L. Douglas Kiel and Euel Elliott, Editors
http://www.press.umich.edu/titleDetailDesc.do?id=14623
The University of Michigan Press, 1996

CHAPTER 3

Measuring Chaos Using the Lyapunov Exponent

Thad A. Brown

At some point in time, all scientific things must be measured and calibrated—even chaos. The Lyapunov characteristic exponent, λ, is the clearest measure to prove the existence and to quantify chaos in a dynamical system or time series. Chaos exists when λ is positive, and this indicates the system under investigation is sensitive to initial conditions. A dictionary would suggest chaos is a state of complete confusion, or one lacking any organization. In the field of nonlinear dynamics, however, the term seems to imply specific properties of turbulence in a nonlinear system. But what is chaos? How can we test for it in an empirical manner, and precisely, thereby separating it from other complicated yet fully predictable phenomena? As it turns out, the Lyapunov exponent, λ, is a powerful measurement tool for separating chaotic behavior from the stable and predictable. Before turning to a discussion of this measure, let us review some of the main features of chaos.

Chaos

Chaos exists when the long-term prediction of a system is impossible. Chaos occurs when the uncertainty of a system's initial state grows exponentially fast. In time-series language, chaos fits the criterion that the autocorrelation function of the time signal goes to zero in finite time. Because trajectories are unstable, errors of estimation of initial conditions or parameters, however small, can later accumulate into substantial errors. Thus, forecasts of future behavior based on the past become problematic as current memory of the past fails. This has both positive and negative effects on information about the underlying system.

Nonlinearity is a necessary, but not sufficient, condition for chaotic motion. The observed chaotic behavior is due neither to external noise nor to an infinite number of degrees of freedom. The source of irregularity is the nonlinear system's property of separating initially close trajectories exponentially fast. Forecasting long-term behavior of chaotic systems may be more akin to art than to statistics (Peitgen, Jürgens, and Saupe 1992). Does this imply that it is impossible to predict the long-time behavior of a political or economic system (or both)? Such systems are potentially chaotic because their initial conditions can be fixed only with finite accuracy. Error may explode.

Yet the long-time unpredictability exists only at the individual level (or, in the language of dynamics, at the level of individual trajectories). At the level of statistical properties of the time evolution (averaged over different trajectories, say as they evolve from different nearby initial conditions), very definite predictions are possible. Said differently, *when* or *where* chaos exists can only be known by statistical prediction. Trajectories will eventually move only on a small submanifold (chaotic or strange attractor to be described later) of the entire state space, with predictable visitation frequencies of the different parts of the attractor. Is it possible, then, to estimate the time up to a short-term prognosis of an individual trajectory? Yes. It is the coexistence of predictable and unpredictable that leads to the phrase *order in chaos* in the field of nonlinear dynamic systems.

What are the attributes of deterministic chaos? Systems of nonlinear equations with time dependencies are the easiest place to find chaos. They range from simple one-dimensional models such as the quadratic equation $x \rightarrow ax(1 - x)$ to more complex three- and four-dimensional equations (Hale and Koçak 1991). Chaotic systems exhibit perverse or complex motion in that the time series never settles down to simple cycles or equilibrium. Orbits are inconsistent, wandering, and erratic, mixing the properties of the time evolution. The key to chaotic behavior is that the equation or the system of equations of motion exhibits a sensitive dependence on initial conditions. Often a series of states can be seen when the value of a parameter of the model is changed and the system moves from one state to another.

In time-series data where the exact equation of motion can only be presumed, chaos is harder to find. Even more challenging is to infer the microscopic process that underlies a time series where chaos has been detected. To know the equations of motion in politics, economics, and social dynamics, we require just such discoveries of microscopic laws. We will turn to this discussion and point to some interesting approaches at the end of this chapter.

An Attractor

The motion of a dynamical system is represented by an orbit around an attractor. An attractor is a closed invariant set that essentially captures all orbits starting in its domain of attraction.[1] As a region of phase space, an attractor exerts an appeal to a system, a point, or a subset of points toward which any dynamical path will converge. When the behavioral history of a system is examined, the nature of the attractor is the core of inquiry. When a rat, voter, or business firm is dropped into a completely different environment, a period of erratic behavior is expected as adjustment takes place. This adjustment is transient behavior that gives way to observed behavior in some equilibrium state.

An equilibrium state is called an attractor. If the attractor is a single point, asymptotically it represents stability because all trajectories that begin in some neighborhood of phase space approach the attractor as time passes. Of course, there are other attractors besides equilibria. Limit cycles may exist. Limit cycles occur when the orbit is drawn toward a cyclic path, rather than toward a point fixed in space. Attractors can also be complex in that they may be quasi-periodic or strange. An example of a quasi-periodic attractor is an orbit on a torus, a doughnut generated by the cosines of a pair of incommensurate frequencies.

Chaos Characteristics

Chaos is measured by the characteristics of its dynamics: fractal structure of the attractor; metric K entropy; and the Lyapunov characteristic exponent, λ_i. We will briefly explore each.

Chaotic Attractor

A chaotic attractor is characterized by exponential divergence away from any point within the attractor (Ruelle and Takens 1971). The exponential diverging movement on a compact attractor means that a trajectory is constantly curling back on the attractor. The strange attractor causes the "flow" to contract the volume in some directions and stretch it in other directions. Remaining within bounded dominion, the volume and backfolding mixing process produces a chaotic motion of the trajectory at the strange attractor. It is this folding motion which produces the fractal structure attributed to chaotic motion.

The properties of a strange attractor are: (1) the attractor set is such that it cannot be decomposed into smaller closed invariant sets (Ruelle 1980, 1989; Guckenheimer and Holmes 1983; Stanišić 1988). (2) There is a noticeable regularity in the physical structure or dimension of the attractor. A point, of course, has no dimension; a line, one dimension; a surface, two dimensions; and so on. By contrast, a strange attractor has a noninteger dimension. A strange attractor usually violates Euclidean geometry. After a long series of iterations, a fractal geometry stems from long-term stretching and folding of trajectories on a manifold. Trajectories on the attractor end up like a Cantor set (the mathematical equivalent of a croissant).

The fractal dimension has a fantastic interpretation. It delimits a minimum number of degrees of freedom necessary to describe chaotic behavior.

The lower the dimensionality, the fewer independent variables are needed to describe the chaotic behavior. Of course, in a chaotic system the dimension of the attractor is less than the dimension of the space within which it is embedded (see Feder 1988).

And, finally, (3) the attractor is sensitive to initial conditions. Sensitivity refers to the amplification of any arbitrarily small interval of values by iteration. Sensitivity to initial conditions is the nub of chaos theory, and indeed may be the point where chaos theory touches science. When nonlinear deterministic models—which are absolutely correct—are sensitive to initial conditions, they cannot be used to predict very far into the future. Lorenz (1963), in his famous experiments on weather patterns, found that the higher iterations appeared statistically independent of earlier values. The smallest error in the n^{th} decimal place becomes magnified and amplified under the numerical left-shift action of the iterative function.

We assume that the volume of an attractor is much smaller than the volume of phase space. In a dissipative system the irregular stretching in some directions and contraction in other directions produces a final motion (once transients are over) that may be unstable within the attractor. Points that are arbitrarily near at the initial time step become exponentially separated at the attractor for sufficiently long times. This leads to positive Kolmogorov entropy, the exponential separation that occurs in the direction of stretching.

A Fourier analysis of motion on a strange attractor reveals a continuous power spectrum. Under conditions of a linear model, a typical interpretation would be that there are an infinite number of Fourier modes or oscillators. Under the conditions of linear theory, such a reasoning has to take place in an infinite-dimension phase space. An interpretation can be either that the system is searching an infinite number of dimensions in phase space or that the system is evolving in a nonlinear fashion on a finite dimensional attractor. Recognition of nonlinearities has provoked a fundamental alteration from thinking about Fourier modes to thinking about the information dimension. The tool used to measure this information dimension is ergodic theory (Billingsley 1965).

K Entropy

A second measure of chaos is K entropy. K entropy is also referred to as the Kolmogorov-Sinai entropy (Farmer 1982; Schuster 1988). Entropy measures the rate of information production and the growth of uncertainty. In a system increasing in entropy, the number of possible system states that evolve from some initial distribution over time also increases. In a chaotic system, information about the system decreases over time when measured against the initial state. If subsequent measurements are not made, indeed, the observer knows less and less over time. With a second set of measures, however, we end up with far more information about the system, even about the initial state, than would have been possible at the beginning (Shaw 1981).

Consider that conditions where entropy increases are similar to a party system where voters were limited to two or three parties but suddenly are allowed to vote for twice as many. The disorder in the system increases because voters are no longer limited or confined to their prior political space. The increase in disorder is coupled with an increase in our ignorance about the state of the system. Before the sudden doubling of the number of parties we knew more about the position of voters.

Precisely then entropy, S, is proportional to $-\Sigma_i P_i \log P_i$, where P_i is the set of probabilities of finding the system in state $\{i\}$ and measures the information needed to locate the system in a specific state $\{i\}$. Thus S captures our ignorance about the system.

Lyapunov Characteristic Exponents

The Lyapunov characteristic exponents, λ_i, of dynamical systems measure the average rate by which the distance between close points becomes stretched or compressed after one iteration. Lyapunovs are generalized eigenvalues over an entire attractor in that they give the average rate of contraction or expansion of trajectories on an attractor (Wiggins 1990). The λ_i allows further resolution by allowing us to see the (orthogonal) direction of the growth or compression. For an attractor to be strange, it must have at least one positive Lyapunov exponent.

The λ_i is closely related to the other measures of chaos. Ruelle (1983, 1989) suggests that K entropy equals the sum of the positive λ_i, implying that positive entropy exists in the presence of chaos. The λ_i is linked to the information lost and gained during chaotic episodes (Ruelle 1980; Shaw 1981, respectively), and it is closely linked to the amount of information available for prediction. The λ_i was "conjectured" to be related to the fractal character of the attractor (Kaplan and Yorke 1979), and indeed this was proven true (with qualifications) for typical attractors (Russell, Hanson, and Ott 1980). The λ_i itself has a fractal, noncontinuous dimension, and that fractal quality is linked to the information available about the system (Kaplan and Yorke 1979; Young 1982; Ruelle 1989).

In a one-dimensional system there are three alternative variations of the Lyapunov.

$\lambda_1 < 0 \Rightarrow$ orbit is stable and periodic
$\lambda_1 = 0 \Rightarrow$ orbit is marginally stable
$\lambda_1 > 0 \Rightarrow$ orbit is chaotic

A Lyapunov-positive system is by definition chaotic. The positive value of the Lyapunov exponents indicates orbital divergence and a time scale on which state prediction is impossible. Negative Lyapunov exponents set the time scale on which transients or perturbations of the system's state will decay to a periodic orbit. A zero Lyapunov indicates a marginally or neutrally stable orbit. This often occurs near a point of bifurcation.

In a more complex phase space, the Lyapunov spectrum is used to classify and describe the dynamics. The sign of the Lyapunov is reported per k^{th} dimension, again with the negative sign indicating convergence and a positive exponential sign indicating divergence. Hence, for a system defined in two dimensions, there are two Lyapunov exponents for the attractors of the system, λ_1 and λ_2. The transformation from the Lyapunov exponent, λ, to the Lyapunov number, e^{λ}, can be used to give a more precise interpretation. The first Lyapunov number, $e^{\lambda 1}$, indicates the error amplification; the second number, $e^{\lambda 2}$, illustrates the degree of contraction (−) or expansion (+). The sum of the first and second exponents, $e^{\lambda 1+\lambda 2}$, can be interpreted as a measure of how much the area is either reduced (−, −) or expanded (+, +). For a system evolving on a fixed point where $k = 3$, the signs of the Lyapunov are (−, −, −). In such a system there is a continuous convergence on a fixed value. For a limit cycle the Lyapunov characteristic exponents are (0, −, −), where convergence occurs on two coordinates, with a marginal limit on the third. A two-torus attractor is indicated by (0, 0, −) and reflects instances where two dimensions are marginally stable and one converges. Strange attractors exist in instances of at least one positive Lyapunov exponent (+, 0, −) and reflect instances of simultaneous convergence and divergence. A purely chaotic state (+, +, +) represents utter divergence. Purely chaotic behavior is illustrated by a cantor-like set in a noninteger, fractal dimension. Intuitively, we can define a purely chaotic state where the three Lyapunov exponents are all positive. The sum of the three Lyapunov numbers, $e^{\lambda 1+\lambda 2+\lambda 3}$, approximates the maximal average by which the volume changes per time step.

Since it is increasingly difficult to see the evolution of a chaotic system as the divergence of the trajectories on the attractor becomes more rapid, the Lyapunov provides the most useful indicator of chaos. It should be mentioned that the usual methods of statistical investigation will not work when chaos exists. Spectral analysis and traditional time-series model fitting will not work to reveal the underlying determinism. Running multiple experiments over a long time and comparing results at the smallest level of observation run straight into the consequences of the "butterfly effect." The slightest measurement error will be magnified, invalidating any comparison or goodness of fit measures based on the expected versus observed observations.

Some Mathematical Underpinnings for λ

Schuster (1988) nicely summarized some of the core mathematical properties of the Lyapunov exponent. The Lyapunov exponent is a local quantity averaged over an entire attractor.[2] It measures exponential separation from some initial value, x_0, due to amplified error, ε. The following illustration shows in one dimension these properties:

$$\frac{\varepsilon}{x_0 \; x_0 + \varepsilon} \rightarrow N \text{ iterations} \rightarrow \frac{\varepsilon e^{N\lambda(x_0)}}{f^N(x_0)\; f^N(x_0 + \varepsilon)},$$

where $\varepsilon e^{N\lambda(x_0)} = |f^N(x_0 + \varepsilon) - f^N(x_0)|$. The limits of the previous expression can be formally expressed as

$$\lambda(x_0) = \lim_{N\to\infty} \lim_{\varepsilon\to 0} \frac{1}{N} \log \left| \frac{f^N(x_0 + \varepsilon) - f^N(x_0)}{\varepsilon} \right|,$$

which reduces to

$$\lambda(x_0) = \lim_{N\to\infty} \frac{1}{N} \log \left| \frac{df^N(x_0)}{dx_0} \right|.$$

Schuster showed that $e^{\lambda(x_0)}$ is the average distance between closely adjacent points after one iteration.

The Lyapunov also provides an unambiguous measure of the presence of chaos and gives us a measure of the average information gain or loss per iteration. The Kaplan-Yorke conjecture proposed that it was possible to predict the dimension of a strange attractor based upon a transformation of the Lyapunov exponents of the selfsame dynamical system (Kaplan and Yorke 1979; Russell, Hanson, and Ott 1980; Frederickson et al., 1983). In essence the units of λ_i are bits of information per iteration from the base of the initial measurement. By an application of the chain rule applied to the following equation, we get

$$\frac{d}{dx} f^2(x) \Big|_{x_0} = \frac{d}{dx} f\,[f(x)] \Big|_{x_0} = f'[f(x_0)]\, f'(x_0)$$

$$= f(x_1) f'(x_0) \text{ where } x_1 \equiv f(x_0).$$

Therefore, to evaluate λ, we calculate

$$\lambda_{x_0} = \lim_{N\to\infty} \frac{1}{N} \log \left| \frac{d}{dx_0} f^N(x_0) \right| = \lim_{N\to\infty} \frac{1}{N} \log \left| \prod_{i=0}^{N-1} f'(x_i) \right|$$

$$= \lim_{N\to\infty} \frac{1}{N} \sum_{i=0}^{N-1} \log \left| f'(x_i) \right| .$$

By Oseledec's multiplicative ergodic theorem (1968), this limit exists for almost all x_0. The average amount of stretching depends upon the initial error as well as on the initial value of the system. Oseledec's theorem has been used in nearly all explications of Lyapunovs to refine the definition of their mathematical underpinning (Eckmann and Ruelle 1985).

The implications are as follows: (1) the Lyapunovs are invariant of the dynamics; (2) because they are independent of the position of the trajectory on the orbit, the Lyapunov can be used to classify the attractor; (3) they also define the limits of predictability for any model. This implies a key interpretation of the Lyapunov, since the largest Lyapunov governs the size of any perturbation to the orbit around the attractor. When the system is Lyapunov-positive, the perturbation grows and predictability is lost. We can precisely say now that $\sum_{i>0}^{k} \lambda_i$ is the Kolmogorov-Sinai entropy, which dictates the rate of loss of predictability. Hence the computation of Lyapunovs becomes crucial since the system with them defines its own predictability (Eckmann and Ruelle 1985; Abarbanel 1992).

A map can illustrate how the change in the level of information may occur after one iteration. Dividing the closed interval [0, 1] on x into n equal intervals, x_0 can occur with a probability of $1/n$. By knowing which interval contains x_0, we gain the following information:

$$I_0 = -\sum_{i=1}^{n} \frac{1}{n} \log_2 \frac{1}{n} = \log_2 n.$$

The linear map, $f(x)$, changes the length of an interval by a known factor, $a = |f'(0)|$. The loss of precision leads to a loss of information:

$$\Delta I = -\sum_{i=1}^{n/a} \frac{a}{n} \log_2 \frac{a}{n} + \sum_{i=1}^{n} \frac{1}{n} \log_2 \frac{1}{n} = -\log_2 a$$

L. Douglas Kiel and Euel Elliott, Editors
http://www.press.umich.edu/titleDetailDesc.do?id=14623
The University of Michigan Press, 1996

Hence,

$$\overline{\Delta I} = -\lim_{N\to\infty} \frac{1}{N} \sum_{i=0}^{N-1} \log_2 |f'(x_i)|.$$

The value of $|\overline{\Delta I}|$ is proportional to the Lyapunov, in that

$$\lambda(x_0) = \log_2 \cdot |\overline{\Delta I}|.$$

The larger the exponent, the greater the information loss. The gain or loss of information reflects the amount of uncertainty in the dynamical process. With a loss of information, the amount of predictability decreases and the certainty in the future is diminished. Thus the Lyapunov exponent allows both a calculation of a chaotic signal and an interpretation.

Calculating λ

The calculation of the Lyapunov exponents depends upon measuring the rate by which close or nearby trajectories split in phase space after small changes in initial conditions. When the equations of motion are known, as in a computer simulation, estimating the Lyapunov is not difficult. When data are used, difficulties arise. As a rule, chaotic systems contain at least one positive Lyapunov exponent. The trick is to calculate the Lyapunov within the smallest possible D-dimensional hypersphere that contains the attractor within a specific number of time steps. There are three control parameters: the embedded dimension, the number of time steps over which pairs of points are evaluated, and the precision of the estimate. A k-dimensional volume segment grows by the average factor of $\lambda_1 + \lambda_2 + \ldots + \lambda_m$ at each time step. Again, a $\lambda > 0$ indicates that the attractor is chaotic and that a very small error or perturbation grows in the short run exponentially fast.

Two algorithms have been used to calculate Lyapunov characteristic exponents: Wolf's algorithm (1985, 1986) and the Eckmann-Ruelle algorithm (1985). The Wolf algorithm tracks a pair of arbitrarily close points over a trajectory to estimate the accumulated error per time step. The points are separated in time by at least one orbit on the attractor. The trajectory is defined by the fiducial and test trajectories.[3] They are tracked for a fixed time period or until the distance between the two components of the trajectory exceeds some specific value. In sequence, another test point near the fiducial trajectory is selected and estimation proceeds. The end product is that the stretching and squeezing are averaged.

Figure 3.1 shows a representation of the Wolf computation of a

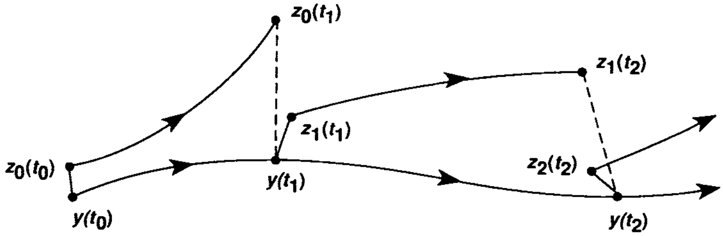

Fig. 3.1. Schematic representation of the Wolf algorithm to compute λ_1. (From Vastano and Kostelich 1986. Reprinted with permission.)

Lyapunov as presented by Vastano and Kostelich (1986, 102). The initial data point, $y(t_0)$, and its nearest neighbor, $z_0(t_0)$, are L_0 units apart. Over Δt, a series of time steps from t_0 to t_k, the two points y and z evolve until their distance, L_0', is greater than some arbitrarily small ε. The y value at t_k becomes $y(t_1)$ and a new nearest neighbor, $z_1(t_1)$, is selected.[4] This procedure continues until the fiducial trajectory **y** reaches the end of the time series. The replacement of the old point by its substitute point and the replacement of the error direction by a new directional vector constitutes a renormalization of errors along the trajectory.

The largest estimated Lyapunov exponent of the attractor is then

$$\hat{\lambda}_1 = \frac{1}{N\Delta t} \sum_{i=0}^{M-1} \log_2 \frac{L_i'}{L_i},$$

where M is the number of replacement steps (where some arbitrarily small ε was exceeded) and N is the total number of time steps that the fiducial trajectory **y** progressed. Wolf et al. (1985) provide the code for two programs: one for estimating the Lyapunov spectrum for systems of differential equations and a second for estimating the Lyapunov from a time series. The code is relatively easy to implement in Fortran or to convert to Turbo C++. The IBM PC computer packages developed by Schaffer (1988) and Sprott and Rowlands (1992) have implemented the Lyapunov exponent by means of the Wolf algorithm. Figure 3.2 illustrates the Lyapunov exponent using the Wolf algorithm for an iterative logistic map. Within the areas of the chaotic region, the Lyapunovs are positive. Note that the boundary of the Lyapunov reflects the erratic nature of chaos within the logistic map. The geometric properties of the Lyapunov exponent are known to be chaotic.

There are several folk wisdoms in estimating the Lyapunov exponent on data using the Wolf algorithm. Schaffer et al. (1988) suggest standardizing the

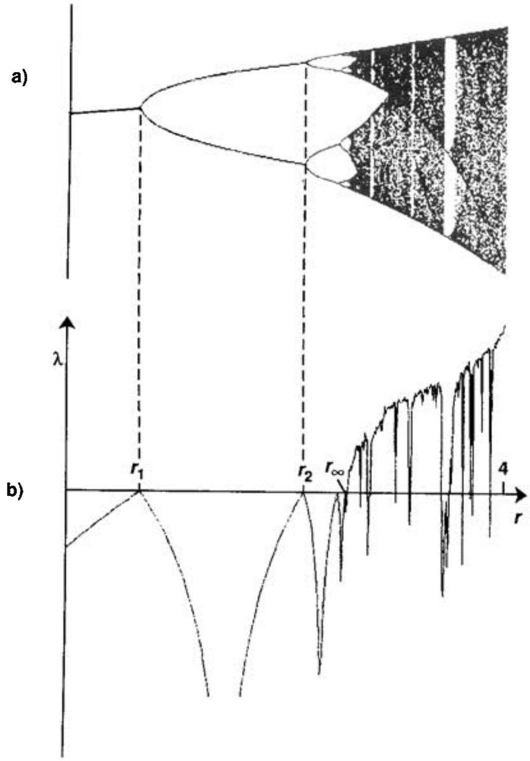

Fig. 3.2. Iterations of the logistic map (a) and accompanying Lyapunov exponents (b). (From Schuster 1988. Reprinted with permission of the author.)

data into interval [0, 1]. Such transformations allow for comparisons between different time series and permit a more ready interpretation of the length scales. Second, Fraser and Swinney (1986) suggest verifying that the estimated Lyapunov exponent is stable over a range of parameter choices. Increasing the maximum admissible length scale between fiducial and test trajectories eventually induces the value of the Lyapunov exponent to drop. Wolf et al. (1985) suggest that the maximum length scales should not exceed 10 percent of the extent of the attractor.

The Eckmann-Ruelle algorithm offers some marginal gain over the Wolf approach and is considered by some the "current state of the art" (Peitgen, Jürgens, and Saupe 1992). It uses a least-squares approximation of the derivative matrices, an approach that was independently proposed by Sano and Sawada (1985). Figure 3.3 illustrates the Eckmann-Ruelle algorithm.

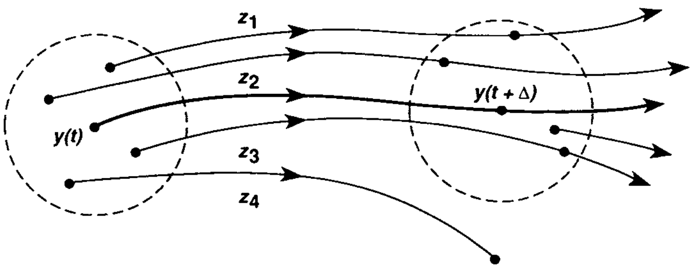

Fig. 3.3. Schematic representation of the Eckmann-Ruelle method. (From Vastano and Kostelich 1986. Reprinted with permission.)

The Eckmann-Ruelle method calculates all positive Lyapunov exponents of the attractor. The fiducial trajectory, **y**, begins with $y(t)$ and k nearest neighbors, $z_i(t) \mid i = 1 - k$, all of which are within ε of $y(t)$. Each point evolves through a specific period of time, Δt, such that $z_i(t)$ and $z_i(t + \Delta t)$ are tracked along with $y(t)$ and $y(t + \Delta t)$. A least-squares estimate is possible of the spatial derivative $Df(y(t))$ of the series using $y(t)$ and the k points. At each time step, all points within ε of $y(t)$ are used to calculate the best possible fit until all the data are used.[5]

The Eckmann-Ruelle algorithm provides a useful way to calculate the entire range of Lyapunov exponent values, while recovering more than one positive Lyapunov exponent with the Wolf algorithm is substantially more complicated.

Conclusion

It might be that chaos at some time in the future will be considered a sideshow of nonlinear dynamics. Issues that allow full investigation of self-organization may simply assume that chaotic states and the nature of prediction, and indeed theoretical understanding, must take chaos into consideration.

The ideas behind chaos may also become a common heuristic device, a way of understanding complexity, criticality, and nuance largely put aside by the earlier ideas underlying the behavioralism in the social sciences. There is an intuitive appeal of nonlinear dynamical ideas in certain areas of social and psychological sciences, where difficulty in measurement and systematic observations often yield qualitative judgment. As a colleague who studies Eastern Europe suggested, we have to take the countries and people as they are, with all their warts (complexity) and goofiness (unreliability).

Chaos may, however, find a new home and more precise use and measurement in more positive, prescriptive areas of social inquiry. At the end of a draft manuscript on infinitely repeated games of incomplete information, McKelvey and Palfrey suggest: ". . . the resulting solutions have properties that are reminiscent of chaotic dynamical systems, and one wonders whether such a solution describes in either a normative or positive fashion the behavior individuals would or should adapt in such games" (McKelvey and Palfrey 1992). Hübler and Pines (1993) investigate in game-theoretic logic how players attempt to model, control, and predict future states within chaotic environments. Indeed, they even suggest that competition between two players can lead to a chaotic state. In each instance, the ability of a player to recognize the presence of chaos presents a strategic advantage in adapting behavior to that chaos. In each instance, it will be critical to identify the degree of chaos, distinguishing between the "edge of chaos" and full-blown chaos.

Dynamical nonlinear systems are thought to predict different things than linear systems. Included are a set of invariant measures such as the Lyapunov exponents or attractor dimensions. Determining their precise interpretation in the social sciences will provide some interesting dissertations. As May (1992, 452) joked: "chaos at last reconciled the Calvinist's God and his foreordained world with the illusion of free-will that we enjoy; sensitivity to initial conditions makes our world and fate appear unpredictable, even though it is determined." In low-dimensional political, social, and economic Lyapunov-positive environments, where does this leave choice theory, free will, and individual responsibility?

NOTES

Figures 3.1 and 3.3 are reprinted from "Comparison of Algorithms for Determining Lyapunov Exponents from Experimental Data" by J. A. Vastano and E. J. Kostelich, in *Dimensions and Entropies in Chaotic Systems*, ed. G. Mayer-Kress (Berlin: Springer-Verlag, 1986), by permission of the authors, Elsevier Science, and Springer-Verlag.

1. Formally, a closed invariant set $A \subset \mathbf{R}^n$ is an attracting set if there is some neighborhood $U(A)$ such that $\Phi(x,t) \in U$ for $t \geq 0$ and $\Phi(x,t) \rightarrow A$ as $t \rightarrow \infty$ for every $x \in U \cdot \bigcup_{t \leq 0} \Phi(x,t)$ is the domain of attraction of A.

2. Some research work concerns local Lyapunov exponents that specifically aim at short-term forecasts of orbital instabilities (Abarbanel 1992; MaCaffrey et al. 1992; Nychka et al. 1992; Wolff 1992). These approaches attempt to describe orbital instabilities for a fixed number of time steps ahead rather than over an infinite number of time steps ahead such as with global Lyapunov exponents. These approaches are especially interesting in instances where estimating the Lyapunov exponent is from experimental data. The difficulty with data is that we may not know the exact equation of motion nor the correct embedding dimension for the phase space. Hence we will not know the Jacobian matrices necessary to calculate the global Lyapunov exponents.

Local Lyapunov exponents may hence provide a useful approach to detecting chaos in some cases.

3. By definition a fiducial distribution makes probability statements about unknown parameter values.

4. It should be noted that every replacement data point should be in the direction of the previous point on the test trajectory. The replacement point, $z_1(t_1)$, should be located as near as possible to the value $y(t_1)$ in the direction $z_0(t_1)$. The search for such a value is restricted to a cone of width Θ and height η about the fiducial trajectory **y**. The angular-width value of Θ is adjusted as needed. Wolf (1985) has suggested that the resulting errors in the estimation of the Lyapunov are not sensitive to the choice of Θ or to approximate values of ε. Usually, $\Theta = .349$, the value of $\pi/9$. See Vastano and Kostelich 1986 for an excellent discussion of the implementation of both the Wolf and Eckmann-Ruelle methods.

5. Using a least-squares estimate of the Jacobian at $y(t)$, given a flow $\dot{x} = f(x)$, the variation equations $\dot{u} = Df(x)u$ are integrated using Benettin's algorithm to compute the Lyapunov.