

BENEATH THE NOISE, CHAOS¹

BY STEVEN P. LALLEY

Purdue University

The problem of extracting a signal x_n from a noise-corrupted time series $y_n = x_n + e_n$ is considered. The signal x_n is assumed to be generated by a discrete-time, deterministic, chaotic dynamical system F , in particular, $x_n = F^n(x_0)$, where the initial point x_0 is assumed to lie in a compact hyperbolic F -invariant set. It is shown that (1) if the noise sequence e_n is Gaussian then it is impossible to consistently recover the signal x_n , but (2) if the noise sequence consists of i.i.d. random vectors uniformly bounded by a constant $\delta > 0$, then it is possible to recover the signal x_n provided $\delta < 5\Delta$, where Δ is a separation threshold for F . A filtering algorithm for the latter situation is presented.

1. Introduction. Physical and numerical experiments carried out over the past 30 or more years suggest that the phenomenon of *deterministic chaos* is ubiquitous in physical systems. Experience has shown that inference of the mathematical objects (the differential equations, equilibrium measures, Lyapunov exponents, etc.) governing the dynamics of such systems from time series data is a delicate problem even when this data is uncorrupted by noise. See [5] and [6] for an extensive review and bibliography. Inference from *noisy* data is therefore bound to be doubly difficult. Although various ad hoc “noise reduction” algorithms have been proposed (some seemingly quite effective when tested on computer-generated data from low-dimensional chaotic systems, e.g., [15] and [10]), their theoretical properties are largely unknown.

The purpose of this paper is to address the following basic question: Is it possible to consistently recover a “signal” $\{x_n\}_{n \in \mathbb{Z}}$ generated by an Axiom A system from a time series of the form

$$(1) \quad y_n = x_n + e_n,$$

where e_n is observational noise? A positive answer would essentially reduce the problem of inference from *noisy* time series data to that of inference from *nonnoisy* data. The following scenario for the signal will be considered here:

$$(2) \quad x_n = F^1(x_{n-1}) = F^n(x_0),$$

where F is a C^2 diffeomorphism and x_0 is a point lying in a hyperbolic invariant set or in the basin of attraction of a hyperbolic attractor (see Section 2 for definitions and examples). Our main result is that the possibility

Received January 1998; revised December 1998.

¹ Supported in part by NSF Grant DMS-96-26590.

AMS 1991 subject classifications. Primary 62M20; secondary 58F15.

Key words and phrases. noise reduction, nonlinear filter, Axiom A, hyperbolic attractor.

of consistent signal extraction depends on the nature and amplitude of the noise. If the noise e_n is uniformly bounded and the bound is below a certain threshold Δ then consistent signal extraction is possible, but if the noise is *unbounded*, in particular Gaussian, then consistent signal extraction is impossible (even when the L^2 -norm of e_n is far below the threshold Δ).

In a companion paper [12] we shall consider a different but related scenario for the signal x_n , which is technically (and perhaps also conceptually) more difficult but probably of greater practical importance. In this scenario, the underlying dynamical system is a topologically mixing Axiom A flow F^t , but observations on the orbit $x_t = F^t(x_0)$ are made only at integer times n . It will be shown that the dichotomy between bounded and unbounded noise persists, and an algorithm for noise removal (more complicated than that given in this paper) will be presented.

We must emphasize at the outset that the results of this and the companion paper, and in particular the type of asymptotics considered, may not be relevant or appropriate for all signal extraction problems connected with noisy data from chaotic dynamical systems. In various circumstances, more or less will be known a priori about the dynamical system than we assume here. In many circumstances, inference about the dynamics F^t and/or the basic set Λ will be of greater importance than extraction of the signal x_n itself. Finally, when dealing with flows F^t rather than diffeomorphisms, the experimenter may sometimes be able to control the frequency of observation.

2. Background: attractors, hyperbolicity and Axiom A.

2.1. Invariant sets and attractors. The model for a smooth discrete-time dynamical system is a C^2 diffeomorphism F of a phase space M . (A C^2 diffeomorphism is a bijective mapping F such that both F and F^{-1} are twice continuously differentiable; see [8]). For simplicity, we take M to be an open subset of \mathbb{R}^d . The *orbits* of the system are the sequences $\{x_n\}$ such that $x_{n+1} = F(x_n) \forall n$. A compact subset Λ of the phase space will be called *F-invariant* if $F^{-1}(\Lambda) = \Lambda$, so that the restriction $F|_{\Lambda}$ of F to Λ is a homeomorphism of Λ . Especially important among the invariant sets are *attractors*, which arise when the phase space contains a relatively compact open set Ω such that $\text{closure}(F\Omega) \subset \Omega$. If there exists such a set Ω , the set $\Lambda = \bigcap_{n>0} F^n\Omega$ is a nonempty F -invariant compact set, called an *attractor* for the diffeomorphism, and Ω is contained in its *basin of attraction*. All orbits $x_n = F^n(x_0)$ beginning at points $x_0 \in \Omega$ converge to Λ .

2.2. Example: Smale's solenoid mapping. The following example, Smale's solenoid mapping, shows that attractors may have a complex structure. The set Ω is a solid torus in \mathbb{R}^3 centered at the origin that may be parametrized by a real coordinate $\theta \in [0, 2\pi)$ and a complex coordinate $z \in \{|z| < 1\}$. (Picture the torus as a solid of revolution obtained by rotating the solid disc

($|z| < 1$) once around the origin.) Fix $\alpha \in (0, \frac{1}{2})$, and define

$$(3) \quad F_\alpha(\theta, z) = (2\theta, \alpha z + e^{i\theta}/2),$$

where 2θ is reduced mod 2π if $\theta \geq \pi$. In geometric terms, the mapping F_α is obtained as follows: (1) Cut the torus and unroll it to get a solid cylinder. (2) Stretch the cylinder lengthwise by a factor of two, then compress the resulting cylinder in the directions orthogonal to its length by a factor of α . (3) Wrap the resulting long, thin cylinder twice around the origin and place it so that it is entirely inside the original solid torus and reattach the two ends. See Figure 1. (Note that $\alpha < 1/2$ keeps the two “branch” cylinders from intersecting, and the centering term $e^{i\theta}/2$ allows the branch cylinder to “roll” completely around once as θ varies from $-\pi$ to $+\pi$.)

For each α , the diffeomorphism F_α has an attractor $\Lambda \subset \Omega$ whose intersection with any “slice” $\Omega_\beta = \{(e^{i\theta}, z) : \theta = \beta\}$ is a Cantor set; see Figure 2. For each $\xi \in \Lambda \cap \Omega_\beta$ there is a smooth curve γ_ξ through ξ transverse to Ω_β that is contained in Λ . The homeomorphism $F_\alpha|_\Lambda$ multiplies distances locally along each γ_ξ by 2, and multiplies distances in $\Omega_\beta \cap \Lambda$ by α . See [4], Section 2.5, for helpful diagrams and further details.

2.3. *Hyperbolicity and orbit separation.* A compact invariant set Λ is called *hyperbolic* if at every point $\xi \in \Lambda$ there is a splitting of \mathbb{R}^d as a direct sum $E^u \oplus E^s$ of subspaces in such a way that the splitting depends on ξ continuously, and for all $n \geq 1$,

$$(4) \quad \|DF^n v\| \geq c_u \lambda^n \|v\| \quad \forall v \in E^u,$$

$$(5) \quad \|DF^n v\| \leq c_s \lambda^{-n} \|v\| \quad \forall v \in E^s,$$

with suitable constants $0 < c_s, c_u < \infty$ and $\lambda > 1$. Here DF^n denotes the matrix of first partial derivatives of F^n . Roughly speaking, the spaces E^u (E^s) consist of those directions in which F^n expands (shrinks) distances locally for large n . The solenoid attractor is hyperbolic: for $\xi \in \Lambda$, E^s is the two-dimensional space of vectors pointing into the slice Ω_β containing ξ , and




FIG. 1.




FIG. 2.

E^u is the one-dimensional space of vectors pointing in the direction of the curve γ_ξ . For our purposes, hyperbolicity is important only insofar as it implies the *orbit separation property*: orbits of nearby points diverge rapidly. In particular, there exist constants $1 > \Delta > 0$ (which we shall call a *separation threshold*) and $C > 0$ such that if $0 < |x - x'| < \Delta$ for two points $x, x' \in \Lambda$ then

$$(6) \quad |F^n(x) - F^n(x')| > \Delta \quad \text{for some } |n| \leq -C \log|x - x'|.$$

Note: The existence of a separation threshold follows from the Hirsch–Pugh theorem ([7], Theorem 5.2.8), which asserts the existence of stable and unstable manifolds at each point of a hyperbolic invariant set. The exponential rates of orbit separation are the *Lyapunov exponents* of the system; see, e.g., [5] for a thorough discussion. The Lyapunov exponents will play no role in our results (but may be important in designing more efficient filters; see the preliminary discussion in Section 3.1.5).

2.4. Smale's Axiom A. A compact hyperbolic invariant set Λ will be called an *Axiom A basic set* if (i) periodic points are dense in Λ , and (ii) there exists $x \in \Lambda$ such that for every $m \geq 0$ the forward orbit $\{F^n(x)\}_{n \geq m}$ is dense in Λ . (See [2] for the standard definition. A *periodic point* x is a point whose orbit is periodic, i.e., such that for some $n \in \mathbb{N}$, $F^n(x) = x$.) It is called *topologically mixing* if for every pair U, V of open subsets U, V there exists n_* sufficiently large that for every $n > n_*$,

$$F^n(U) \cap V \neq \emptyset.$$

It is not difficult to verify that the solenoid is a topologically mixing Axiom A attractor.

There are some celebrated theoretical results in dynamical systems theory that suggest the importance of Axiom A systems, and which we take as (partial) justification for focusing our attention on these. First, Axiom A systems are structurally stable (see, e.g., [16], Corollary 8.24); a small perturbation of a diffeomorphism F with an Axiom A basic set Λ results in another diffeomorphism F^* with an Axiom A basic set Λ^* near Λ , and the restriction of F^* to Λ^* is topologically conjugate to the restriction of F to Λ , that is, there is a homeomorphism $\phi: \Lambda \rightarrow \Lambda^*$ such that

$$F^* \circ \phi = \phi \circ F.$$

Second, according to the Birkhoff–Smale theorem, if a diffeomorphism F has a hyperbolic periodic point whose stable and unstable manifolds intersect transversally, then this point must be contained in an Axiom A basic set (a *horseshoe*). (See [7], Section 5.3 for definitions and a precise statement; also [14], Chapter 2, and [17].) It is known that such hyperbolic periodic points occur in a number of physically important dynamical systems (see, e.g., [14]).

The ergodic theory of Axiom A basic sets and attractors is well understood; see [2] for a thorough exposition. Of special importance in the study of Axiom A attractors is the existence of a (unique) strongly mixing F -invariant probability measure μ_* , the so-called SRB measure (for Sinai, Ruelle and Bowen), that is supported by Λ and has the following property: for every continuous function $\varphi: \Omega \rightarrow \mathbb{R}$ and for a.e. $x \in \Omega$ (relative to Lebesgue measure on Ω),

$$(7) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^n \varphi(F^k(x)) = \int \varphi d\mu_*.$$

It is the SRB measure that one would expect to “see” in time series data. For our purposes, the essential fact about the SRB measure is that it is a Gibbs state in the sense of [2], Chapter 1. More background on Axiom A diffeomorphisms, Gibbs states and SRB measures, of a more technical nature, is given in the Appendix. This additional material is needed for the proofs, but not the statements, of the results in the following section.

3. Signal extraction.

3.1. *Bounded noise.* Consider now the problem of reconstructing an orbit $\{x_n\}$ from a noise-corrupted time series $y_n = x_n + e_n$. The sequence x_n is generated by (2), and we assume that the initial point x_0 is either an element of a (compact) hyperbolic invariant set or in the basin of attraction of a hyperbolic attractor. We first consider the problem of noise removal under the hypothesis that the noise is uniformly bounded.

HYPOTHESIS 1. *Conditional on the sequence $\{x_n\}$ (equivalently, conditional on x_0) the random vectors e_n are independent, uniformly bounded by a constant $\delta > 0$, and have expectations*

$$(8) \quad E(e_n | x_0) = 0.$$

3.1.1. *Smoothing Algorithm D.* This algorithm is designed for time series produced by a diffeomorphism (hence the D), with noise satisfying Hypothesis 1, and assumes that a suitable bound $\delta > 0$ for the noise is known a priori. The algorithm takes as input a finite sequence $\{y_n\}_{0 \leq n \leq m}$ and produces as output a sequence $\{\hat{x}_n\}_{0 \leq n \leq m}$ of the same length that will approximate the unobservable signal $\{x_n\}_{0 \leq n \leq m}$. Let κ_m be an increasing sequence of integers such that

$$(9) \quad \lim_{m \rightarrow \infty} \kappa_m = \infty \quad \text{and} \quad \lim_{m \rightarrow \infty} \frac{\kappa_m}{\log m} = 0;$$

for example, $\kappa_m = \log m / \log \log m$. For each integer $\kappa_m < n < m - \kappa_m$, define A_n to be the set of indices $\nu \in \{0, 1, \dots, m\}$ such that

$$(10) \quad \max_{|j| \leq \kappa_m} |y_{\nu+j} - y_{n+j}| < 3\delta,$$

with the convention that $|y_j - y_i| = \infty$ if either of i or j is not in the range $[0, m]$. Observe that $n \in A_n$, so A_n is nonempty. For $n \leq \kappa_m$ or $n \geq m - \kappa_m$, define A_n to be the singleton $\{n\}$. In rough terms, A_n consists of the indices of those points in the time series whose orbits “shadow” the orbit of x_n for κ_m time units. In Lemma 1 we will show that $\nu \in A_n$ implies that $|x_\nu - x_n|$ is small. Thus, even though the values x_j are unobservable, neighboring points may still be picked out by virtue of having similar orbits. Now define

$$(11) \quad \hat{x}_n = \frac{1}{|A_n|} \sum_{\nu \in A_n} y_\nu.$$

THEOREM 1. *Assume that x_0 is either an element of a compact hyperbolic invariant set Λ or an element of the basin of attraction of a compact hyperbolic attractor Λ and assume that the noise sequence e_n satisfies Hypothesis 1. Let Δ be a separation threshold for the invariant set. If $5\delta < \Delta$, then for every $\varepsilon > 0$,*

$$(12) \quad \lim_{m \rightarrow \infty} P \left\{ m^{-1} \sum_{n=0}^m \mathbf{1}\{|\hat{x}_n - x_n| \geq \varepsilon\} \geq \varepsilon \right\} = 0.$$

Theorem 1 is valid for every orbit $\{x_n\}_{n \geq 0}$ contained in Λ , but the conclusion is only a weak convergence statement. For “generic” orbits of an Axiom A basic set, the conclusion can be strengthened to an a.s. convergence statement.

THEOREM 2. *Assume that x_0 is chosen at random from a Gibbs state μ_* supported by an Axiom A basic set Λ , and assume that the noise sequence e_n satisfies Hypothesis 1. Let Δ be a separation threshold for Λ . If $5\delta < \Delta$ then with probability 1,*

$$(13) \quad \lim_{m \rightarrow \infty} \max_{\kappa_m < n < m - \kappa_m} |\hat{x}_n - x_n| = 0.$$

The most important case (probably) is when Λ is an Axiom A attractor and μ_* is the SRB measure. In practice, when dealing with an attractor, the initial point might be chosen at random from an absolutely continuous distribution on the basin of attraction Ω , and an initial segment of the orbit would then be discarded. Theorem 2 remains valid under this hypothesis.

Since the almost sure convergence statement (13) holds for points x_0 chosen at random from *any* F -invariant Gibbs state and since Gibbs states are dense in the space of ergodic F -invariant probability measures on Λ , one might at first suspect that Theorems 1 and 2 might be strengthened to the stronger statement that (13) holds for *every* x_0 in Λ . This is false: it can be shown that *every* Axiom A basic set contains orbits for which (13) fails.

Theorems 1 and 2 will be proved in Sections 5 and 6 below, respectively. In both cases, only the proofs for orbits x_n contained in Λ will be given, as the proofs for orbits initiated in the basin of attraction are nearly identical. The proof of Theorem 1 is relatively elementary, but that of Theorem 2 requires deeper results from the ergodic theory of Axiom A basic sets, which are collected in the Appendix.

3.1.2. *Other noise reduction schemes.* The problem of noise reduction has been studied by a number of authors; see [1], Chapter 7, [3], [9] and [10] for reviews and further pointers to the recent literature. The methods proposed in these papers can be partitioned into two broad classes: (1) those that attempt to estimate F using local linear (or polynomial) maps, and then replace the series y_n by a nearby orbit of the estimated map and (2) those that use a principal-components decomposition of the autocovariance matrix, usually removing the smaller principal components. Smoothing Algorithm D does not fit into either category. It seems to be the first proposed method that directly exploits the orbit separation property. (It is also the first proposed method for which rigorous results concerning convergence properties are known.) It may naturally be expected that the usefulness of Smoothing Algorithm D will be limited to those dynamical systems with sensitive dependence on initial conditions, as it depends on orbit separation to “recognize” nearby points. However, these are precisely the systems that many experimenters expect to see. Comparison of the performance of Smoothing Algorithm D with the performance of various other methods will be a worthwhile and interesting project.

3.1.3. *Implementation.* One might expect to use filters of the type described above on time series of length $m = 10^6$ or more, and so from a practical standpoint *computational* efficiency may be as important as *statistical* efficiency. Although implementation of Smoothing Algorithm D in the form described above may require $O(m^2)$ comparisons, there are simple modifications that can be implemented by $O(m \log m)$ -step algorithms. In perhaps the simplest such modification, the indices $n \in [1, m]$ are sorted into bins B_v indexed by integer vectors v , using the following rule: $n \in B_v$ if and only if v is the integer vector whose entries are the integer parts of the

entries of $2y_n/3\delta$. (Note that if two indices n, n_* are in the same bin, then, by the triangle inequality, $|y_n - y_{n_*}| < 3\delta$.) The indices $n \in [1, m]$ are then resorted into bins B_w^* indexed by arrays w of integer vectors of length $2\kappa_m$, with $n \in B_w^*$ if and only if for each $|j| \leq \kappa_m$ the index $n + j \in B_{v_j}$, where v_j is the j th entry of w . The n th entry \tilde{x}_n of the smoothed sequence is then gotten by averaging the vectors y_ν over the indices ν in the bin B_w^* containing n . (Note that the sets B_n are, in general, not the same as the sets A_n used in Smoothing Algorithm D above.)

How should one choose the window size κ_m and the parameter δ ? For a time series of length $m = 10^6$, the condition (9) suggests that κ_m should be no larger than $\ln 10^6 \approx 14$. There are not many integers between 1 and 14, so it will usually be easy to run the algorithm for each possible choice, starting at $\kappa = 1$, and stopping when κ gets so large that the bins have fewer than five or ten (?) indices each. One should attempt to use $\delta \approx \Delta/5$, where Δ is a separation threshold, and hope that the noise values e_n are really smaller than this. For the separation threshold, one could use half the diameter of the attractor, because if x_n, x'_n are orbits with x_0, x'_0 independently chosen from a mixing, invariant probability measure with support Λ , then

$$(14) \quad \lim_{m \rightarrow \infty} \max_{0 \leq n \leq m} |x_n - x'_n| = \text{diam}(\Lambda)$$

with probability 1. If the noise values e_n really are smaller than $\text{diam}(\Lambda)/10$, then for any x_0 chosen from an ergodic, invariant probability measure with support Λ ,

$$(15) \quad \frac{5}{6} \lim_{m \rightarrow \infty} \max_{0 \leq n, \nu \leq m} |y_n - y_\nu| \leq \text{diam}(\Lambda)$$

with probability 1, by the ergodic theorem. This suggests that one might choose δ to be $5/12$ times the diameter of the point set $\{y_n\}_{0 \leq n \leq m}$.

3.1.4. *Consequences for Axiom A attractors.* By the ergodic theorem, it is almost surely the case that the empirical distribution of the points x_1, x_2, \dots, x_m converges weakly to the Gibbs state μ_* . Therefore, by Theorem 2, the empirical distribution of the points $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_m$ converges weakly to μ_* . Since F is continuous and the support of μ_* is dense in Λ , the set of ordered pairs $(\hat{x}_n, \hat{x}_{n+1})$, where $\kappa_m < n < m - \kappa_m$, converges in Hausdorff metric to the graph of $F| \Lambda$. Thus, one can in effect reconstruct the basic set Λ and the mapping $F| \Lambda$. Various other quantities of dynamical importance, for example, certain measures of fractal dimension of Λ , may also be consistently estimated using the filtered sequence \hat{x}_n . We shall leave a thorough discussion of such estimation problems to a subsequent paper.

3.1.5. *Second stage smoothing.* There is, obviously, a bias-variance trade-off in the choice of the sequence κ_m used in the smoothing algorithm of Section 3.1.1. Decreasing the rate of growth of κ_m increases the number of points in A_n and therefore decreases the variance of the average (11);

however, the values of x_ν included in the average will then tend to be farther from x_n , therefore increasing the bias. But there is an even larger impediment to the accuracy of the algorithm that derives from the fact that the rate of orbital separation depends on the direction of the difference between initial points. In particular, the dynamical distance between orbit segments $\{F^\nu(x)\}_\nu$ and $\{F^\nu(x')\}_\nu$ will tend to be smaller when $x' - x$ points approximately in a “Lyapunov direction” corresponding to a smaller Lyapunov exponent. Thus, for most n it will be the case that the points $\{x_\nu\}_{\nu \in A_n}$ will lie in a (very) thin ellipsoid and that many ν for which $|x_\nu - x_n|$ is relatively small will be excluded from A_n .

This peculiarity might, in principle, be exploited to obtain more accurate estimates of the points x_n . Fix $\beta \in (0, 1)$, and for each $1 \leq n \leq m$ let B_n be the set consisting of those m^β integers $\nu \in [1, m]$ for which $|\hat{x}_\nu - \hat{x}_n|$ is smallest. Define

$$(16) \quad \tilde{x}_n = m^{-\beta} \sum_{\nu \in B_n} y_\nu.$$

We conjecture that, with a suitable choice of β , use of this second-stage filter might considerably improve the accuracy of estimation of x_n .

3.2. *Gaussian noise.* Hypothesis 1 is quite a bit more stringent than one would like. However, if the errors are unbounded, even Gaussian, then it is impossible to consistently reconstruct the signal x_n , or even a part of it, from a long stretch of the time series y_n . In fact, it is impossible to infer even a single value x_0 of the signal from the entire two-sided time series $y_n = x_n + e_n$.

HYPOTHESIS 2. *Conditional on the sequence $\{x_n\}$ (equivalently, conditional on x_0) the random vectors e_n are independent and Gaussian with mean-vector 0 and nonsingular covariance matrix Σ .*

THEOREM 3. *Assume that x_0 is chosen at random from a Gibbs state μ_* supported by an Axiom A basic set Λ . If the errors e_n satisfy Hypothesis 2 then there is no measurable function $\xi_* = \xi_*(\{y_n\}_{n \in \mathbb{Z}})$ such that*

$$(17) \quad x_0 = \xi_* \quad \text{with probability 1.}$$

The proof, which will be given in Section 4, will show that orbit reconstruction is impossible even if the macroscopic features of the dynamics (the diffeomorphism F , the attractor Λ and the SRB measure μ_*) are known a priori. Furthermore, it should be clear from the proof that the result extends to a large class of error distributions. We shall refrain, however, from trying to state and prove an extremely general form of the result.

Although it is not possible to consistently recover the signal $\{x_n\}$ from the time series y_n when the noise e_n is Gaussian, it is nevertheless possible to consistently estimate important features of the dynamics provided the covariance matrix Σ is known. In particular, Birkhoff’s ergodic theorem implies

that for every polynomial $g(x)$ in d variables,

$$(18) \quad \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{i=1}^m g(y_i) = \int_{\mathbb{R}^d} \int_{\Lambda} g(\xi + \zeta) d\mu_*(\xi) \varphi_{0, \Sigma}(\zeta) d\zeta,$$

where $\varphi_{0, \Sigma}$ is the Gaussian density with parameters $0, \Sigma$. This implies that the moments of μ_* can be consistently estimated; since μ_* has compact support, it is determined by its moments, and so μ_* can be consistently estimated (in the weak topology). Similarly, the joint distribution of $(X, F(X))$, where $X \sim \mu_*$, may also be consistently estimated. Since the support of this latter distribution is the graph of $F|_{\Lambda}$, this too may be recovered.

Unfortunately, proving the existence of consistent estimators is not the same as the construction of good or useful estimators. The substantial problem of inference about the dynamics of F from time series data $y_n = x_n + e_n$ when the noise e_n is Gaussian will be left to another paper.

4. Proof of Theorem 3.

PROOF. The proof that there is no such ξ_* uses the existence of homoclinic pairs; see Section A.4 in the Appendix. By Proposition 2 of the Appendix, on some probability space are defined random vectors x_0 and x'_0 , each with marginal distribution μ_* , such that (a) with positive probability, $x'_0 \neq x_0$ and (b) with probability 1, x_0 and x'_0 constitute a homoclinic pair; that is, for some $\alpha > 0$,

$$(19) \quad \lim_{|n| \rightarrow \infty} (1 + \alpha)^{|n|} |x_n - x'_n| = 0,$$

where $x_n = F^n(x_0)$ and $x'_n = F^n(x'_0)$. We may assume that the probability space also accommodates a sequence e_n of Gaussian random vectors that are jointly independent of x_0 and x'_0 . Define $y_n = x_n + e_n$ and $y'_n = x'_n + e_n$; then conditional on the values of x_0 and x'_0 , the sequences $\mathbf{y} = \{y_n\}_{n \in \mathbb{Z}}$ and $\mathbf{y}' = \{y'_n\}_{n \in \mathbb{Z}}$ have Gaussian distributions with the same autocovariance and mean vector sequences $\{x_n\}_{n \in \mathbb{Z}}$, $\{x'_n\}_{n \in \mathbb{Z}}$ satisfying (19). Since (19) implies that $\sum_{n \in \mathbb{Z}} |x_n - x'_n|^2 < \infty$, a theorem of Kakutani (see, for instance, [11], Section II.2, Theorem 2.1 and Example 3) implies that the conditional distributions of the sequences \mathbf{y} and \mathbf{y}' , given x_0 and x'_0 , are mutually absolutely continuous. Consequently, for any Borel measurable function $\xi_* : (\mathbb{R}^d)^{\mathbb{Z}} \rightarrow \mathbb{R}^d$, the conditional distributions of the random vectors $\xi_*(\mathbf{y})$ and $\xi_*(\mathbf{y}')$, given x_0 and x'_0 , are also mutually absolutely continuous. If there were a function $\xi_* = \xi_*(\mathbf{y})$ such that $x_0 = \xi_*(\mathbf{y})$ almost surely, then it would also be the case that $x'_0 = \xi_*(\mathbf{y}')$ almost surely, and so the mutual absolute continuity of the conditional distributions would then imply that $x'_0 = x_0$ almost surely, a contradiction. \square

5. Proof of Theorem 1. In essence, the proof of Theorem 1 consists of showing (1) that the sets A_n are large (so that averaging over A_n will get rid of the errors); (2) that the sets A_n contain only indices ν such that $|x_n - x_\nu|$

is small and (3) that although the sets A_n and the error random vectors e_ν are not a priori independent (since the sets A_n are defined using the values y_ν), the dependence may be circumvented in the averaging. It is only for task (2) that hyperbolicity of the invariant set Λ is needed.

LEMMA 1. *There exists a constant $C > 0$ such that if $\nu \in A_n$ then*

$$(20) \quad |x_n - x_\nu| \leq \exp\{-\kappa_m/C\}.$$

PROOF. This is a consequence of the orbit separation property, which in turn follows from the hyperbolicity of Λ . By hypothesis, $5\delta < \Delta$, where Δ is a separation threshold for the attractor [see inequality (6)], and by Hypothesis 1, $|e_n| < \delta$. Consequently, if $\nu \in A_n$ [i.e., if inequality (10) holds], then

$$\max_{|j| \leq \kappa_m} |x_{n+j} - x_{\nu+j}| < 5\delta < \Delta.$$

But this cannot hold unless (20) is true, by the orbit separation property (6) [the constant C being the same as the constant C in (6)]. Thus, $\nu \in A_n$ implies (20). \square

LEMMA 2. *For every $\varepsilon > 0$,*

$$(21) \quad \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{n=0}^m \mathbf{1}\{|A_n| \leq m^{1-\varepsilon}\} = 0.$$

PROOF. This follows from the hypothesis (9) that $\kappa_m = o(\log m)$ as $m \rightarrow \infty$, by a routine counting argument. Since Λ is compact, it has a finite subset B that is $\delta/2$ -dense. Since $\kappa_m = o(\log m)$, the cardinality N_m of the set $B^{2\kappa_m+1}$ of length- $(2\kappa_m + 1)$ sequences with entries in B satisfies

$$(22) \quad N_m = o(m^\varepsilon) \quad \text{as } m \rightarrow \infty$$

for every $\varepsilon > 0$. If B is $\delta/2$ -dense in Λ , then for every $x \in \Lambda$, there is at least one element $\xi = (\xi_0, \xi_1, \dots, \xi_{2\kappa_m})$ of $B^{2\kappa_m+1}$ that $\delta/2$ -shadows the orbit segment $\{F^n(x)\}_{-\kappa_m \leq n \leq \kappa_m}$, that is, such that

$$(23) \quad |F^n(x) - \xi_{n+\kappa_m}| < \delta/2 \quad \forall |n| \leq \kappa_m.$$

For each $\xi \in B^{2\kappa_m+1}$, define $B_m(\xi)$ to be the set of all indices $\nu \in \{0, 1, 2, \dots, m\}$ such that (23) holds with $x = x_\nu$. Every index ν is contained in at least one of the sets $B_m(\xi)$. If two indices n, ν both lie in the same set $B_m(\xi)$, then by (23) and the triangle inequality, $|x_{n+j} - x_{\nu+j}| < \delta$ and hence $|y_{n+j} - y_{\nu+j}| < 3\delta$ for all $|j| \leq \kappa_m$; thus, $\nu \in A_n$. Therefore, to prove (21) it suffices to show that for large m most of the indices ν lie in sets $B_m(\xi)$ with at least $m^{1-\varepsilon}$ elements. However, by (22), the number of sets $B_m(\xi)$ is $o(m^\varepsilon)$. Consequently, the number of indices ν that are contained in sets $B_m(\xi)$ of cardinality less than $m^{1-\varepsilon}$ cannot be larger than $o(m^\varepsilon)m^{1-\varepsilon} = o(m)$. \square

PROOF OF THEOREM 1. The estimate \hat{x}_n is obtained by averaging the vectors y_ν over the indices $\nu \in A_n$ (11). Since $y_\nu = x_\nu + e_\nu$, we have

$$(24) \quad \hat{x}_n = x_n + \frac{1}{|A_n|} \sum_{\nu \in A_n} e_\nu + \frac{1}{|A_n|} \sum_{\nu \in A_n} (x_\nu - x_n).$$

Lemma 1 implies that the latter average converges to zero uniformly for $\kappa_m < n < m - \kappa_m$ as $m \rightarrow \infty$. Thus, it suffices to show that for most of the indices n the average of the errors e_ν for $\nu \in A_n$ is small, with probability approaching 1 as $m \rightarrow \infty$. If the random vectors e_ν were independent of the index sets A_n then in view of Lemma 2, the result would follow immediately from the Chebyshev inequality. However, the random vectors e_ν are *not* independent of the index sets A_n ; thus, some delicacy is required.

For each index n , define A_n^* to be the set of all indices ν such that $\nu \in A_n$ and $|n - \nu| \leq 2\kappa_m$; note that $|A_n^*|$ is no larger than $4\kappa_m + 1 = o(\log m)$, so on the event that $|A_n| > m^{3/4}$ the indices $\nu \in A_n^*$ have a negligible effect on the average $\sum_{\nu \in A_n} e_\nu / |A_n|$. For each index n and each integer $i \in [1, 2\kappa_m + 1]$, define A_n^i to be the set of all indices $\nu \notin A_n^*$ such that $\nu \in A_n$ and $\nu \equiv i \pmod{2\kappa_m + 1}$. Obviously, the sets $A_n^*, A_n^1, A_n^2, \dots, A_n^{\kappa_m}$ are pairwise disjoint, and

$$(25) \quad A_n = A_n^* \cup \left(\bigcup_{i=1}^{2\kappa_m+1} A_n^i \right).$$

For each integer $i \in [1, 2\kappa_m + 1]$, the set A_n^i is independent of the collection of random vectors $\{e_\nu\}$ indexed by integers $\nu \equiv i \pmod{2\kappa_m + 1}$. To see this, consider an integer $\nu \equiv i \pmod{2\kappa_m + 1}$. The event $\nu \in A_n^i$ is completely determined by the values of y_{n+j} and $y_{\nu+j}$ for $|j| \leq \kappa_m$; furthermore, no other event $\nu' \in A_n^i$, where $\nu' \neq \nu$, is influenced by the values of $y_{\nu+j}$ for $|j| \leq \kappa_m$ (this is the point of partitioning the indices ν into blocks of size $2\kappa_m + 1$). Moreover, the event $\nu \in A_n^i$ is not affected by the value of e_ν , because if $|y_{\nu+j} - y_{n+j}| < 3\delta$ for all $1 \leq |j| \leq \kappa_m$, then by the same argument as in the proof of Lemma 1, $|x_n - x_\nu| < \delta/2$ (provided m is large) and so $|y_n - y_\nu| < 3\delta$ regardless of the values of e_n and e_ν . Thus, the composition of the set A_n^i can be determined without reference to the values of the random vectors $\{e_\nu\}$ indexed by integers $\nu \equiv i \pmod{2\kappa_m + 1}$.

For each index n , the sets A_n^i may be partitioned as $\mathcal{S} \cup \mathcal{T}$, where \mathcal{S} consists of the special index $*$ and those indices i for which $|A_n^i| < \sqrt{m}$, and \mathcal{T} consists of the remaining indices. For each $i \in \mathcal{T}$, Chebyshev's inequality implies that for any $\varepsilon > 0$,

$$(26) \quad P \left(\left| \left[\sum_{\nu \in A_n^i} e_\nu \right] / |A_n^i| \right| > \varepsilon \middle| A_n^i \right) \leq \delta^2 / |A_n^i| \varepsilon^2 \leq \delta^2 / (\sqrt{m} \varepsilon^2),$$

since the random vectors e_ν indexed by $\nu \in A_n^i$ are independent of A_n^i , by the preceding paragraph. Since there are no more than $4\kappa_m + 2$ elements of \mathcal{T} ,

and $|A_n^i| < \sqrt{m}$ for each $i \in \mathcal{I}$,

$$(27) \quad \left| \sum_{i \in \mathcal{I}} \sum_{\nu \in A_n^i} e_\nu \right| \leq (4\kappa_m + 2)\sqrt{m} \delta.$$

Consequently, if $|A_n| \geq m^{3/4}$ and m is sufficiently large that $(4\kappa_m + 2)/m^{1/4} < \varepsilon/\delta$ then the event $|\sum_{\nu \in A_n} e_\nu|/|A_n| > 2\varepsilon$ is contained in the union over $i \in \mathcal{I}$ of the events $|\sum_{\nu \in A_n^i} e_\nu| > \varepsilon|A_n^i|$. It therefore follows from inequality (26) that

$$P\left(\left[\left|\sum_{\nu \in A_n} e_\nu\right|/|A_n| > 2\varepsilon \mid |A_n| \geq m^{3/4}\right]\right) \leq (2\kappa_m + 1)\delta^2/(\sqrt{m}\varepsilon^2).$$

Together with Lemma 2, this implies that

$$(28) \quad \sum_{n=0}^m P\left(\left[\left|\sum_{\nu \in A_n} e_\nu\right|/|A_n| > 2\varepsilon\right]\right) = o(m),$$

which, in view of Lemma 1, proves (12). \square

6. Proof of Theorem 2. The proof of Theorem 2 differs from that of Theorem 1 in two respects: (1) Lemma 2 must be replaced by the stronger statement that the cardinality of A_n is large for every index n between κ_m and $m - \kappa_m$, and (2) Chebyshev’s inequality must be replaced by an exponential large deviations probability inequality. The latter change is relatively minor; the former, however, requires hard results from the ergodic theory of Gibbs states on Axiom A basic sets. See the Appendix for a resume of the most important definitions and facts, and [2] for a detailed exposition of the theory.

Assume that Λ is an Axiom A basic set for F , that μ_* is a Gibbs state for F supported by Λ (see Section A.3 for the definition and basic properties), and that the initial point x_0 of the orbit x_n is distributed in Λ according to μ_* .

LEMMA 3. For every $\varepsilon > 0$, all sufficiently large m and all integers $n \in (\kappa_m, m - \kappa_m)$,

$$(29) \quad P(|A_n| \leq m^{1-4\varepsilon}) \leq \exp\{-m^\varepsilon\}.$$

PROOF. The basic set Λ admits a Markov partition \mathcal{M} of diameter less than δ (see Section A.2 below). Let $z_0, z'_0 \in \Lambda$ be points with orbits $z_j = F^j(z_0)$ and $z'_j = F^j(z'_0)$ and itineraries $\{i_j\}$ and $\{i'_j\}$ (relative to the Markov partition \mathcal{M}), respectively. If $i_j = i'_j$ for all $|j| \leq \kappa_m$, then $|z_j - z'_j| < \delta$ for all $|j| \leq \kappa_m$, since the diameters of the sets G_i of \mathcal{M} are less than δ . Consequently, if x_n and x_ν are two points on the orbit of $x = x_0$ with itineraries $\{i_j\}, \{i'_j\}$ that coincide for $|j| \leq \kappa_m$, then $|y_{n+j} - y_{\nu+j}| < 3\delta$ for all $|j| \leq \kappa_m$, and so $\nu \in A_n$. Thus, to prove the inequality (29) it suffices to prove that for every finite itinerary $\mathbf{i} = \{i_j\}_{|j| \leq \kappa_m}$ of length $2\kappa_m + 1$, the probability that fewer than $m^{1-4\varepsilon}$ of the points $\{x_n\}_{1 \leq n \leq m}$ share the itinerary \mathbf{i} is smaller than $\exp\{-m^\varepsilon\}$.

Let I be the (doubly infinite) itinerary of a random point of Λ with distribution μ_* . Because μ_* is a Gibbs state, there exists a constant $\beta > 0$ and an integer L , both independent of m , such that the following is true [see inequalities (38) and (39) of the Appendix]: for any infinite itinerary \mathbf{i} and any finite itinerary \mathbf{i}^* of length $2\kappa_m + 1$,

$$(30) \quad P(I_{L+n} = i_n^* \forall 1 \leq n \leq 2\kappa_m + 1 \mid I_n = i_n \forall n \leq 0) \geq \beta^{2\kappa_m+1}.$$

Thus, if the random itinerary I is broken up into segments of length $L + 2\kappa_m + 1$, each segment will provide an opportunity for the letters \mathbf{i}^* to occur with success probability at least $\beta^{2\kappa_m+1}$. Hence, if $N(\mathbf{i}^*)$ is the number of times that the finite string \mathbf{i}^* occurs in the first m entries of I , then $N(\mathbf{i}^*)$ stochastically dominates the sum of $k = \lfloor m/(L + 2\kappa_m + 1) \rfloor$ i.i.d. Bernoulli random variables with success parameter $\beta^{2\kappa_m+1}$. Since $\kappa_m = o(\log m)$, for sufficiently large m , this success probability is, for any $\varepsilon > 0$, eventually larger than $m^{-\varepsilon}$, and furthermore $k \geq m^{1-\varepsilon}$. It follows that the expectation of the sum is larger than $m^{1-2\varepsilon}$. Consequently, by a very crude probability inequality for sums of independent Bernoulli random variables,

$$(31) \quad P\{N(\mathbf{i}^*) \leq m^{1-4\varepsilon}\} \leq \exp\{-m^\varepsilon\}. \quad \square$$

LEMMA 4. *With probability 1,*

$$(32) \quad \lim_{m \rightarrow \infty} \max_{\kappa_m < n < m - \kappa_m} \frac{1}{|A_n|} \sum_{\nu \in A_n} e_\nu = 0.$$

PROOF. The proof will use the following standard large deviations probability estimate for sums of independent random variables: if ξ_1, ξ_2, \dots are independent random variables (or vectors) uniformly bounded by a constant $\delta < \infty$ and if $E\xi_j = 0$ for every j , then for every $\eta > 0$ there exists $\gamma = \gamma(\eta, \delta) > 0$ such that for all sufficiently large n ,

$$(33) \quad P\left\{\frac{1}{n} \left| \sum_{j=1}^n \xi_j \right| \geq \eta\right\} \leq \exp\{-n\gamma\}.$$

As in the proof of Theorem 1, the set A_n may be decomposed as the disjoint union of the sets A_n^* and A_n^i ; see (25). Recall that for each i the set A_n^i is independent of the collection of random vectors $\{e_\nu\}$ indexed by integers $\nu \equiv i \pmod{2\kappa_m + 1}$. Recall also that the indices $*, i$ may be partitioned as $\mathcal{S} \cup \mathcal{T}$, where \mathcal{S} consists of the special index $*$ and those indices i for which $|A_n^i| < \sqrt{m}$, and \mathcal{T} consists of the remaining indices. For each $i \in \mathcal{T}$, the large deviations inequality (33) implies that for any $\varepsilon > 0$ and all sufficiently large m ,

$$(34) \quad P\left(\frac{1}{|A_n^i|} \left| \sum_{\nu \in A_n^i} e_\nu \right| > \varepsilon \mid A_n^i\right) \leq \exp\{-\gamma|A_n^i|\} \leq \exp\{-\gamma\sqrt{m}\}$$

for a constant $\gamma > 0$ depending on ε and δ but not on m . Now for sufficiently large m ,

$$(35) \quad \left\{ \left| \sum_{\nu \in A_n} e_\nu \right| > 2\varepsilon |A_n| \right\} \subset \{|A_n| \leq m^{3/4}\} \cup \left(\bigcup_{i \in \mathcal{I}} \left\{ \left| \sum_{\nu \in A_n^i} e_\nu \right| > \varepsilon |A_n^i| \right\} \right).$$

Consequently, by Lemma 3 and inequality (34), for all large m and $\kappa_m < n < m - \kappa_m$,

$$(36) \quad P \left(\frac{1}{|A_n|} \left| \sum_{\nu \in A_n^i} e_\nu \right| > 2\varepsilon \right) \leq (2\kappa_m + 1) \exp\{-\gamma\sqrt{m}\} + \exp\{-m^{1/16}\}.$$

Since the series $\sum_m m e^{-am^\alpha}$ is summable for any values of $a > 0$ and $\alpha > 0$, the result (32) follows from the Borel–Cantelli lemma. \square

Theorem 2 follows immediately from Lemmas 3 and 4.

APPENDIX

Markov partitions for Axiom A basic sets.

A.1. *Example: Smale’s solenoid.* In this example there is a simple Markov partition, and the resulting “symbolic dynamics” is relatively transparent. Partition the attractor Λ (or its basin of attraction Ω) into two sets,

$$G_0 = \{(\theta, z) : 0 \leq \theta \leq \pi\},$$

$$G_1 = \{(\theta, z) : \pi \leq \theta \leq 2\pi\}.$$

(This isn’t really a partition in the usual sense of the word, since the sets have nonempty intersection, nor would Markov understand why his name is attached, but it is called a Markov partition anyway.) For any point $x \in \Lambda$, define an itinerary of x to be a doubly infinite sequence $\mathbf{i} = \{i_n\}_{n \in \mathbb{Z}}$ of 0’s and 1’s such that $F^n(x) \in G_{i_n}$ for each integer n . Observe that if \mathbf{i} is an itinerary of $x = (\theta, z)$ then $i_0 i_1 i_2 \dots$ is a binary expansion of $\theta/2\pi$; moreover, if $x \in \Lambda_\beta$ for some particular cross-sectional slice Λ_β then the value of i_{-1} indicates which of the two “first generation” circles (see Figure 2) contains x , and $i_{-n} i_{-n+1} \dots i_{-1}$ determines which of the 2^n “ n th generation” circles contains x . With this in mind, it is not difficult to see that (1) every infinite sequence of 0’s and 1’s is an itinerary of a unique $x \in \Lambda$, and (2) for μ_* -a.e. x there is only one itinerary. The projection from sequence space to Λ (semi)conjugates the forward shift operator on sequence space to the solenoid mapping F_α . (In fact, Smale invented the solenoid mapping for just this reason.) See [4], Chapter 2, for further details concerning this example.

A.2. *Markov partitions and symbolic dynamics.* Every Axiom A basic set admits *Markov partitions* of arbitrarily small diameter, but in general neither the partitions nor their construction are simply described. See [2],

Chapter 3, or [16], Chapter 10 for the precise definition and construction. A Markov partition \mathcal{M} consists of finitely many closed sets G_1, G_2, \dots, G_r whose union contains Λ , and such that for μ_* -a.e. z_0 , every point $z_n = F^n(z_0)$ in the orbit of z_0 lies in only one of the sets G_i . The diameter of the partition is the maximum of the diameters of its constituent sets. For any point $z_0 \in \Lambda$, define an itinerary of z_0 to be a two-sided sequence $\dots i_{-1}i_0i_1\dots$ such that for each $n, z_n \in G_{i_n}$; note that for μ_* -a.e. z_0 , there is only one itinerary. If the diameter of \mathcal{M} is sufficiently small, then no two distinct points $x, x' \in \Lambda$ may share the same itinerary, since this would entail a violation of the orbit separation property mentioned in Section 2.3.

Let Σ be the space of all doubly infinite itineraries, and let σ be the forward shift operator on Σ . Since distinct points of Λ may not share the same itinerary, there is a projection $\pi: \Sigma \rightarrow \Lambda$ that maps each itinerary \mathbf{i} to the unique point $x \in \Lambda$ with itinerary \mathbf{i} . It is not difficult to see that π is continuous (and even Hölder continuous with respect to the appropriate metric on Σ ; see [2] or [16]). Clearly, $F \circ \pi = \pi \circ \sigma$, and so σ is a homeomorphism of Σ , since F is a homeomorphism of Λ . Not every sequence \mathbf{i} need be an element of Σ ; however, the Markov property of the partition \mathcal{M} implies that the space Σ of all doubly infinite itineraries, together with the forward shift operator σ , is a topologically mixing shift of finite type (see [2], Lemma 1.3 and Proposition 3.19). A shift (Σ, σ) is of *finite type* if there exists a finite set \mathcal{F} of finite words from the alphabet $\mathcal{A} = \{1, 2, \dots, r\}$ such that for any doubly infinite sequence \mathbf{i} with entries in \mathcal{A} , \mathbf{i} is an element of Σ if and only if \mathbf{i} contains none of the words in \mathcal{F} . In general, a shift (Σ, σ) of finite type is topologically mixing if there exists an integer $M < \infty$ such that for every pair $\omega, \omega' \in \Sigma$ there exists a finite word w of length M such that the concatenation

$$\dots \omega_{-2}\omega_{-1}\omega_0w_1w_2 \dots w_M\omega'_1\omega'_2 \dots$$

is an element of Σ . For shifts constructed from Markov partitions for topologically mixing Axiom A basic sets, $M = 1$.

A.3. *Gibbs states.* A Gibbs state μ_* on Λ is defined to be an invariant probability measure whose pullback to a shift-invariant probability measure $\bar{\mu}_*$ on the sequence space Σ has the Gibbs property described in [2], Chapter 1 (see [2], Chapter 4, for the proof). In particular, $\bar{\mu}_*$ must satisfy a system of inequalities,

$$(37) \quad C_1 \leq \frac{\bar{\mu}_*\{\mathbf{w} \in \Sigma: w_j = i_j \forall 0 \leq j \leq n\}}{\exp\{-\lambda n + \sum_{j=0}^n \varphi(\sigma^j \mathbf{i})\}} \leq C_2,$$

valid for all itineraries \mathbf{i} and all integers $n \geq 0$, for constants $0 < C_1 < C_2 < \infty$ independent of n and of the itinerary \mathbf{i} . Here φ is a real-valued, Hölder continuous function on the space of all doubly infinite sequences \mathbf{i} , σ is the forward shift operator and $\lambda \in \mathbb{R}$ is a constant called the *pressure*. See [2], Section 1.4, for details. Note that (37) implies that there exists a constant

$\beta > 0$ such that for any finite itinerary $i_1 i_2 \cdots i_n$,

$$(38) \quad \bar{\mu}_* \{ \mathbf{w} \in \Sigma : w_j = i_j \forall 1 \leq j \leq n \} \geq \beta^n.$$

The SRB measure μ_* for an Axiom A attractor is a Gibbs state (see [2], Chapter 4 for a proof). For Smale’s solenoid mapping F_α , the measure $\bar{\mu}_*$ is the product Bernoulli-1/2 measure, that is, the measure that makes the coordinate random variables i.i.d. Bernoulli-1/2. In general, Gibbs states enjoy very strong mixing properties, among which the following, concerning the conditional distribution of the future given the past, is perhaps the most useful.

PROPOSITION 1. *There exist constants $\rho_k > 0$ satisfying $\lim_{k \rightarrow \infty} \rho_k = 1$ and such that for every infinite itinerary $\mathbf{i} = \cdots i_{-1} i_0 i_1 \cdots \in \Sigma$ and every finite itinerary $\mathbf{i}^* = i_1^* i_2^* \cdots i_n^*$ (of any positive length),*

$$(39) \quad \begin{aligned} \bar{\mu}_*(w_{j+M+k} = i_j^* \forall 1 \leq j \leq n \mid w_j = i_j \forall j \leq 0) \\ \geq \rho_k \bar{\mu}_* \{ w_j = i_j^* \forall 1 \leq j \leq n \}, \end{aligned}$$

where $M \geq 1$ is the integer in the definition of topological mixing.

See [13] for a proof.

Equations (38) and (39) have the following consequence: there is a constant $\beta > 0$ such that for any finite itinerary $\mathbf{i}^* = i_1^* i_2^* \cdots i_n^*$ (of any positive length), the conditional probability, given the past, that the next $M + n$ steps of the itinerary will end in $i_1^* i_2^* \cdots i_n^*$ is at least β^n .

A.4. *Homoclinic pairs.* One of the important features of Axiom A (and, more generally, hyperbolic) systems is the existence of homoclinic pairs. Two distinct points x and x' are said to be a homoclinic pair if, for some $\varepsilon > 0$,

$$(40) \quad \lim_{|n| \rightarrow \infty} (1 + \varepsilon)^{|n|} |F^n(x) - F^n(x')| = 0;$$

in words, x, x' are distinct but their orbits approach each other exponentially fast both forwards and backwards in time. In Axiom A systems, homoclinic pairs are dense; in particular, for any points $\xi, \xi' \in \Lambda$ and any $\delta > 0$, there exists a homoclinic pair of points such that $|x - \xi| < \delta$ and $|x' - \xi'| < \delta$.

This may be proved using the existence of Markov partitions of small diameter. Let \mathbf{i} and \mathbf{i}' be itineraries of ξ and ξ' , respectively. By the separation of orbits property, there exists an integer k such that if the itinerary \mathbf{i}'' of a point $x \in \Lambda$ satisfies $i_j'' = i_j$ for all $|j| \leq k$, then $|x - \xi| < \delta$, and similarly, if $i_j'' = i_j'$ for all $|j| \leq k$, then $|x - \xi'| < \delta$. But topological mixing (see Section A.2) guarantees that itineraries may be spliced together to obtain itineraries \mathbf{i}^* and \mathbf{i}^{**} so that (1) $i_j^* = i_j$ for all $|j| \leq k$; (2) $i_j^{**} = i_j'$ for all $|j| \leq k$ and (3) $i_j^* = i_j^{**}$ for all $|j| > M + k$. If x and x' have itineraries \mathbf{i}^* and \mathbf{i}^{**} , respectively, then $|x - \xi| < \delta$ and $|x' - \xi'| < \delta$, by (1) and (2), and x, x' are a homoclinic pair, by (3) and the orbit separation property.

The foregoing argument may be adapted to prove the following proposition, which is the key to Theorem 3.

PROPOSITION 2. *On some probability space there exist random vectors X', X'' valued in Λ such that:*

- (a) *Each of X' and X'' has marginal distribution μ_* .*
- (b) *With probability 1, X' and X'' are a homoclinic pair.*
- (c) *With positive probability, $X' \neq X''$.*

PROOF. The probability space should be large enough to accommodate a random vector X with distribution μ_* and several independent uniform-(0, 1) random variables. Let $\mathbf{I} = \cdots I_{-1}I_0I_1 \cdots$ be the itinerary of X . Construct new itineraries $\mathbf{I}', \mathbf{I}''$ as follows: for some large integer k , set $I'_j = I''_j = I_j$ for all $|j| > k$ and choose the random vectors (I'_{-k}, \dots, I'_k) and (I''_{-k}, \dots, I''_k) independently from the conditional distribution of (I_{-k}, \dots, I_k) given $\{I_j\}_{|j| > k}$. (This is possible if the underlying probability space supports uniform random variables independent of \mathbf{I} .) By construction, each of \mathbf{I}' and \mathbf{I}'' will be an itinerary. Define X' and X'' to be the unique points with itineraries \mathbf{I}' and \mathbf{I}'' , respectively. Clearly, each of X' and X'' has the same marginal distribution as X . Moreover, since the itineraries of X' and X'' coincide except in finitely many entries, X' and X'' must be a homoclinic pair. Finally, Proposition 1 implies that if k is large, then the joint distribution of (X', X'') approximates the product measure $\mu_* \times \mu_*$. Since under $\mu_* \times \mu_*$ there is positive probability that the coordinates are not equal, the same is true for the joint distribution of (X', X'') . \square

REFERENCES

- [1] ABARBANEL, H. (1996). *Analysis of Observed Chaotic Data*. Springer, New York.
- [2] BOWEN, R. (1975). *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms. Lecture Notes in Math.* **470**. Springer, Berlin.
- [3] DAVIES, M. (1992). Noise reduction schemes for chaotic time series. *Phys. D* **79** 174–192.
- [4] DEVANEY, R. (1986). *An Introduction to Chaotic Dynamical Systems*. Benjamin and Cummings, Menlo Park, CA.
- [5] ECKMANN, J.-P. and RUELLE, D. (1985). Ergodic theory of chaos and strange attractors. *Rev. Modern Phys.* **57** 617–656.
- [6] ECKMANN, J.-P., KAMPHORST, S. O., RUELLE, D. and CILIBERTO, S. (1986). Lyapunov exponents from times series. *Phys. Rev. A* **34** 4971–4979.
- [7] GUCKENHEIMER, J. and HOLMES, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer, New York.
- [8] HIRSCH, M. (1976). *Differential Topology*. Springer, New York.
- [9] KOSTELICH, E. and SCHREIBER, T. (1993). Noise reduction in chaotic time-series data: a survey of common methods. *Phys. Rev. E* **48** 1752–1763.
- [10] KOSTELICH, E. and YORKE, J. (1990). Noise reduction—finding the simplest dynamical system consistent with the data. *Phys. D* **41** 183–196.
- [11] KUO, H. H. (1975). *Gaussian Measures in Banach Spaces. Lecture Notes in Math.* **463**. Springer, Berlin.
- [12] LALLEY, S. (1997). More noise, more chaos. Unpublished manuscript.

- [13] LEDRAPPIER, F. (1976). Sur la condition de Bernoulli faible et ses applications. *Théorie Ergodique. Lecture Notes in Math.* **532** 152–159. Springer, Berlin.
- [14] MOSER, J. (1973). *Stable and Random Motions in Dynamical Systems*. Princeton Univ. Press.
- [15] SAUER, T. (1992). A noise reduction method for signals from nonlinear systems. *Phys. D* **58** 193–201.
- [16] SHUB, M. (1987). *Global Stability of Dynamical Systems*. Springer, New York.
- [17] SMALE, S. (1963). Diffeomorphisms with many periodic points. In *Differential and Combinatorial Topology* (S. S. Cairns, ed.) Princeton Univ. Press.

DEPARTMENT OF STATISTICS
PURDUE UNIVERSITY
WEST LAFAYETTE, INDIANA 47907
E-MAIL: lalley@stat.purdue.edu