

THE DISCRETE CORRELATION FUNCTION: A NEW METHOD FOR ANALYZING UNEVENLY SAMPLED VARIABILITY DATA

R. A. EDELSON

Center for Astronomy and Space Astrophysics, University of Colorado

AND

J. H. KROLIK

Department of Physics and Astronomy, Johns Hopkins University

Received 1987 December 29; accepted 1988 April 12

ABSTRACT

A method of measuring correlation functions without interpolating in the temporal domain, "the discrete correlation function," is introduced. It provides an assumption-free representation of the correlation measured in the data and allows meaningful error estimates. This method avoids the problem of spurious correlations at zero lag due to correlated errors. It is shown that physical interpretation of the cross-correlation function of two series believed to be related by a convolution requires knowledge of the input function's fluctuation power spectrum. In the case of AGN line-continuum cross-correlation functions, the interpretation also involves model dependence in the form of symmetry assumptions and must take into account intrinsic scale bias. Application to published data for Akn 120 and NGC 4151 illustrates this method's capabilities. No correlation was found for the optical data for Akn 120, but the ultraviolet NGC 4151 data show a strong correlation, indicating that the broad C IV feature emanates from a region whose size is greater than 1.2 and less than 20 light-days. These bounds on the size of the line-emitting region in NGC 4151 are in good agreement with the predictions of photoionization models.

Subject headings: galaxies: individual (NGC 4151, Akn 120) — galaxies: Seyfert — numerical methods — quasars — radio sources: variable

I. INTRODUCTION

There are a great many astronomical problems in which two signals are observed to vary in time, and the goal is to determine whether they are correlated and, if so, how. Although it is not the most general relationship between two time series, many can be described in terms of an input function which varies in time, driving an output function which therefore also varies in some related way. In the simplest version of such systems, the response is linear and may be described mathematically by the convolution

$$b(t) = \int_{-\infty}^{\infty} d\tau \Psi(\tau) a(t - \tau), \quad (1)$$

where a is the input, b is the output, and Ψ is called the "transfer" or "response" function. Typically, one measures b at certain sample times, and either knows Ψ and wishes to infer a , or else measures a at a set of sample times and wishes to infer Ψ . One of the principal problems in implementing this program in astronomical contexts is that it is very difficult to control the sampling times: most objects are only observable during certain times of the year and certain phases of the Moon, and observing schedules are also perturbed by the depredations of both weather and time allocation committees. In this paper we present a method which takes account of these difficulties. Although it is easiest to interpret the results when the convolution relation (eq. [1]) applies, this statistical method can be used to determine correlations between any two time series.

In order to illustrate this method's capabilities, we will apply it to a problem in the study of active galactic nuclei (AGNs) which has attracted much effort over the past 20 years, namely,

the determination of the physical character and size of the region responsible for producing the broad emission lines (the so-called "BLR"). Arguments based on photoionization modeling (cf., Davidson and Netzer 1979) suggest that the line-emitting region is $\sim 1 \times L_{46}^{1/2}$ pc across (L_{46} is the ionizing luminosity measured in units of 10^{46} ergs s^{-1}). Consequently, it is natural to expect that these lines would vary in flux on observable time scales. Furthermore, since it is generally thought that the line emission is powered by the observed ionizing continuum, fluctuations in the continuum strength should be reflected in fluctuations in the lines. In the language of equation (1), if the lines respond linearly to the continuum, the continuum luminosity is the function a , the line luminosity is b , and the geometrical structure of the region is encoded in Ψ since the delays are due entirely to light-travel time effects (see § IV).

Blandford and McKee (1982) proposed a program by which Ψ for the BLR could be inferred from coordinated observations of continuum and line variations. Although this method is extremely powerful, it has yet to be applied because it requires extremely good temporal sampling and also extremely low noise in the data.

For these reasons, other methods, which are less ambitious, but also less stringent in their prerequisites, have been used to try to measure a characteristic size of the BLR. For instance, the autocorrelation function of a single line contains information on the size of the BLR (Alloin, Boisson, and Pelat 1987). Size information can also be derived by cross-correlating continuum fluxes with line fluxes (Gaskell and Sparke 1986; Gaskell and Peterson 1987). However, these techniques have been used in a way that requires interpolating the data between observed points to form a continuous function. When, as is

usually the case, the fluctuation power spectrum has substantial amplitude at frequencies above the mean sampling rate, interpolation is dangerous. Furthermore, correlated errors in the measurement of continua and lines often cause a spurious peak at zero lag (see § III). A simpler application of the cross-correlation method has also been tried, in which the maximum in the cross-correlation function is determined "by sight" (Clavel *et al.* 1987). This method suffers from the same problems as the interpolation method, as well as the possibility of introducing subjective biases.

In this paper, a new method for determining auto- and cross-correlation functions from unevenly sampled data is introduced. Analogous methods have been applied to other problems, but they have been left in a comparatively undeveloped state (e.g., Mayo, Shay, and Ritter 1974; Hjellming and Narayan 1986). In this method, the correlation function is defined only for those lags for which measured data exist, so that no interpolation, hence no "invention" of data, is required. It also yields reliable error estimates and avoids spurious features caused by correlated errors. The only price paid for these advantages is that we are only able to evaluate the correlation function at a set of discrete sample points. This technique is described in detail in the next section and is tested using simulated data in § III. We discuss the physical interpretation of cross-correlation functions in § IV before applying both our new method and the interpolation method to two sets of published data (for Akn 120 and NGC 4151) in § V. A concluding discussion is presented in § VI.

II. THE DISCRETE CORRELATION FUNCTION

For two continuous, statistically stationary stochastic functions, $a(t)$ and $b(t)$, the classical correlation function is defined as

$$CF(\tau) = \frac{E\{[a(t) - \bar{a}][b(t + \tau) - \bar{b}]\}}{\sigma_a \sigma_b}, \quad (2)$$

where $E\{f\}$ is the expectation value of the function f , \bar{f} is its mean, and σ_f is its standard deviation (Oppenheim and Schaffer 1975). The autocorrelation function is produced when $a(t) = b(t)$, and the cross-correlation is measured when they are different. This function is normalized such that the autocorrelation function at $\tau = 0$ is unity.

In practice, the expectation value in equation (2) is computed by sampling $a(t)$ and $b(t)$ at a discrete set of points $\{t_j; j = 1, \dots, N\}$. If these points were separated by a constant spacing Δt , one could accurately compute $CF(\tau)$ at a large number of lags using only real data, provided these lags were all multiples of Δt . For any $\tau = m \Delta t$, there are $N - m$ points t_k such that $t_k + \tau$ is also in the set of sample points, so the expectation value is well defined so long as $N - m \gg 1$. Unfortunately, astronomical data are only rarely evenly spaced, and therefore in general there are no values of τ such that t_k and $t_k + \tau$ are both in the set $\{t_j\}$ for more than one t_k .

In order to apply the classical technique to astronomical data, one or both of the observed data trains have been interpolated in time. Since this interpolation method weights all temporal points evenly, the result is dominated by interpolated data in at least one of the data trains. Unless there are good reasons for a particular choice of interpolation, the result must be regarded as highly uncertain. Furthermore, the interpolation method gives no indication of the uncertainty in the calculated cross-correlation, while the irregularity in sampling quality over the different possible lags must lead to a consider-

able variation in reliability of the correlation values as a function of lag.

Correlated measurement errors are another problem with cross-correlation functions determined with the interpolation method. Line and continuum fluxes are typically measured from the same spectra, and some errors, such as those in the overall flux level, will affect both continua and lines in the same fashion. It is well known that such correlated errors can lead to spurious contributions to the cross-correlation function at zero lag (Gaskell and Peterson 1987).

To remedy these problems, we suggest the use of what we call the Discrete Correlation Function (DCF). It is defined in the following manner: For two discrete data trains, a_i and b_j , we collect the set of unbinned discrete correlations

$$UDCF_{ij} = \frac{(a_i - \bar{a})(b_j - \bar{b})}{\sqrt{(\sigma_a^2 - e_a^2)(\sigma_b^2 - e_b^2)}}, \quad (3)$$

for all measured pairs (a_i, b_j) . Each of these is associated with the pairwise lag $\Delta t_{ij} = t_j - t_i$. The parameter e_f is the measurement error associated with the data set f . For noisy data, it is necessary to replace the $\sigma_a \sigma_b$ in equation (2) with $[(\sigma_a^2 - e_a^2)(\sigma_b^2 - e_b^2)]^{1/2}$ to preserve the proper normalization (A. Lawrence, private communication). We emphasize that every point represents real information.

Binning this result in time allows the directly useful function $DCF(\tau)$ to be measured. Averaging over the M pairs for which $\tau - \Delta\tau/2 \leq \Delta t_{ij} < \tau + \Delta\tau/2$,

$$DCF(\tau) = \frac{1}{M} \sum UDCF_{ij}. \quad (4)$$

[The $DCF(\tau)$ is not defined for a bin with no points.] Choosing the right bin size is governed by a tradeoff between the desire for high accuracy in the mean defined by equation (4) and a countervailing desire for resolution in the description of the cross-correlation curve. The former consideration argues for large bins, the latter for small. Our simulations (see § III) show that the results depend only weakly on the specific bin size chosen.

It is true that this binning implicitly means that one must accept a modest amount of interpolation in the correlation function (although not in the time series). This procedure is justified in the frequently encountered circumstance in which the correlation is smoother than at least one of the constituent time series. A smoothly varying phase for the product of the Fourier transforms of $a(t)$ and $b(t)$, i.e., $\hat{a}(f)\hat{b}(f)$ is all that is required. This is usually the case with astronomical data.

In order to eliminate correlated errors, lags $t_j - t_i$ for which $i = j$ can be excluded. This procedure automatically removes zero-lag correlated error (the most common kind) and does so at the very limited cost of reducing the number of points contributing to the zero-lag bin. All other bins are automatically free of correlated error even without this cleansing of the zero-lag bin. Note that, when one subtracts the variance due to measurement error from the total variance, one *must* eliminate the zero-lag pairs if their errors are correlated in order to keep the normalization correct. Indeed, requiring the autocorrelation at zero lag to be within 1σ of unity amounts to a crude check of the error estimate.

Unlike the interpolation correlation function, for which the errors are difficult and probably impossible to define, it is straightforward to define a standard error for the DCF. If the individual $UDCF_{ij}$ within a single bin were totally uncor-

related, then the standard error in the determination of their mean would be simply

$$\sigma_{\text{DCF}}(\tau) = \frac{1}{M-1} \left\{ \sum [\text{UDCF}_{ij} - \text{DCF}(\tau)]^2 \right\}^{1/2}. \quad (5)$$

However, many of the UDCF_{ij} within a single bin may have mutual correlations if one of the series is strongly autocorrelated over widths greater than the bin size. This happens, for example, if the relation between the two time series is through convolution, as in equation (1), and the transfer function is nonzero over a range of lags. It is then necessary to replace $M-1$ in the denominator of equation (5) by $[(M-1)(M'-1)]^{1/2}$, where M' is the number of uncorrelated UDCF values within the bin, that is, the number of different measurement times t_i for the series a_i . This definition of the standard error is quite reasonable: where the point-by-point UDCF has a large scatter, or where there are few independent points contributing to a bin, σ_{DCF} is large; where there is little scatter and many points, σ_{DCF} is small.

With a quantitative error estimate in hand, it makes sense to compare the magnitude of the correlation found with that which might arise from random fluctuations in the correlation of two causally unrelated time series. If $a(t)$ and $b(t)$ are independent Gaussian processes with zero mean and unit standard deviation, the probability distribution for their product γ is

$$P(\gamma)d\gamma = \frac{d\gamma}{2\pi} \int da \int db e^{-(a^2+b^2)/2} \delta(\gamma - ab). \quad (6)$$

The integrals may be easily evaluated using the δ -function and the method of steepest descents:

$$P(\gamma)d\gamma \approx 0.370 d\gamma \begin{cases} 1, & |\gamma| \ll \frac{1}{2}, \\ \frac{e^{-|\gamma|}}{|\gamma|^{1/2}}, & |\gamma| \gg \frac{1}{2}. \end{cases} \quad (7)$$

Note that the probability distribution is a Gaussian in $|\gamma|^{1/2}$ for $|\gamma| \gg \frac{1}{2}$. The DCF of a and b is independent of lag (since the statistical properties of both a and b are independent of time), so its expectation value at any lag is the mean of γ , i.e., zero. However, the standard deviation of $P(\gamma)$ is $\sim 1.07M^{1/2}/(M-1)$, where M is the number of measured pairs in a bin. Thus, even completely uncorrelated time series are expected to produce values of the cross-correlation of $\sim \pm M^{-1/2}$.

III. COMPUTER SIMULATIONS

In order to test the DCF and interpolation methods, we have applied them to simulated data. The following example was chosen to mimic the physics of AGNs as well as the exigencies of ground-based observing, while at the same time remaining as conceptually simple as possible.

To that end, we choose the power spectrum of continuum fluctuations to be inversely proportional to frequency, i.e., flicker noise. We assume that the phases of different frequency components are independent and random. No good studies of the actual fluctuation power spectrum of the ionizing continuum over the temporal frequency range of interest (days to years) have yet been done, but studies of the X-ray continuum at higher frequencies (minutes to hours) indicate power-law behavior of this character (McHardy and Czerny 1987; Lawrence *et al.* 1987).

We also assume that the continuum is radiated isotropically and excites a spherical shell of isotropically radiating line-

emitting material which responds linearly to continuum fluctuations. The radius of this shell is taken to be 200 light-days, a number which is within the range predicted on the basis of photoionization models for the Seyfert galaxy Akn 120, whose actual time series we analyze in § Va.

All cross-correlation methods work well for finely spaced, evenly sampled, noise-free data. The goal of this paper is to understand how to work with coarsely spaced, unevenly sampled, noisy data; that is, the kind encountered in real life. Toward this end, we chose a sampling pattern which exactly duplicates the intervals between observations reported for the Akn 120 monitoring by Peterson *et al.* (1983, 1985) and Gaskell and Peterson (1987). They measured the flux in that object 58 times with a mean spacing of 38 days, but in a highly uneven pattern.

In another step toward verisimilitude, we add two different kinds of measurement error to the continuum and line light curves: each is altered by its own, independent Gaussian process with standard deviation of 7%, and then both are further disturbed by a different Gaussian process of the same standard deviation which is identical for the two light curves. The net result is to introduce a 1σ measurement error of 10% which is evenly divided between random and correlated error.

We define the variability parameter, R , as the ratio of the total rms fractional fluctuation in the time series to the estimated rms fractional error. Thus defined, the minimum value of R is unity. This ratio is more important than the measurement signal-to-noise ratio in determining the quality of the resulting correlation functions. Again, in order to match the condition of the Akn 120 data, we constructed the simulation to also have similar values of R : 2.8 for the continuum, and 1.6 for the line. Figure 1 shows the line and continuum light curves which result from this procedure. It is apparent that the continuum is very badly represented by this sampling pattern and that a smooth interpoint interpolation would be extremely misleading.

The resulting correlation functions are presented in Figure 2. The solid line is the true correlation function, measured by applying equation (2) to the continuous light curves shown in Figure 1. As expected, the cross-correlation is large near 200 days. However, the peak of the curve is very broad, its half-maximum points falling at -50 and $+400$ days. Even the 90% maximum level spans the range from $+35$ to $+275$ days. The dotted line is the cross-correlation function using the noisy, irregularly sampled data shown in Figure 1 and calculated by the interpolation method. The individual points with error bars are the result of the DCF method applied to the noisy, irregularly sampled data of Figure 1 using a bin size of 50 days. The difference between Figures 2a and 2b is that the noise level was doubled for the data used to construct Figure 2b.

At a crude level, the interpolation method and the DCF do comparably well in reproducing the ideal cross-correlation. Both are large in roughly the right place, and are generally small in magnitude elsewhere. Both also deviate from the ideal curve in similar fashion. In detail, they differ in several respects. The interpolation curve has three peaks: at 0, 110, and 400 days. With no additional information, any one of them might be chosen as the true peak. Formally, the DCF has a single peak at 150 days, but the large error bars at this bin and at 200 days bespeak the lack of pairs at intervals of half a year imposed by seasonal constraints on ground-based observations.

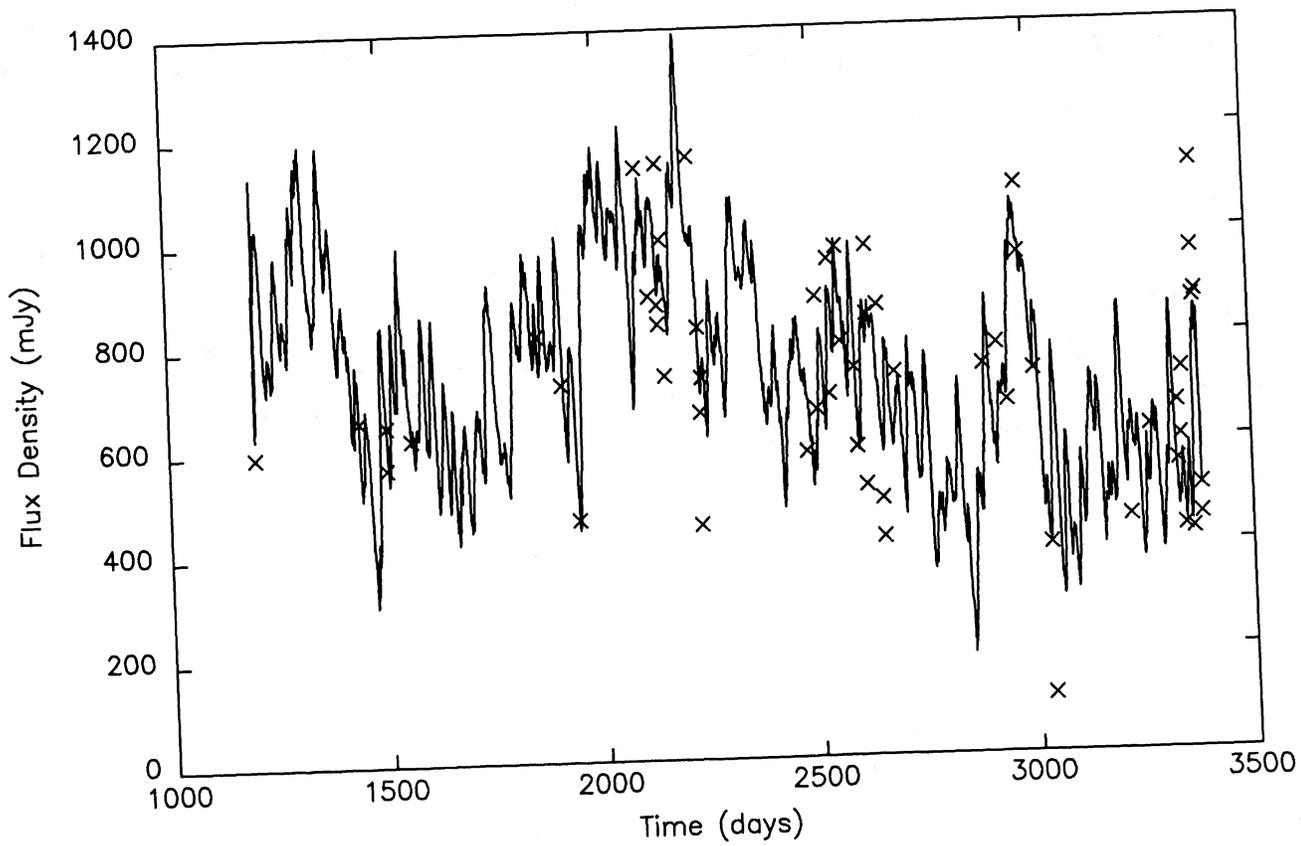


FIG. 1a

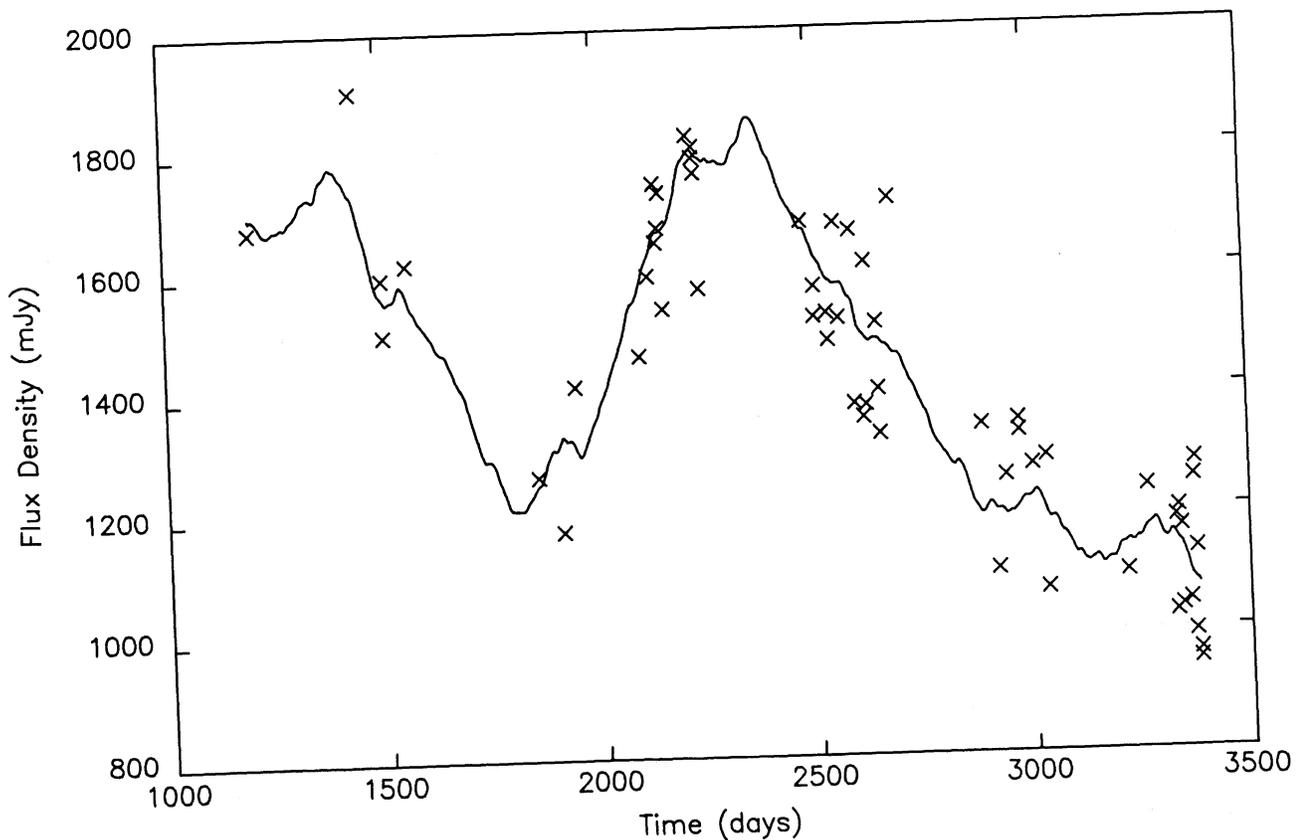


FIG. 1b

FIG. 1.—(a) Simulated continuum light curve with flicker noise power spectrum. Crosses represent this light curve sampled at intervals corresponding to those in the Akn 120 monitoring, and with 10% noise added as described in the text. (b) Simulated line light curve driven by the continuum shown in (a) according to the model described in the text, sampled as in (a).

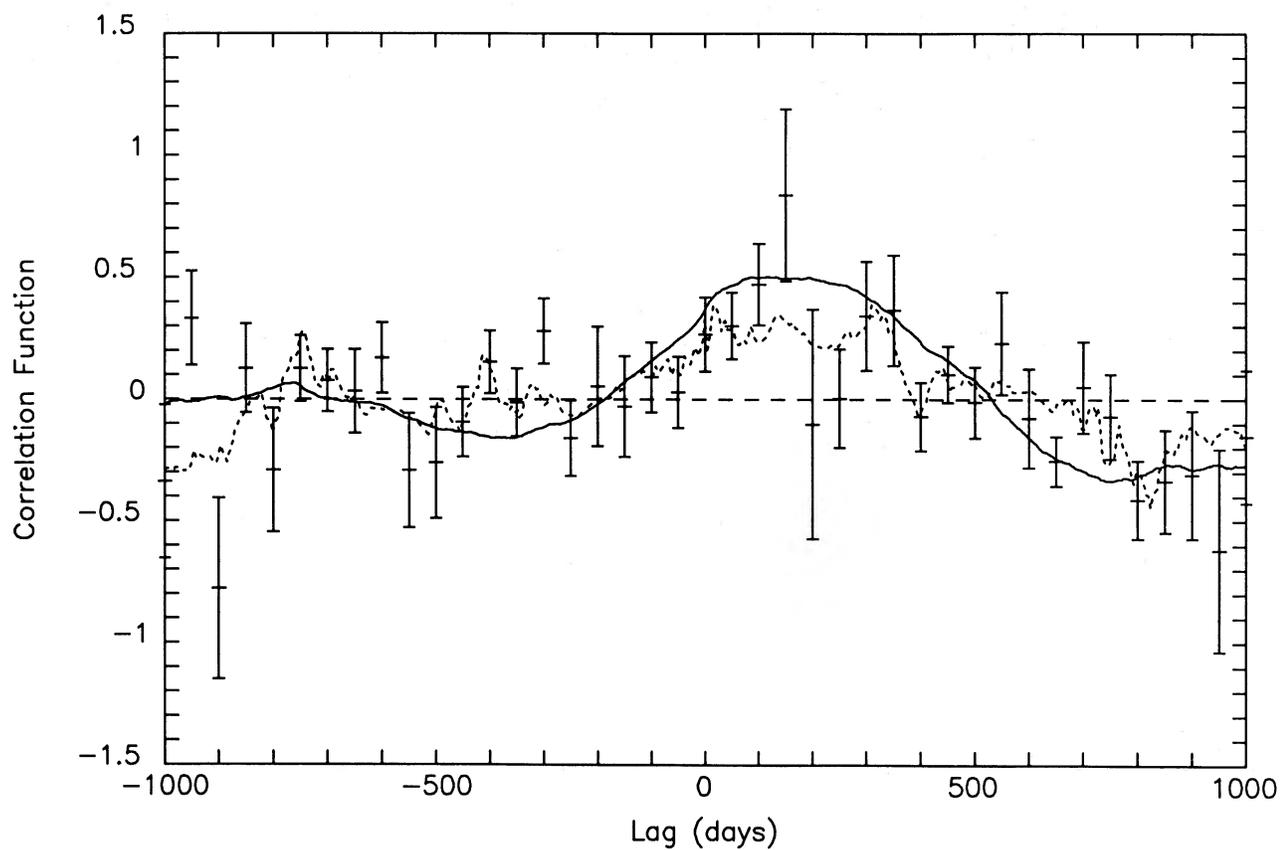


FIG. 2a

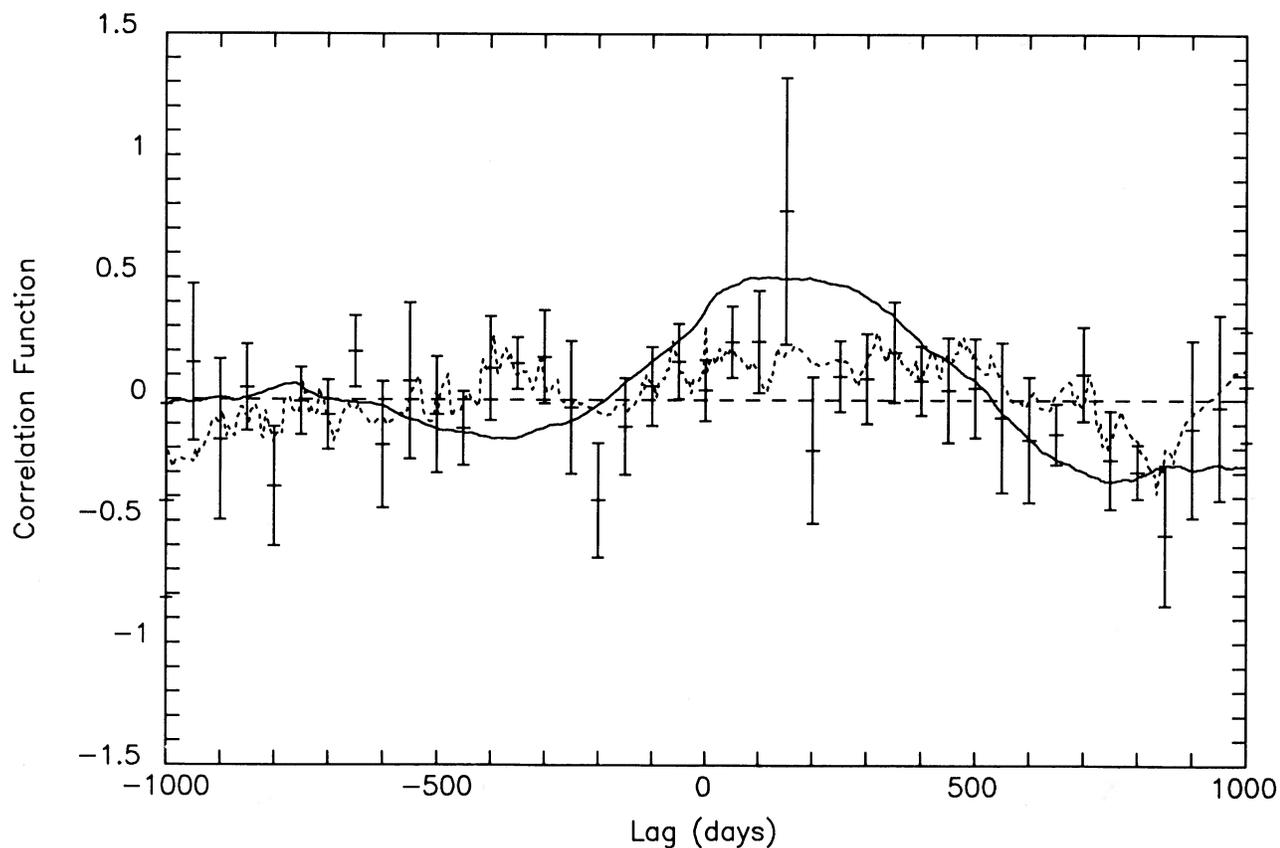


FIG. 2b

FIG. 2.—(a) Cross-correlation functions for the simulated data. Dotted line is the cross-correlation function measured with the interpolation method, sampled every 10% of the mean spacing. The points with error bars are the discrete correlation function, binned in intervals of 50 days. (b) Same as (a), except that the measurement noise was doubled to 20%. Note that this increase in noise effectively wipes out the correlation.

Their most important contrast is in the error bars provided by the DCF method. Both the DCF method and the interpolation method show a dip near +200 days. Although there would be no way to judge its reality in terms of the interpolation method, the large error bar on the DCF point shows that in fact the true level at that lag cannot be distinguished from the level obtained between 0 and 300 days. Its spurious character is made even clearer by simulations done with the same noise level, but regular sampling with the same average interval, which show that this feature is due to the peculiarities of the sampling pattern. We also comment that the fact that the true cross-correlation function lies within the error bars $\sim 60\%$ of the time demonstrates that the error bars are genuine measures of the 1σ uncertainty.

In this simulation, the correlated error produces a clear peak precisely at zero lag in the interpolation method curve. Other simulations (e.g., Fig. 2*b*) demonstrate that at slightly smaller signal-to-noise ratios, this zero-lag peak grows in prominence in the interpolation method cross-correlation as true correlation disappears from the rest of the curve. It would be easy to remove this feature, which is the result of correlated errors, if its influence were felt *only* at zero lag. However, the interpolation method spreads its effects over a characteristic width which is related (in a way made rather complicated by the irregular sampling) to the mean spacing (in this case 38 days). Correlated error, as we have previously remarked, is very easily removed in the DCF method.

The adequacy of the sampling for cross-correlation analysis depends strongly on whether the autocorrelations in the individual time series are significant at time differences comparable to the typical spacing. The interpolation method does as well as it does in this simulation because the line flux has substantial autocorrelation up to ~ 400 days, and even the continuum flux is significantly autocorrelated at large lags because of the properties of flicker noise.

The specific value of R required to obtain a meaningful result is a function of both the sampling pattern and the physical model assumed. Within those limitations, a minimum R can be determined by simulation. For instance, Figure 2*b* shows the result of doubling the error level imposed on the same underlying light curves as shown in Figure 1. The amplitude of the cross-correlation found by the interpolation method is very nearly cut in half by the additional noise, and the clearest feature remaining is the zero lag peak created by the correlated error. Even the cross-correlation curve produced by the DCF method has almost no real features remaining. For this case, then, the minimum R is around 2.2 for the continuum and 1.4 for the line.

Simulations show that other sampling patterns and other physical models require different values of R to yield a meaningful result. For the physical model discussed in this section, sampling at regular intervals of 38 days would allow much larger error levels: R values as low as 1.5 for the continuum and 1.15 for the lines could be tolerated. On the other hand, as we show in the following section, even noise-free data sampled at very small, regular intervals produce a very weak cross-correlation when $a(t)$ has strong high-frequency components.

IV. INTERPRETATION

Once a cross-correlation curve has been computed, and its reliability assessed, the next task is to use it to derive physical information. Here we devote ourselves to the case in which $b(t)$ and $a(t)$ are believed to be related by a convolution, and we

seek information about $\Psi(\tau)$. For there to be any usable features in the cross-correlation function, either Ψ or a must have at least one characteristic time scale. In this section, we determine which features of the cross-correlation curve can be used to infer possible characteristic time scales embedded in the response function Ψ , and how their interpretation depends on both the continuum power spectrum $|\hat{a}(f)|^2$ and symmetry assumptions.

Plausible, but erroneous, statements about this program abound in the literature. For example, if there is a one-to-one functional relationship between characteristic time scales and characteristic length scales (e.g., by light ray kinematics), it is frequently assumed that the principal characteristic length scale of the system can be directly identified with the time lag at which the cross-correlation has its maximum. Unfortunately, this is *not* true in general. Furthermore, the position of the peak is quite commonly *not* the most prominent feature associated with the principal characteristic time scale. The specific features which *do* reveal that scale are determined by both the statistical character of the fluctuations in $a(t)$ and the form of $\Psi(\tau)$.

These problems are illustrated clearly if the cross-correlation function is rewritten (in the continuous limit) in terms of Fourier transforms:

$$CF(\tau) = \frac{1}{\sigma_a \sigma_b (T - \tau)} \int df e^{-2\pi i f \tau} |\hat{a}(f)|^2 \hat{\Psi}(f), \quad (8)$$

where T is the total duration of the run of data and \hat{X} is the Fourier transform of the time-dependent function X . The function $\hat{\Psi}(f)$ may be regarded as a frequency-dependent gain function, so the importance of different fluctuation frequencies to the cross-correlation depends *jointly* on the continuum power spectrum $|\hat{a}(f)|^2$ and the gain $\hat{\Psi}$.

Once again, our illustrative example is based on AGN broad emission lines. The delay between continuum fluctuations and their associated line fluctuations is entirely due to geometry because the local response time of line-emitting material to fluctuations in the continuum strength is extremely short (< 1 hr) compared to the expected light travel times across the region (days to years). Although lines do not respond in a truly linear fashion to continuum fluctuations (Gaskell and Sparke 1986), typically $|\partial \log F_l / \partial \log F_c| \approx 1$ (T. R. Kallman, private communication), and the fluctuation level is rarely more than 50%. Thus, the linear approximation implicit in the convolution description is generally adequate.

For simplicity, we restrict our attention, for the most part, to cases in which $|\hat{a}(f)|$ has no characteristic scales, and Ψ has no more than two. These constraints are appropriate because (as was remarked in § III) AGN continuum power spectra appear to be broad band and smoothly varying, and because it is generally thought that the dynamic range in radius over which the emission lines are made is no more than a factor of a few (Mathews and Capriotti 1985).

A particularly simple example to study is that of a spherical shell of isotropically emitting material at radius r , responding to a central source of isotropic radiation. In that case, the gain is

$$\hat{\Psi}(f) = -\frac{ic}{4\pi r f} (e^{4\pi i r f/c} - 1). \quad (9)$$

The gain is unity at frequencies $f \ll c/r$, and falls off roughly as f^{-1} for higher frequencies. This behavior is, in a qualitative

way, quite general. Whenever the emitting region has a finite size, $\Psi(f) \rightarrow 1$ as $f \rightarrow 0$. Physically, this is because the line flux can track perfectly continuum fluctuations with frequencies $f \ll \tau_0^{-1}$, where τ_0 is the light-crossing time. Similarly, if most of the line power comes from a region with characteristic light-crossing time τ_0 , $\Psi(f) \rightarrow 0$ as $f \rightarrow \infty$; i.e., fluctuations faster than the light travel time are smeared out.

We immediately see, then, that in order to produce any recognizable signal of a characteristic time scale τ_0 in $\Psi(\tau)$, there must be significant fluctuation power in the continuum at frequencies corresponding to roughly τ_0^{-1} . Examination of equation (8) then shows that if the only fluctuation power is at frequencies $f \ll \tau_0^{-1}$, the cross-correlation reduces to the auto-correlation of the continuum, while if the only fluctuation power is at frequencies $f \gg \tau_0^{-1}$, the magnitude of the cross-correlation is very small. Thus, if $|\hat{a}(\tau_0^{-1})| \ll 1$, it is entirely possible for the peak of the cross-correlation to fall at $\tau = 0$, independent of the value of τ_0 .

The exact criteria required for the continuum power spectrum to produce a useful line-continuum cross-correlation depend on the details of the line-emitting geometry. For example, for the particular $\Psi(f)$ given in equation (9) (e.g., an isotropic continuum source, with isotropic radiation from a spherical shell of radius r), there must be significant power above $f_{\min} \approx (24)^{1/2}(4\pi)^{-1}(c/r)$ in order to distinguish the line-continuum cross-correlation from the continuum autocorrelation. On the other hand, there must be significant power below $f_{\max} \approx (4\pi)^{-1}(c/r)$ in order to maintain a nonzero cross-correlation. Other geometries would change these coefficients of c/r by factors of order unity.

a) Intrinsic Scale Bias

Scale biases (Gaskell and Sparke 1986) are very difficult to avoid whenever the continuum power spectrum has enough power at high frequencies to probe the shortest time scales in $\Psi(\tau)$. When that is the case, the cross-correlation amplitude at $\tau = r/c$ due to material with characteristic scale r can be approximated by:

$$CF_r(\tau) \approx \frac{2}{\sigma_a \sigma_b (T - \tau)} \int_0^{c/(2\pi r)} df |\hat{a}(f)|^2, \quad (10)$$

where we have used the fact that the gain function has an effective cutoff at frequencies above $\sim c/(2\pi r)$ and the oscillation in the inverse transform adds to the suppression of higher frequency contributions. Clearly, as r decreases, a wider range of frequencies contributes to the integral and its value grows. For example, if $|\hat{a}(f)|^2 \propto f^{-1}$ above some lower cutoff frequency f_0 , to the accuracy of this approximation, the correlation function is proportional to $\ln [c/(2\pi r f_0)]$. The quantitative dependence on r would change if the power spectrum were different or if the responding material's configuration were changed, but the basic sense is always in the same direction, emphasizing the smallest sizes.

An example of this effect is presented in Figure 3. It shows the cross-correlation of arbitrarily finely sampled, noise-free simulated data produced by a model in which an isotropic continuum source with a flicker noise power spectrum excites two spherical shells with equal line luminosities, one at $r = 50$ lt-days, the other at $r = 200$ lt-days. The weak shoulder at $\tau \approx 200$ days would be lost in the noise with real data.

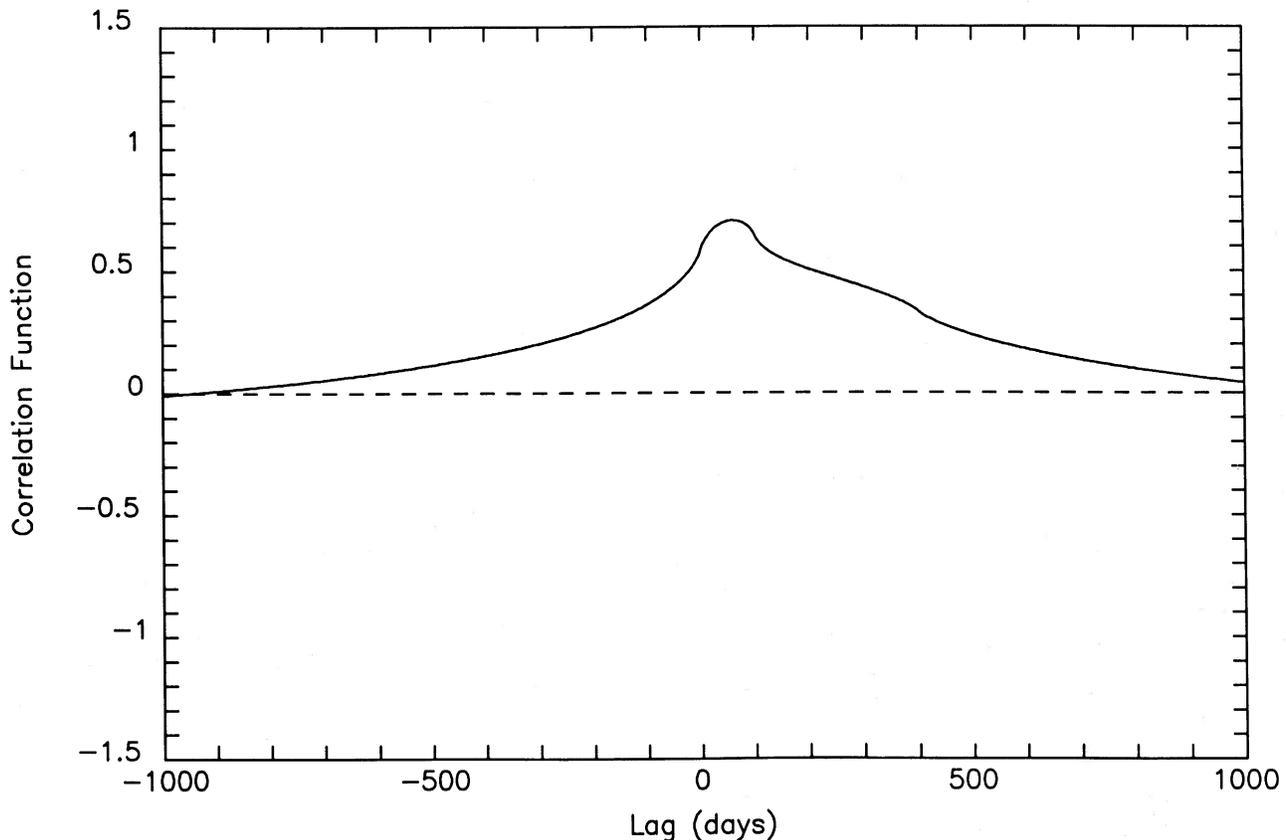


FIG. 3.—The cross-correlation produced by two shells of equal power but different distance from the center. Only the peak due to the smaller shell is apparent.

b) *Dependence on Input Function Power Spectrum*

As is readily seen in equation (8), the *shape* of the cross-correlation function is strongly influenced by the character of the power spectrum $|\hat{a}(f)|^2$. White noise provides a particularly striking example of this influence: if $|\hat{a}(f)|^2$ is a constant, the shape of the cross-correlation curve is identical to that of the response function $\Psi(\tau)$.

“Blue” fluctuation spectra (i.e., spectra with more power toward higher frequencies) are genuinely pathological. In general, their cross-correlations are of very small magnitude except for possible divergences at a finite number of points related to characteristic time scales of Ψ (e.g., $|\hat{a}(f)| \propto f$ operating on a spherical shell, which produces a cross-correlation of zero everywhere except for δ -functions at $\tau = 0$ and $\tau = 2r/c$).

However, the fluctuation spectra of most continuum sources are “red.” These spectra are easier to work with, producing more robust curves with greater magnitude correlations. Figure 4 shows the “ideal” cross-correlation, computed by evaluating equation (8), for a variety of continuum power spectra $|\hat{a}(f)|^2$ operating on the same spherical shell. For fixed Ψ , the peaks tend to narrow and the overall magnitude of the cross-correlation increases as the fluctuation spectrum becomes “redder.” Both the greater magnitude and the sharper peak stem from the growing relative importance of slow fluctuations, which any structure can track precisely. Unfortunately, this trend, which makes interpretation easier, also makes measurement harder: at constant variance, the large

fluctuations which are required for real events to stand out above measurement noise move farther apart.

In many cases, the positions of the “shoulders” of the cross-correlation are better indicators of the characteristic scale than the position of the peak. The peak is often very low and broad, while the places of greatest derivative, or alternatively, the positions of half maximum, may be more easily discerned. For power spectra which produce vague peaks, the shoulders lie near $\tau = 0$ and $\tau = 2r/c$. Thus, at least for these “red” spectra, there is always a feature which can be used, but the best choice depends on the continuum power spectrum (assuming, of course, that the system has a characteristic time scale).

c) *Symmetry Assumptions*

Our comments so far have focused on line emission from isotropically radiating, spherically symmetric distributions. If the true distribution is less symmetric, the most prominent features for a given input power spectrum change their character and their position, even if the underlying scale stays the same. Consider two examples, the first having spherical geometry but anisotropic radiation, the second having isotropic radiation but aspherical geometry. If the emitting material is confined to a spherical shell, but radiates predominantly in the direction toward the ionizing source (as it would if the line is very optically thick and the outside edge of the emitting material is colder than the inside edge), then the peak is sharper than it would be if the same continuum fluctuations illuminated an isotropic radiator, but it also moves toward

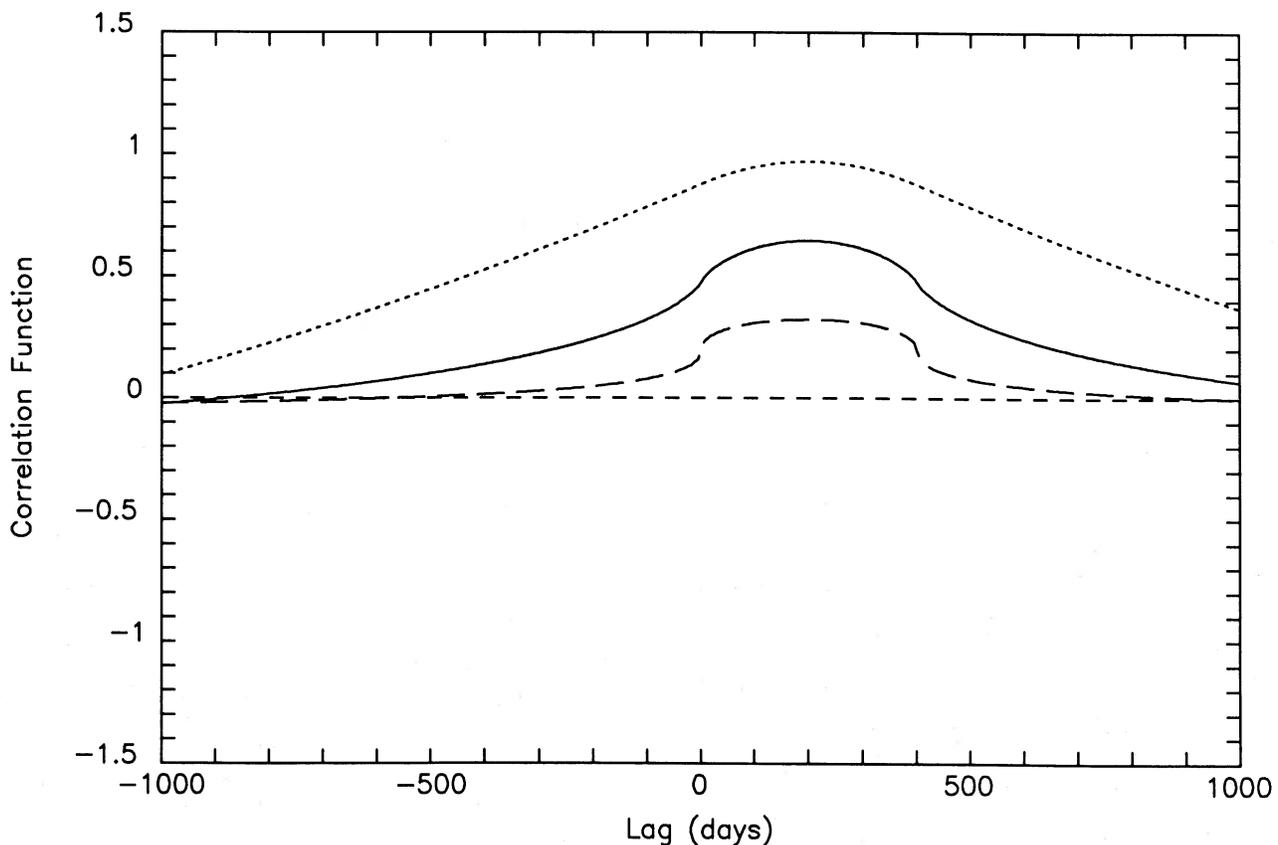


FIG. 4.—“Ideal” cross-correlations for isotropically emitting spherical shells irradiated by continua with a variety of power spectra. The power spectrum corresponding to the dotted line is proportional to f^{-2} , that corresponding to the solid line is proportional to f^{-1} , and that corresponding to the broken line is proportional to $f^{-1/2}$.

$\tau = 2r/c$. On the other hand, rings of isotropically emitting material tend to amplify the cross-correlation near $\tau = (r/c)(1 \pm \sin \alpha)$, where α is the viewing angle between the ring axis and the line of sight. Both flattened cross-correlation curves (when the input power spectrum is very "red") and double-peaked curves (when the input power spectrum is close to that of white noise) may result.

Clearly, a great many considerations must be taken into account before even a well-determined cross-correlation function can be interpreted physically. Before drawing conclusions about geometrical structure from these functions, it is always necessary to clarify the symmetry assumptions which have been made, and to study the nature of the continuum fluctuations.

V. APPLICATIONS

As a second test, we apply the DCF and interpolation methods to published data from two variability studies: optical continuum and $H\beta$ line flux observations of Akn 120 (Peterson *et al.* 1983, 1985; Gaskell and Peterson 1987), and *IUE* monitoring of the continuum and broad component of the C IV $\lambda 1550$ line in NGC 4151 (Clavel *et al.* 1987).

a) Akn 120

Figure 5 is a plot of the cross-correlation between broadband optical flux and $H\beta$ line flux, measured with both the DCF and interpolation methods. Gaskell and Peterson (1987) have also analyzed these data by the latter technique, and claim to find a broad peak whose most likely position is at +7 days, but which is consistent with being anywhere from 0 to +30 days. Our application of the DCF method is performed with two different bin widths in order to test two proposed sizes for the emission line region: 50 day bins to test the photoionization prediction of ~ 150 days; and 5 day bins to test the peak near 7 days claimed by Gaskell and Peterson.

The interpolation method curve in Figure 5 shows a peak at 0 days lag, not +7 days; this peak is entirely due to correlated error (see § III). There is also a broad shoulder in the interpolation curve that extends out to +180 days, and an equally strong positive feature near +400 days.

The DCF undercuts the reliability of *any* feature in the interpolation cross-correlation curve. It shows *no* clear peak; indeed, the (unphysical) negative lag side has as many positive features as the positive lag side, and there is only the merest hint of positive correlation between -10 and $+25$ days. Moreover, in both binnings, 60% of the points lie within 1σ of zero; this is, of course, consistent with a true correlation which is identically zero. Near lags which are odd multiples of half a year, the error bars are very large because few independent points contribute. It appears that it is impossible to reach any firm conclusions regarding the characteristic size or structure of the BLR in Akn 120 on the basis of these data.

There are a number of possible reasons why a correlation was not detected, and yet the hypothesis of photoionization driving of the lines would still be valid: (1) the source may be either much larger ($\gg 1$ pc) or much smaller (\ll a few light-days) than sizes sampled by the data; (2) the gas may be distributed in an unfavorable geometry which smears out correlations; (3) if the true shape of the continuum power spectrum is closer to white noise than flicker noise, the amplitude of any cross-correlation is diminished (see § IV); (4) the optical continuum may not be well correlated with the ionizing continuum ($H\beta$ responds to the entire continuum from 1 Ry up to

a few keV; Krolik and Kallman 1988); or (5) the measurement errors may have been underestimated—our simulations show that if the true error level is 15%, rather than the quoted value of 10%, R would probably be below the minimum useful level. With regard to the last point, the large width of the correlated error spike suggests that there were extra errors correlated on time scales of a few days (due, for example, to calibration errors which persisted throughout a run).

It is particularly unfortunate that this monitoring data has so few pairs of points separated by half a year, because the best guess photoionization model estimate of the characteristic size of the broad line region in this galaxy is $r \approx 150(L_{\text{ion}}/L_{\text{UV}})^{1/2}(\Xi/0.2)^{-1/2}h^{-1}$ lt-days. This estimate assumed that the ionization parameter Ξ and pressure in the broad line gas are the "typical" values found, e.g., by Kwan and Krolik (1981). The line ratio data of Wu, Boggess, and Gull (1983) are consistent with this assumption, or possibly with an ionization parameter up to a factor of 2 or so larger. The parameter L_{UV} is defined as the monochromatic luminosity λL_{λ} at 1450 \AA : the values shown by Alloin, Boisson, and Pelat (1988) were combined with those given by Wu, Boggess, and Gull (1983) to estimate its value. The symbol h is the usual abbreviation for the Hubble constant in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$, so a smaller value would imply a larger emitting region. Combining the uncertainties in these scaling factors yields an overall uncertainty in the photoionization prediction for the characteristic scale of at least a factor of several in either direction.

b) NGC 4151

The results of the second study, which sought a correlation between ultraviolet continuum spectral flux density (average of the values at 1450 \AA and 1710 \AA) and (broad) C IV line flux in NGC 4151 (Clavel *et al.* 1987), are plotted in Figure 6. Clavel *et al.* have claimed, on the basis of fits "by sight" of selected stretches of data, that the cross-correlation is greatest for lags around +5 days. Gaskell and Sparke (1986), using part of these data, found a similar result. These claims and the photoionization prediction of ~ 15 days characteristic scale (see discussion below) can all be tested by applying the DCF method with a bin width of 5 days. We also plot the result of binning in 50 day segments in order to demonstrate that there is no significant correlation at large lags. The continuum autocorrelation is plotted in Figure 7.

In this case, the DCF method confirms the overall shape of the interpolation method curve, except for the peak at precisely zero lag in the latter curve. More strikingly, the cross-correlation as measured by the DCF method is very nearly symmetrical around zero lag, and has a large positive amplitude between approximately -300 and $+300$ days. In fact, the cross-correlation is almost identical to the continuum *autocorrelation*. The only place where they differ significantly (the deviations near ± 200 days are less than 1σ because the error bars are so large) is between -15 and $+15$ days, where the autocorrelation is ~ 0.3 larger. Thus, except for a narrow region around zero lag, fluctuations in the broad component of the C IV emission line are statistically indistinguishable from fluctuations in the neighboring continuum.

One possible explanation for this identity is that the fluctuations in both the line and the continuum are due to a correlated error of duration ~ 300 days. It is more likely, however, that we are seeing an effect alluded to in § IV: if there is little continuum power at frequencies high enough to probe the structure of $\Psi(\tau)$, the cross-correlation simply reduces to the

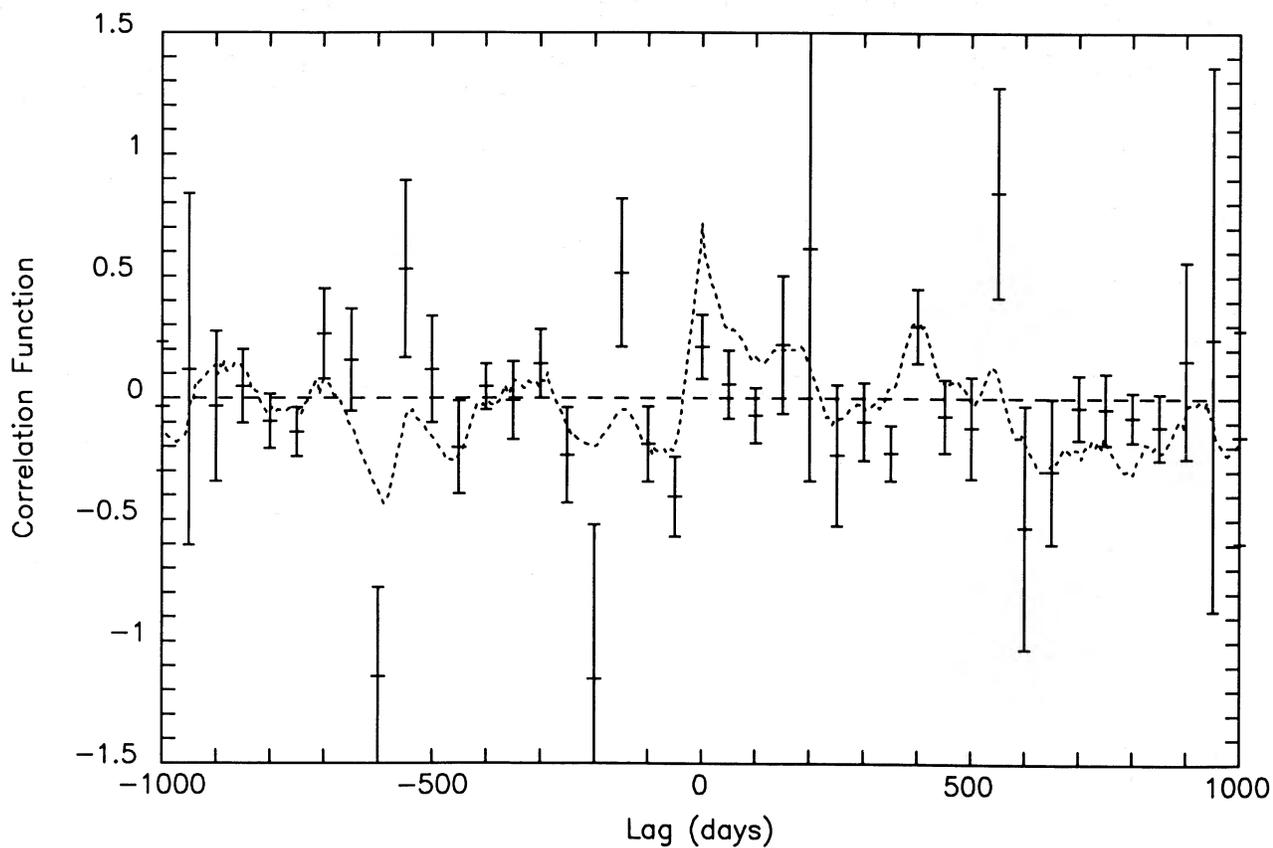


FIG. 5a

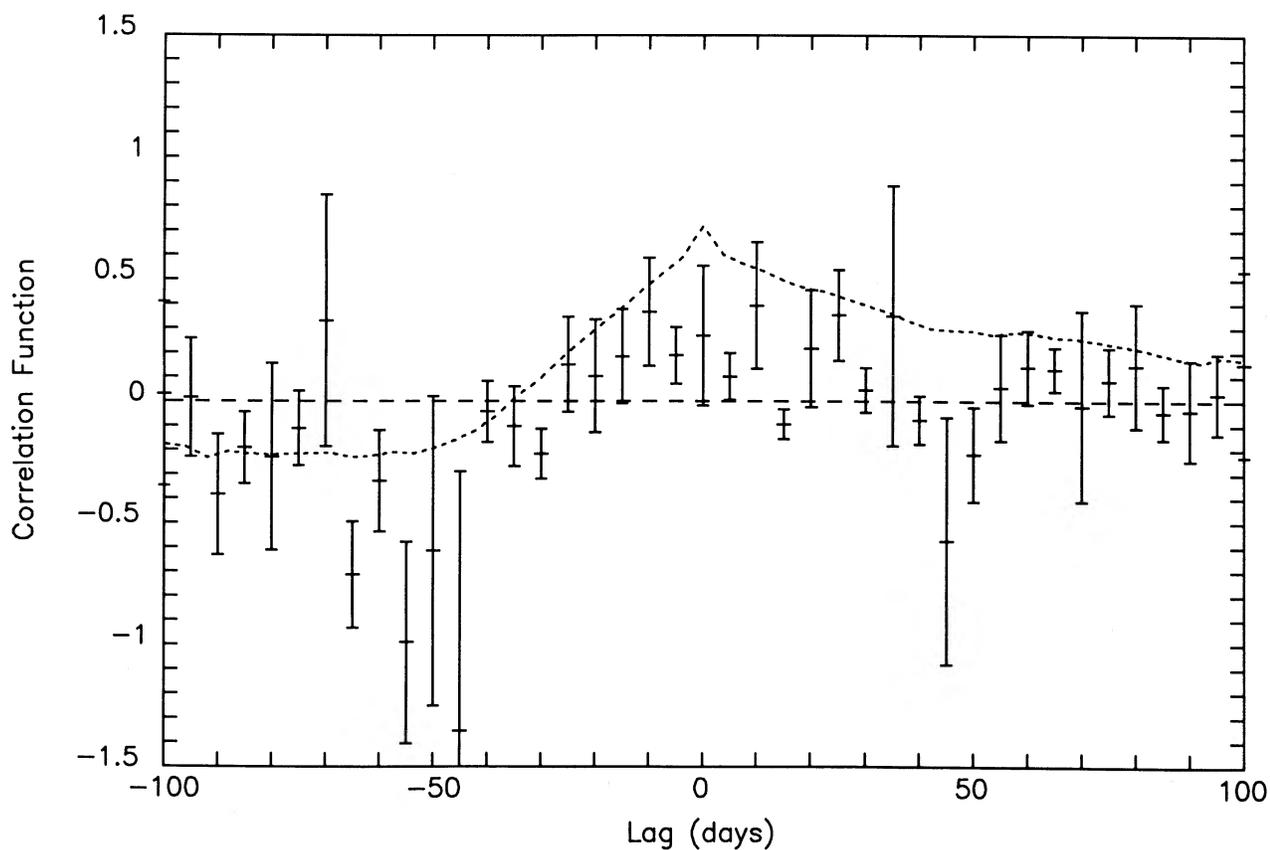


FIG. 5b

FIG. 5.—(a) Cross-correlation functions for the Akn 120 data. Dotted line is the cross-correlation function measured with the interpolation method, sampled every 10% of the mean spacing. Points with error bars are the discrete correlation function, binned in intervals of 50 days. (b) Same as (a), but plotted on a larger scale to show the region from -100 to $+100$ days more clearly and with the DCF binned every 5 days.

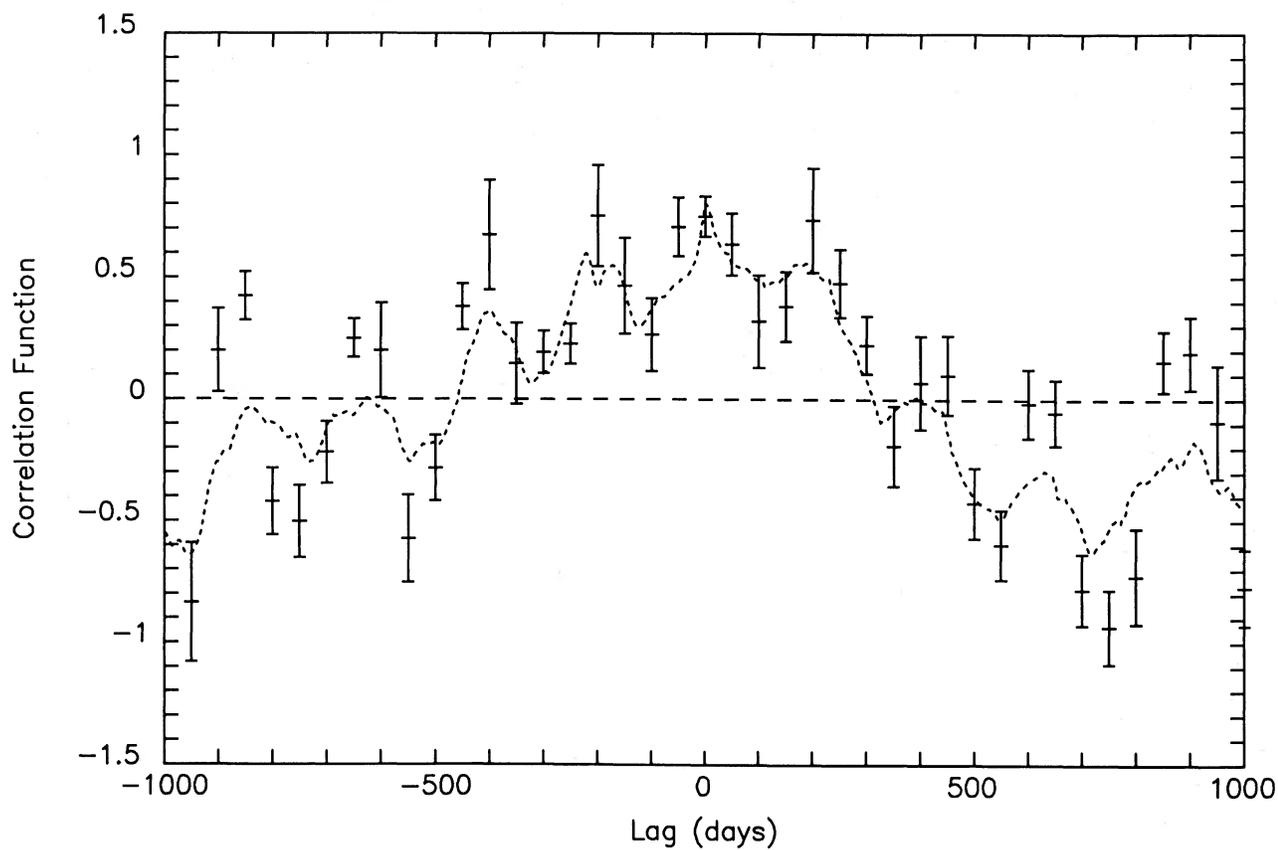


FIG. 6a

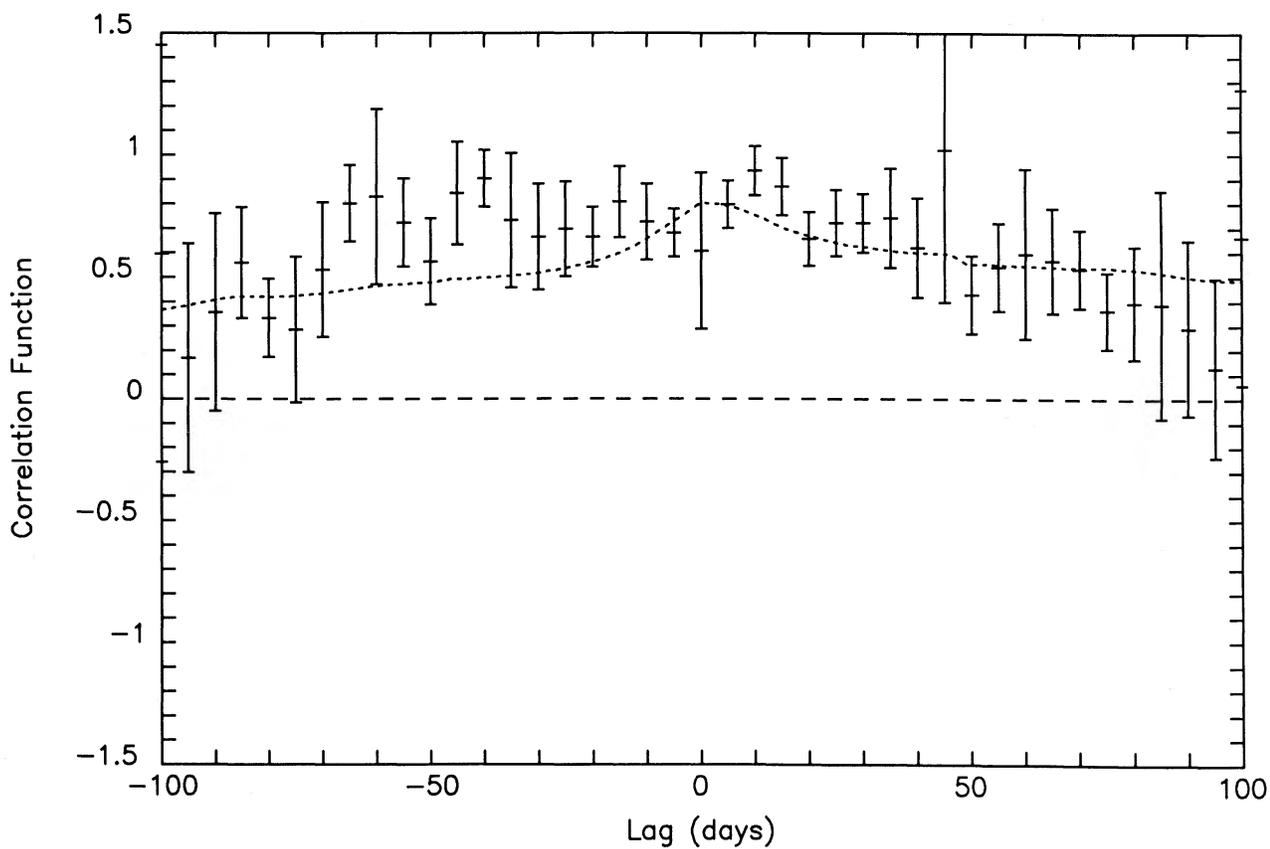


FIG. 6b

FIG. 6.—Cross-correlation functions for the NGC 4151 data. Dotted line in both panels is the result of using the interpolation method, sampled every 10% of the mean spacing. Points with error bars are the discrete correlation function, binned (a) every 5 days and (b) every 50 days.

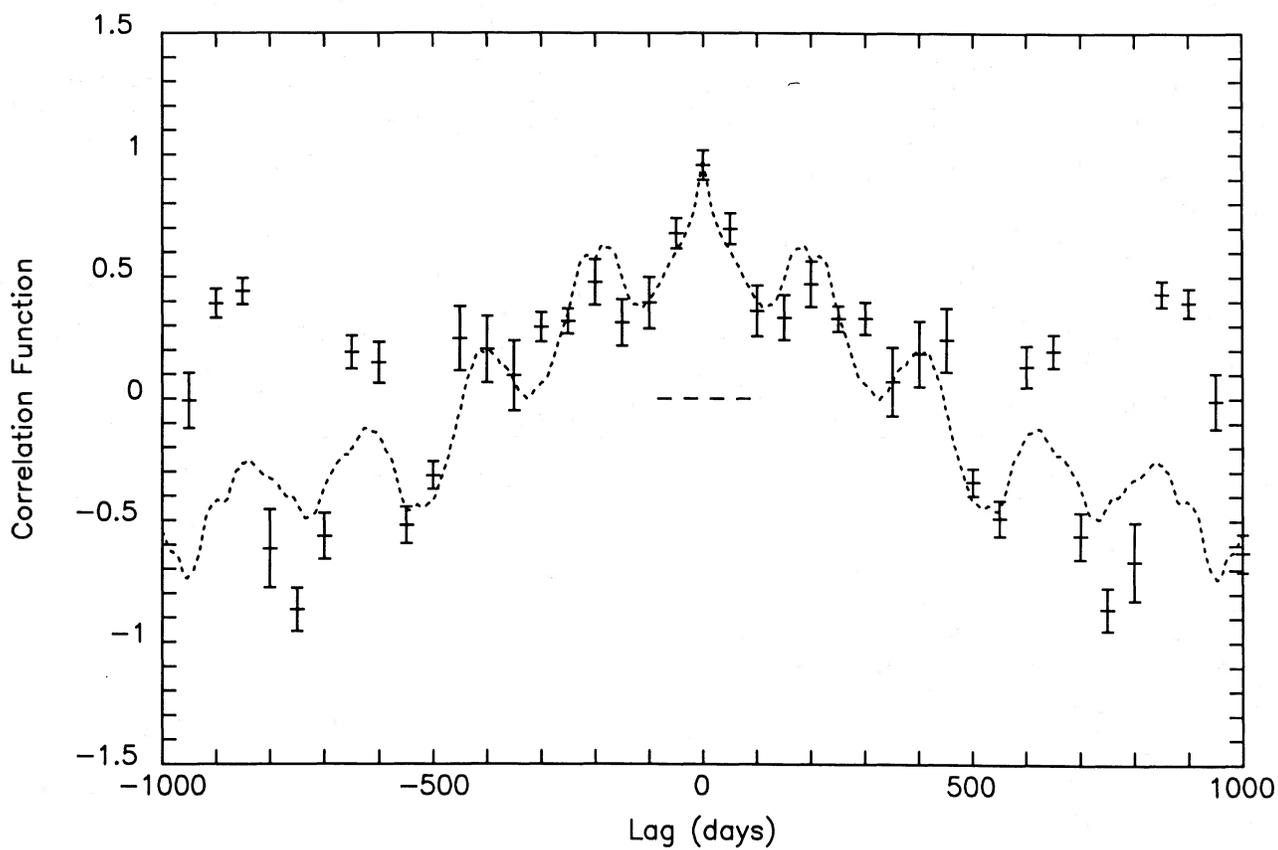


FIG. 7a

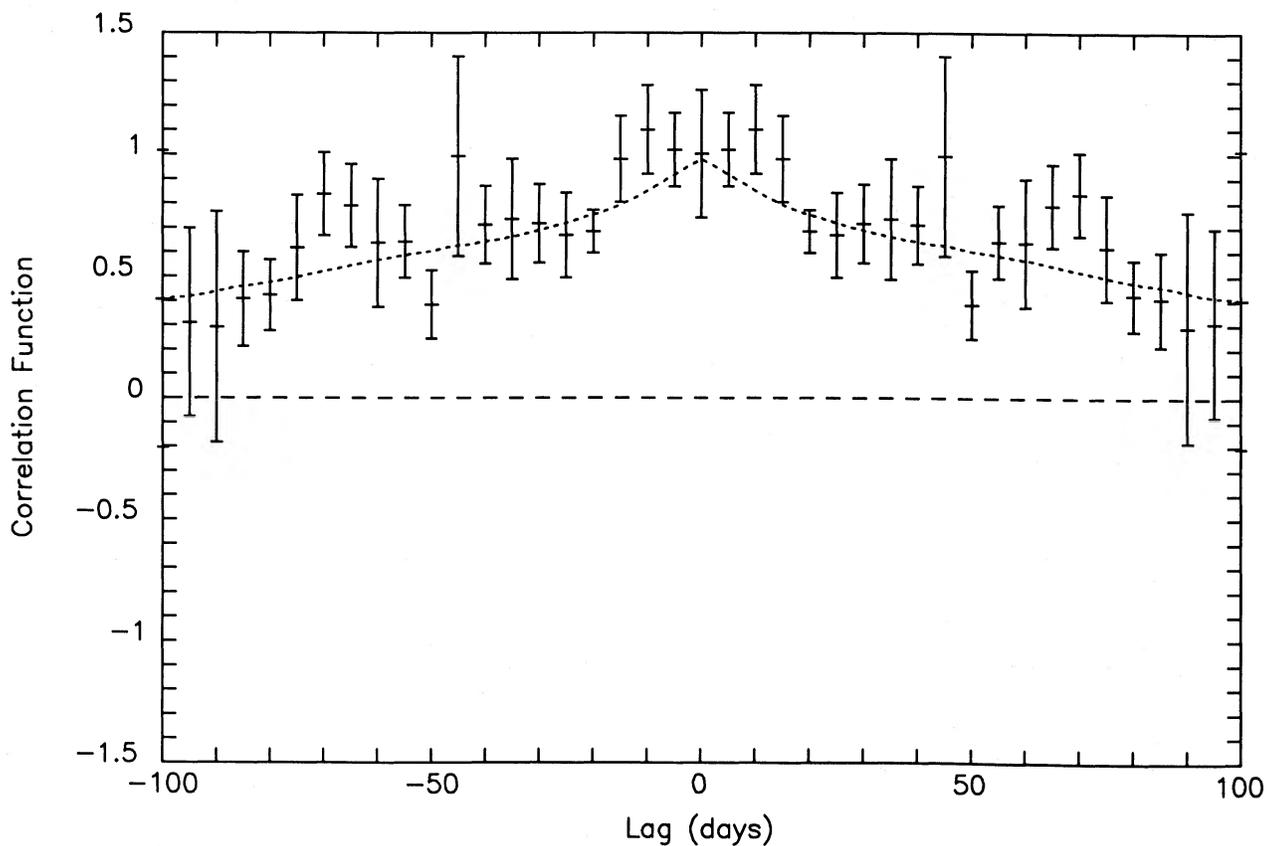


FIG. 7b

FIG. 7.—The continuum autocorrelation function for the NGC 4151 data computed by the DCF method for a bin width of (a) 50 days and (b) 5 days.

continuum autocorrelation. Although the power spectrum is not very well determined by this data, a Fourier transform of the autocorrelation function does show that it drops sharply at frequencies around $f \approx 0.02 \text{ d}^{-1}$. (This is also why the interpolation method does relatively well with this data.) The analytic arguments of § IV indicate that the close tie between the cross-correlation and the autocorrelation places an upper limit on the size of the emitting region which would be $r \approx 20$ lt-days if the simple isotropic spherical shell model applies. We emphasize, however, that this limit is uncertain by at least a factor of 2 due to our ignorance of the true symmetry of the emission region, and by at least another factor of 2 due to the imprecision of our estimate of the power spectrum. On the other hand, the clear failure of the cross-correlation to reproduce the continuum autocorrelation within 15 days of zero lag (computing the correlation functions with 30 day bins demonstrates that the two functions are $\sim 2.5 \sigma$ apart in this range) indicates that the small amount of power at periods shorter than 15 days finds little response in the line. If the criterion suggested in § IV for the simple isotropic spherical shell model is applied to all frequencies 0.066 d^{-1} and higher, a lower limit on the size of the line-emitting region of ~ 1.2 lt-days is found. Again, this number is uncertain by at least a factor of 2 due to uncertainty in the true geometry and by at least another factor of 2 due to the roughness of our criteria for frequency matching.

These results are entirely consistent with photoionization models. The very large amplitude of the cross-correlation certainly argues for a close relation between continuum excitation and line emission. In addition, estimates of the likely length scale, made in the same manner as for Akn 120, suggest a size of $r \approx 15(\Xi/0.2)^{-1/2}(L_{\text{ion}}/L_{\text{UV}})^{1/2} h^{-1}$ lt-days. L_{UV} is defined as for AKN 120, but the data are taken from Clavel *et al.* (1987). If anything, the measured line ratios would indicate a somewhat larger value of Ξ than we have chosen as our fiducial value, further diminishing the predicted size of the line-emitting region, but a smaller value of H_0 would increase it.

It should not be too surprising that these measurements provide a clearer signal than the optical monitoring of Akn 120. The ionizing continuum which powers the C IV line is much closer in wavelength to the measured continuum than is the case for the optical estimates of the continuum responsible for H β ; variability amplitude tends to increase with photon frequency in active galaxies (Cutri *et al.* 1985); and the more uniform sampling made possible with IUE, all combine to make it easier to detect any intrinsic correlation. For instance, if we assume that the measurement errors are 10% in this case

(no error estimate is given in Clavel *et al.* 1987), R is 7.4 for the continuum and 6.8 in the line, since the rms variations are almost an order of magnitude larger than in the optical Akn 120 data.

VI. CONCLUSIONS

We have presented a new method for analyzing correlations in time series data. The discrete correlation function has the advantage of taking a very conservative approach to the data: no interpolation, that is, no "invention" of data, is required. Moreover, this method permits the easy elimination of spurious effects due to correlated errors. Its final major advantage is that, unlike other methods used to compute correlations, it gives a quantitatively meaningful error estimate.

We also present a discussion of the problems of physical interpretation of cross-correlations. We have raised a number of previously ignored issues, such as the influence of the continuum power spectrum on the shape of the cross-correlation curve and how to choose which features of that curve to identify with the characteristic time scales of the underlying system. Bias in favor of small scales and symmetry assumptions also color interpretation of these correlation functions.

Finally, we applied this new technique to two sets of Seyfert galaxy line and continuum monitoring data. In one case, Akn 120, we were unable to confirm the presence of any significant cross-correlation. This is in good agreement with simulations, which show that a high ratio of intrinsic variability to measurement error is required to derive significant information from unevenly sampled observations. In the other case, NGC 4151, strong positive cross-correlation is present, but there is so little high-frequency variability in the continuum that only very approximate limits may be placed on the size of the line-emitting region. Nonetheless, these limits—between ~ 1.2 and ~ 20 lt-days—are good enough to demonstrate quite satisfactory agreement with the predictions of the photoionization model.

The authors would like to thank A. Lawrence, A. Hamilton, J. M. Shull, A. Uomoto, and R. White for useful discussions. They would also like to thank the organizers of two meetings, F. Córdova and J. Clavel, for provoking their interest in the subject, and, in the latter case, for providing an opportunity for a first discussion of the ideas presented in this paper. In addition, the referee, J. Scargle, made helpful suggestions. The work of R. A. E. was supported by NASA grants NAG-5-193 and NAG-8-651, while that of J. H. K. was partially supported by NASA grant NAGW-1017.

REFERENCES

- Alloin, D., Boisson, C., and Pelat, D. 1988, *Astr. Ap.*, submitted.
 Blandford, R. D., and McKee, C. 1982, *Ap. J.*, **255**, 419.
 Clavel, J., *et al.* 1987, *Ap. J.*, **321**, 251.
 Cutri, R. C., *et al.* 1985, *Ap. J.*, **296**, 423.
 Davidson, C., and Netzer, H. 1979, *Rev. Mod. Phys.*, **51**, 715.
 Gaskell, C. M., and Peterson, B. M. 1987, *Ap. J. Suppl.*, **65**, 1.
 Gaskell, C. M., and Sparke, L. S. 1986, *Ap. J.*, **305**, 175.
 Hjellming, R. M., and Narayan, R. 1986, *Ap. J.*, **310**, 786.
 Krolik, J. H., and Kallman, T. R. 1988, *Ap. J.*, **324**, 714.
 Kwan, J., and Krolik, J. H. 1981, *Ap. J.*, **250**, 478.
 Mathews, W. G., and Capriotti, E. R. 1985, in *Astrophysics of Active Galaxies and Quasi-Stellar Objects*, ed. J. S. Miller (Mill Valley: University Science Books), p. 185.
 Mayo, W. T., Jr., Shay, M. T., and Riter, S. 1974, in *Proc. 2nd Internat. Conf. Laser Velocimetry*, Vol. 1 (W. Lafayette: Purdue University Press), p. 16.
 McHardy, I., and Czerny, B. 1987, *Nature*, **325**, 694.
 Lawrence, A., *et al.* 1987, *Nature*, **325**, 143.
 Peterson, B. M. 1988, *Pub. A.S.P.*, in press.
 Peterson, B. M., *et al.* 1983, *A.J.*, **88**, 926.
 ———, 1985, *Ap. J.*, **292**, 164.
 Oppenheim, A. V., and Schaffer, R. W. 1975, *Digital Signal Processing* (Prentice Hall: New Jersey).
 Wu, C., Boggess, A., and Gull, T. 1983, *Ap. J.*, **266**, 28.

Note added in proof.—After this paper was accepted and preprints distributed, Linda Sparke communicated the results of her application of the DCF to the Clavel *et al.* (1987) ultraviolet data for NGC 4151. Her analysis showed that, while the source is in a state of increased activity, the cross-correlation function shows a clear peak at a lag of +10 days. The authors are indebted to Dr. Sparke for her contribution to their understanding of this object.

During the Clavel *et al.* observations (JD 2,443,568 to 2,445,658), the short time scale (< 1 yr) variability properties of NGC 4151 appear to be correlated with the phase of long time scale (≥ 1 yr) variations, in the sense that during times of consistently high continuum flux there is comparatively more power in high-frequency fluctuations. The data are naturally divided into three segments, with the source being in a "low" state in the middle, during the period from JD 2,444,449 to 2,445,026 inclusive (mean continuum flux $\bar{F} = 2.5 \times 10^{-14}$ ergs cm^{-2} s^{-1} \AA^{-1}), and in a "high" state for the observations on either side (mean continuum flux $\bar{F} = 8.9 \times 10^{-14}$ ergs cm^{-2} s^{-1} \AA^{-1}). Both the magnitude of the fluctuations and the continuum autocorrelation functions of the two states are dramatically different. During the high state, the rms fractional fluctuation is 0.74, and the autocorrelation function is virtually constant from lags of 0 to 15–20 days, beyond which it quickly goes to zero. On the other hand, the rms fractional fluctuation in the low state is only 0.47, and the autocorrelation function is different from zero only at essentially zero lag, and again between 80 and 100 days. Thus, during the high state there is substantial power in fluctuations having time scales up to ~ 20 days, while in the low state, the only time scale for which there is much power is ~ 90 days.

This contrast is mirrored in the cross-correlation functions. In the high state, there is a strong peak at a lag of about +10 days. There is also a peak at +10 days for the low state data, but it is much weaker than that in the high state, and there are also peaks of comparable statistical significance near -60 and $+80$ days.

Following the analysis in § IV, these results indicate that, at least during the high state, the C IV line-emitting region has a characteristic scale of ~ 10 lt-days. This result is completely consistent with our previously determined range of 1.2 to 20 lt-days, and with the order of magnitude photoionization estimate of ~ 15 lt-days. Evidently, using only part of the data allows a finer estimate than using all of it because the continuum variations probe the line-emitting region more effectively in the high state than in the low state. As we demonstrated in § IV, interpretable structure in the cross-correlation function appears only when there is significant power in the continuum fluctuations on time scales matching the natural time scales of the emission-line region. Thus, a region with a characteristic scale of ~ 10 lt-days is much more effectively probed by continuum fluctuations with power on time scales shorter than 20 days (i.e., the high state) than by variations with power on ~ 90 day time scales (the low state).

It should be emphasized that the arbitrary elimination of data is always extremely risky. In this case, it apparently works because the fractional variations on the natural time scale of the system during the high state are greater than those in the low state. Confirmation of these results requires further monitoring or other evidence to justify the contention that the short time scale continuum fluctuation power spectrum in this object depends upon the phase of its long time scale variations, and to determine if the size and structure of the line-emitting region remain constant with time.

R. A. EDELSON: Center for Astronomy and Space Astrophysics, Campus Box 391, University of Colorado, Boulder, CO 80309-0391

J. H. KROLIK: Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218