Annals of Biomedical Engineering, Vol. 24, pp. 269–293, 1996
Printed in the USA.

0090-6964/96 $10.50 + .00

Fractal Character of the Electrocardiogram: Distinguishing Heart-Failure and Normal Patients

ROBERT G. TURCOTT*,† and MALVIN C. TEICH*,‡

*Columbia University, Department of Electrical Engineering, New York, NY; †Stanford University, School of Medicine, Stanford, CA; ‡Boston University, Department of Electrical, Computer & Systems Engineering and Department of Biomedical Engineering, Boston, MA

Abstract—Statistical analysis of the sequence of heartbeats can provide information about the state of health of the heart. We used a variety of statistical measures to identify the form of the point process that describes the human heartbeat. These measures are based on both interevent intervals and counts, and include the interevent-interval histogram, interval-based periodogram, rescaled range analysis, the event-number histogram, Fano-factor, Allan Factor, and generalized-rate-based periodogram. All of these measures have been applied to data from both normal and heart-failure patients, and various surrogate versions thereof. The results show that almost all of the interevent-interval and the long-term counting statistics differ in statistically significant ways for the two classes of data. Several measures reveal $1/f$-type fluctuations (long-duration power-law correlation). The analysis that we have conducted suggests the use of a conveniently calculated, quantitative index, based on the Allan factor, that indicates whether a particular patient does or does not suffer from heart failure. The Allan factor turns out to be particularly useful because it is easily calculated and is jointly responsive to both short-term and long-term characteristics of the heartbeat time series. A phase-space reconstruction based on the generalized heart rate is used to obtain a putative attractor's capacity dimension. Though the dependence of this dimension on the embedding dimension is consistent with that of a low-dimensional dynamical system (with a larger apparent dimension for normal subjects), surrogate-data analysis shows that identical behavior emerges from temporal correlation in a stochastic process. We present simulated results for a purely stochastic integrate-and-fire model, comprising a fractal-Gaussian-noise kernel, in which the sequence of heartbeats is determined by level crossings of fractional Brownian motion. This model characterizes the statistical behavior of the human electrocardiogram remarkably well, properly accounting for the behavior of all of the measures studied, over all time scales.

Acknowledgment—We are grateful to Ary Goldberger and David Rigney (Beth Israel Hospital, Boston) for providing us with the data used in this paper, which is part of the BIH Congestive Heart-Failure Database. Steven Schiff shared a great deal of insight with us in connection with the use of surrogate data sets. We thank James B. Bassingthwaighte, Ary Goldberger, Conor Heneghan, Daniel Kaplan, Steven B. Lowen, H. Eugene Stanley, and C.-K. Peng for valuable discussions. This work was supported by the Office of Naval Research under Grant N00014-92-J-1251.

Address correspondence to M. C. Teich, Boston University, Department of Biomedical Engineering, 44 Cummington Street, Boston, MA 02215.

(Received 20Jun94, Revised 28Sep94, Revised 31Aug95, Accepted 7Dec95)

Keywords—Electrocardiogram, Heart failure, Fractal, Point process, $1/f$ Noise, Chaos, Time series

INTRODUCTION

Scale-invariant fluctuations and power-law correlation have been demonstrated in the sequence of heartbeats (6,15–17,37,41,42,49,63). Different heartbeat power spectral densities are associated with various pathological conditions, such as diabetic autonomic neuropathy (12,33), uncomplicated essential hypertension (21), sudden infant death syndrome (29), potential for sudden cardiac death (15–17), severe heart disease (41,42,63), and myocardial infarction (6). The studies that have been conducted focus, for the most part, on the relative power in various spectral peaks as indicators of health.

In this paper, we examine normal and heart-failure electrocardiograms (ECGs) in the context of more general analyses designed to quantify multiscale fluctuations in point processes. We also examine a measure (the capacity dimension) used to quantify the behavior of a nonlinear dynamical system, but we do so using a count-based rather than an interval-based representation. Our approach enables us to distinguish between normal and heart-failure data and therefore may provide a more general way of determining the degree of health of the heart. We also present a stochastic model that characterizes the statistical properties of the electrocardiogram remarkably well. A preliminary report of portions of this work was presented at the SPIE meeting on Chaos in biology and medicine in 1993 (63).

As an indicator of the potential roles of scale-invariant fluctuations and fractals in biology and medicine, we direct the reader's attention to the long-duration, power-law correlation that has been found to be present in the sequence of action potentials generated by primary auditory (VIII-nerve) neurons in the cat (35,55–57), chinchilla (43), and chicken (44); in retinal ganglion cells and in the lateral geniculate nucleus of the cat (58); in striate-cortex neurons in the cat (59); in mesencephalic reticular-formation neurons in the cat (19); and in a visual interneuron in the locust (64). At all these loci, power-law

correlation extends to long times, with the upper limit of observed correlation time often imposed by the duration of the recording.

Fractal fluctuations, which give rise to $1/f$-type power spectra, are readily apparent in estimates of the heart rate (49,63). In Fig. 1, a rate estimate is formed by counting the number of contractions (QRS complexes) in successive counting windows of duration T. It might be expected that the fluctuations exhibited by the estimate would decrease rather rapidly as T is increased, and this is indeed the case for nonfractal signals (e.g., the heart data with randomly reordered interevent intervals shown in Fig. 1a). The fluctuations of fractal signals, on the other hand, either converge slowly or persist at the same magnitude as the counting time is increased, as is the case for the normal heart data presented in Fig. 1b. Despite the different convergence properties, the point processes on which Figs. 1 a and b are based have identical interevent-interval histograms. Similar scale invariance is observed for the sequence of interevent intervals (16,17,30,41). The statistical measures we present quantify this phenomenon.

Previous studies have, for the most part, focused on analyses based on the sequence of intervals between heartbeats, in which the abscissa of the time series is the interval number (30,41,42,53). Our analysis joins this approach with information obtained from count or rate esti-

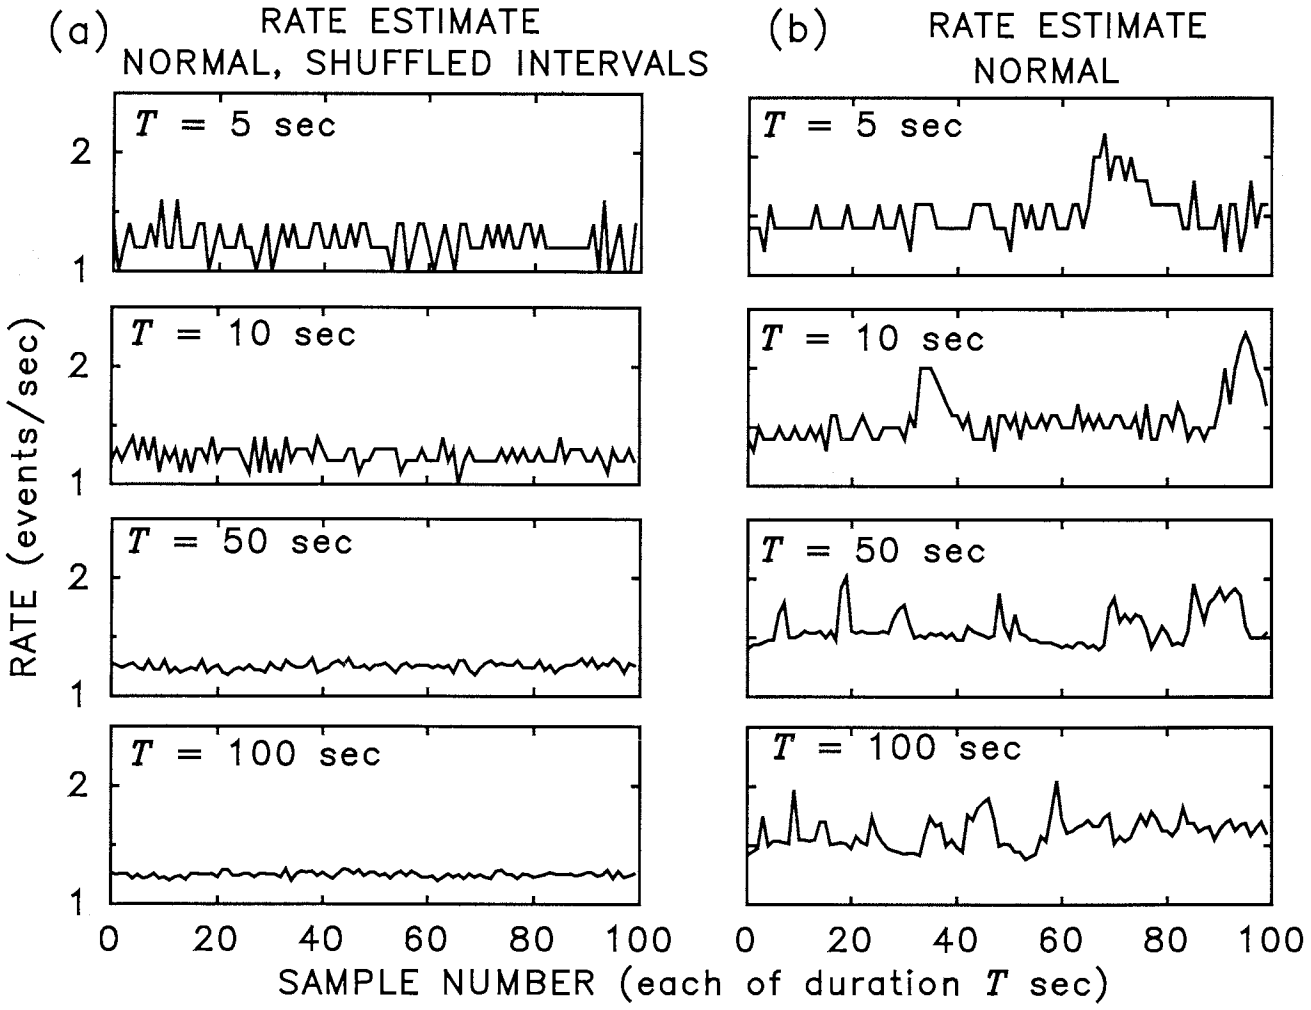

FIGURE 1. Rate estimates formed by counting the number of beats in successive counting times. (a) The data with shuffled intervals is nonfractal. For nonfractal signals, the fluctuations decrease rapidly with increasing counting time T. The magnitude of fluctuations can be quantified by the standard deviation of the rate. (b) The magnitude of fluctuations in fractal signals converge with a slower (power-law) dependence on counting time, if they converge at all. For each of the estimates in (b), fluctuations of the order of 0.5 beat/sec persist, even though the counting time is increased by a factor of 20 from the top panel to the bottom. The convergence properties are characterized by several measures presented in later figures. The data in (b) are from a long-duration recording of a normal heart electrocardiogram (data set 16265). The data in (a) are from the same recording, but with the intervals randomly reordered (shuffled), which maintains the same relative frequency of the intervals (the IIH), but destroys long-term correlations arising from other sources. These same two data sets (16265) are denoted in subsequent figures as 'normal', and 'normal, shuffled intervals', respectively.

mates (54,56) of the heartbeat sequence, which have the merit of preserving real time along the abscissa. An advantage of the rate approach over the interval approach is that it therefore allows a direct interpretation of the observed correlation (8).

Indications that the heart exhibits chaos over short time scales (1,3,46), and under special experimental conditions (13,14,20,28,66,68), have been set forth previously. To examine the possibility that heartbeat-rate variability over long time scales in normal and heart-failure patients has a nonlinear-dynamical origin, we have calculated the dimension of the associated attractor using a counting paradigm. Though the results are consistent with what would be expected from a low-dimensional chaotic attractor, surrogate-data analysis suggests that simple temporal correlation in a stochastic system leads to the same results.

In the *Methods* section we present the theoretical background for the statistical measures we utilize. The subsequent two sections present the results and a discussion of how the various statistical measures differ for normal patients and those with heart failure. A stochastic model which characterizes the behavior of the ECG remarkably well is presented subsequently, and the principal conclusions are set forth in the summary.

METHODS

Point Processes

The statistical behavior of the sequence of heartbeats can be studied by replacing the complex waveform of an individual heartbeat recorded in the electrocardiogram (QRS-complex) with the time of occurrence of the contraction (R-phase), a single number (8). In mathematical terms, the heartbeat sequence is then modeled as an unmarked point process. This simplification greatly reduces the computational complexity of the problem and permits us to use the substantial methodology that treats stochastic point processes.

The occurrence of a contraction at time t_i is therefore simply represented by an impulse $\delta(t - t_i)$ at that time, so that the sequence of heartbeats is represented by

$$s(t) = \sum_i \delta(t - t_i) \qquad (1)$$

As illustrated in Fig. 2a, a realization of a point process is specified by the set of occurrence times $\{t_i\}$ of the events. A single realization of the data is often all that is available to the observer, so that the identification of the point process, and the elucidation of the mechanisms that underlie it, must be gleaned from this one realization.

One way in which the information in an experimental point process can be made more digestible is to reduce the data into a statistic that emphasizes a particular aspect of the data at the expense of other features. These statistics fall into two broad classes (7,54) which have their origins, respectively, in the sequence of interevent intervals $\{\tau_i\}$ represented in Fig. 2b, and in the sequence of counts $\{N_i\}$ represented in Fig. 2c. Measures based on interval and counting statistics are discussed below, in turn.

We first consider the homogeneous Poisson point process (HPP), which is the simplest of all stochastic point processes (22). It is memoryless: the occurrence of an event at any time t_0 is independent of the presence (or absence) of events at other times $t \neq t_0$. Because of this property, both the intervals $\{\tau_i\}$ and counts $\{N_i\}$ form sequences of independent, identically distributed (iid) random variables. The HPP interval process is therefore completely characterized by the interevent-interval distribution (which is exponential) or by the event-number distribution (which is Poisson). The HPP serves as a benchmark against which other point processes are measured, and therefore plays the role that the Gaussian process plays in the realm of continuous-time stochastic processes.

A related point process is the nonparalyzable dead-time-modified Poisson point process (DTMP) (47), a close cousin of the HPP that differs only by the imposition of a dead-time (refractory) interval after the occurrence of each event, during which other events are prohibited from occurring.

Interevent-Interval Measures of a Point Process

We employ three statistical measures to characterize the discrete-time stochastic process $\{\tau_i\}$, which is a sequence of positive real-valued random variables, as illustrated in Fig. 2b. These are the interevent-interval histogram (IIH), rescaled range analysis (R/S), and the interval-based periodogram (IBP).

Interevent-Interval Histogram (IIH). The interevent-interval histogram displays the relative frequency of occurrence $p_\tau(\tau)$ of an interval of size τ; it is an estimate of the probability density function of interevent-interval magnitude [see Fig. 2b]. It is, perhaps, the most commonly used of all statistical measures of point processes in the life sciences. The IIH provides information about the underlying process over time scales that are of the order of the mean interevent interval. Its construction involves the loss of interval ordering, and therefore dependencies among intervals; a reordering of the sequence $\{\tau_i\}$ does not alter the IIH since the order plays no role in the relative frequency of occurrence.

Some point processes exhibit no dependencies among their interevent intervals at the outset, in which case the sequence of interevent intervals forms a sequence of iid random variables and the point process is completely specified by its IIH. Such a process is called a renewal process, a definition motivated by the replacement of failed parts

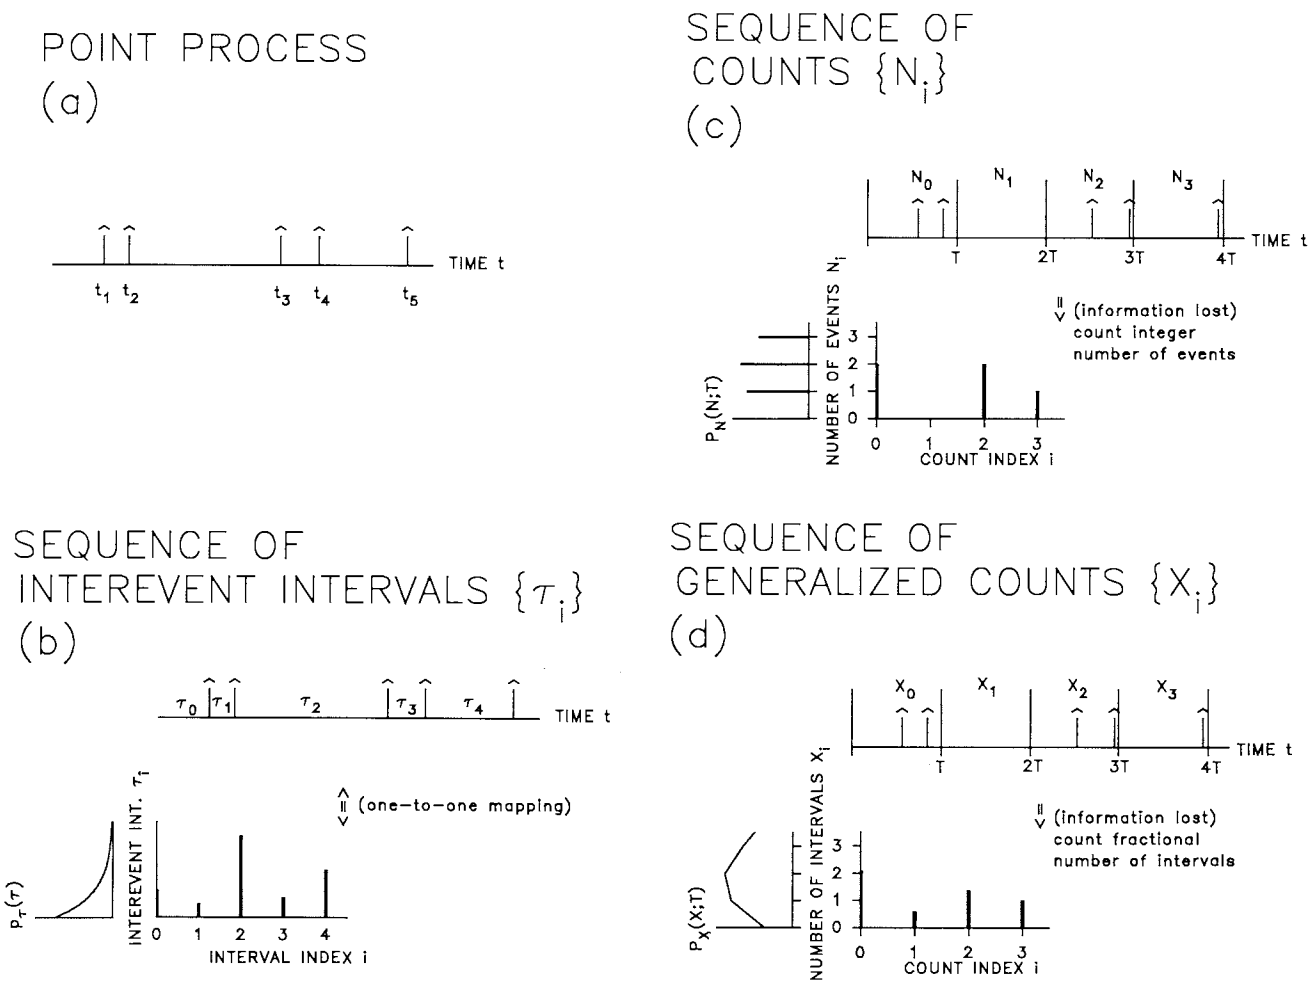

FIGURE 2. (a) The electrocardiogram is reduced to a set of event occurrence times $\{t_i\}$ that form a point process. (b) A sequence of interevent intervals $\{\tau_i\}$ is formed from the time between successive events, resulting in a discrete-time, positive, real-valued stochastic process. All information contained in the original point process is preserved in this representation, but the discrete-time axis of the sequence of interevent intervals is distorted relative to the real time axis of the point process. (c) The sequence of counts $\{N_i\}$, a discrete time, non-negative, integer-valued stochastic process, is formed from the point process by recording the number of events in successive counting windows of duration T. Information is lost in mapping the point process to the sequence $\{N_i\}$, but the amount lost can be made arbitrarily small by reducing T. (d) The sequence of generalized counts $\{X_i\}$, a discrete-time, positive, real-valued stochastic process, is formed from the point process by recording the fractional number of intervals in successive counting windows of duration T. Information is lost in mapping the point process to the sequence $\{X_i\}$, but in this case the amount lost cannot be made arbitrarily small by reducing T.

(such as light bulbs), each replacement of which forms a renewal of the point process.

The HPP and the DTMP are both renewal processes. The interevent-interval probability density function for the HPP assumes the exponential form

$$p_\tau(\tau) = \lambda e^{-\lambda\tau}, \tag{2}$$

where λ is the mean number of events per unit time. The interevent-interval mean and variance are readily calculated to be $\langle\tau\rangle = \int_0^\infty \tau p_\tau(\tau)d\tau = 1/\lambda$ and $\text{var}(\tau) = (\langle\tau^2\rangle - \langle\tau\rangle^2) = 1/\lambda^2$, respectively. The interevent-interval probability density function for the DTMP exhibits the same exponential form as for the HPP, but is truncated at short interevent intervals as a result of the dead time:

$$p_\tau(\tau) = \begin{cases} 0 & \tau < \tau_d \\ \lambda e^{-\lambda(\tau-\tau_d)} & \tau \geqslant \tau_d \end{cases} \tag{3}$$

Here τ_d is the dead time, and λ is the rate of the process before dead time is imposed.

If a process is nonrenewal, so that dependencies exist among its interevent intervals, then the IIH does not completely characterize the process. In that case, measures that reveal the nature of the dependencies provide information that is complementary to that contained in the IIH. The heartbeat time series is such a nonrenewal process.

Rescaled Range Analysis (R/S). Rescaled range analysis (R/S) provides information about correlations among blocks of interevent intervals. This widely used measure is generally assumed to be well suited to processes that exhibit long-term correlation or have a large variance (11,23,24), but it appears not to be robust since it exhibits large systematic errors and highly variable estimates of the Hurst coefficient for fractal-Gaussian-noise processes (4). In any case, it does not appear to have been applied to heart data previously.

For a block of k interevent intervals, the difference between each interval and the mean interevent interval is obtained and successively added to a cumulative sum. The normalized range $R(k)$ is the difference between the maximum and minimum values that the cumulative sum attains, divided by the standard deviation of the interval size. $R(k)$ is plotted against k. Information about the nature and degree of correlation in the process is obtained by fitting $R(k)$ to the function k^H. For $H > 0.5$, positive correlation exists among the intervals, while $H < 0.5$ indicates the presence of negative correlation. For negatively correlated intervals, an interval that is larger than the mean tends, on average, to be preceded or followed by one smaller than the mean.

Interval-Based Periodogram (IBP). Fourier transform methods provide a way of quantifying the correlation properties of a stochastic process. The averaged interval-based periodogram $S_\tau(f)$, for the sequence of intervals $\{\tau_i\}$, is obtained by calculating the average of the individual periodograms

$$S(f) \equiv \frac{1}{M}|\tilde{\tau}(f)|^2, \tag{4}$$

where $\tilde{\tau}(f)$ is the discrete Fourier transform (DFT) of the sequence of intervals and M is the length of the DFT (40,45). Information can be inferred about the correlation among intervals from $S_\tau(f)$ since it is simply the Fourier transform of the interval-process autocorrelation function. The IBP bears a close relationship to the rate-based periodogram (discussed subsequently) for point processes whose interevent-interval coefficient of variation is relatively small (8), in which case the information contained in both is essentially equivalent.

Event-Number Measures of a Point Process

It is advantageous to study some characteristics of a point process in terms of the sequence of event numbers (counts) $\{N_i\}$ rather than via the sequence of intervals $\{\tau_i\}$ (7,54,63).

Figure 2c illustrates how the sequence of counts is obtained. The time axis is divided into equally spaced, contiguous time windows, each of duration T sec, and the (integer) number of events in the ith window is counted and denoted N_i. This sequence $\{N_i\}$ forms a discrete-time random counting process of nonnegative integers. In general, information is lost in the transformation from the point process to the counting process since the specific occurrence times of the events within each window are ignored. Nevertheless, for regular point processes, the information loss can be rendered arbitrarily small by using counting windows that are sufficiently short. Closely related to the sequence of counts is the sequence of rates (events/sec), which is obtained by dividing each count N_i by the counting time T (this measure was used in Fig. 1).

We employ a number of statistical measures to characterize the counting process $\{N_i\}$: the event-number histogram (ENH); the Fano factor (FF); the normalized coincidence rate (NCR); and the Allan factor (AF). We also consider the generalized count and rate; the latter conveniently serves as a basis for the rate-based periodogram (RBP) and for a rate-based phase-space reconstruction.

Event-Number Histogram (ENH). Just as the IIH provides an estimate of the probability density function of interevent-interval magnitude, the event-number histogram $P_N(N;T)$ provides an estimate of the probability mass function of the number of events N. Construction of the ENH, like the IIH, involves loss of information, in this case the ordering of the counts. However, whereas the time scale of information contained in the IIH is the mean interevent interval, which is intrinsic to the process under consideration, the ENH reflects behavior occurring on the time scale of the counting window T. Since this time is externally specified by the observer, the character of the process at arbitrary time scales can be examined by use of this measure (54).

For the HPP, the probability mass function is the Poisson distribution

$$P_N(N;T) = \frac{(\lambda T)^N e^{-\lambda T}}{N!} \tag{5}$$

The event-number distribution for the DTMP (47) is considerably more complex than Eq. 5 because the possible overlap of a dead-time interval across the boundary of adjacent counting windows results in correlation between the numbers of events in these intervals. As a result, the counts $\{N_i\}$ are no longer independent, though they become approximately so for counting times much greater than the dead-time interval.

Additional information pertaining to a point process can also be revealed by particular characteristics of the ENH. For example, a sawtooth-like form for the ENH, revealing higher probabilities for even than odd numbers in a counting time T, implies that events tend to occur in pairs separated by less than T sec.

The moments of the histogram, such as the event-number variance and mean, and their ratio, provide suc-

cinct and useful information about the process, as sketched below.

Fano Factor (FF). The Fano factor, $F(T)$, is defined as the event-number variance divided by the mean:

$$F(T) = \frac{\text{var}[N_i(T)]}{\langle N_i(T)\rangle}. \tag{6}$$

This quantity provides an abbreviated way of describing correlation in a sequence of events. It indicates the degree of event clustering or anticlustering in a point process relative to the benchmark HPP, for which $F(T) = 1$ for all T. This latter result is readily derived by calculating the count mean $\langle N_i(T)\rangle = \sum_{N=0}^{\infty} N P_N(N;T)$ and count variance $\text{var}[N_i(T)] = (\langle N_i^2(T)\rangle - \langle N_i(T)\rangle^2)$ which, with the help of Eq. 5 for the Poisson distribution, leads to $\text{var}[N_i(T)] = \langle N_i(T)\rangle = \lambda T$. Thus $F(T) = 1$ for all counting times T for the HPP.

In fact, the FF must approach unity at sufficiently small values of the counting time T for *any* regular point process because only zero or one event can be registered in an arbitrarily short counting window for such processes. The sequence of counts then becomes a sequence of Bernoulli random variables, with a value of 0 or 1, and with a mean event number equal to the probability p of observing an event in the counting window. The variance of the Bernoulli distribution is simply $p(1 - p)$ so that

$$\lim_{T\to 0} F(T) = \lim_{T\to 0} \frac{p(1-p)}{p} = 1, \tag{7}$$

since $p \to 0$ as $T \to 0$.

For the DTMP, the dead time imposes anticlustering (more regularity) on the point process for all but the shortest counting times. This anticlustering reduces the variance relative to the mean which suppresses $F(T)$ so that it lies below unity. The asymptotic result for an HPP subject to nonparalyzable fixed dead time, valid in the limit of large T, is (56)

$$F(T) \approx (1 - \mu\tau_d)^2, \tag{8}$$

where μ represents the post-dead-time event rate. When $\lambda\tau_d$ is appreciable, the imposition of dead time produces a nearly periodic series of events because one is always available immediately after the termination of the dead-time interval. The process is therefore anticlustered relative to the HPP, and as a consequence has low count variance and Fano factor.

In general, a Fano factor less than unity indicates that a point process is more regular than the HPP at the particular time scale T, whereas an excess over unity indicates increased clustering at the given time scale (34,54–58). This measure is sometimes called the index of dispersion; it appears to have been first used by Fano in 1947 (10) for characterizing the statistical fluctuations of the number of ions generated by individual fast charged particles. Equivalent measures are provided by the variance–time curve (7) and relative dispersional analysis (51).

Normalized Coincidence Rate (NCR). The normalized coincidence rate $g^{(2)}(\tau)$ is defined as (7,55,56)

$$g^{(2)}(\tau) \equiv \frac{\Pr\{\mathscr{E}(t,t+dt) \text{ and } \mathscr{E}(t+\tau,t+\tau+dt)\}}{\Pr\{\mathscr{E}(t,t+dt)\}\ \Pr\{\mathscr{E}(t+\tau,t+\tau+dt)\}}, \tag{9}$$

where $\mathscr{E}(x,y)$ denotes the occurrence of an event in the interval (x,y), and τ is a delay time. The NCR is sometimes called the autocorrelogram, and plays the role of the correlation function for point processes. For an HPP, $g^{(2)}(\tau) = 1$ for all τ.

The Fano factor $F(T)$ and the normalized coincidence rate $g^{(2)}(\tau)$ have a unique relation for an arbitrary stationary point process (7,34,55,56):

$$\begin{aligned} F(T) &= 1 + 2\mu \int_0^T \left(1 - \frac{\tau}{T}\right)[g^{(2)}(\tau) - 1]d\tau, \\ g^{(2)}(\tau) &= 1 + \frac{1}{2\mu}\frac{\partial^2}{\partial T^2}[TF(T)]\big|_{T=\tau}, \end{aligned} \tag{10}$$

where μ is the mean rate of the point process.

In particular, a power-law dependence of the form $F(T) \sim T^{\alpha}$ $(0 < \alpha < 1)$ in the long-counting-time limit implies that the underlying point process has a power-law normalized coincidence rate $g^{(2)}(\tau) \sim |\tau|^{\alpha-1}$ and a rate-based power spectral density that behaves as $S_R(f) \sim f^{-\alpha}$ (34–36,55,56,65). Processes that behave in this way are called fractal stochastic point processes (FSPPs) (34). The parameter α is identified as the fractal exponent (or scaling exponent) of the point process. It is ambiguously related to the Hurst exponent (4,24) since some authors have used the quantity H to index fractal Gaussian noise (FGN) whereas others have used the same value of H to index the integral of FGN [which is fractional Brownian motion (FBM)]. The relationship between the quantities is $\alpha = 2H - 1$ for FGN and $\alpha = 2H + 1$ for FBM (51). We avoid this ambiguity by using α rather than H.

Even though the correlation between a single pair of events is typically rather small in FSPPs, the value of $F(T)$ at a particular counting time T can become quite large because the FF integrates the many correlations from different pairs of events within the counting window T. As a result, even weak correlation in $g^{(2)}(\tau)$ can lead to dramatic departures of the Fano factor from unity (55–59). Consequently, the FF exhibits particular sensitivity to correlation in a point process; it is also superior to the coincidence rate as an estimator (34).

Allan Factor (AF). Though the FF can detect the presence of self-similarity even when it cannot be discerned in a

visual representation of a sequence of events, mathematical constraints prevent it from increasing with the counting time faster than $\sim T^1$. It therefore provides a suitable measure only for fractal exponents in the range $0 < \alpha < 1$ (34,35,58).

The estimation of a fractal exponent that assumes a value greater than unity requires the use of a measure whose increase is not constrained in this way. In this section we define a measure called the Allan factor (AF), which is the ratio of the event-number Allan variance to twice the mean:

$$A(T) = \frac{\langle [N_{i+1}(T) - N_i(T)]^2 \rangle}{2\langle N_i(T) \rangle}. \tag{11}$$

The Allan variance is defined in terms of the variability of successive counts; it was first introduced in connection with the stability of atomic-based clocks (2).

Like the FF, the AF is a useful measure of the degree of event clustering (or anticlustering) in a point process relative to the benchmark HPP, for which $A(T) = 1$ for all T. In fact, the AF is simply related to the Fano factor by (35,50)

$$A(T) = 2F(T) - F(2T) \tag{12}$$

so that, in general, both quantities vary with the counting time T. For an FSPP, the AF exhibits a power-law dependence that varies with the counting time T as $A(T) \sim T^\gamma$ $(0 < \gamma < 3)$; it can rise as fast as $\sim T^3$ and can therefore be used to estimate fractal exponents over the expanded range $0 < \gamma < 3$.

For an FSPP with $0 < \alpha < 1$, the FF and the AF both vary as $\sim T^\alpha$, with the same fractal exponent $\gamma = \alpha$, over a large range of counting times T. Thus, a doubly logarithmic plot of the AF for such a process will yield an estimate γ of the fractal-exponent that is similar in value to the estimate α.

For a DTMP with nonparalyzable fixed dead time, Eq. 12 shows that the asymptotic formula for the AF valid in the limit of large T is identical to that for the FF given in Eq. 8, that is, $A(T) = F(T) \approx (1 - \mu\tau_d)^2$.

Wavelet-based measures can also be used for estimating the fractal exponent of a point process (58). The Wavelet Fano Factor (WFF) and Wavelet Allan Factor (WAF) serve in this capacity as natural generalizations of the Fano factor and the Allan factor, respectively.

Sequence of Generalized Rates. A generalization of the positive integer-valued sequence of counts $\{N_i\}$ is obtained by counting the (fractional) number of intervals that appear in each counting window, resulting in a positive real-valued discrete-time sequence $\{X_i\}$. An illustrative example is provided in Fig. 2d, where it is seen that approximately 60% of the interevent interval extending across the counting window $[T, 2T]$ falls within that window. The generalized count for this window is therefore assigned the *real* value 0.6, rather than the *integer* value 0, as it is in Fig. 2c. As the mean count increases, the $\{X_i\}$ approach the $\{N_i\}$.

Dividing each generalized count X_i by the counting time T results in a sequence of generalized rates $\{R_i\}$. This measure is useful for spectral analysis and for dynamical systems analysis, as discussed below. For the FF, the ordinary sequence of counts $\{N_i\}$ is preferred to the generalized sequence because the former offers a direct and simple comparison with the benchmark HPP.

Generalized-Rate-Based Periodogram (RBP). The rate-based periodogram is an estimate of the power spectral density of a point process, revealing how the power is concentrated across frequency. Following Berger *et al.* (5), the periodogram $S_R(f)$ is formed from the sequence of generalized rates $\{R_i\}$. The data set is divided into contiguous segments of equal length $\mathcal{T}$. Within each segment, a process is formed by dividing $\mathcal{T}$ into M equal bins, so that the binwidth is $\mathcal{T}/M$. A periodogram is then formed for each segment using

$$S(f) \equiv \frac{1}{M} |\tilde{R}(f)|^2, \tag{13}$$

where $\tilde{R}(f)$ is the discrete Fourier transform (DFT) of the sequence and M is the length of the DFT (40,45). The $S(f)$ are averaged together to form the final averaged periodogram $S_R(f)$ which estimates the PSD in the frequency range from $1/\mathcal{T}$ to $M/2\mathcal{T}$ Hz. This quantity provides direct undistorted information about the time correlation of the underlying point process because the count index increases by unity for every $\mathcal{T}/M$ seconds, in proportion to the real time of the point process. Unlike the FF and AF, the fractal exponent of the RBP has no upper bound.

Although $S_R(f)$ is related to the interval-based periodogram $S_\tau(f)$ for point processes that are relatively regular (8), the latter suffers from artifacts associated with the distortion of real time that are inherent in its construction (5,8). This arises from the fact that the interval index increases by unity at the termination times of successive interevent intervals, rather than in step with real time.

Using the generalized rate, rather than the generalized count, eliminates the noise floor of the count-based periodogram at high frequencies (40). This is because the sequence of counts behaves as a two-state Bernoulli process for short counting times, resulting in excess autocorrelation at zero and, consequently, a noise floor in the count-based periodogram. In contrast, the generalized rate is a smooth process for short counting times and therefore lacks the high-frequency noise floor.

Generalized-Rate-Based Phase-Space Reconstruction (PSR). It is important to determine the conditions under which the fractal fluctuations in the heartbeat are stochastic and when they might be accounted for in terms of a low-dimensional attractor of a deterministic nonlinear dynamical system. One way of approaching this question is in terms of a phase-space reconstruction and the use of an algorithm to estimate one or more of its fractal dimensions, D_q (3,51,60). The box-counting algorithm (32) provides a method for estimating the capacity (or box-counting) dimension of the attractor, D_0 (3,9,60).

The phase-space reconstruction approach operates on the basis that the topological properties of the attractor for a dynamical system can be determined from the time series of a single observable (18,31,67). An m-dimensional vector $\vec{Y}_m(t_n) \equiv [y(t_n), y(t_n + l), \ldots, y(t_n + (m-1)l)]$ is formed from the discrete time series of the observable $y(t_n)$. The parameter m is the embedding dimension and l is the lag (usually taken to be the location of the first zero crossing in the autocorrelation function of the time series). As the time t_n elapses, the vector $\vec{Y}_m(t_n)$ traces out a trajectory in the m-dimensional embedding space.

The observable that we investigate is the generalized rate so that $y(t_n) = R(t_n)$. It is advantageous to use the generalized version of the rate or count because the sequence of integer counts maps all vectors of the generalized count to nearby integer lattice points, resulting in a loss of fine structure in the phase-space trajectory. For $m = 1$ and T less than the dead-time interval, for example, the trajectory traced by the sequence of integer counts then collapses to two points in the embedding space: zero and one. The generalized rate, on the other hand, provides a continuum of real numbers for the reconstructed trajectory.

The box-counting algorithm estimates the capacity dimension of this trajectory. For uncorrelated noise, the capacity dimension D_0 continues to increase as the embedding dimension m is increased. For an attractor in a deterministic system, in contrast, D_0 saturates as m becomes larger than $2D_0 + 1$ (9). Such saturation, however, is not a definitive signature of deterministic dynamics; temporal correlation in a stochastic process can also be responsible for what appears to be underlying deterministic dynamics (39,52,61). Surrogate data analysis, discussed below, is useful in establishing the underlying cause.

Analysis of Surrogate Data

Information about the nature of the fractal fluctuations in the heartbeat time series may be obtained by applying the various statistical measures detailed above to surrogate data sets. These are point processes constructed from the original sequence of heartbeats in ways designed to preserve certain characteristics of the original data while eliminating (or modifying) others. Surrogate data analysis provides a way of determining whether a given result arises from a particular property of the data set.

We make use of three kinds of surrogate data sets: shuffled intervals, rescaled intervals, and randomized phases. In particular, we compare statistical measures calculated from both the original data and its various surrogates to distinguish those properties of the data that arise from correlation among intervals (such as from long-term rate fluctuations) from those properties inherent in the form of the IIH.

Shuffled Intervals. One of our sets of surrogate data is formed by shuffling (randomly reordering) the sequence of interevent intervals $\{\tau_i\}$ of the original data set. Such random reordering destroys dependencies among the intervals, and therefore the long-term correlation properties of the data, while exactly preserving the interevent-interval histogram. It provides a method for generating a renewal point process with an identical IIH as that of the original data set.

Rescaled Intervals. The second set of surrogate data is formed by rescaling, but not changing the order of, the sequence of interevent intervals $\{\tau_i\}$ of the original data set. This modifies the interevent-interval histogram while essentially preserving the long-term correlation properties of the data. The rescaled-intervals surrogate is, in a sense, the opposite of the shuffled-intervals surrogate (which maintains the IIH but destroys the long-term correlation properties). This surrogate set is used, for example, for calculating various normalized Fano and Allan factors. We consider four versions of it:

(A) Normalized counting time: $F(T)$ is calculated in the usual way, in accordance with Eq. (6), and the counting time T is then divided by the mean interevent interval $\langle\tau\rangle$ for each data set. The shape of the FF is not altered when plotted on a logarithmic abscissa. **(B) Normalized interval mean:** The mean interevent interval $\langle\tau\rangle$ of the point process is normalized to unity before the FF is constructed. The transformed point process exhibits unity mean rate. The result is the same as normalizing the counting time, as considered above. **(C) Normalized interval mean and variance:** The intervals of the point process are linearly transformed so that the mean and variance of the interevent intervals for each data set are normalized to $\langle\tau\rangle = 1.0$ sec and $\text{var}(\tau) = 0.04$ sec^2, respectively (the value of the variance was arbitrarily selected to make the standard deviation 20% of the mean). The shape of the FF is altered. **(D) Normalized interval histogram:** Normalizing the moments of the interval histogram can be generalized to normalizing the entire histogram to a specified form. First, a sequence of exponentially distributed random variables with unity mean is generated and they are ranked from smallest to largest. The relative position from the mean of each interval in the original data set is

noted, and the interval is then replaced by the interval which has the same relative rank in the ordered sequence of exponential random variables. Imposing this ordering on the heartbeat intervals essentially allows the long-term interevent-interval correlations to persist in the transformed point process while imposing on it an exponential IIH. This approach is similar in spirit to the normalizing transform used by Theiler *et al.* (61), but their express purpose was to transform the *amplitude* of a time series to Gaussian form, whereas ours is to transform the interevent-*interval* histogram to exponential form.

Randomized Phases. The third class of surrogate data we consider is obtained by Fourier transforming the generalized rate, and then randomizing the phases while leaving the periodogram magnitude intact. The modified function is then inverse transformed to return to a time-domain representation of the generalized rate. This rate is, in turn, converted into a point process through an integrate-and-fire mechanism as discussed subsequently. This technique preserves the second-order temporal correlation properties of the point process while removing other temporal structure essential for phase-space reconstruction.

RESULTS

The interevent-interval and event-number measures described in *Methods* were applied to 12 electrocardiogram records from normal patients and 15 records from heart-failure patients. Three of the heart-failure patients also suffered from atrial fibrillation (Table 1). The recordings were made with a Holter monitor (Del Mar Avionics, Model 445, Irvine, CA), digitized at 250 Hz. The beat-to-beat (R-R) intervals were measured automatically with a computer program (38). The data were provided to us by Ary Goldberger and David Rigney of Beth Israel Hospital (BIH), Boston, as part of the BIH Congestive Heart-Failure Database. Various characteristics of the data sets are summarized in Table 1.

TABLE 1. Characteristics of the data sets. The parameters were generally obtained from the entire duration of the recording.

File Number	Number of Intervals	Duration (sec)	Mean Rate (sec^{-1})	$\langle\tau\rangle$ (sec)	var(τ) (sec^2)	FF Exponent α	AF Exponent γ	RBP Exponent β	IBP Exponent δ	CD D_0
					NORMAL PATIENTS					
16265	100460	80061.9	1.255	0.7970	0.0291	0.95	0.95	0.93	1.04	2.66
16272	93177	84395.5	1.104	0.9058	0.0202	0.87	1.14	1.21	1.24	2.64
16273	89846	74348.6	1.208	0.8275	0.0212	0.87	1.06	1.20	1.23	2.75
16420	102081	77761.0	1.313	0.7618	0.0102	0.90	1.15	1.01	1.23	2.41
16483	104338	76099.5	1.371	0.7294	0.0079	0.94	1.04	0.87	1.29	2.58
16539	108331	84669.3	1.279	0.7816	0.0225	0.87	1.07	1.07	0.97	2.81
16773	82160	78141.1	1.051	0.9511	0.0600	0.96	0.93	1.00	0.95	2.84
16786	101630	84051.4	1.209	0.8270	0.0134	0.93	1.15	1.11	1.20	3.04
16795	87061	74734.7	1.165	0.8584	0.0448	0.94	1.24	1.21	1.37	2.45
17052	87548	76399.6	1.146	0.8727	0.0251	0.84	1.18	1.19	1.20	2.98
17453	100674	74482.0	1.352	0.7398	0.0106	0.92	0.95	0.86	1.12	2.94
c4	88140	71398.7	1.234	0.8101	0.0172	0.87	1.28	1.30	1.47	2.84
average			1.224	0.8219	0.0238	0.91	1.10	1.08	1.19	2.75
std. dev.			0.097	0.0667	0.0153	0.04	0.11	0.15	0.15	0.20
					HEART-FAILURE PATIENTS					
6796	75821	71940.9	1.054	0.9488	0.0081	0.85	1.03	1.40	1.48	1.70
7257*	118376	71166.4	1.663	0.6012	0.0013	0.88	1.09	0.93	0.94	2.44
8519	80878	71941.4	1.124	0.8895	0.0040	0.78	1.50	1.48	1.63	1.81
8552*	111826	71827.0	1.557	0.6423	0.0039	0.77	1.30	1.76	1.71	1.99
8679	119094	71180.1	1.673	0.5977	0.0026	0.94	1.52	1.35	1.45	2.37
8988*	118058	71140.3	1.660	0.6026	0.0081	0.94	1.27	1.28	1.23	2.44
9049	92497	71964.8	1.285	0.7780	0.0033	0.96	1.28	1.35	1.58	2.44
9377	90644	71965.0	1.260	0.7939	0.0033	0.87	1.31	1.24	1.31	2.13
9435	114959	71196.7	1.615	0.6193	0.0008	0.94	1.55	1.42	1.32	1.60
9643	148111	72015.6	2.057	0.4862	0.0002	0.92	1.62	1.44	1.34	2.03
9674	115542	71976.3	1.605	0.6229	0.0071	0.92	1.22	1.57	1.53	2.29
9706	115064	71320.0	1.613	0.6198	0.0103	0.96	1.34	1.21	1.22	2.53
9723	115597	71999.9	1.606	0.6229	0.0003	0.91	1.46	1.75	1.80	2.19
9778	93607	71946.3	1.301	0.7686	0.0051	0.96	1.60	1.69	1.89	2.17
9837	115205	71947.0	1.601	0.6245	0.0043	0.95	1.26	1.62	1.69	2.37
average			1.512	0.6812	0.0042	0.90	1.36	1.43	1.47	2.17
std. dev.			0.259	0.2158	0.0031	0.06	0.18	0.22	0.25	0.29
p values			<0.002	<0.005	<0.001	not significant	<0.001	<0.001	<0.005	<0.001

Asterisks indicate heart-failure patients who also suffered from atrial fibrillation. The average values of all parameters (except α) show a statistically significant difference ($p < 0.005$) between normal and heart-failure patients.

The analysis presented in this section relies on two representative data sets: normal set 16265 and heart-failure set 6796 (see Table 1). These files are, by-and-large, typical of their respective classes of data. In the following section, we consider the behavior of collections of normal and heart-failure data sets.

In Fig. 3, we present interevent-interval histograms for the normal data (solid curve) and the heart-failure data (dotted curve). For these particular data files the mean interevent interval is greater for the heart-failure than for the normal (*i.e.*, the normal has a higher heart rate), though on average the opposite is true (Table 1). The narrow width of the histogram for the heart-failure patient is reflected in its small interval variance [var(τ) = 0.0081 sec^2 for data set 6796]. By comparison, the normal patient generates an excess of both shorter and longer intervals [var(τ) = 0.0291 sec^2 for data set 16265]; this is typical. Since random reordering of the intervals does not alter the relative frequency with which they occur, the IIHs of the shuffled data sets are the same as those of the originals.

A rescaled range analysis for the two data sets is presented in Fig. 4. The open circles represent the reference function $k^{1/2}$. Positive correlation among interevent intervals is present in both cases since $R(k)$ grows more rapidly than $k^{1/2}$. On the whole, this measure reveals no substantial difference between data from normal subjects and those with heart failure. Shuffling the intervals (lower two curves) results in a dependence quite close to $k^{1/2}$, as expected for renewal point processes (4,11,23,24).

The averaged interval-based periodograms $S_\tau(f)$ generated from the sequence of intervals $\{\tau_i\}$ are presented in Fig. 5a for the normal subject (solid curve) and for the heart-failure patient (dotted curve). The approximately straight-line behavior of the curves for sufficiently low frequencies on these doubly logarithmic coordinates indicates the presence of power-law correlation among the sequence of intervals. The IBP is reasonably well de-

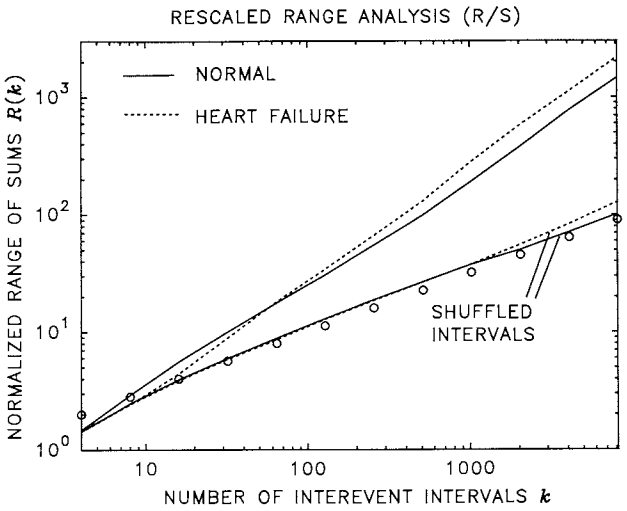

FIGURE 4. Rescaled range analysis (R/S) for the data from a normal subject (upper solid curve) and for a patient with heart failure (upper dotted curve). The results are nearly indistinguishable and reveal the presence of positive correlation. Shuffling the intervals (lower two curves) results in a dependence very close to $k^{1/2}$, as expected for sequences of independent random variables.

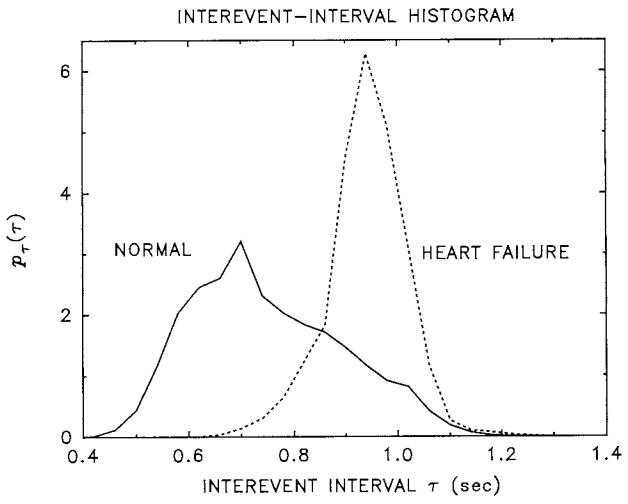

FIGURE 3. Interevent-interval histograms for the normal data (solid curve) and for the heart-failure data (dotted curve). For the particular data sets portrayed here, the mean interevent interval is greater for the heart-failure than for the normal (i.e., the normal has a higher heartrate), though the reverse is generally true (see Table 1). The interevent-interval variance of the heart-failure data is substantially smaller than that of the normal data, as is generally the case.

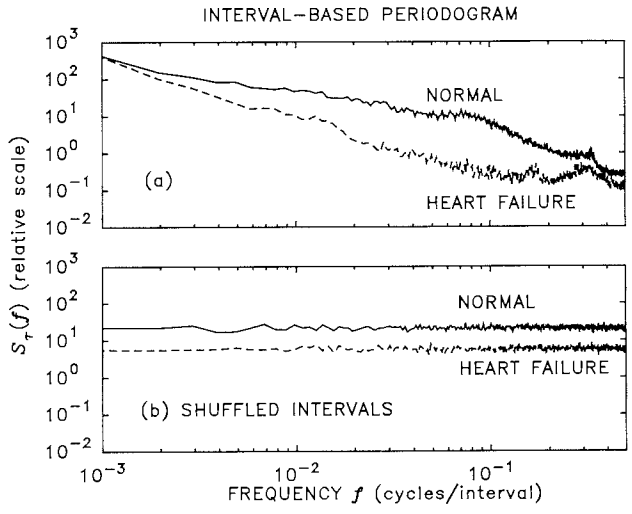

FIGURE 5. (a) Averaged interval-based periodograms (IBPs) $S_\tau(f)$ for normal subject (solid curve) and heart-failure patient (dashed curve) generated from the sequence of interevent intervals $\{\tau_i\}$. The IBP was calculated by partitioning the data into successive blocks of 1024 intervals. The periodogram of each block was calculated and the average was obtained. (b) The IBPs for the sequence of shuffled intervals are flat, corresponding to uncorrelated intervals.

scribed by the function $S_\tau(f) = af^{-\delta}$ over the range $0.0001 < f < 0.005$ cycles/interval, where a is a constant (the data is illustrated in the figure only down to 0.001 cycles/interval, however). From Table 1 it is seen that $\delta = 1.04$ for the normal subject, whereas $\delta = 1.48$ for the heart-failure patient. The integral of the IBP is somewhat greater for the normal data set; this reflects its larger interevent-interval variance (Fig. 3). However the steeper negative slope of the IBP for the heart-failure patient indicates that the relative proportion of low-frequency power, corresponding to slow time fluctuations, is greater than for the normal patient. The IBPs for the sequence of shuffled intervals (Fig. 5b) are flat since the intervals are uncorrelated after shuffling.

The event-number histogram for normal (lower solid curve) and heart-failure data (lower dotted curve), obtained with a counting time of $T = 10.0$ sec, is shown in Fig. 6. The normal ENH is substantially broader than the heart-failure histogram, indicating greater count variance arising from the clustering of events in the underlying point process. This implies greater rate variability in the normal data. This variability cannot be attributed solely to the interevent-interval variance, since the ENHs of the shuffled data sets (higher solid and dotted curves) both show reduced count variance, indicating that the count variance of the original process arises at least in part from the ordering of the intervals.

The Fano factors for the same two data sets are shown as the upper solid curves in Fig. 7a. For short counting

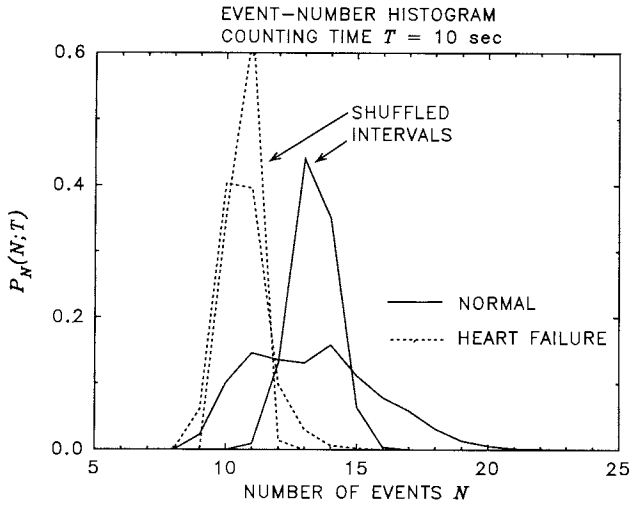

FIGURE 6. Event-number histogram (ENH) for normal (lower solid curve) and heart-failure data (lower dotted curve) using a counting time of $T = 10.0$ sec. The normal ENH is wider than the heart-failure histogram, indicating its greater count variance on this time scale. Shuffling the intervals eliminates long-duration correlation while leaving the IIH unaffected. The substantially narrower ENHs for the shuffled data sets demonstrate the presence of long-duration correlation in the original sequence of interevent intervals.

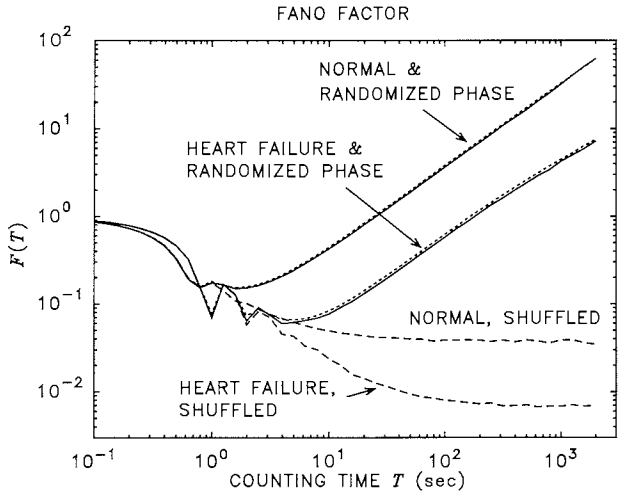

(a)

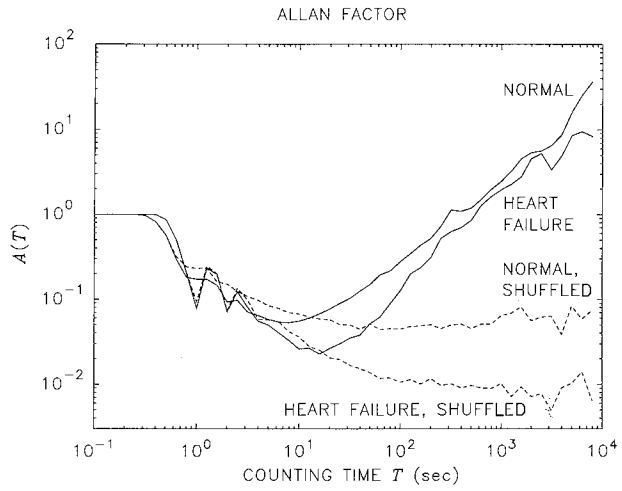

(b)

FIGURE 7. (a) Fano factor (FF) for normal and heart-failure data (solid curves). The FFs of the phase-randomized surrogate data (dotted curves) were obtained by carrying out a 65536-point DFT of the rate and then assigning independent and uniformly distributed values of the phase over the interval $[0, 2\pi)$. The sequence was then inverse transformed to provide a rate function $R(t)$ with correlation properties identical to those of the original data but with other temporal structure removed; this rate was converted into a point process by using an integrate-and-fire construct. The phase-randomized surrogate FFs are nearly identical to those of the original data, indicating that second-order correlation is sufficient for understanding the behavior of the FF. The FFs for the shuffled surrogate processes (dashed curves) always fall below unity, since long-duration correlation has been eliminated, leaving only the anticorrelation arising from the short-term regularity of the heartbeat intervals. The surrogate data presented is the mean value of $F(T)$ calculated from an ensemble of 10 realizations. Error bars are not included because the thickness of the plotted curves is greater than a standard deviation for the randomized-phase surrogates, and there is only a small (but discernible) standard deviation for the shuffled surrogates, but only for $T > 100$ sec. (b) Allan factor (AF) for normal and heart-failure data (solid curves). The AFs and FFs for the shuffled surrogate processes (dashed curves) are essentially the same.

times ($T \rightarrow 0$), the Fano Factor approaches unity, in accordance with Eq. (7). For time scales where the underlying events are highly regular, and thereby exhibit low variance, the FF dips well below unity. The oscillations of the FF in the vicinity of $T = 1$ sec in Fig. 7a (the dips are particularly pronounced for heart-failure patients), arise from the clocklike regularity of the heartbeats on this time scale. When the counting time T is such that the number of beats per counting time is almost always the same, variability is greatly diminished and the FF dips to a low value. A slight increase in T then occasionally admits an extra event, resulting in a greater variance-to-mean ratio and a rise in the curve. This behavior is considered more thoroughly in the subsequent discussion surrounding the normalized Fano factor shown in Fig. 13a.

For counting times greater than approximately 10 sec, the Fano factor increases above unity and grows as a fractional power-law function proportional to T^{α}, consistent with the behavior of an FSPP. The observed values of α are all rather similar and lie below unity, as shown in Table 1. The appearance of the power-law behavior (fractal onset time) occurs when the counting window is sufficiently large so as to allow fractal event clustering to overcome the anticlustering imposed by the pause between heartbeats.

The phase-randomized surrogate data yields an FF that is nearly identical to that of the original data. Since this surrogate retains the same periodogram magnitude as the generalized heartrate, but not the details of its trajectory in phase space, this indicates that the power-law increase in the FF can be explained in terms of the second-order temporal correlation properties of the data. This issue is discussed subsequently. The shuffled surrogates (lower dashed curves), on the other hand, lack the long-term, positive correlation in the sequence of heartbeats, resulting in FFs whose values at long counting time reflect only the magnitudes of var(τ).

The Allan factors shown in Fig. 7b have similar features which, for the most part, share the same explanations. There are, however, two important distinctions that ultimately render the AF more suitable than the FF as a clinical diagnostic measure: first, the AF curves dip to lower values as a result of delayed fractal onset time relative to the FF, as can be understood from Eq. 12 (the shuffled versions of the FF and AF are essentially indistinguishable); and second, the power-law exponent γ that describes the growth of $A(T)$ for large counting times may assume a wider range of values than α, and they almost always exceed unity (see Table 1). The differences between the normals and heart failures shown in Fig. 7, for both the AFs and the FFs, are typical of the data sets examined.

In Fig. 8a we present the rate-based periodograms $S_R(f)$

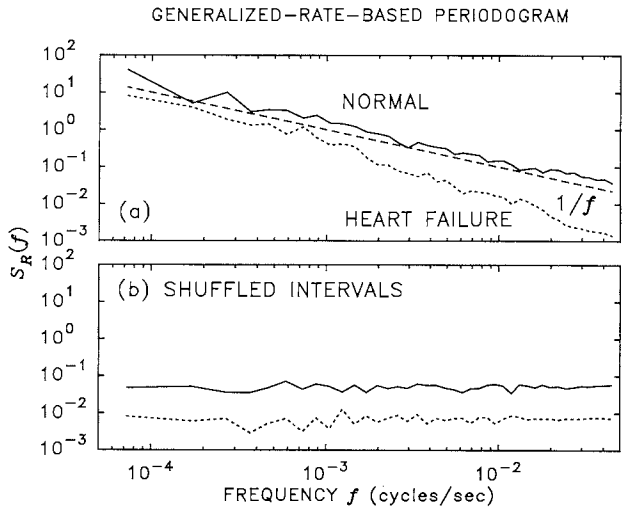

FIGURE 8. (a) Generalized-rate-based periodograms $S_R(f)$ for normal subject (solid curve) and heart-failure patient (dotted curve) generated from the sequence of generalized rates $\{R_i\}$. A periodogram with 4096 points was formed from a segment of data using a binwidth of 10 sec. The function $1/f$ is included for comparison (dashed curve). (b) After shuffling, the RBPs are flat, indicating that the rates are uncorrelated, as expected for a renewal process with $\mathcal{T}/M \gg \langle\tau\rangle$.

for the normal subject (solid curve) and heart-failure patient (dotted curve), generated from the sequence of generalized rates $\{R_i\}$ using a binwidth $\mathcal{T}/M = 10.0$ sec. The straight-line behavior of the curves on these doubly logarithmic axes indicates the presence of power-law correlation among the sequence of rates. The RBP is well described by the function $S_R(f) = cf^{-\beta}$ over the range $0.0002 < f < 0.05$ Hz (cycles/sec) for the normal data, and over the range $0.0002 < f < 0.01$ Hz for the heart-failure data (the dashed line represents $1/f$ behavior and is shown for comparison). The decreasing power-law behavior of the RBPs is consistent (8) with that of the IBPs evidenced in Fig. 5a. Indeed, it is apparent that the values $\beta = 0.93$ and $\delta = 1.04$ from Table 1 compare favorably for the normal subject, as do $\beta = 1.40$ and $\delta = 1.48$ for the heart-failure patient. The decreasing power-law form is also consistent with the power-law growth of the AFs for large counting times, as discussed subsequently.

As shown in Fig. 8b, the RBPs become flat after shuffling the intervals, reflecting the absence of correlation among the values of the generalized rate. The sequence of generalized rates is therefore a white-noise process for this binwidth; however in general, it is only flat for $\mathcal{T}/M \gg \langle\tau\rangle$. For smaller values of the binwidth, the RBP reveals correlation in the point process associated with the form of the interevent-interval histogram, even though the intervals have been shuffled. This correlation occurs because an interevent interval that extends beyond a binwidth affects the values of all succeeding binwidths across

which it falls, for all point processes other than the Poisson.

The interval-based periodograms (IBPs) of the shuffled intervals shown in Fig. 5b are also white, as expected for a renewal process. However, as discussed above, a white sequence of intervals does not imply a white sequence of rates. The two periodograms are not equivalent measures in general (8), and indeed have different units on their abscissas. Nevertheless, both periodograms show that the total power in the heart-failure patient record is lower at all frequencies than in the normal record, reflecting the lower variances of the heart-failure histograms in Figs. 3 and 6.

In Fig. 9 we display the capacity dimension D_0 obtained from a phase-space reconstruction based on generalized rate with $T = 10$ sec, for a normal subject (solid curve in upper panel) and a heart-failure patient (solid curve in lower panel). The value of D_0 for the data (solid curves) increases with the embedding dimension until m reaches about 5. The behavior shown in Fig. 9 is typical of the other data sets: curves from normal patients (Fig. 9a) saturate at a slightly higher value of D_0 (approximately 2.7 in this case) than curves from heart-failure patients (Fig. 9b) (approximately 2.4 in this case).

The saturation of D_0 with increasing m is one hallmark of a low-dimensional dynamical system, but its presence is not sufficient to conclude that such dynamics are present (39). Indeed, both sets of surrogate data reflecting purely stochastic behavior also show saturation. The greater temporal correlation in the phase-randomized surrogate data (dotted curves) yields a smaller value of D_0 than for the shuffled intervals (dashed curves). Neither surrogate provides a perfect replica of the results obtained from the original data, which lie 1–3 standard deviations away from the mean values of the surrogates. It is unclear why D_0 for the phase-randomized surrogate should lie below the value for the raw data. The heart data may contain a form of temporal structure (other than second-order) that imbues it with a larger apparent capacity dimension than its second-order correlations alone would generate. Nevertheless, the general behavior of these stochastic surrogates is similar enough to that of the original data to suggest that the saturation of D_0 with increasing m arises from temporal correlation in the heartbeat time series rather than from underlying low-dimensional deterministic dynamics. This observation is supported by the fact that the model simulations presented in the final section of this paper emulate the D_0 vs m data very well (Fig. 26).

The relationship between $F(T)$ and $g^{(2)}(\tau)$ was presented in Eq. (10) and discussed in *Methods*. Similarly, we can forge a connection between the behavior of the FF and the saturation of the capacity dimension D_0. This connection is illustrated in Fig. 10, where the reconstructed trajectories of the generalized rate (see *Methods*) in a two-dimensional ($m = 2$) embedding space are shown for a

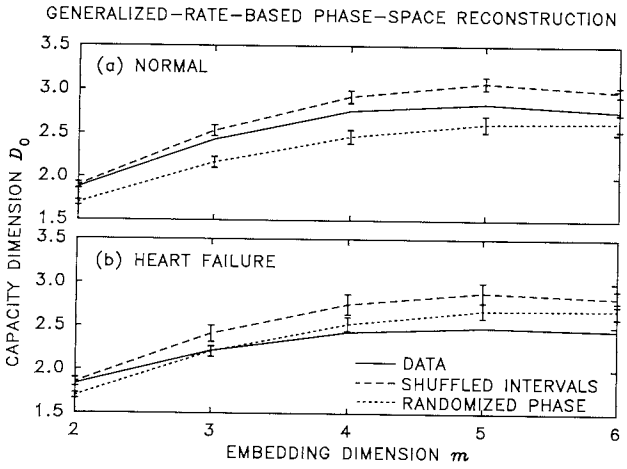

FIGURE 9. The capacity-dimension estimate D_0 of the attractor obtained from the sequence of generalized rates, with a counting time $T = 10$ sec, as a function of the embedding dimension m. The data presented are from a subset of the record and are representative of their respective classes, with the normal data [(a), solid curve] typically attaining a slightly larger value of D_0 than the heart-failure data [(b), solid curve]. For both sets of data, D_0 reaches its maximum at about $m = 5$. Also shown is the average dimension estimated from ensembles of 10 simulations of the shuffled-interval surrogate (dashed curves) and of the phase-randomized surrogate (dotted curves), along with their $\pm\sigma$ error bars. D_0 was calculated using the same segment of data that was used in Fig. 8.

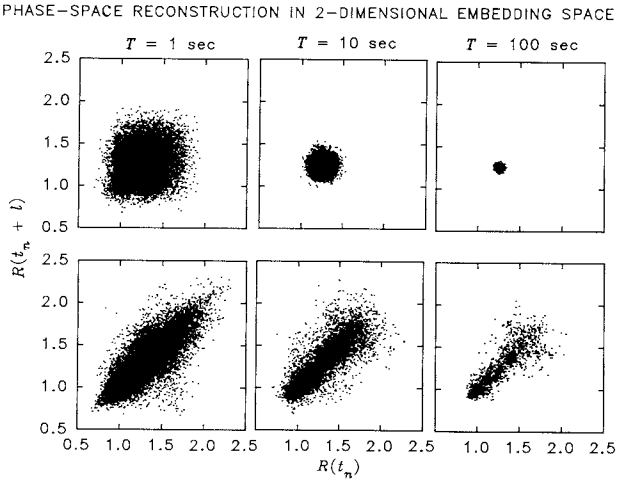

FIGURE 10. Reconstructed trajectories of the generalized rate in a two-dimensional ($m = 2$) embedding space for a range of counting windows T. Normal heart data (data set 16265) are displayed in the lower panels, whereas their shuffled surrogates are represented in the upper panels. The lag l was chosen to be 1 count sample for all counting times. For the original fractal process (normal patient, unshuffled intervals, lower panels), the spatial extent of the trajectory is relatively unaffected by increases in T, whereas for the nonfractal renewal process (normal patient, shuffled intervals, shown in the upper panels), the extent of the trajectory contracts far more dramatically with increasing T.

range of counting windows T. The lag (on the ordinate) was chosen to be $l = 1$ count sample. Normal heart data (data set 16265) are displayed in the lower panels, whereas their shuffled surrogates are represented in the upper panels.

For a particular value of T and m, D_0 provides a measure of the fractal dimension of the set of trajectory points. The Fano factor $F(T)$, on the other hand, provides a measure of the spatial extent of the trajectory in phase space, which is independent of the embedding dimension m. This can be demonstrated by considering the construction of $F(T)$ from a rate trajectory, such as that shown in Fig. 10. The coordinates of each trajectory point are multipled by T (to convert to count). The relative frequency of count value for any dimension is equivalent to the event-number histogram with counting time T, which allows the event-number variance-to-mean ratio to be calculated.

Indeed, when the FF increases in power-law fashion due to power-law decaying temporal correlation (Fig. 7a), the reconstructed phase space will necessarily contract slowly, in power-law fashion, with increasing counting time (lower panels in Fig. 10), as with the heartrate estimate displayed in Fig. 1b. For the shuffled-interval surrogate, on the other hand, the FF fails to increase with increasing counting time (Fig. 7a), the spatial extent of the phase space contracts substantially with increasing counting time (upper panels in Fig. 10), and the fluctuations in the heartrate estimate decrease substantially (Fig. 1a).

Clearly, long-duration correlation affects the dependence of D_0 on m. In particular, it prevents D_0 from increasing with m as it would for uncorrelated noise.

COMPARISON OF RESULTS FROM NORMAL SUBJECTS AND HEART-FAILURE PATIENTS

Reduced variability of the point process representing the ECG is evident in heart failure. The relative regularity of the heartbeat of these patients manifests itself in short-term measures such as the reduced variance of the interevent-interval histogram (Fig. 3), and in measures such as the reduced variance of the event-number histogram (Fig. 6). However, aside from the lower overall power, the spectral distribution is different in heart failure: there is relatively more power at low frequencies than in normal patients. Quantitative values for all of the measures that we have investigated are reported in Table 1, both for the individual patients and as averages over collections of normal and heart-failure data sets separately.

The results are striking. Nearly every parameter measured shows a statistically significant difference between the two classes of data (see p values in Table 1). In comparison with normal patients, heart-failure patients display, on average, a significantly smaller mean interevent interval ($p < 0.005$), interevent-interval variance ($p < 0.001$), and capacity dimension ($p < 0.001$); and a significantly larger mean heartrate ($p < 0.002$), and fractal exponent [as estimated from the IBP ($p < 0.005$), RBP ($p < 0.001$), and AF ($p < 0.001$)], using the Student's t-test of significance. The only exception is the fractal exponent α estimated from the Fano factor, which shows no statistically significant difference between the two classes of data. This is expected since the FF is constrained to fractal exponents less than unity.

Ideally, a clinical diagnostic should completely separate the ill from the well. Though the average values of the two classes of patients are well separated, there are nevertheless regions in each of the parameter spaces in which the membership of individual patients in one or the other group is ambiguous. This is displayed in Fig. 11, in which several interevent-interval and event-number measures displayed in Table 1 are graphically presented. One measure of clinical significance may be developed by choosing a threshold value of the parameter at hand such that all normal patients are so identified (100% specificity); the sensitivity is then defined as the proportion of heart-failure patients that lie to the opposite side of this threshold and would therefore be correctly identified as suffering from heart failure. As is evident in Fig. 11a, the variance of the

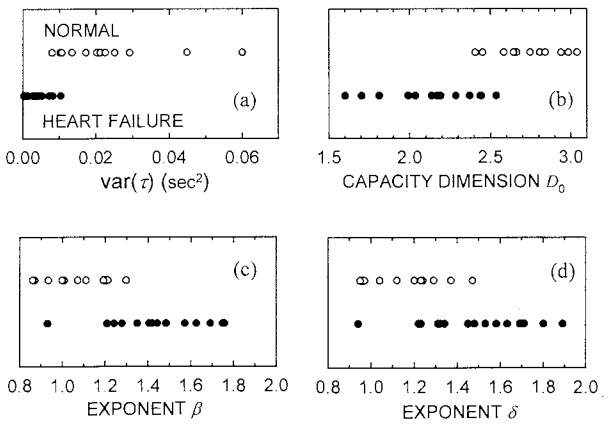

FIGURE 11. (a) The variance of the interevent intervals is generally smaller for heart-failure patients (solid circles) than that for normal subjects (open circles). (b) The capacity dimension D_0 is also typically smaller for heart-failure patients than that for normal subjects. (c) The RBPs were fit by a power-law function of the form $S_R(f) = cf^{-\beta}$ over the range $0.0002 < f < 0.05$ Hz (cycles/sec) for the normal data, and over $0.0002 < f < 0.01$ Hz for the heart-failure data. The best-fitting values of β are shown for all data sets. The heart-failure RBPs generally exhibit a larger exponent, so this parameter also separates the two classes of data on average, but not quite as well as does var(τ) and D_0. (d) The IBPs were fit by a power-law function of the form $S_\tau(f) = af^{-\delta}$ over the range $0.0001 < f < 0.005$ cycles/interval. The best-fitting values of δ are shown for all data sets. This parameter separates the data less well than var(τ), D_0, and β.

interevent intervals [var(τ)] for heart-failure patients (solid circles) is, on average, significantly smaller than that for normal subjects (open circles). However, three heart-failure sets fall among the normal sets. Thus, choosing the threshold just to the left of the leftmost open circle (at var(τ) = 0.0079), under this criterion three heart-failures to the right of this value would fail to be diagnosed with heart failure, leading to a sensitivity of (15-3)/15 = 80%. The capacity dimension D_0 of the phase-space trajectory, an event-number measure, behaves similarly, as shown in Fig. 11b. In this case, four heart-failure sets fall among the normal sets (sensitivity = 73%). The decaying power-law exponent β of the generalized-rate-based periodogram (RBP) is, on average, significantly larger for heart-failure patients, which separates the two classes of data as shown in Fig. 11c, but it is not as sensitive a measure as var(τ) since four heart-failure sets fall among the normals (sensitivity = 73%). The power-law exponent of the interval-based periodogram (IBP), like the RBP, is also larger for heart-failure patients on average, but in this case it is even more difficult to properly assign a patient to one or the other groups; there are seven heart-failure patients among the normals (sensitivity = 53%).

Though the Fano factor exponent α is not useful for this purpose because it saturates at unity, the magnitude of the Fano factor for large counting times provides a reasonable measure of separation between the normals and the heart-failures, as shown in Fig. 12. Two heart-failure sets fall among the normals (sensitivity = 87%) so, at least for the data at hand, this measure discriminates the two classes of data a bit better than the variance of the interevent intervals.

The underlying basis for the separation provided by the FF is elicited by examining this count-based measure for various surrogate data sets, as shown in Figs. 13 and 14. Normalizing the counting time (Fig. 13a), or the interval mean (Fig. 13b), does not affect the ability of the FF to separate the two classes of data, thereby demonstrating that the origin of the separation does not lie in differences in heartrate. However, normalizing the interval mean and variance (Fig. 13c), or the interval histogram (Fig. 13d), seriously impairs the ability of the FF to discriminate between normals and heart failures, demonstrating that the separation provided by the FF has its principal origin in the underlying interval statistics.

The interval statistics are exhibited not only in interval-based measures, but are also manifested in event-number (count-based) measures such as the FF, in which they can persist to arbitrarily long counting times. The behavior of the shuffled-surrogate FF readily illustrates this point. Shuffling preserves the IIH while destroying long-term correlations. The shuffled-FF, shown in Fig. 14, reveals three heart-failures among the normals. For 100% specificity, the shuffled-FF magnitude at large counting times therefore has the same sensitivity (80%) as the interevent-interval variance. Renewal theory shows that $F(\infty) = \text{var}(\tau)/\langle\tau\rangle^2$, providing the link between these two measures (7). Thus, it is principally differences in the interevent-interval statistics that are responsible for the separation of normals and heart failures by the FF, and by surrogates that leave the IIH statistics intact.

The Allan factor is more successful at separating the two classes of data, as shown in Fig. 15a, because the AF is jointly responsive to both short- *and* long-term characteristics of the heartbeat time series. The AF curves descend to lower values than the FF curves because the AF fractal onset time is delayed, as can be understood from Eq. 12. This delay permits the two classes of data to become more distinguishable at intermediate times than can be achieved using the FF. The smaller interevent-interval variances for the heart-failure patients result in their curves reaching lower levels than the normal curves, thereby enhancing distinguishability. The asymptotic values of the AF and FF for the shuffled data, determined by the contributions from the interevent-interval variance through the relation $A(\infty) = F(\infty) = \text{var}(\tau)/\langle\tau\rangle^2$ and valid for any renewal point process, are approached until the fractal characteristics begin to dominate and the AF curves turn upward. As T increases, the AF curves rise more aggressively for the heart-failure patients because they have a relatively greater proportion of low-frequency spectral components (long-time-scale fluctuations). Since the fractal exponent γ for the AF is not pinned at unity, as α is for the FF, the range of steepness with which the curves rise differs for the two classes of data. The net result is two sets of AF curves, one for the normals and another for the heart-failures, that differ in character. Both sets exhibit dips, but their shapes are sufficiently different

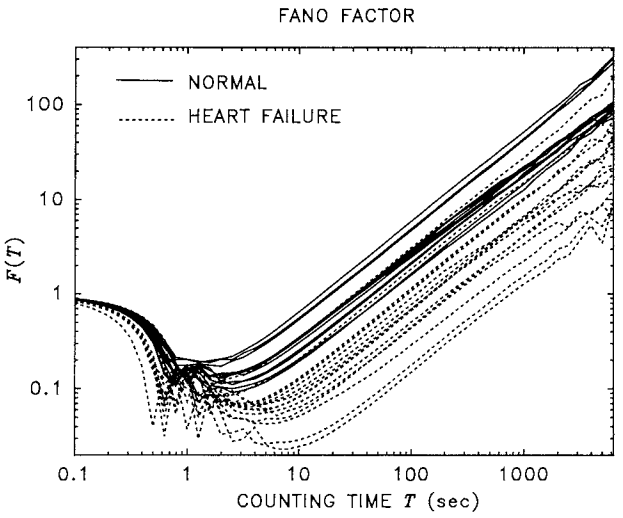

FIGURE 12. Fano factors for all data sets presented in a single plot. Two heart-failure sets fall among the normals. The FF therefore is slightly more effective than the variance of the interevent intervals in separating the two classes of data.

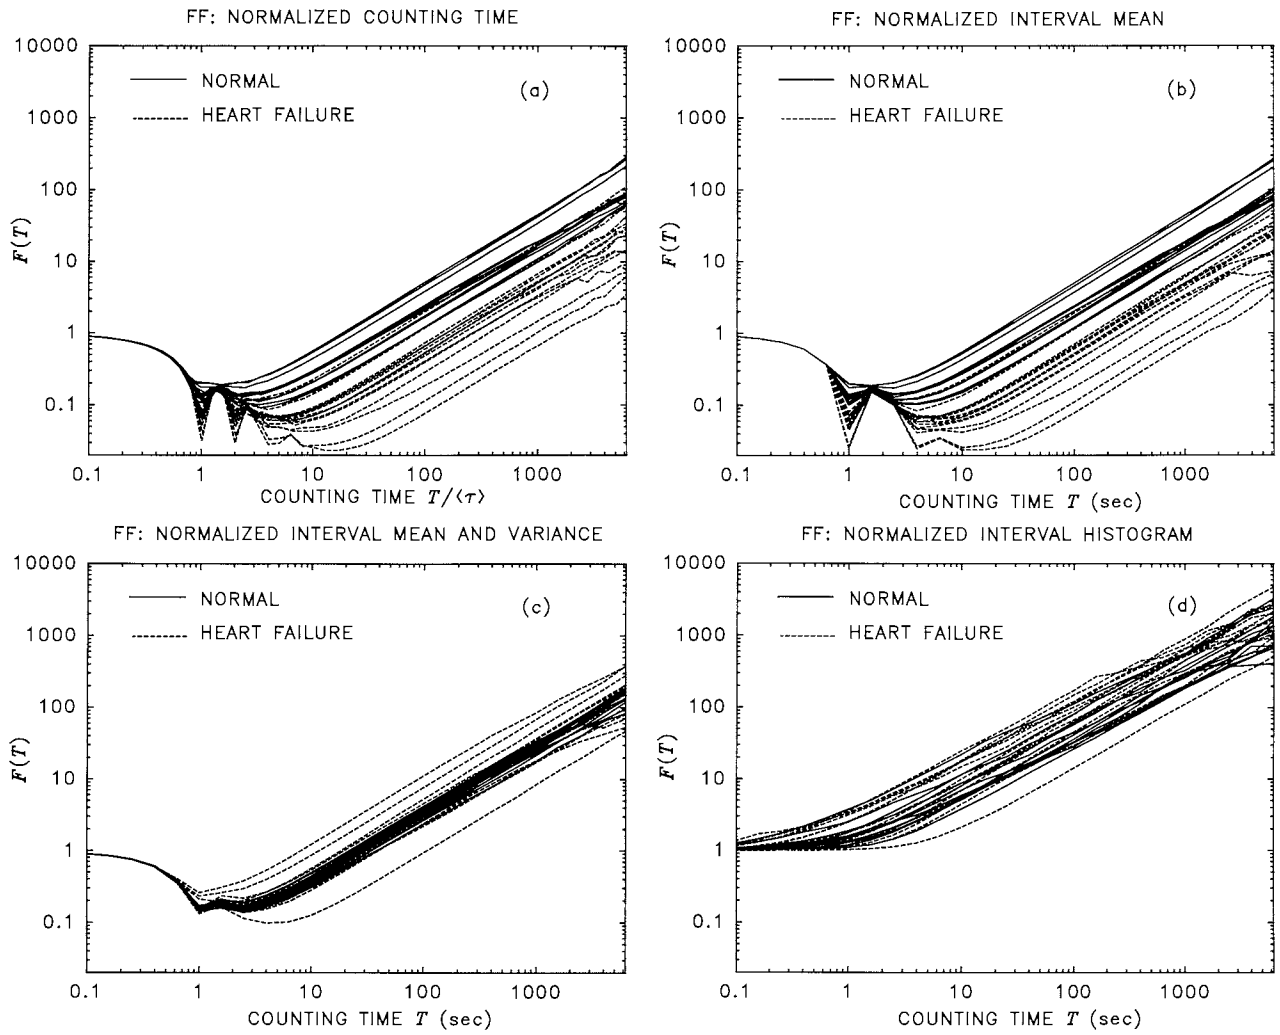

FIGURE 13. Fano factors for surrogate data sets: (a) ***Normalized counting time:*** Because the variable T on the abscissa in Fig. 12 is plotted on a logarithmic scale, dividing it by the constant $\langle\tau\rangle$ simply serves to shift the entire curve horizontally without changing its shape. It is evident from the figure that all of the curves normalized in this way first attain a local minimum at $T/\langle\tau\rangle = 1$, viz., when the counting time T is equal to the mean interevent interval $\langle\tau\rangle$. This can be intuitively understood by considering a nearly periodic point process such as the ticks of an imperfect clock. For $T = \langle\tau\rangle$, the counting time will almost always admit exactly one tick and the count variance (and thereby the Fano factor) will dip to a very small value (for a perfect clock the Fano factor would dip precisely to zero). For $T = 1.5\langle\tau\rangle$, on the other hand, two ticks will be admitted just about as frequently as a single one, depending on where the counting window falls in relation to the occurrence times of the ticks. For a yet larger counting time, $T = 2\langle\tau\rangle$, two ticks will almost always occur within the counting time, which again results in a dip in the Fano factor. Minima in the FF will therefore occur at integer multiples of $\langle\tau\rangle$. The imperfect ticking (along with the logarithmic spacing of the FF samples chosen for the plot) cause the dips to become less crisp as $T/\langle\tau\rangle$ increases. Of course, the less periodic the ticking, the less dramatic the dips. The net result is that normalization of the counting time by the mean interevent interval results in the first few minima of the FF coming into register. This particular normalization leaves intact the reasonable effectiveness of the FF in separating normals and heart-failure patients. (b) ***Normalized interval mean:*** The curves are essentially the same as those shown in (a); the only distinction is the effective difference in the data sampling strategies. This normalization therefore also leaves intact the ability of the FF to separate the two classes of data. (c) ***Normalized interval mean and variance:*** As with the counting-time and interval-mean normalizations discussed above, the mean interevent interval of all data sets are the same, so that the first minimum of the FFs are in register. However, the normalization of the interval variance causes the shapes of the individual FFs to change so that the normals and heart-failures can no longer be separated by the FFs. This illustrates that it is the interevent-interval statistics, and not the rate, that enables the FFs to distinguish them. (This procedure resulted in a few events (<10) for which $\tau < 0$; these intervals were arbitrarily set to 1 msec with no significant effect on the results.) (d) ***Normalized interval histogram:*** The IIH is exponentialized (see ***Methods***). Those properties of the FF that depend on details of the interspike-interval histogram are modified, whereas those that depend on the long-term correlation properties of the point process remain virtually unchanged. The Fano factor is always greater than unity because the anticlustering arising from the regularity in the point process that causes the FF to dip below unity in the original data has been eliminated by the transformation. As in the case of the normalized interval mean and variance considered in (c), the normal and heart failure data are not segregated. The normalization of the IIH has, as a byproduct, resulted in a normalization of the first two moments which, as demonstrated in (c), is sufficient for eliminating the segregation. In short, the two classes of data remain separated in (a) and (b), but not in (c) and (d), demonstrating that the separation provided by the FF has at its root the interevent-interval statistics.

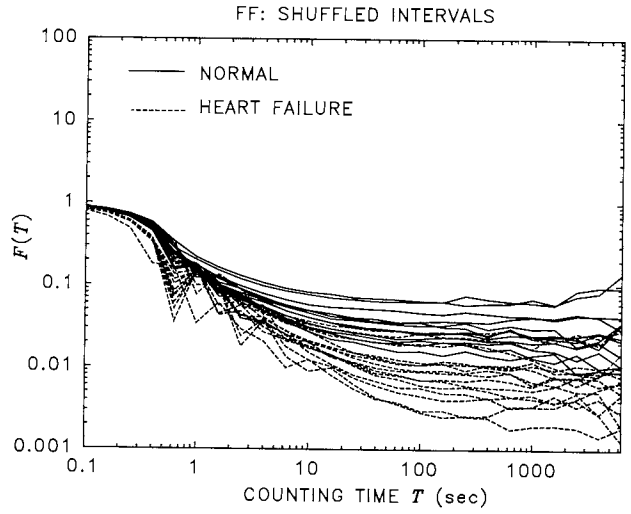

FIGURE 14. Fano factors for shuffled surrogate data sets. The ability of the shuffled FFs to separate the two classes of data is the same as that of the interevent-interval variance shown in Fig. 11a (three heart-failures among the normals). This confirms that it is principaplly, but not exclusively, differences in the underlying interevent-interval statistics that are responsible for the ability of all forms of the FF to separate normals from heart failures.

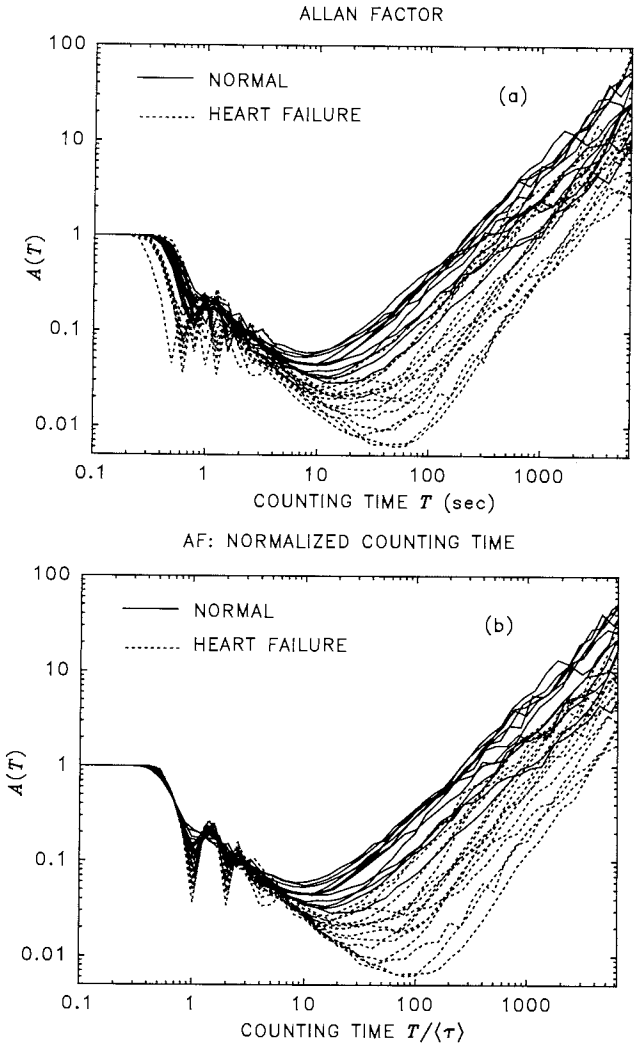

FIGURE 15. (a) The AFs of all data sets presented in a single plot. In the vicinity of $T = 10$ sec, the magnitude of the Allan Factor falls into different ranges for heart-failure patients and normal subjects, providing a measure with 100% sensitivity and specificity. (b) The AF with normalized counting time. This normalization leaves intact the effectiveness of the AF for separating normals and heart-failures.

that there is a region of counting time, in the vicinity of $T = 10$ sec, for which the magnitude of the AF falls into a different region for each class of data. Thus, a threshold can be chosen that yields both 100% sensitivity and 100% specificity.

The AF is the most successful measure we have constructed to date for separating heart-failure from normal patients. A surrogate AF with normalized counting time, shown in Fig. 15b, behaves quite similarly to the original AF, also allowing the two classes of data to be neatly divided. The shuffled-surrogate AF (not shown) is barely distinguishable from the shuffled-surrogate FF displayed in Fig. 14, as expected. The decreases in the AF and FF then result purely from short-term effects that are completely characterized by var(τ) and $\langle\tau\rangle$, and the asymptotic value $A(\infty) = F(\infty)$ is fully realized. The failure of this surrogate to completely separate the two classes of data confirms that long-term correlations contribute to the ability of the AF to discriminate as well as it does.

From the foregoing discussion, it is apparent that the FF fractal-exponent estimator α will generally be smaller than the AF fractal-exponent estimator γ for heart data. This limitation is shown explicitly in Fig. 16a, where it is clear that α is pinned below unity, whereas γ seldom lies below unity. Furthermore, the experimental values of the AF fractal-exponent estimator γ generally agree with those of the RBP fractal-exponent estimator β, as shown in Fig. 16b. Finally, in Fig. 16c, we demonstrate that β is in accord with the IBP fractal-exponent estimator δ. All of the results we observe are therefore in accord with expectations (8,58).

It is of interest to examine which aspects of the AF curves are the most suitable for discriminating the two classes of data. In Fig. 17, we graphically display the discriminability provided by several different features of the Allan-factor curves. The mean heartrate μ, shown in the upper-left panel, fails to separate the classes of data very well. The AF fractal-exponent estimator γ (upper right-hand panel), like the RBP estimator β and the IBP estimator δ (see Fig. 11), provides only partial separation, exhibiting a sensitivity of 60% at 100% specificity. The most successful indicators are the value of the AF at $T = 10$ sec (lower-left panel), and the minimum value of

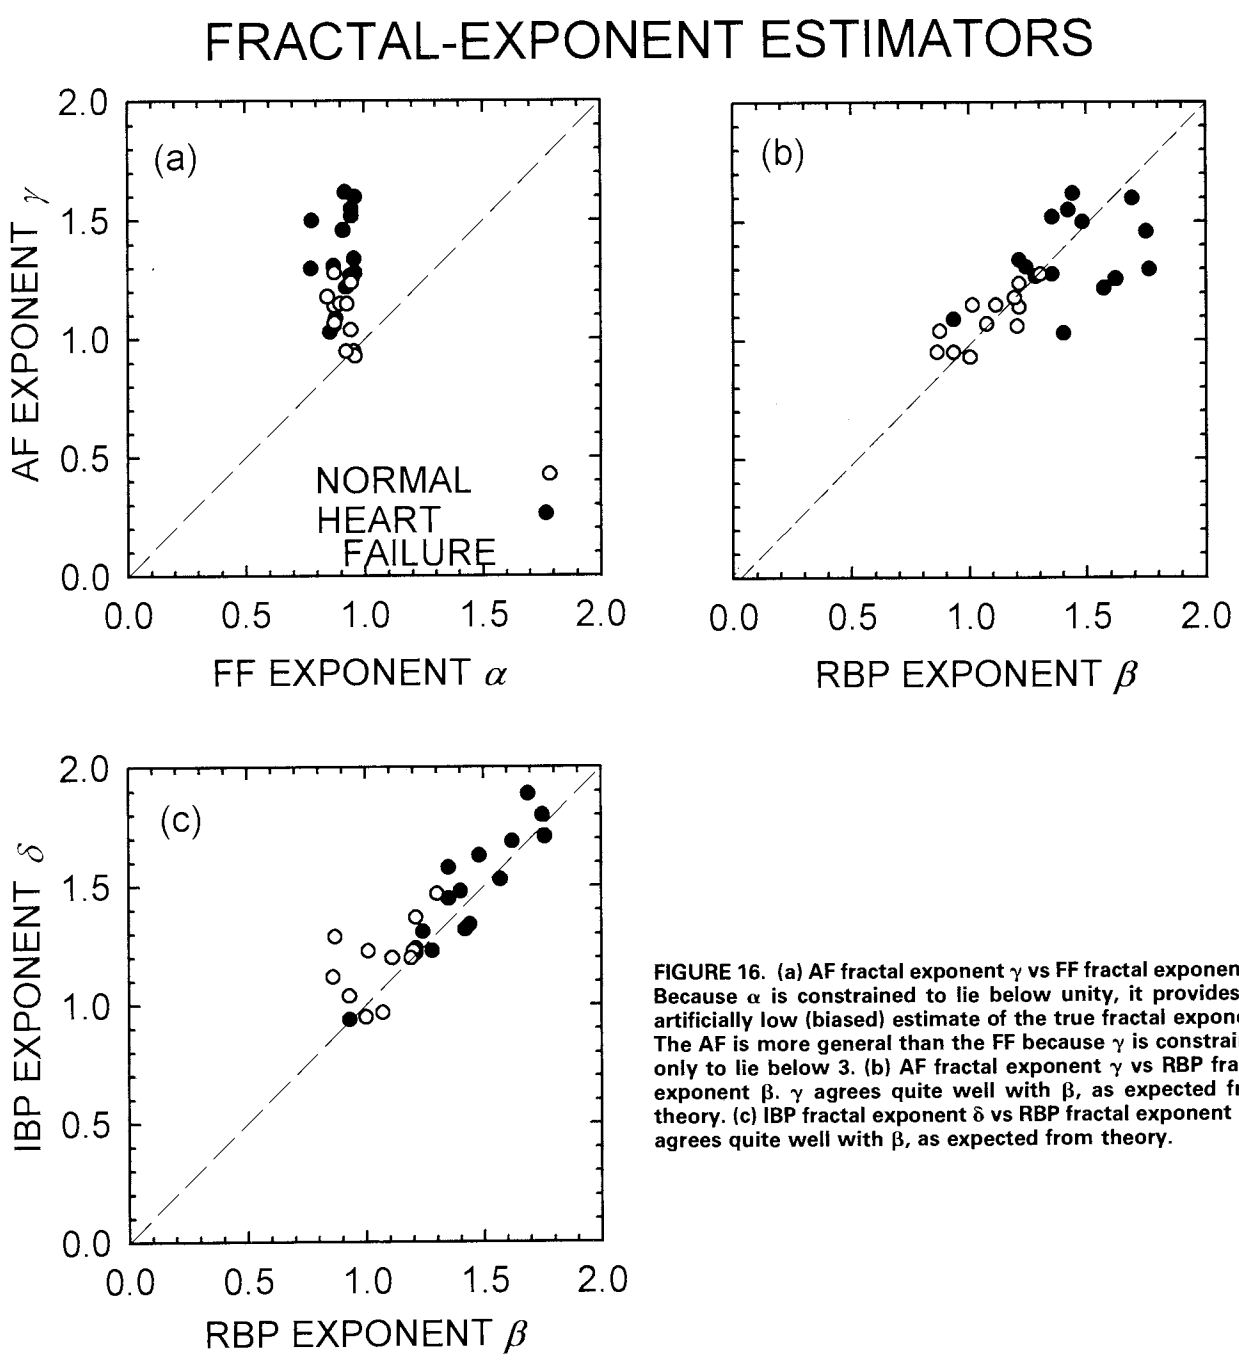

FIGURE 16. (a) AF fractal exponent γ vs FF fractal exponent α. Because α is constrained to lie below unity, it provides an artificially low (biased) estimate of the true fractal exponent. The AF is more general than the FF because γ is constrained only to lie below 3. (b) AF fractal exponent γ vs RBP fractal exponent β. γ agrees quite well with β, as expected from theory. (c) IBP fractal exponent δ vs RBP fractal exponent β. δ agrees quite well with β, as expected from theory.

the AF (lower-right panel). Indeed, $A(T = 10$ sec) completely separates the heart failures from the normals, exhibiting both a sensitivity and a specificity of 100%.

Finally, it is useful to compare our results with those of Peng *et al.* (41), who assumed that fractional Brownian motion (36,65) provides a suitable model for the sequence of heartbeat intervals. These authors examined the correlation properties of the increment process of the heartbeat intervals $I(i) \equiv \{\tau_{i+1} - \tau_i\}$, and obtained the exponent of the associated power-law periodogram (they denoted this quantity as β; in the following we denote it as β′ to avoid confusion with our RBP fractal exponent). Because $I(i)$ is generated from the original time series by passing the interval sequence through a simple differencing digital filter, however, it is straightforward to relate β′ directly to δ, the fractal exponent of the interval-based periodogram $S_\tau(f)$ (see *Methods*). In the frequency domain we have $S_I(f) = |H(f)|^2 S_\tau(f) = 2(1 - \cos 2\pi f) S_\tau(f)$, where $H(f)$ is the transfer function of the differencing filter. For sufficiently small f ($\lesssim 0.2$ cycles/interval), a Taylor-series ex-

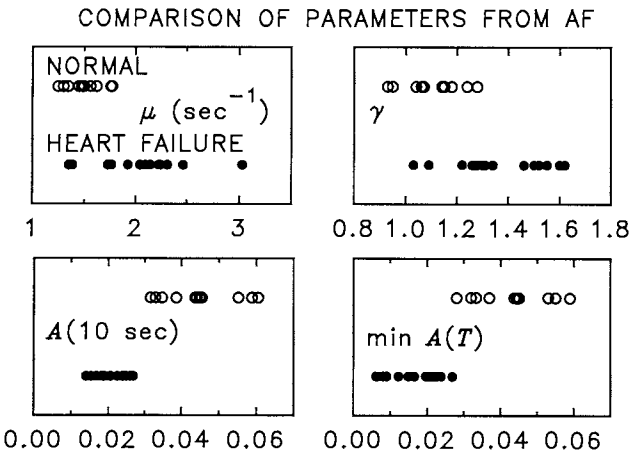

FIGURE 17. Comparison of AF parameters associated with all of the data sets. The mean heartrate μ (upper-left panel) fails to separate the classes of data very well. The AF fractal-exponent estimator γ (upper-right panel), like the RBP and IBP fractal-exponent estimators [see Fig. 11], provides only partial separation. The most successful indicators are the value of the AF at T = 10 sec (lower-left panel) and the minimum value of the AF (lower-right panel). Indeed, $A(T = 10$ sec) completely separates the heart failures from the normals, providing a sensitivity and specificity of 100%. The AF is a successful measure for separating the two classes of data because it is jointly responsive to both short-term and long-term statistical characteristics of the heartbeat time series.

pansion of the cosine function leads to $S_I(f) \sim f^2 S_\tau(f)$ which, using $S_I(f) \sim f^{+\beta'}$ and $S_\tau(f) \sim f^{-\delta}$, yields $\beta' = 2 - \delta$.

Peng *et al.* (41) calculated the average value of β' for ten normal data sets, obtaining $\beta' = 1.01 \pm 0.16$ (mean ± s.d.); for ten heart-failure data sets they obtained $\beta' = 0.54 \pm 0.25$. The corresponding values of $\delta = 2 - \beta'$ would therefore be 0.99 ± 0.16 and 1.46 ± 0.25, respectively. These values are, in fact, in good agreement with the values of δ we obtained directly from the IBPs (see Table 1), using independent calculations: 1.19 ± 0.15 for the average of 12 normal data sets (including the 10 that they used), and 1.47 ± 0.25 for the average of 14 heart-failure data sets (including the 10 that they used). Further evidence that δ and β' are equivalent is provided by the two data sets illustrated in Fig. 3 of their paper (41): in Fig. 3a they reported β'(data set 17453) = 0.93 for a normal patient, corresponding to $\delta = 1.07$; in Fig. 3b they reported β'(data set 9778) = 0.14 for a heart-failure patient, corresponding to $\delta = 1.86$. Again, these values closely match those we obtained directly from the IBPs (Table 1): δ(data set 17453) = 1.12 and δ(data set 9778) = 1.89.

Effectively, therefore, Peng *et al.* (41) propose that normal and heart-failure patients be distinguished by making use of the parameter δ, the exponent of the interval-based periodogram. Indeed δ does distinguish the two classes of data, on average, with good statistical significance ($p < 0.005$ from Table 1), as discussed earlier. However, this measure has limited ability for properly classifying individual heart-failure patients, as was illustrated in Fig. 11d. Comparing the four panels in Fig. 11 reveals that the sensitivity of Peng *et al.'s* measure β', which is the same as that of δ, *i.e.* 53%, is substantially inferior to that achievable by using var(τ), D_0, or β.

The magnitude of the Allan factor at $T = 10$ sec, because it is determined by short-term as well as long-term characteristics of the heartbeat series, is a superior measure to all of these, inasmuch as it completely separates the patients into two distinct classes and therefore exhibits both a sensitivity and specificity of 100%.

In a more recent contribution, Peng *et al.* (42) used detrended-fluctuation analysis to identify crossover behavior arising from differences in scaling over short versus long time scales, and introduced a stochastic model for pathologic data. However, in the following section, we show that a stochastic integrate-and-fire model, comprising a fractal-Gaussian-noise kernel, provides results in accord with *all* of the experimental measures that we have investigated, including the IBP and the AF.

INTEGRATE-AND-FIRE MODEL

Integrate-and-fire (IAF) models are widely used in neurophysiology (62) and in cardiology (5,25,26,48). IAF models are attractive, in part, because they capture known physiology in a simple way. The integration of a rate can represent the cumulative effect of neurotransmitter on the postsynaptic membrane of a neuron, or of the currents responsible for the pacemaker potential in the sino-atrial node of the heart, with the firing of an action potential or a heart contraction occurring when the integrated rate crosses a preset threshold.

In the context of the heartbeat, Berger *et al.* (5) considered an IAF model in which an underlying rate function $R(t)$ is integrated until it reaches a fixed threshold θ, whereupon a point event is triggered and the integrator is

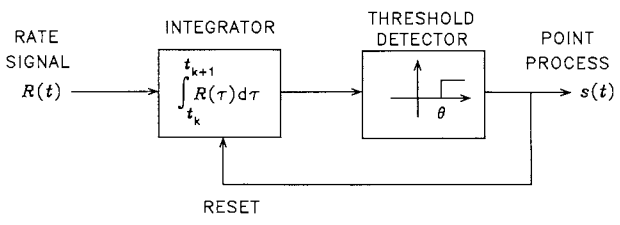

FIGURE 18. Schematic representation of the integrate-and-fire (IAF) mechanism for modeling the human heartbeat. An underlying rate process $R(t)$, assumed to be bandlimited fractal Gaussian noise (FGN), is integrated until it reaches a fixed threshold θ, whereupon a point event is generated and the integrator reset. The continuous rate process is thereby converted into a point process $s(t)$. The IAF approach was also used for generating the phase-randomized surrogate data (see *Methods*).

reset. As shown in Fig. 18, the occurrence time of the $(k + 1)$st event is then implicitly obtained from

$$\theta = \int_{t_k}^{t_{k+1}} R(\tau)d\tau. \tag{14}$$

In general, the rate $R(t)$ comprises a constant plus a time-varying component:

$$R(t) = R_0 + M(t); \tag{15}$$

R_0 is the mean heart rate. Heart rate variability is introduced through $M(t)$. In the limiting case when $M(t) = 0$ and $\theta = 1$, a regularly spaced deterministic firing pattern, with rate R_0, emerges.

Modeling the stochastic part of the rate function as bandlimited fractal Gaussian noise, and using a constant threshold $\theta = 1$, gives rise to a heartbeat sequence that is in remarkably good accord with observations, as demonstrated below. Ideal FGN is a continuous stochastic process of infinite bandwidth, with Gaussian amplitude statistics and a power-law decaying correlation function (36,65); its integral is fractional Brownian motion (FBM). However, since real data has a low-frequency cutoff A imposed by the duration of the signal under study, and a high-frequency cutoff B imposed by the time resolution of the measurement, bandlimited FGN (34) may be used, with no loss of generality, to carry out the model calculations more simply. A convenient point of departure is the rate-based power spectral density, which is characterized by four parameters (β, c, A, B) (see Fig. 19):

$$S(f) = cf^{-\beta}, \quad A \leq f \leq B. \tag{16}$$

The procedures used for selecting model parameter values are as follows. Two of the parameters of the model (the lower and upper cutoff frequencies, A and B respectively) are imposed implicitly by the duration and resolution of the simulated sample. The remaining three free parameters (the mean heartrate R_0, the coefficient c, and the RBP fractal exponent β) are readily obtained from the data. The threshold θ specifies the class of models; in the case considered here, θ is a fixed constant equal to unity. The power-law exponent β is estimated directly from the RBP. The coefficient c can be obtained from the data in one of two ways: directly from the RBP when the latter is calculated in absolute terms, or from β and the rate variance (which is the integral of the RBP). Because frequency components beyond the cutoffs A and B imposed by the statistical analysis contribute to the variance of the real data, however, the latter method will generally give rise to a larger apparent value of c (see schematic illustration in Fig. 19).

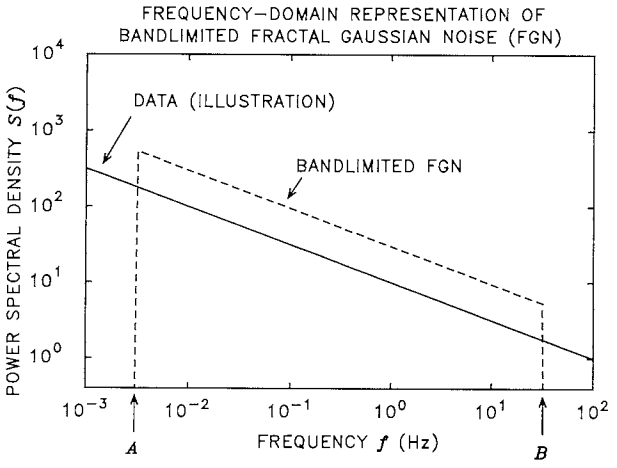

FIGURE 19. Frequency-domain representation of bandlimited fractal Gaussian noise. Frequency components lying beyond the bandwidth admitted by the model contribute to the rate variance of the data. Because of this, the power-law coefficient c of the FGN is correspondingly larger than that of the data over the frequency range of the model.

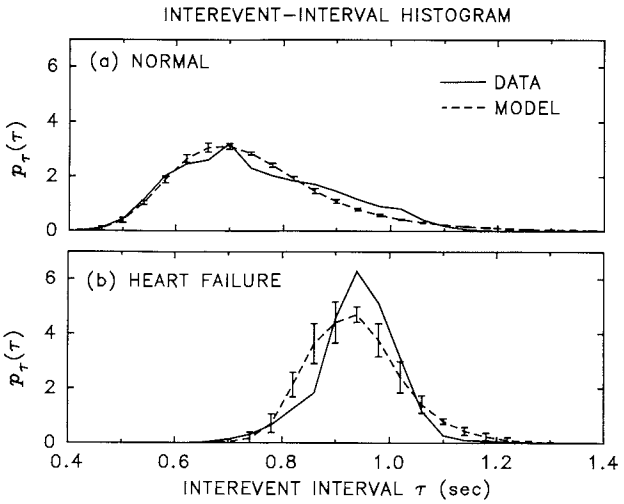

FIGURE 20. Comparison of interevent-interval histograms (IIHs) obtained from the heart data (solid curves, same as Fig. 3) and from the FGN-IAF model (dashed curves). The model simulations closely follow the general features of the data. The mean and standard deviation of the model results (indicated by error bars) were estimated from 10 computer simulations. The agreement is representative of the 27 data sets examined.

For each of the 27 data sets, the best-fitting model parameters were obtained and 10 computer simulations of the model point process were constructed. The mean and standard deviation (indicated by error bars) of various statistical measures were obtained from these 10 simulations. Figures 20–26 present both the fractal-Gaussian-noise/integrate-and-fire (FGN-IAF) model results (dashed curves) and the heart data (solid curves). The heart data shown in Figs. 20–26 (data set 16265 for normal and data set 6796 for heart failure) are identical to those displayed in Figs. 3 and 5–9. The degree of agreement between the data and model in Figs. 20–26 is representative of the 27 sets of heart data examined.

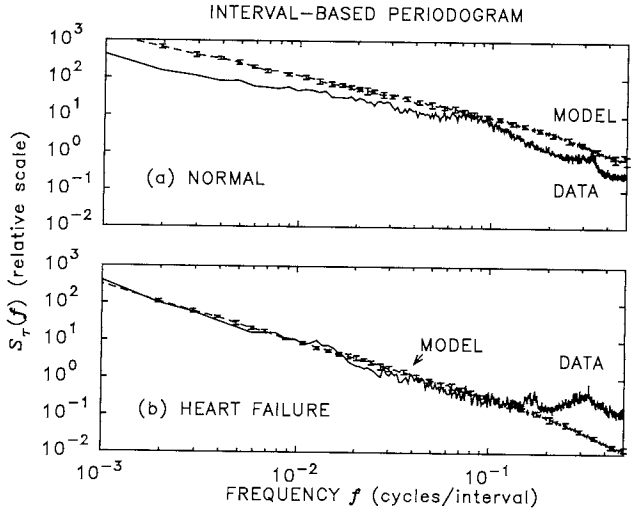

FIGURE 21. Interval-based periodograms (IBPs) calculated on the basis of the FGN-IAF model (dashed curves for mean, error brackets for s.d.) follow the general form of the data (solid curves) well, though they lack the spectral features associated with cyclic processes (*e.g.*, respiration) since these are not included in the model. The data are the same as those presented in Fig. 5. The fit of the model to the data is representative of all the data sets examined.

The interevent-interval histograms are shown in Fig. 20. The FGN-IAF model histograms (both normal and heart failure) are somewhat smoother than those of the heart data (in part as a result of the averaging of simulations), but they follow the general shapes of the data curves well. Indeed, the degree of similarity is particularly gratifying in light of the fact that the model parameters were extracted from the long-time-scale properties of the

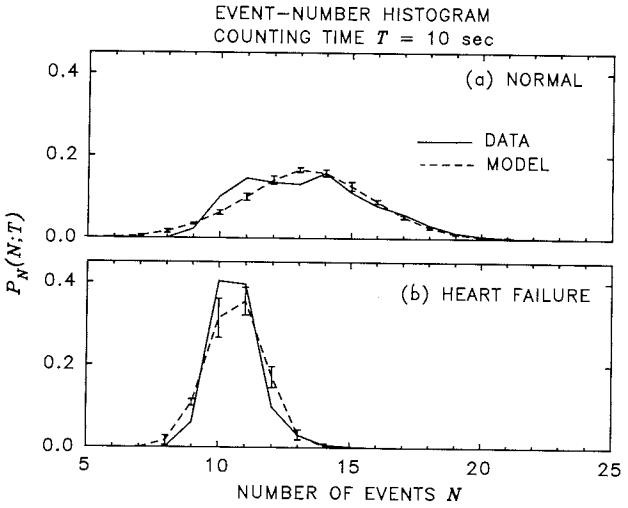

FIGURE 22. Event-number histograms (ENHs) with a counting time of T = 10 sec, calculated from simulations using the model (dashed curves for mean, error brackets for s.d.), accord well with the data (solid curves, same as Fig. 6). The agreement between the model curves and the data is representative of the 27 data sets examined.

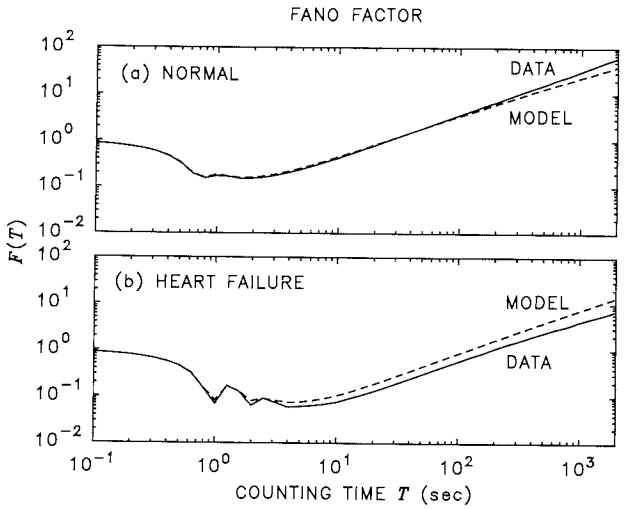

FIGURE 23. Fano-factors (FFs) of the data [solid curves, same as Fig. 7(a)] and model (dashed curves) are in good agreement. Since the standard deviation associated with the model curves is less than the thickness of the curve itself, error bars are not shown. The degree of divergence exhibited by these particular data sets is typical.

sequence of heartbeats, and no explicit limitations on the interevent-interval statistics (such as refractoriness or maximum interval size) were imposed.

As shown in Fig. 21, the interval-based periodograms calculated from the model also follow the general form of the data quite well. Cyclic processes such as respiration, which give rise to peaks in the periodogram, were not included in the model so the simulations lack the fine detail of the data.

The model event-number histograms for a counting time T = 10 sec, presented in Fig. 22, are approximately bell-shaped. Both the mean number of events and the

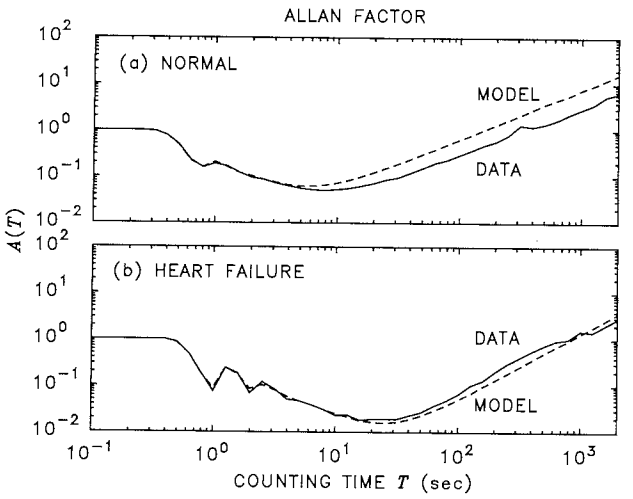

FIGURE 24. Allan-factors (AFs) of the data [solid curves, same as Fig. 7(b)] and model (dashed curves) accord quite well. The level of agreement is typical of all the data sets.

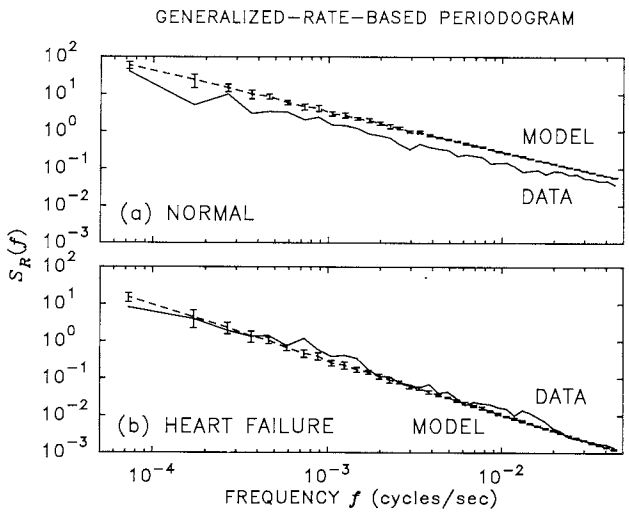

FIGURE 25. Generalized rate-based-periodograms (RBPs) computed with a binwidth of 10 sec. The model curves (dashed for mean, error brackets for standard deviation) agree reasonably well with the data (solid curves, same as Fig. 8).

event-number variance match the data well. Consequently the Fano and Allan factors, which are both calculated from the ENH, must be similar for the model and data at $T = 10$ sec.

Figures 23 and 24 show that the Fano factors and Allan factors for the model and data agree well not just at $T = 10$ sec, but over a broad range of counting times.

The generalized rate-based periodograms, computed with a binwidth of 10 sec, are shown in Fig. 25. The shapes of the model curves are quite similar to those of the data. As described above, however, the method of obtaining the coefficient c can overestimate its value in the band-limited region (see Fig. 19). This is apparent in Fig. 25a, where the model RBP has a greater magnitude than the data. While this overestimation was common, it was not universal; as shown in Fig. 25b for heart failure, the data and simulated RBPs are quite similar in magnitude. The degree of accord of the magnitudes depends on the spectral content of the data outside the bandwidth being simulated.

Finally, the capacity dimension is presented in Fig. 26. As with the surrogate data for this measure (Fig. 9), the FGN-IAF model, though it does not describe the data perfectly, succeeds rather well in emulating the general trend of the data, including saturation with increasing embedding dimension m. This confirms that such behavior is not a definitive indicator of underlying nonlinear deterministic dynamics but also can arise from temporal correlation in a stochastic process.

In short, the IAF model with a FGN kernel is remarkably successful at predicting all of the interevent-interval- and event-number-based measures that we have examined.

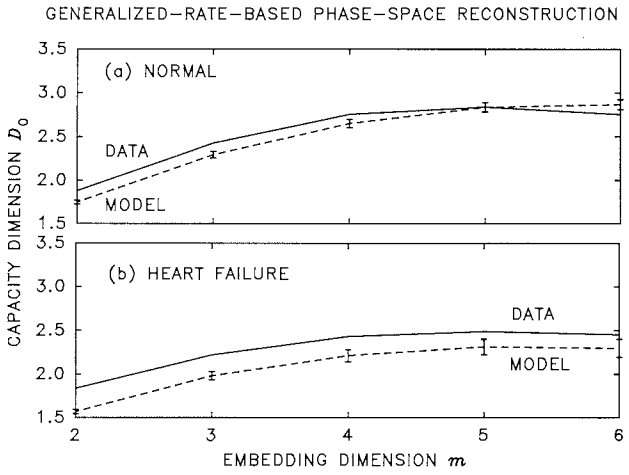

FIGURE 26. Capacity dimension vs. embedding dimension for data (solid curves, same as Fig. 9) and FGN-IAF model (dashed curves for mean, error brackets for standard deviation). The results predicted by the model level off with increasing embedding dimension m, mimicing quite nicely the way the data behave. Since the model is purely stochastic, this shows that such behavior need not be associated with nonlinear deterministic dynamics but can instead arise from temporal correlation in a stochastic process.

SUMMARY

We have studied the interevent-interval and event-number statistics of the sequence of human heartbeats from normal subjects and heart-failure patients. Both short-term and long-term measures reveal that, on average, the two classes of data behave differently.

We examined three measures generated from the sequence of interbeat intervals. The interevent-interval histogram is generally narrower for heart-failure patients, as determined by its variance. Rescaled range analysis manifests a strong correlation in the sequence of intervals, but it does not reveal a significant difference in the degree of correlation between the two classes of data. The interval-based periodogram also shows strong correlation in the sequence of intervals, and on average reveals a larger fractal exponent for heart-failure than for normal patients.

Event-number-based measures also exhibit differences between the two classes of data. The event-number histogram, computed for a sufficiently large counting time, reveals that the rate fluctuations are generally greater in the normal data than in the heart-failure data. The Fano factor, Allan factor, and rate-based periodogram all show power-law correlation in the sequence of counts for both classes of data; the fractal exponent for the heart-failure group is generally larger. For heart-failure patients there is less overall power in the ECG but relatively more power at low frequencies.

The behavior of the capacity dimension D_0, estimated from the reconstructed generalized-rate-based phase

space, is consistent with a putative low-dimensional attractor of a deterministic dynamical system. It is generally larger for the normal data. However, randomized-phase surrogate-data analysis gives similar results for D_0, suggesting that temporal correlation in the heartbeat time series, rather than the presence of an underlying attractor, is responsible for the effect. A similar conclusion has been reached by Kanters *et al.* (27), who based their conclusions on a dynamical-systems analysis of interbeat intervals and surrogates thereof.

We have developed a measure we call the Allan factor, which is jointly responsive to both short-term and long-term statistical characteristics of the heartbeat time series. Its success in properly identifying all 27 patients as normals or heart-failures (100% specificity and 100% sensitivity) leads us to consider it as a possible diagnostic tool for clinically distinguishing between the two groups of patients. We expect that it will be possible to devise other measures responsive to both the short- and long-term statistical characteristics of the ECG to separate the two classes of data. After all, the AF is an excellent measure whereas the FF is less useful; yet they contain the same information and are in fact related by the simple transformation provided in Eq. 12. Wavelet-based versions of the Allan factor are good candidates (58). Various combinations of interval-based measures, such as the intereventinterval histogram and the interval-based periodogram, should be examined as well.

For both normal and heart-failure patients, the sequence of heartbeats is well modeled by the point process generated by an integrate-and-fire mechanism with a fractal-Gaussian-noise kernel. It is remarkable that such a simple stochastic model so successfully characterizes all of the measures we have examined (save cyclic processes which are not incorporated in the model), particularly since they comprise both short- and long-time-scale behavior. One implication is that measures that exhibit differences between the two classes of data are simply reflecting differences in the free parameters used in the model.

Finally, we note that a more general class of point processes can be formulated in terms of the IAF model. For example, the threshold for firing θ need not be constant, but can be a random variable or a stochastic process. Randomness imparted to the threshold can then be used to represent variability in biological elements of the system, such as the amount of neurotransmitter released per synaptic vesicle or the duration of ion channel openings. The homogeneous Poisson point process with rate R_0 would then be a special case within this construct, generated when $M(t) = 0$ and θ is an independent, exponentially distributed random variable associated with each interevent interval. An inhomogeneous Poisson point process would result when the condition $M(t) = 0$ is relaxed. Models such as these have been used extensively in characterizing neural firing patterns (62).

REFERENCES

1. Babloyantz, A., and A. Destexhe. Is the normal heart a periodic oscillator? *Biol. Cybern.* 58:203–211, 1988.
2. Barnes, J. A. and D. W. Allan. A statistical model of flicker noise. *Proc. IEEE* 54:176–178, 1966.
3. Bassingthwaighte, J. B., L. S. Liebovitch, and B. J. West. *Fractal Physiology*. New York: American Physiological Society/Oxford, 1994, 364 pp.
4. Bassingthwaighte, J. B. and G. M. Raymond. Evaluating rescaled range analysis for time series. *Ann. Biomed. Eng.* 22:432–444, 1994.
5. Berger, R. D., S. Akselrod, D. Gordon, and R. J. Cohen. An efficient algorithm for spectral analysis of heart rate variability. *IEEE Trans. Biomed. Eng.* BME-33:900–904, 1986.
6. Bigger, J. T., Jr., J. L. Fleiss, R. C. Steinman, L. M. Rolnitzky, R. E. Kleiger, and J. N. Rottman. Frequency domain measures of heart period variability and mortality after myocardial infarction. *Circulation* 85:164–171, 1992.
7. Cox, D. R. and P. A. W. Lewis. *The Statistical Analysis of Series of Events*. London: Methuen, 1966, 285 pp.
8. DeBoer, R. W., J. M. Karemaker, and J. Strackee. Comparing spectra of a series of point events particularly for heart rate variability data. *IEEE Trans. Biomed. Eng.* BME-31:384–387, 1984.
9. Ding, M., C. Grebogi, E. Ott, T. Sauer, and J. A. Yorke. Plateau onset for correlation dimension: when does it occur? *Phys. Rev. Lett.* 70:3872–3875, 1993.
10. Fano, U. Ionization yield of radiations. II. The fluctuations of the number of ions. *Phys. Rev.* 72:26–29, 1947.
11. Feller, W. The asymptotic distribution of the range of sums of independent random variables. *Ann. Math. Stat.* 22:427–432, 1951.
12. Freeman, R., J. P. Saul, M. S. Roberts, R. D. Berger, C. Broadbridge, and R. J. Cohen. Spectral analysis of heart rate in diabetic autonomic neuropathy. *Arch. Neurol.* 48: 185–190, 1991.
13. Garfinkel, A., M. L. Spano, W. L. Ditto, and J. N. Weiss. Controlling cardiac chaos. *Science* 257:1230–1235, 1992.
14. Glass, L., A. Shrier, and J. Bélair. Chaotic cardiac rhythms. In *Chaos,* edited by A. V. Holden. Princeton: Princeton University Press, 1986, pp. 237–256.
15. Goldberger, A. L., D. R. Rigney, J. Mietus, E. M. Antman, and S. Greenwald. Nonlinear dynamics in sudden cardiac death syndrome: heartrate oscillations and bifurcations. *Experientia* 44:983–987, 1988.
16. Goldberger, A. L. Nonlinear dynamics, fractals and chaos: applications to cardiac electrophysiology. *Ann. Biomed. Eng.* 18:195–198, 1990.
17. Goldberger, A. L., D. R. Rigney, and B. J. West. Chaos and fractals in human physiology. *Sci. Am.* 262(2):42–49, 1990.
18. Grassberger, P. and I. Procaccia. Measuring the strangeness of strange attractors. *Physica D* 9:189–208, 1983.
19. Grüneis, F., M. Nakao, Y. Mizutani, M. Yamamoto, M. Meesmann, and T. Musha. Further study on $1/f$ fluctuations observed in central single neurons during REM sleep. *Biol. Cybern.* 68:193–198, 1993.
20. Guevara, M. R. and A. Shrier. Rhythms produced by high-

amplitude periodic stimulation of spontaneously beating aggregates of embryonic chick ventricular myocytes. *Ann. N.Y. Acad. Sci.* 591:11–22, 1990.
21. Guzzetti, S., E. Piccaluga, R. Casati, S. Cerutti, F. Lombardi, M. Pagani, and A. Malliani. Sympathetic predominance in essential hypertension: a study employing spectral analysis of heart rate variability. *J. Hypertens.* 6:711–717, 1988.
22. Haight, F. A. *Handbook of the Poisson Distribution.* New York: Wiley, 1967, 168 pp.
23. Hoop, B., H. Kazemi, and L. Liebovitch. Rescaled range analysis of resting respiration. *Chaos* 3:27–29, 1993.
24. Hurst, H. E. Long-term storage capacity of reservoirs. *Trans. Amer. Soc. Civil Eng.* 116:770–808, 1951.
25. Hyndman, B. W. and R. K. Mohn. A pulse modulator model for pacemaker activity. *Dig. 10th Int. Conf. Med. Biol. Eng.* (Dresden, Germany) p. 223, 1973.
26. Hyndman, B. W. and R. K. Mohn. Model of the cardiac pacemaker and its use in decoding the information content of cardiac intervals. *Automedica* 1:239–252, 1975.
27. Kanters, J. K., N.-H. Holstein-Rathlou, and E. Agner. Lack of evidence for low-dimensional chaos in heart rate variability. *J Cardiovasc. Electrophysiol.* 5:591–601, 1994.
28. Kaplan, D. T. and R. J. Cohen. Searching for chaos in fibrillation. *Ann. N.Y. Acad. Sci.* 591:367–374, 1990.
29. Kluge, K. A., R. M. Harper, V. L. Schechtman, A. J. Wilson, H. J. Hoffman, and D. P. Southall. Spectral analysis assessment of respiratory sinus arrhythmia in normal infants and infants who subsequently died of sudden infant death syndrome. *Pediatr. Res.* 24:677–682, 1988.
30. Kobayashi, M. and T. Musha. $1/f$ fluctuation of heartbeat period. *IEEE Trans. Biomed. Eng.* BME-29:456–457, 1982.
31. Liebovitch, L. S. Introduction to the properties and analysis of fractal objects, processes, and data. In: *Advanced Methods of Physiological System Modeling,* Vol. 2, edited by V. Z. Marmarelis. New York: Plenum Press, 1989, pp. 225–239.
32. Liebovitch, L. S. and T. Tóth. A fast algorithm to determine fractal dimensions by box counting. *Phys. Lett. A* 141:386–390, 1989.
33. Lishner, M., S. Akselrod, V. Mor Avi, O. Oz, M. Divon, and M. Ravid. Spectral analysis of heart rate fluctuations. A non-invasive, sensitive method for the early diagnosis of autonomic neuropathy in diabetes mellitus. *J. Auton. Nerv. Sys.* 19:119–125, 1987.
34. Lowen, S. B. and M. C. Teich. Estimation and simulation of fractal stochastic point processes. *Fractals* 3:183–210, 1995.
35. Lowen, S. B. and M. C. Teich. The periodogram and Allan variance reveal fractal exponents greater than unity in auditory-nerve spike trains. *J. Acoust. Soc. Am.* 99:in press, 1996.
36. Mandelbrot, B. B. and J. W. van Ness. Fractional brownian motions, fractional noises and applications. *SIAM Rev.* 10:422–437, 1968.
37. Meesmann, M., F. Grüneis, P. Flachenecker, and K.-D. Kniffki. A new method for analysis of heart rate variability: counting statistics of $1/f$ fluctuations. *Biol. Cybern.* 68:299–306, 1993.
38. Moody, G. B. and R. G. Mark. Development and evaluation of a two-lead ECG analysis program. *Proc. IEEE Comp. in Cardiol. Conf.* 9:39–44, 1982.
39. Osborne, A. R. and A. Provenzale. Finite correlation dimension for stochastic systems with power-law spectra. *Physica D* 35:357–381, 1989.
40. Papoulis, A. *Probability, Random Variables, and Stochastic Processes,* 2nd Ed. New York: McGraw-Hill, 1984, 576 pp.
41. Peng, C.-K., J. Mietus, J. M. Hausdorff, S. Havlin, H. E. Stanley, and A. L. Goldberger. Long-range anticorrelations and non-Gaussian behavior of the heartbeat. *Phys. Rev. Lett.* 70:1343–1346, 1993.
42. Peng, C.-K., S. Havlin, H. E. Stanley, and A. L. Goldberger. Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series. *Chaos* 5:82–87, 1995.
43. Powers, N. L., R. J. Salvi, and S. S. Saunders. Discharge rate fluctuations in the auditory nerve of the chinchilla. Abstracts of the XIV Midwinter Research Meeting, Assoc. for Res. in Otolaryngology, Des Moines, IA; D. J. Lim, Ed., Abstract No. 411, p. 129, 1991.
44. Powers, N. L. and R. J. Salvi. Comparison of discharge rate fluctuations in the auditory nerve of chickens and chinchillas. Abstracts of the XV Midwinter Research Meeting, Assoc. for Res. in Otolaryngology, Des Moines, IA; D. J. Lim, Ed., Abstract No. 292, p. 101, 1992.
45. Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C,* 2nd Ed. Cambridge, U.K.: Cambridge University Press, 1992, 994 pp.
46. Ravelli, F. and R. Antolini. Complex dynamics underlying the human electrocardiogram. *Biol. Cybern.* 67:57–65, 1992.
47. Ricciardi, L. M. and F. Esposito. On some distribution functions for non-linear switching elements with finite dead time. *Kybernetik (Biol. Cybern.)* 3:148–152, 1966.
48. Rompelman, O., J. B. I. M. Snijders, and C. J. van Spronsen. The measurement of heart rate variability spectra with the help of a personal computer. *IEEE Trans. Biomed. Eng.* BME-29:503–510, 1982.
49. Saul, J. P., P. Albrecht, R. D. Berger, and R. J. Cohen. Analysis of long term heart rate variability: methods, $1/f$ scaling and implications. *Proc. IEEE Comput. in Cardiol. Conf.* 14:419–422, 1988.
50. Scharf, R., M. Meesmann, J. Boese, D. R. Chialvo, and K.-D. Kniffki. General relation between variance-time curve and power spectral density for point processes exhibiting $1/f^{\beta}$-fluctuations, with special reference to heart rate variability. *Biol. Cybern.* 73:255–263, 1995.
51. Schepers, H. E., J. H. G. M. van Beek, and J. B. Bassingthwaighte. Four methods to estimate the fractal dimension from self-affine signals. *IEEE Eng. Med. Biol. Mag.* 11:57–71, 1992.
52. Schiff, S. J. and T. Chang. Differentiation of linearly correlated noise from chaos in a biologic system using surrogate data. *Biol. Cybern.* 67:387–393, 1992.
53. Shono, H., M. Yamasaki, M. Muro, M. Oga, Y. Ito, K. Shimomura, and H. Sugimori. Chaos and fractals with $1/f$ spectrum below 10^{-2} Hz demonstrates full-term fetal heart rate changes during active phase. *Early Human Dev.* 27: 111–117, 1991.
54. Teich, M. C. and S. M. Khanna. Pulse-number distribution for the neural spike train in the cat's auditory nerve. *J. Acoust. Soc. Am.* 77:1110–1128, 1985.
55. Teich, M. C. Fractal character of the auditory neural spike train. *IEEE Trans. Biomed. Eng.* 36:150–160, 1989.
56. Teich, M. C. Fractal neuronal firing patterns. In: *Single*

Neuron Computation, edited by T. McKenna, J. Davis, and S. Zornetzer. Boston: Academic Press, 1992, pp. 589–625.
57. Teich, M. C. and S. B. Lowen. Fractal patterns in auditory nerve-spike trains. *IEEE Eng. Med. Biol. Mag.* 12:197–202, 1994.
58. Teich, M. C., C. Heneghan, S. B. Lowen, and R. G. Turcott. Estimating the fractal exponent of point processes in biological systems using wavelet- and Fourier-transform methods. In: *Wavelets in Medicine and Biology*, edited by A. Aldroubi and M. Unser. Boca Raton, FL: CRC Press, 1996.
59. Teich, M. C., R. G. Turcott, and R. M. Siegel. Variability and long-duration correlation in the sequence of action potentials in cat striate-cortex neurons. *IEEE Eng. Biol. Med. Mag.* 14:in press, 1996.
60. Theiler, J. Estimating fractal dimension. *J. Opt. Soc. Am. A* 7:1055–1073, 1990.
61. Theiler, J., S. Eubank, A. Longtin, B. Galdrikian, and J. D. Farmer. Testing for nonlinearity in time series: the method of surrogate data. *Physica D* 58:77–94, 1992.
62. Tuckwell, H. C. *Stochastic Processes in the Neurosciences*. Philadelphia: Society for Industrial and Applied Mathematics, Philadelphia, 1989, 129 pp.
63. Turcott, R. G. and M. C. Teich. Long-duration correlation and attractor topology of the heartbeat rate differ for normal patients and those with heart failure. *Proc. SPIE (Chaos in Biology and Medicine)* 2036:22–39, 1993.
64. Turcott, R. G., P. D. R. Barker, and M. C. Teich. Long-duration correlation in the sequence of action potentials in an insect visual interneuron. *J. Statist. Comput. Simul.* 52: 253–271, 1995.
65. West, B. J. and W. Deering. Fractal physiology for physicists: Lévy statistics. *Physics Reports* 246:1–100, 1994.
66. Witkowski, F. X., K. M. Kavanagh, P. A. Penkoske, R. Plonsey, M. L. Spano, W. L. Ditto, and D. T. Kaplan. Evidence for determinism in ventricular fibrillation. *Phys. Rev. Lett.* 75:1230–1233, 1995.
67. Wolf, A., J. B. Swift, H. L. Swinney, and J. A. Vastano. Determining Lyapunov exponents from a time series. *Physica D* 16:285–317, 1985.
68. Zbilut, J. P., G. Mayer-Kress, P. A. Sobotka, M. O'Toole, and J. X. Thomas, Jr. Bifurcations and intrinsic chaotic and $1/f$ dynamics in an isolated perfused rat heart. *Biol. Cybern.* 61:371–378, 1989.

LIST OF ABBREVIATIONS

AF = Allan factor
CD = Capacity dimension
DFT = Discrete Fourier transform
DTMP = Dead-time-modified Poisson point process
ECG = Electrocardiogram
ENH = Event-number histogram
FBM = Fractional Brownian motion
FF = Fano factor
FGN = Fractal Gaussian noise
FSPP = Fractal stochastic point process
HPP = Homogeneous Poisson point process
IAF = Integrate-and-fire
IBP = Interval-based periodogram
IIH = Interevent-interval histogram
NCR = Normalized coincidence rate
PSD = Power spectral density
PSR = Phase-space reconstruction
RBP = Generalized rate-based periodogram
R/S = Rescaled range analysis
WAF = Wavelet Allan factor
WFF = Wavelet Fano factor