optimset('TolX', 0.0001);
alpha_upper_bound = 5;  %upper bound for alpha (can be changed)
if avgs == 1
   alpha_lower_bound = -2.9;  %if measurements are time-averaged
else
   alpha_lower_bound = -0.9;  %if measurements are point samples
end
alpha = fminbnd(@misfit, alpha_lower_bound, alpha_upper_bound, ...
               options, fs, fc, freq, log_pwr, freq_mask, avgs); %get best-fit alpha

[model_pwr, aliased_pwr, RMSE] = alias(alpha, fs, fc, freq, log_pwr, freq_mask, avgs); %recalculate
spectra for best-fit alpha
filtered_pwr = pwr .* model_pwr ./ aliased_pwr; %perform the filtering here (equation 23)
%end of function alias_filter



function RMSE = misfit(alpha, fs, fc, freq, log_pwr, freq_mask, avgs);
%returns RMSE (Root Mean Square Error) value for misfit of model-plus-alias spectrum for specified alpha
[model_pwr, aliased_pwr, RMSE] = alias(alpha, fs, fc, freq, log_pwr, freq_mask, avgs);
%end of function misfit



function [model_pwr, aliased_pwr, RMSE] = alias(alpha, fs, fc, freq, log_pwr, freq_mask, avgs);
%calculates model spectrum plus aliases (equation 21), and RMSE deviation from real spectrum
model_pwr = model(alpha, fs, fc, freq, avgs); %this is the model spectrum
aliased_pwr = model_pwr; %this is the start of the model-plus-alias spectrum
if avgs == 1    %if sampling is time-averaged, here's where we model its effect
   aliased_pwr = aliased_pwr .* sinc(freq ./ fs) .* sinc(freq ./ fs);  %equation 26
end
for k = 1:1:10 %loop through aliases, building up spectrum of model+aliases: equations 21 and 28
   alias = model(alpha, fs, fc, ((k*fs)-freq), avgs);
           if avgs == 1    %if sampling is time-averaged, here's where we model its effect on aliases
      alias = alias .* sinc(((k*fs)-freq) ./ fs) .* sinc(((k*fs)-freq) ./ fs);  %equation 26
           end
```

```
   aliased_pwr = aliased_pwr + alias; %add alias
   alias = model(alpha, fs, fc, ((k*fs)+freq), avgs);
            if avgs == 1    %if sampling is time-averaged, here's where we model its effect on aliases
      alias = alias .* sinc(((k*fs)+freq) ./ fs) .* sinc(((k*fs)+freq) ./ fs);  %equation 26
            end
   aliased_pwr = aliased_pwr + alias; %add alias
end

%After 10 exact iterations, integrate from k=11 to k=infinity
%alias=integral(S(kfs-f) dk) for k=11 to k=infinity.
%This integration is done by a change of variables, such that z=(k*fs-f)^(-alpha-1)
%and dk=-1/(alpha+1)fs z^(-(alpha+2)/(alpha+1)) dz.
%so the limits of integration are zo=(11*fs-f)^(-alpha-1) and z=0
%and the integrand is almost constant across the entire integral.
%For convenience we define beta=alpha+1.

%For time-averaged sampling, we take the mean of sin^2 over the (rapid) cycles in sinc, so that
%sinc(k*fs-f) is approximated by 1/(2*pi*(k*fs-f)/fs).  Then the entire effect of time-averaged
%sampling is a redefinition of beta and of the leading constant for the integration.
if avgs == 1
   beta = alpha+3;
   const = 1/(2*(pi^2)*beta/fs);
else
   beta = alpha+1;
   const = 1/(beta*fs);
end

zo = (11*fs - freq).^(-beta);
dz = zo ./ 20;
for j=1:1:20
   aliased_pwr = aliased_pwr + const ./ ((j.*dz).^(2/beta) + 1/fc^2).*dz;
end

%now do the same thing over again for the (k*fs+f) aliases -- note the integral is the same, just
%the limits are (slightly) different
zo = (11*fs + freq).^(-beta);
dz = zo ./ 20;
for j=1:1:20
   aliased_pwr = aliased_pwr + const ./ ((j.*dz).^(2/beta) + 1/fc^2).*dz;
end

log_aliased = log(aliased_pwr);

% now calculate prefactor to rescale model+alias spectrum
% prefactor=difference of means in log space; minimizes RMSE deviations from log_pwr
prefactor = sum((log_pwr - log_aliased) .* freq_mask)/sum(freq_mask);

log_aliased = log_aliased + prefactor;
aliased_pwr = aliased_pwr * exp(prefactor);
model_pwr = model_pwr * exp(prefactor);

% now calculate RMSE deviation from measured spectrum, using frequency mask
RMSE = sqrt( sum( (log_aliased - log_pwr) .* (log_aliased - log_pwr) .* freq_mask ) ...
   / sum(freq_mask) );
%end of function alias
```

```
function spectr = model(alpha, fs, fc, freq, avgs);
%calculates model spectrum for a specified set of frequencies
   spectr = (freq .^ (-alpha)) ./ (1 + (freq .* freq) ./ (fc * fc));  %equation 24
%end of function model
```