Phase Space Reconstruction from Time Series Data: Where History Meets Theory

Ray Huffaker

Food and Resource Economics Department, University of Florida, USA
rhuffaker@ufl.edu

1   Introduction

In 'dissipative' dynamical systems, variables evolve asymptotically toward low-dimensional 'attractors' that define their dynamical properties. Unfortunately, real-world dynamical systems are generally too complex for us to directly observe these attractors. Fortunately, there is a method—'phase space reconstruction'—that can be used to indirectly detect attractors in real-world dynamical systems using time series data on a single variable (Broomhead and King, 1985; Schaffer and Kott, 1985; Kott et al, 1988; Williams,1997). Armed with this knowledge, we can formulate more accurate and informative models of real-world dynamical systems.

We begin by introducing the concept of phase space attractors within the context of a dynamic ISLM model. We next demonstrate how phase space reconstruction faithfully reproduces one of the model's attractors.   Finally, we discuss how phase space reconstruction fits into a more general 'diagnostic' modeling approach that relies on historical data to guide and test the deterministic formulation of theoretical dynamical models. As an example of diagnostic modeling, we test how closely the attractor generated by the dynamic ISLM model visually approximates the attractor reconstructed from time series data on real-world interest rates.

2   Phase Space Attractors

Consider the following dynamic ISLM model (Shone, Chapter 10):

(1a) $$\frac{dr}{dt} = \beta \underbrace{\left[m^d \left(r, y; \Omega_1 \right) - m_0 \right]}_{\text{excess demand in money market}}$$

$$\frac{dy}{dt} = \alpha \underbrace{\left[e(r, y; \Omega_2) - y \right]}_{\text{excess demand in goods market}}$$
(1b)

where r is the nominal interest rate, y is real income, $m^d(r, y)$ measures the demand for real money balances, m_0 is the exogenous real money supply, $e(r, y)$ measures real expenditures, and underlying parameters Ω_1 and Ω_2 are constant. The model exhibits a wide range of dynamic behaviors depending on how rapidly the money and goods markets adjust to excess demand. One possible behavior is for interest rates and incomes to oscillate toward an equilibrium that balances demand and supply in each market (Figures 1a,b). Figure 1(c) shows the solution in phase space obtained by plotting $r(t),$ and $y(t)$ for each time t, and connecting the points (a 'scatter' diagram). System

dynamics are characterized by a 'stable-focus' attractor.

(a)   (b)

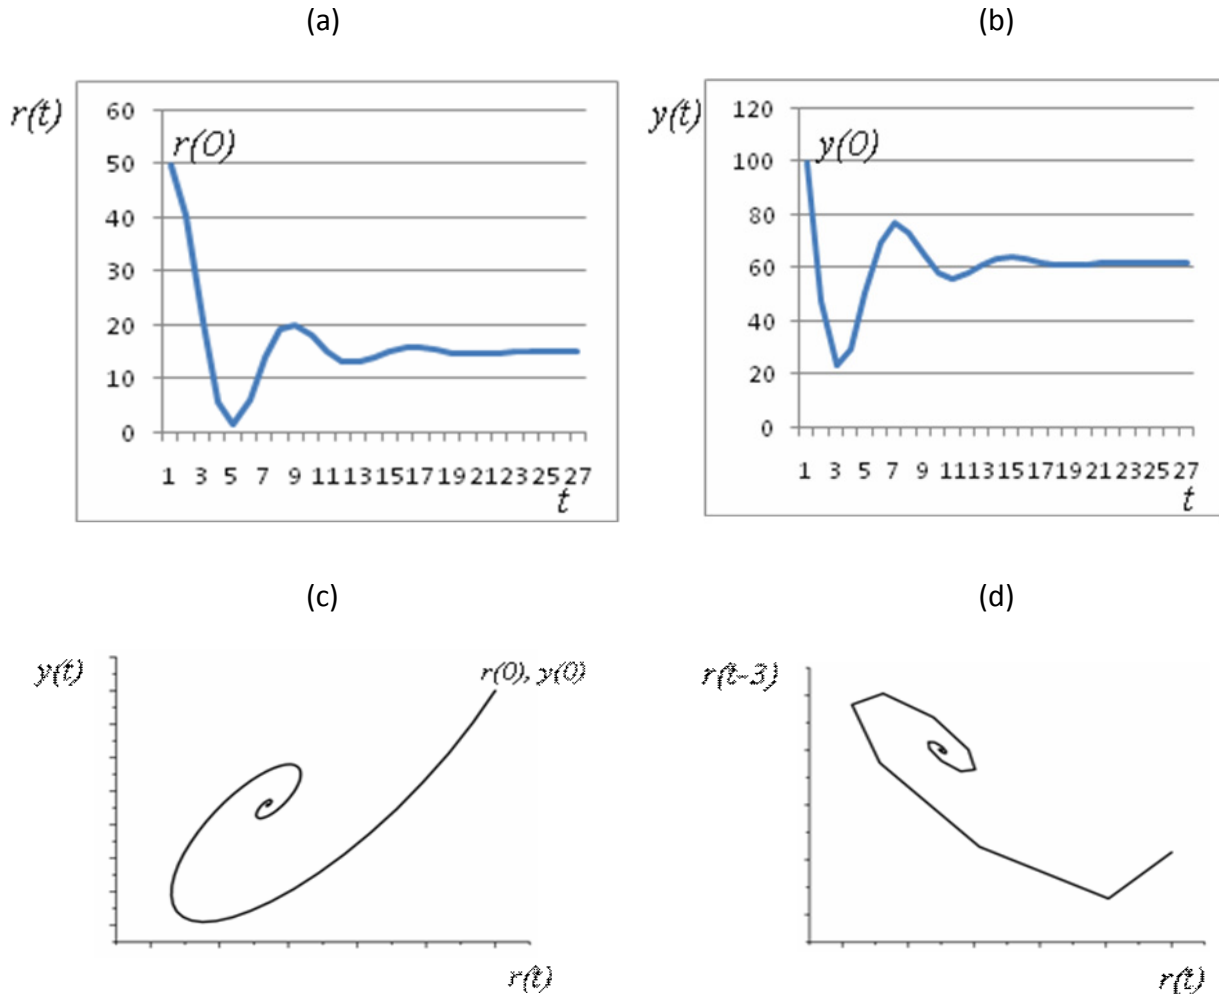(c)   (d)

Figure 1. The spiral-node atrractor in the dynamic ISLM model:
(a) Time series of interest rates; (b) Time series of incomes; (c) Phase diagram solution; (d) Attractor reconstruction. We solved the model numerically in Excel using an Euler approximation.

3   Attractor Reconstruction

To demonstrate how attractor reconstruction works, we reconstruct the stable-focus attractor in the ISLM model using only one of the system variables, in this case, $r(t)$.

3.1 Time-Delay Embedding Procedure

The time-delay embedding procedure of attractor reconstruction is firmly rooted in mathematical topology theory thanks to Takens' theorems (1980). In general, suppose that the solution of a real world dynamical system asymptotically contracts onto an unobservable attractor contained within a lower-dimensional submanifold of unknown dimension *m*. The observed time series, call it *x(t)*, lies on this attractor. The time-delay embedding procedure uses *x(t)* to 'reconstruct' the attractor with a 'scatter' plot of the following lagged coordinate vectors: $x(t), x(t-\tau), x(t-2\tau), ..., x(t-(n-1)\tau)$ where τ is the length of the lag, and *n* is the number of lagged coordinate vectors (called the 'embedding dimension').

For example, assume that we observe the time series $x(t) = (4,2,7,4,9,6,10,3)^T$ [T is the transpose]. Assume a lag length $\tau = 2$ and an embedding dimension n = 3. This generates the following lagged coordinate vectors:

$$x(t) = \begin{pmatrix} 4 \\ 2 \\ 7 \\ 4 \\ 9 \\ 6 \\ 10 \\ 3 \end{pmatrix}, \quad x(t-2) = \begin{pmatrix} 7 \\ 4 \\ 9 \\ 6 \\ 10 \\ 3 \end{pmatrix}, \quad x(t-4) = \begin{pmatrix} 9 \\ 6 \\ 10 \\ 3 \end{pmatrix}$$

The reconstructed attractor is a 'scatter plot' of the elements in the first four rows of the lagged coordinate vectors. The first point on the attractor is composed of the first elements of each vector (4,7,9). Similarly, the second point is (2,4,6); the third point is (7,9,10); and the fourth point is (4,6,3). Elements in the remaining rows are lost in the lagging process. The reconstructed attractor for this example is:

Takens' theorems give the following sufficient condition for a reconstructed attractor to have the same dynamical properties as the unobservable attractor from the real-world system: $n > (2m+1)$. In words, the embedding dimension n has to be at least as large as twice the dimension of the real-world solution manifold m plus 1. Three notes are in order: (1) Since this is a sufficient condition, fewer embedding dimensions may be necessary; (2) Takens' theorems do not imply that the reconstructed attractor is identical to the real-world attractor—just that it has the same dynamical (topological) properties; and (3) Since m is unobservable, the embedding dimension n cannot be directly calculated from the sufficient condition.

3.2 Selecting Lag Length (τ) and Embedding Dimension (m)

Selecting the reconstruction parameters τ and m is like focusing a camera—some pairs result in much clearer reconstructed attractors than others. A 'brute force' approach would simply search over a number of combinations until an attractor comes into focus. Fortunately, the literature suggests various diagnostic procedures to narrow the search (Williams, 1997, pp. 275-284).

The lag length τ is conventionally chosen as the first minimum of the 'mutual information function'—a probabilistic measure of the extent to which $x(t+\tau)$ is related to $x(t)$ at a given τ. The rationale behind the approach is to introduce the needed statistical independence between successive lagged values. For example, the problem with overly short lags is that successive values contain redundant mutual information. Consequently, the reconstructed attractor is restricted to lay on a 45° diagonal in the embedding space. Alternatively, the problem with overly long lags is that successive lagged values contain too little mutual information to reconstruct the attractor.

The embedding dimension *m* is conventionally chosen using the 'false nearest neighbors' method. This method measures the percentage of close neighboring points in a given dimension that remain so in the next highest dimension. The minimum embedding dimension capable of containing the reconstructed attractor is that for which the percentage of false nearest neighbors drops to zero for a given tolerance level. For example, assume that an overcoat is crammed into small box. If you gaze into the top of the box, you see the material, but can't tell if it's part of a coat or just a lump of cloth. Take the coat out of the box, unfurl it into its proper 'dimension', and you see plainly that it's a coat. When the coat is wadded up, the end of the sleeve might be touching the collar. When the coat is unfurled, the end of the sleeve is distant from the collar. Thus, when the coat is wadded up, the end of the sleeve and the collar are false nearest neighbors. When the coat is completely unfurled, these two points (and all other points on the coat) no longer move away from each other—the percentage of false nearest neighbors is zero. Similarly, an attractor embedded into too few dimensions does not have room to fully express itself. Alternatively, an attractor embedded into too many dimensions goes out of focus.

Software is available to compute lag length and embedding dimension. For example, *Visual Recurrence Analysis* (*VRA*) software is free for downloading, use, copying, and distribution at http://visual-recurrence-analysis.software.informer.com/ . We use it below to compute lag length and embedding dimension.

3.3 What Constitutes a 'Good' Reconstruction?

The test of a 'good' reconstruction is largely visual. Does the reconstructed attractor exhibit apparent structure? Recent research in *recurrence quantification analysis* is adding objective numerical measures to identify recurring patterns in data (see, e.g., Webber and Zbilut, 2004).

3.4 Reconstructing the ISLM Stable-Focus Attractor

The stable-focus attractor reconstructed from the single time series $r(t)$ generated in the dynamic ISLM model is shown in Figure 1(d). The first minimum of the mutual information function occurs at lag length $\tau = 3$:

The embedding dimension for which the percentage of false nearest neighbors drops effectively to zero is *m*=2:

The attractor reconstructed with these parameters (Figure 1d) faithfully reproduces the stable-spiral attractor from the original system (Figure 1c).

3.5 Problems in Application: Small and Noisy Data Sets

Attractor reconstruction may fail for a couple of reasons (Williams, 1997, pp. 274-275). First, the real world dynamical system may not evolve toward a low-dimensional attractor that can be visualized in two or three dimensions. In other words, there is no low-dimensional attractor to reconstruct, and the false nearest neighbors test gives a minimum embedding dimension greater than what can be visualized (i.e., > 3 dimensions). Second, attractor reconstruction has difficulty with small and noisy data sets—common characteristics of economic time series. Principal components analysis may remedy both problems (Williams, 1997, pp. 284-287).

Principal components analysis (sometimes called 'singular value decomposition') filters noise from data without distorting the underlying dynamical properties. As a result, the attractor may become more clearly focused in fewer embedding dimensions. A brief outline of the procedural steps follows (Theil, 1971):

1.   Construct the 'embedded date' matrix X from the observed time series $x(t)$ with q observations. Set lag length to $\tau = 1$ and embedding dimension to the level n computed from the false nearest neighbors test:

$$X_{(q-n \times n+1)} =$$

$$\begin{bmatrix} x(t) & x(t-1) & x(t-2) & x(t-3) & \bullet & \bullet & x(t-n) \\ x(t-1) & x(t-2) & x(t-3) & x(t-4) & \bullet & \bullet & x(t-n-1) \\ x(t-2) & x(t-3) & x(t-4) & x(t-5) & \bullet & \bullet & x(t-n-2) \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & x(t-q-n) \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \\ \bullet & \bullet & \bullet & \bullet & \bullet & & \\ \bullet & \bullet & \bullet & x(t-q-3) & & & \\ \bullet & \bullet & x(t-q-2) & & & & \\ \bullet & x(t-q-1) & & & & & \\ x(t-q) & & & & & & \end{bmatrix}$$

2. Compute the eigensystem of $X'X$: $(X'X - \lambda_i I)a_i = 0$ where $\lambda_i \ (i=1,n)$ are eigenvalues and a_i are the associated eigenvectors.

3. Compute the principal components: $\underset{q \times 1}{p_i} = \dfrac{1}{\lambda_i} X a_i \quad i=1,n$

4. Compute the proportion of the variance of the embedded data accounted for by each principal component: $V = \dfrac{\lambda_i}{Trace(X'X)}$.

5. Reconstruct phase space with a scatter plot of the principal components matrix: $P = (p_1, p_2, p_3)$ where p_1, p_2, p_3 are the principal components associated with the three largest eigenvalues and corresponding eigenvectors.

This procedure filters noise by using the principal components explaining the greatest proportion of the variance in the embedded data (the principal components associated with smaller eigenvalues contain more noise). The procedure is better justified as the proportion *V* associated with the three largest eigenvalues approaches 100%. Reducing the noise in the data increases the possibility that a reconstructed attractor can be visualized in two or three dimensions with a scatter plot of the principal components associated with the two or three largest eigenvalues.

4   Diagnostic Modeling Using Real-World Interest Rate Data

Diagnostic modeling is a two-pronged procedure for guiding and testing the specification of dynamic models (Figure 2). The left-ward prong relies on reconstruction of a real-world attractor from historic data on a single observed variable. The right-ward prong formulates a dynamic model—derived from fundamental economic principles—that attempts to explain the dynamic forces characterizing the reconstructed attractor. A useful specification test for the model's attractor is whether it is visually similar to the real-world attractor reconstructed from the historic data.
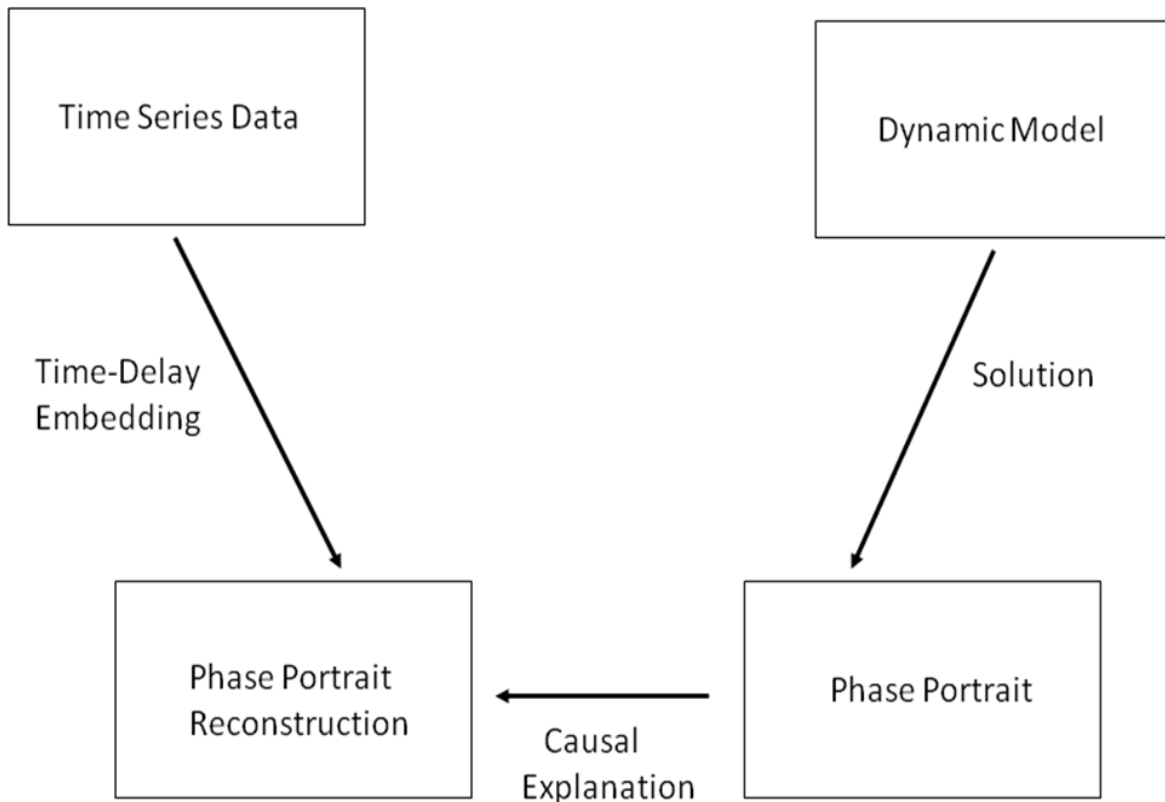
Figure 2: Diagnostic Modeling

We now consider an example of diagnostic modeling. Figure 3 (upper left corner) shows a plot of time series data on U.S. interest rates (Moody's Seasoned AAA Corporate Bond Yields, January 1, 1919 to January 5, 2010, Board of Governors of the Federal Reserve System). Directly below is the graph of the attractor reconstructed from these data using the time-delay method and principle components. The attractor exhibits aperiodic cycling that is more complex than the stable-focus attractor (and other possible attractors) generated by specification (1a,b) of the dynamic ISLM model. This indicates that the model must be re-specified so that it can generate an aperiodic attractor. We re-specified the model by making the marginal propensity to consume in the real expenditure function seasonally forced, i.e., an explicit function of time: $e(r,y,t)$ (upper right corner of Figure 3). Specifically, we modeled the controlling parameter as a periodic cosine function. Similar to the historic data, the attractor generated by the re-specified model exhibits aperiodic cycling (lower right corner of Figure 3).
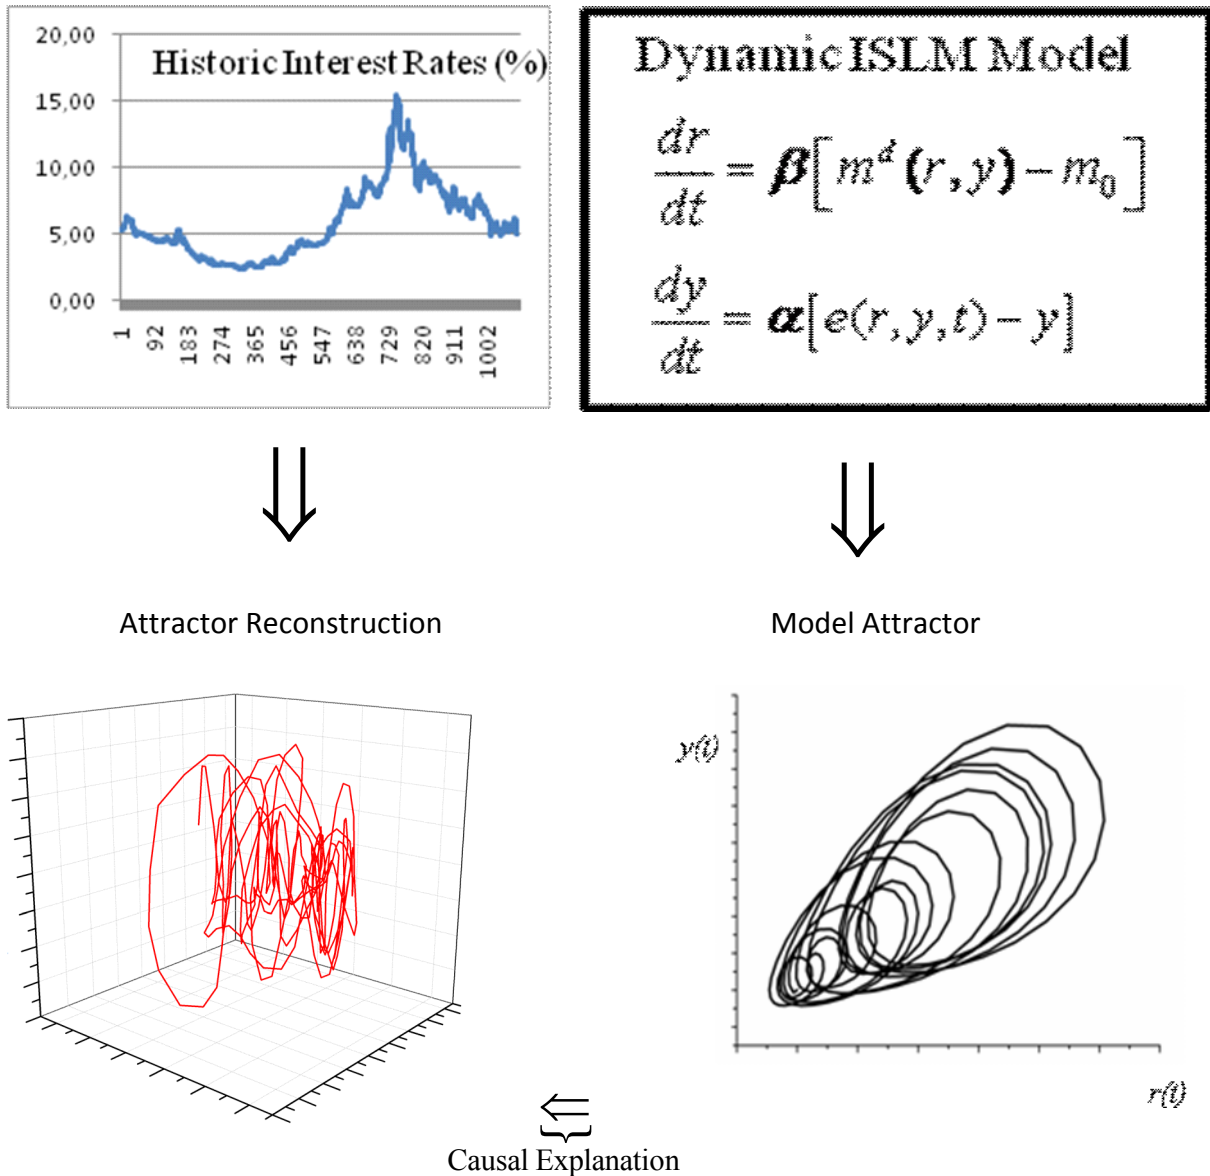
Figure 3. Example of diagnostic modeling with the dynamic IS-LM model: To simulate the aperiodic cycling of the attractor reconstructed from the interest rate data (Moody's seasoned AAA Corporate Bond Yield), the marginal propensity to consumer in the dynamic IS-LM model was assumed to adjust seasonally.

5   Concluding Comments

Phase space reconstruction is a potentially useful method for detecting the dynamical structure of real-world systems that generate the historical data we observe. As a part of diagnostic modeling, it can improve the formulation and performance of dynamical models. The method uncovers surprising structure in apparently random data. This justifies more extensive use of deterministic modeling based on principles of theory—as opposed to 'knee-jerk' application of stochastic methods—when faced with apparently random data.

6   References

Broomhead, D.S., King, G.P. (1985). Extracting Qualitative Dynamics from Experimental Data. Physica 20D: 217-236.

Huffaker, R., Mittelhammer, R., Barkley, P., Folwell, R. (2003). Community Dynamics in a University Environment 7: 181-203.

Kot, M., Schaffer, W., Truty, G., Graser, D., Olsen, L. (1988). Changing Criteria for Imposing Order. Ecological Modeling 43: 75-110.

Schaffer, W., Kot, M. (1985). Nearly One Dimensional Dynamics in an Epidemic. J. Theor. Biol. 112: 403-427.

Shone, R. (2002). Economic Dynamic. Cambridge University Press, Cambridge, UK.

Takens, F. (1980). Detecting strange attractors in turbulence. In: D.A. Rand and L.S. Young (Editors), Dynamical Systems and Turbulence. Lecture Notes in Mathematics, 898. Springer, New York, pp. 366-381.

Theil, H. (1971). Principles of Econometrics. John Wiley & Sons, New York.

Weber, C. and Zbilut, J. (2004). Introduction to recurrence quantification analysis. In: C. Weber (Ed), Recurrence Quantification Analysis of Nonlinear Dynamical Systems. Springer-Verlag, New York.

Williams, G. (1997). Chaos Theory Tamed (Chapter 19: Attractor Reconstruction). Joseph Henry Press, Washington D.C.

USDA/ERS (2009). Nutrient Availability 1909-2004, (http://www.ers.usda.gov/data/FoodConsumption/).

USDA/ERS (2009). Nutrient Availability: Documentation, http://www.ers.usda.gov/Data/foodconsumption/NutrientAvailDoc.htm).