Vectorfield

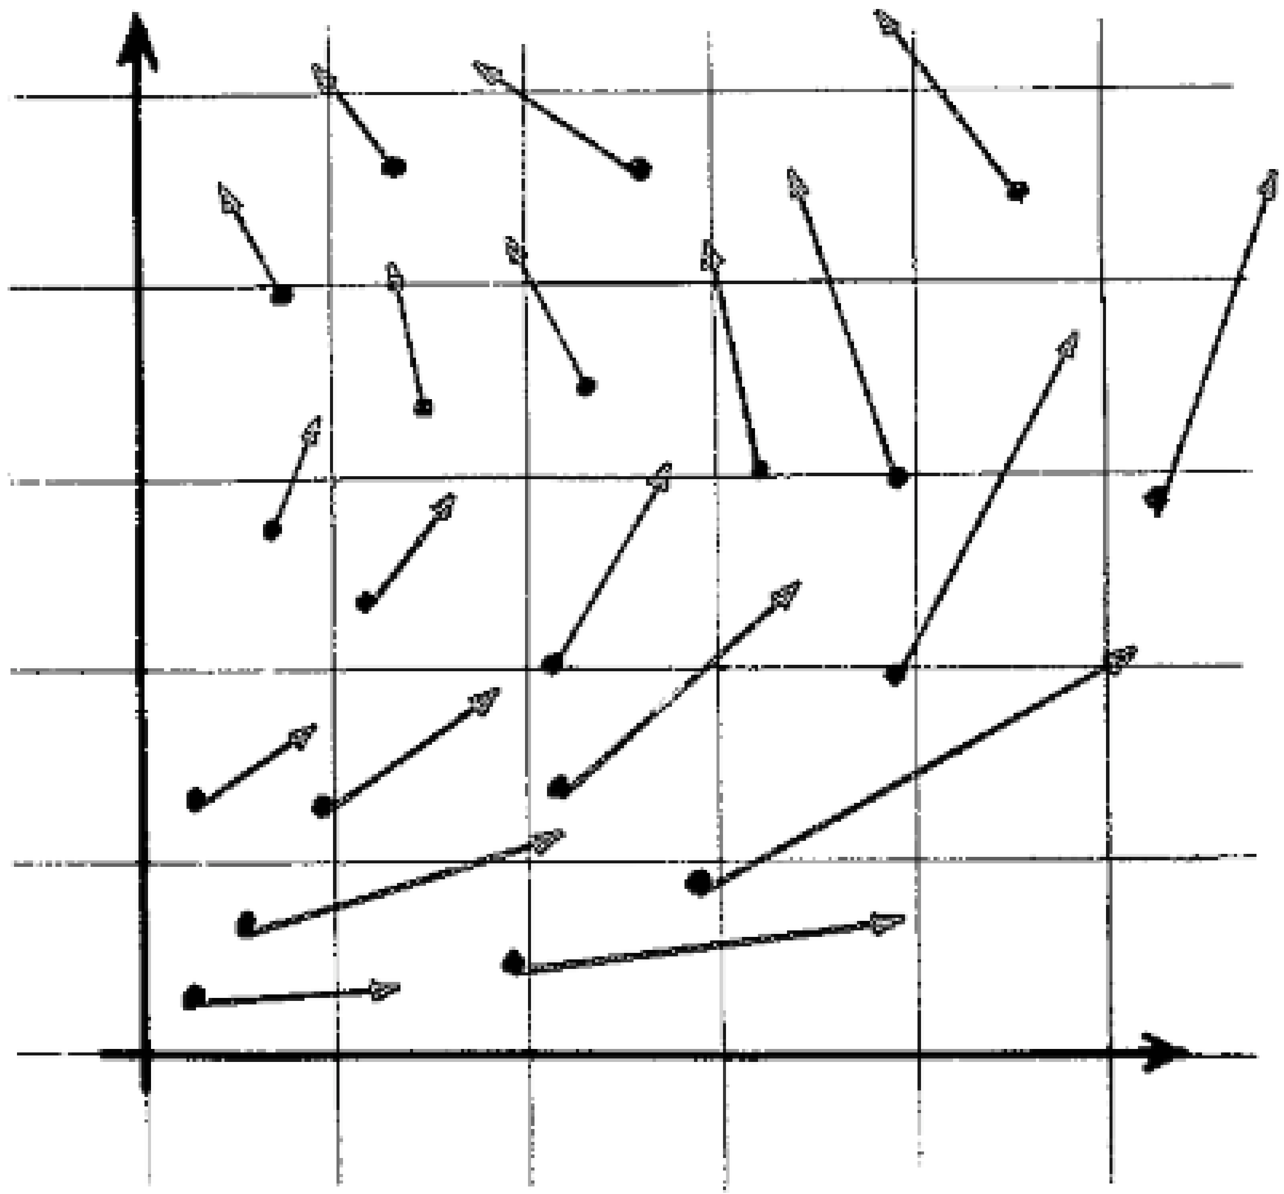

1.2.4. A *vectorfield* is a field of bound vectors, one defined at (and bound to) each and every point of the state space. Here only a few of the vectors are drawn, to suggest the full field.

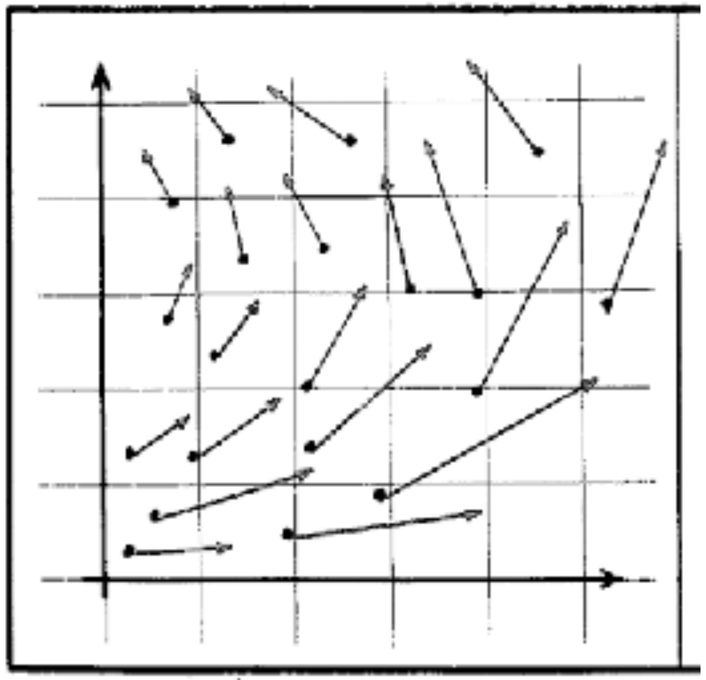

The state space, filled with trajectories, is called the *phase portrait* of the dynamical system. The velocity vectorfield has been derived from the phase portrait by *differentiation*.

We regard this vectorfield as the model for the system under study. In fact, the phrase *dynamical system* will specifically denote this vectorfield.

(Annoyingly, some folk use the term *phase space* where we are using *state space*. Sorry.)

Trajectory:
the evolution
of a *specific* system
over time, from
some initial conditions

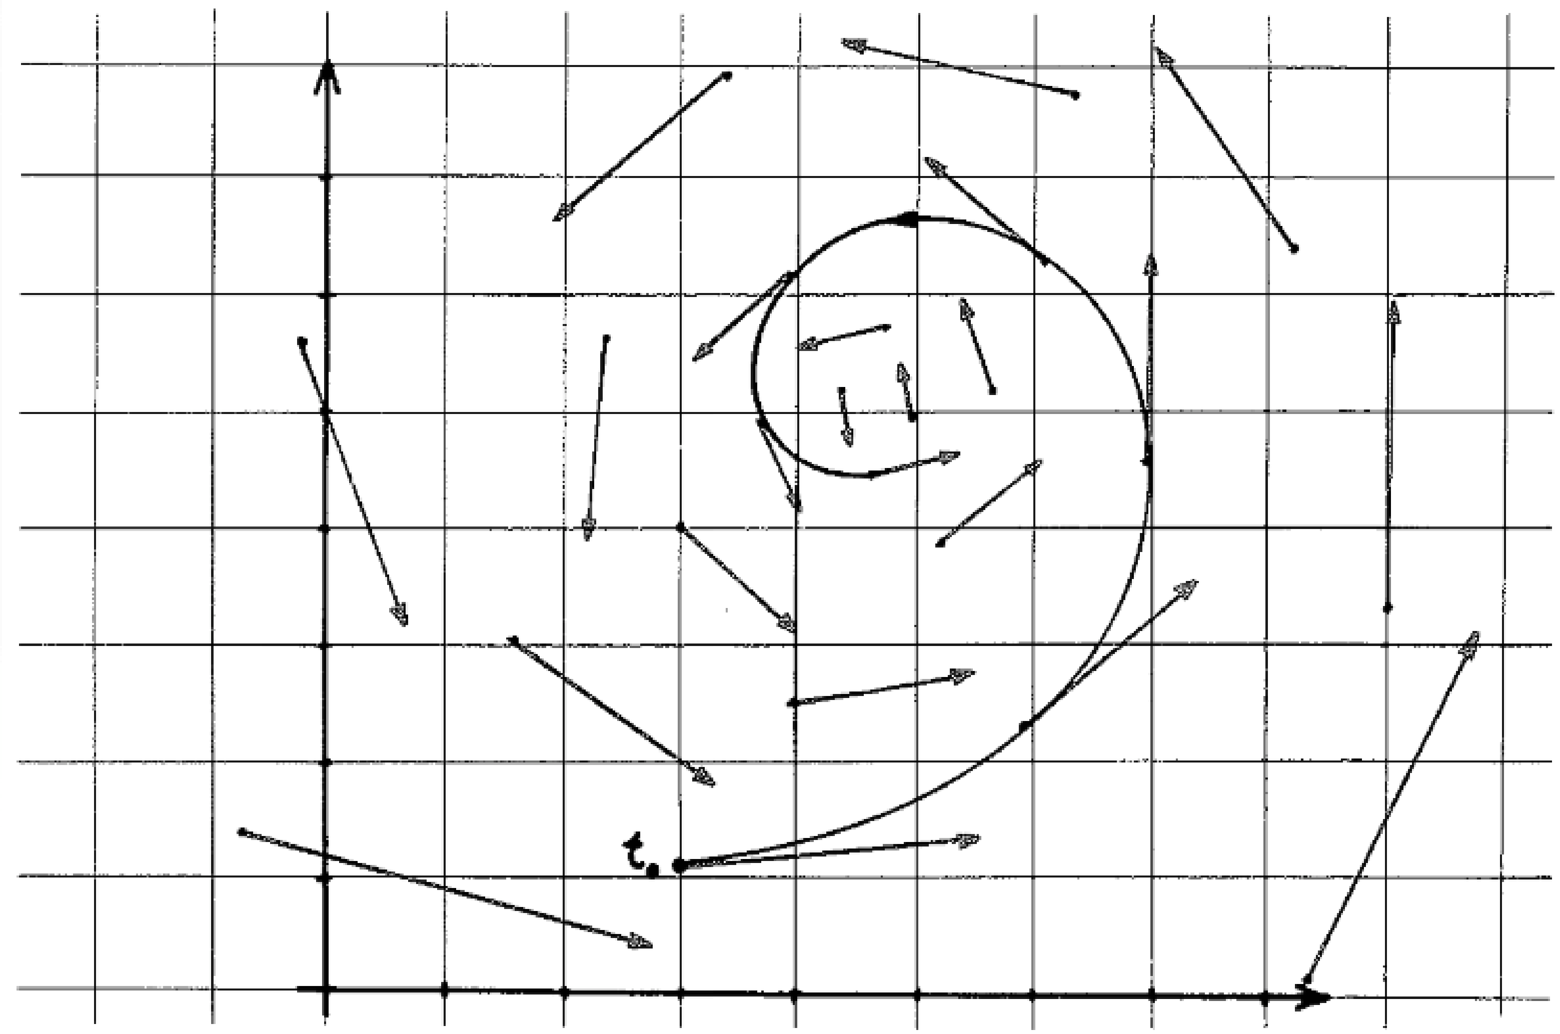

1.2.9. Given a state space and a dynamical system (smooth vectorfield), a curve in the state space is a *trajectory, or integral curve*, of the dynamical system if its velocity vector agrees with the vectorfield at each point along the curve. This means the curve must evolve so as to be tangent to the vectorfield at each point, as shown here. The point on the trajectory corresponding to elapsed time zero, t_0, is the *initial state* of the trajectory.

In a deterministic dynamical system, trajectories never cross. Why?

Parameters and Variables

If we observe a system for a while, the things which change while we observe it are naturally called variables.

Things which characterize the system, but which don't change while we watch it, are parameters of the system.

For some purposes, my bank balance might be a variable, and interest rates might be a system parameter

For other purposes, interests rates might be regarded as a variable.

In a network, activations might be variables, and weights be parameters, but weights change slowly over time, so we need to consider, not one system, but a family of systems.

Parameters:

Serve to pick out a specific instance of a dynamic system from a family of related systems.

Often, we are interested in the whole family: how does the vector field change as the parameters are altered?

E.g. the length of a pendulum
the mass of a planet

A worked example: the simple pendulum

The pendulum may be the most classical example of the dynamical modeling process. It has a two-dimensional state space, and a dynamical system established by Newton.

This model assumes that the rod is very light, but rigid. The hinge at the top is perfectly frictionless. The weight at the lower end is heavy, but very small. It moves in a vacuum. The force of gravity always pulls it straight down.

These idealizations describe the modeling assumptions in this example, called the *simple pendulum*.

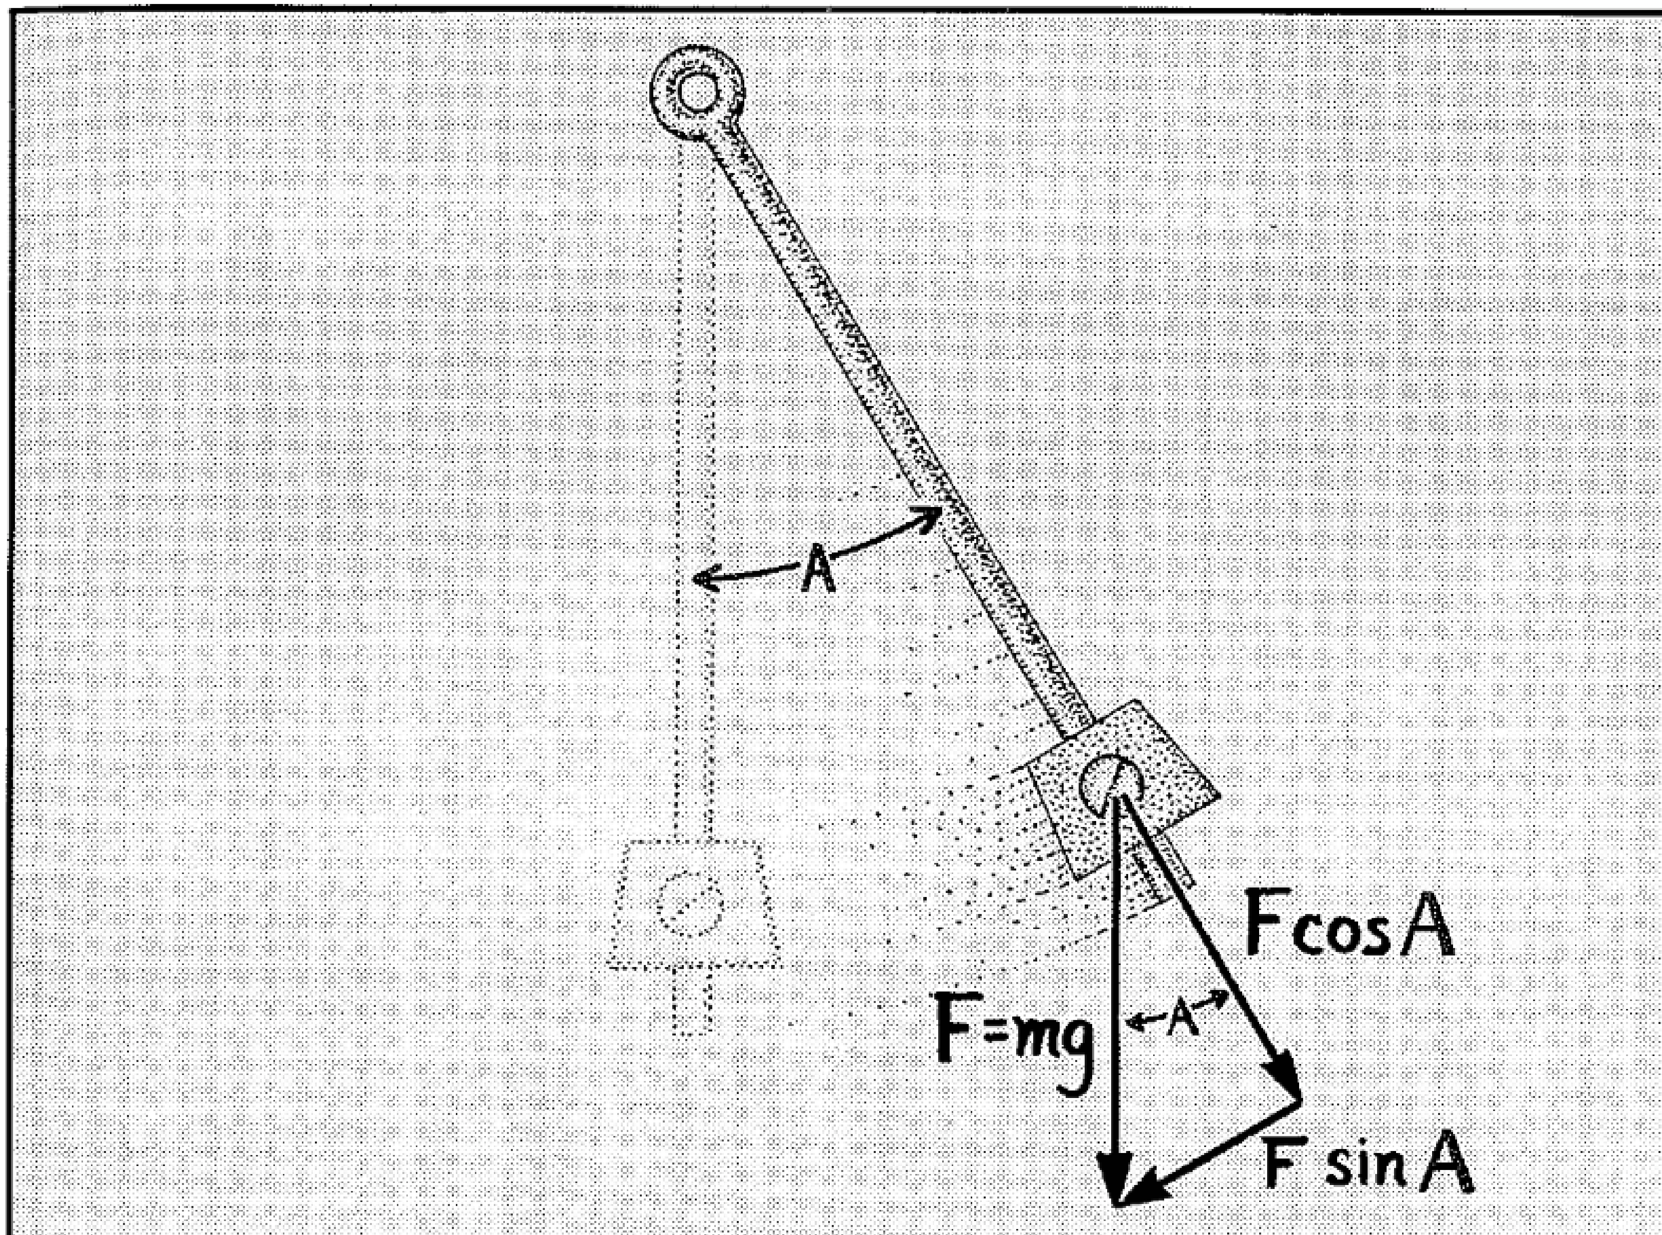

2.1.1. If A denotes the angle of elevation of the pendulum and F the force of gravity, then $F \cos A$ is the pull along the rod, and $F \sin A$ the force turning it, as shown here.

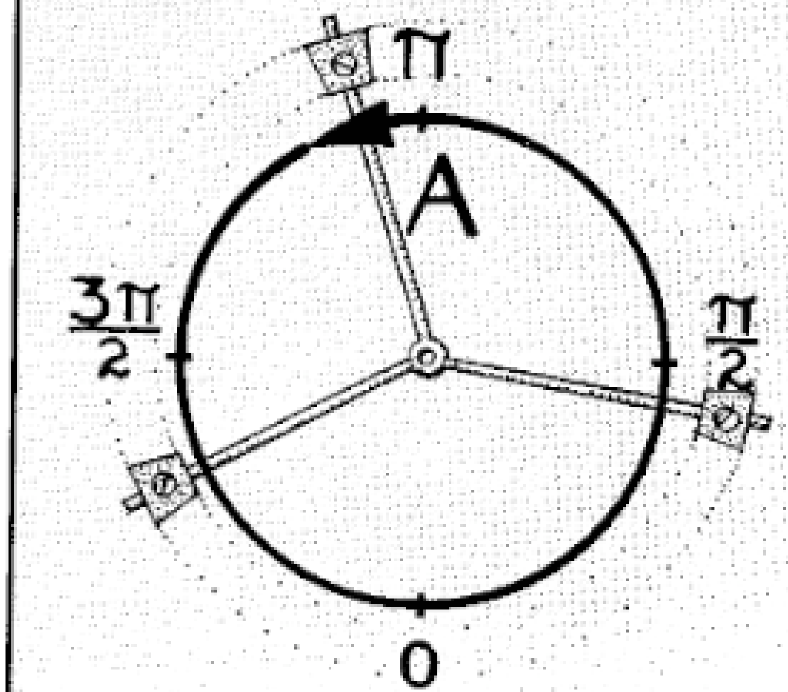

A is the angle of elevation. It is a circular variable

2.1.2. The angle of elevation, A, parameterizes a circle. That is, values of A can be any real number, but $A = 0$ and $A = 2\pi$ denote the same angle. The angle A represents a point of the circle. It is called an *angular variable*.

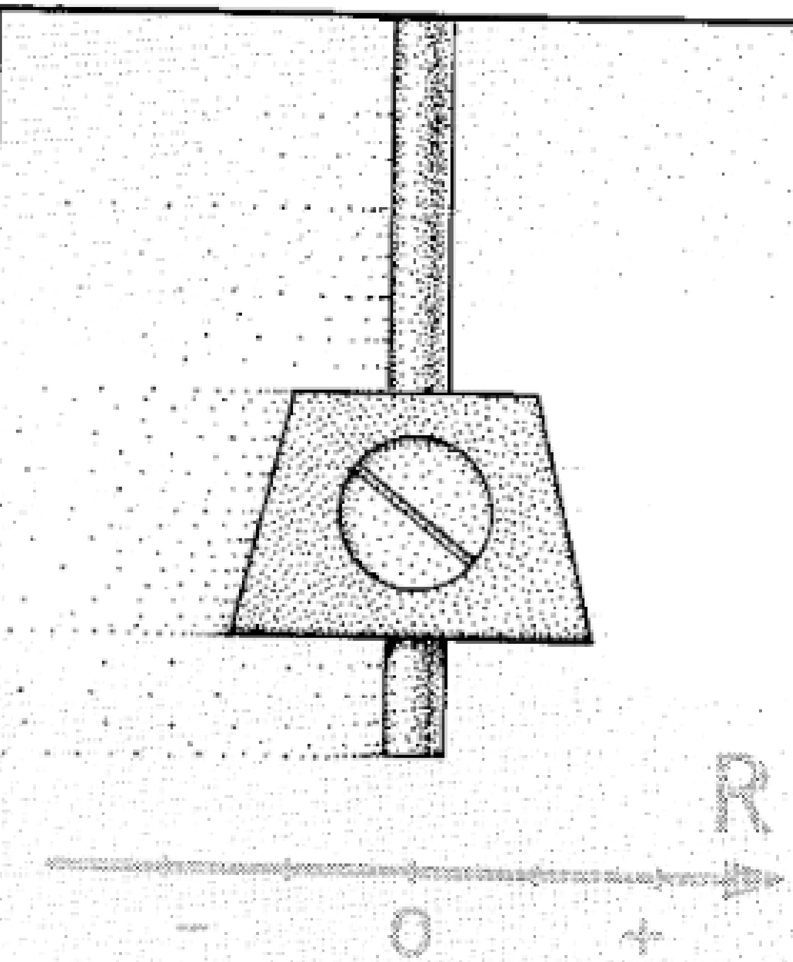

R is the rate of rotation. It can be any real number

2.1.3. Let R denote the rate of rotation of the rod at a given moment. This rate is also observable, by radar for example. In Newton's model, this parameter is included, along with the angle A, as a descriptor of the state of the pendulum. The rate of rotation, R, may have any real number as its value. It represents a point of the real number line.

2.1.4. The two parameters, A and R, together locate a point on a circular cylinder. This is the state space of Newton's model. The vertical circle in the center of this cylinder denotes the states of zero angular velocity, $R = 0$. The straight line from front to back, at the bottom of the cylinder, is the axis of zero inclination, $A = 0$, where the pendulum is lowest.

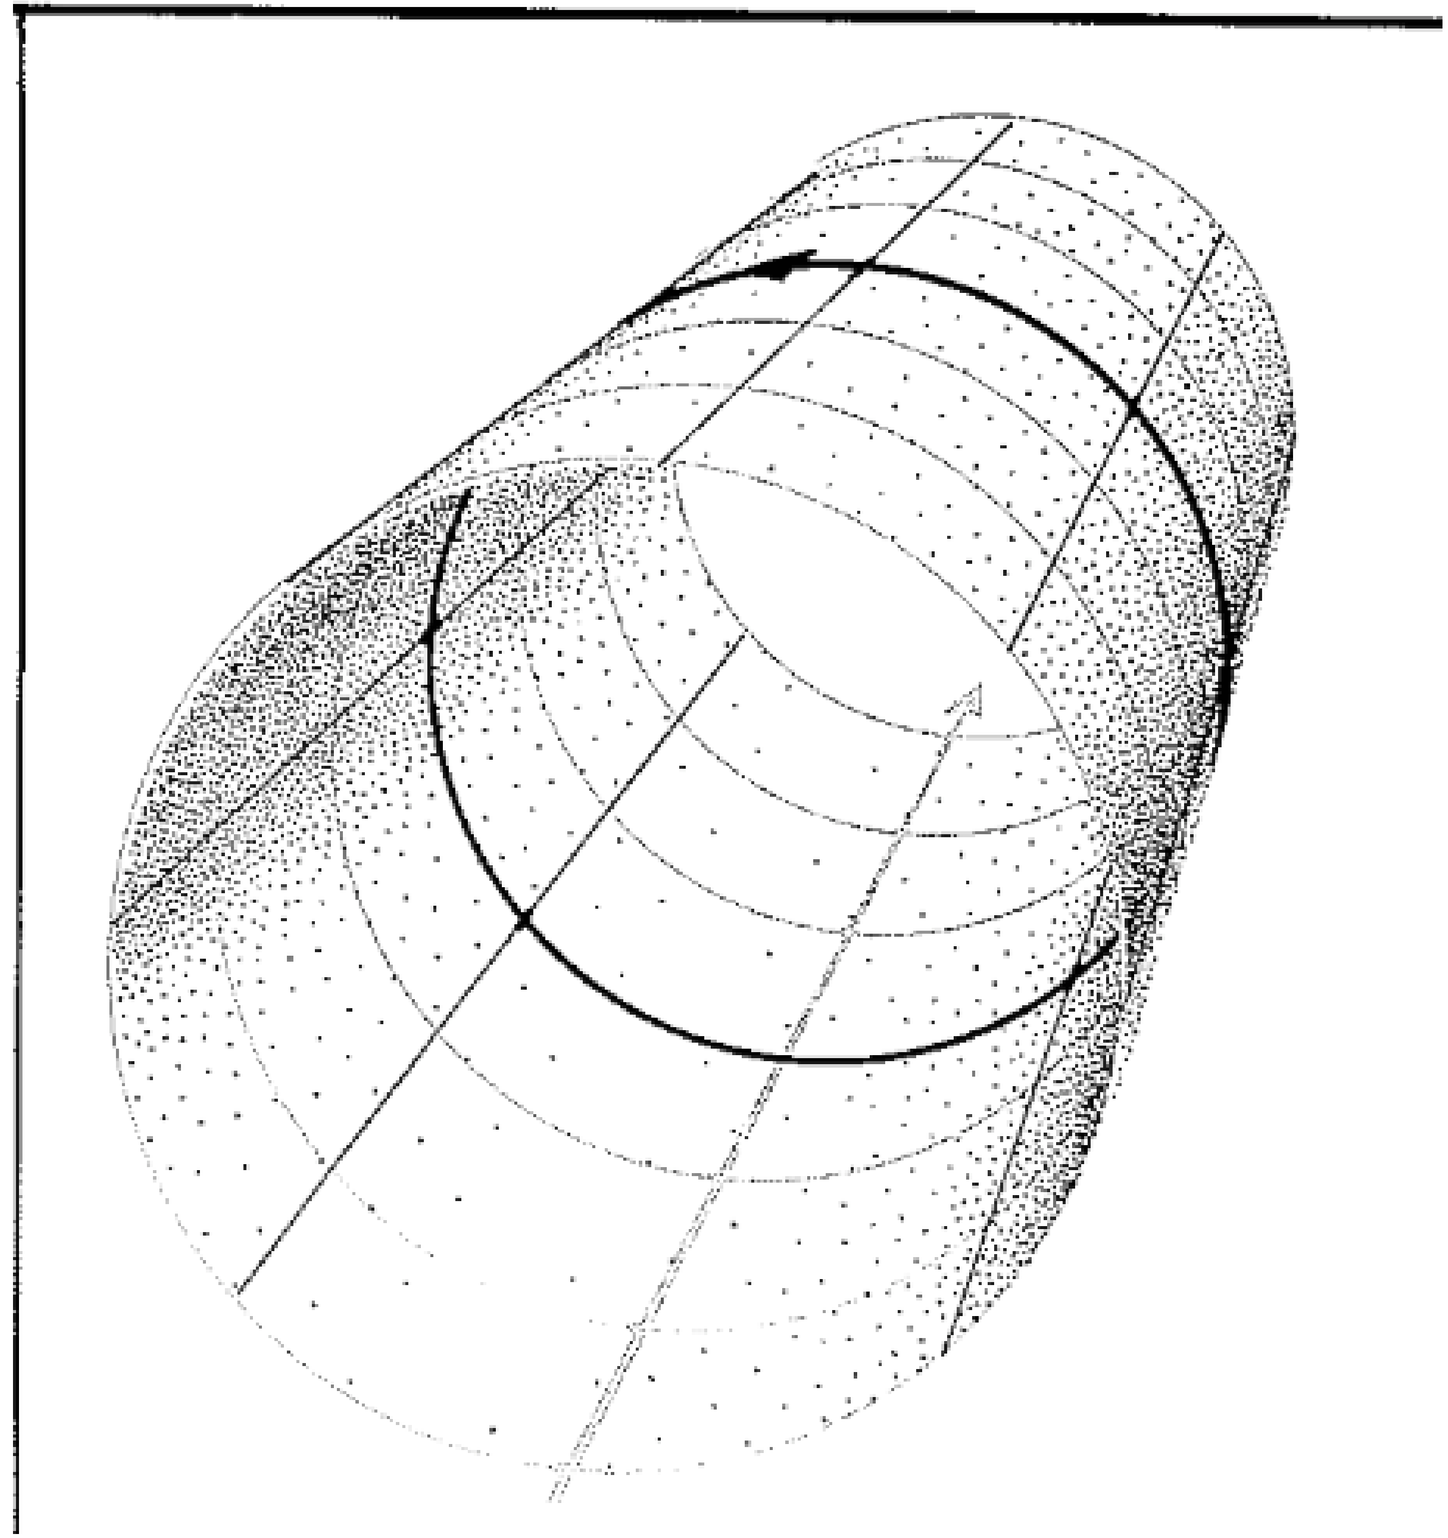

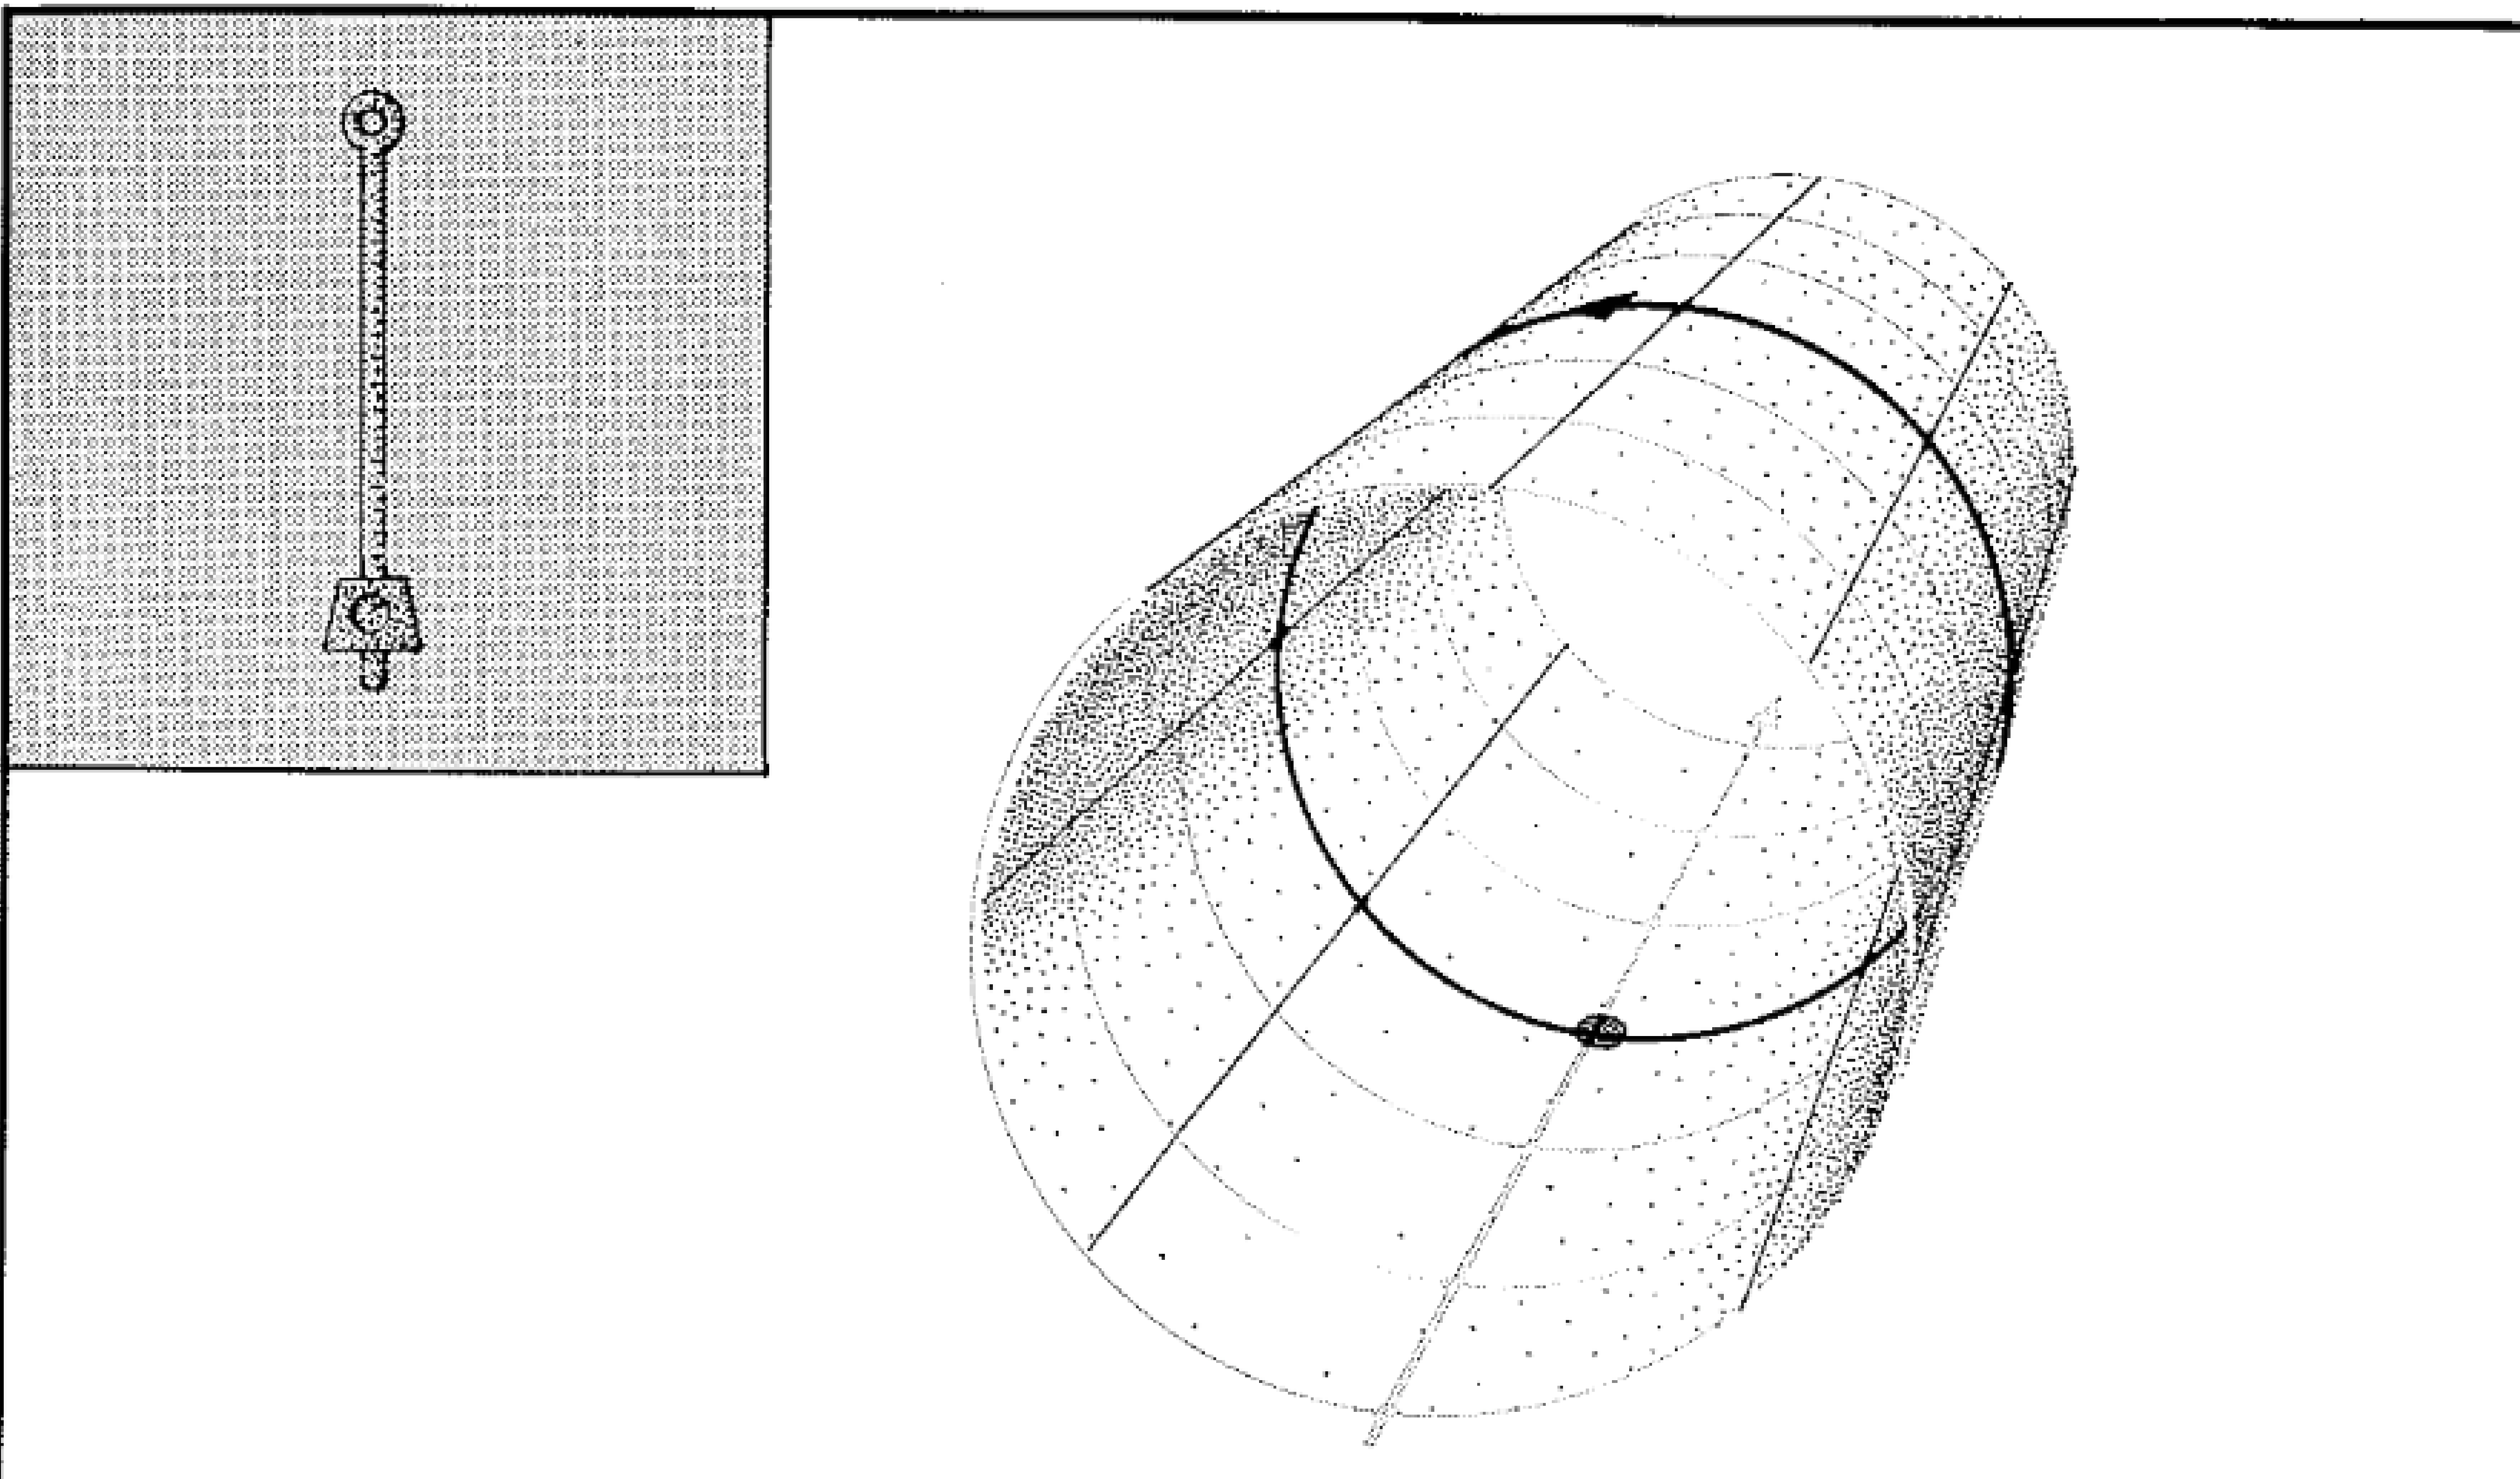

2.1.5. At the *origin*, defined by $(A,R) = (0,0)$, the pendulum is at rest at its lowest position.

Moving the pendulum a little to the left and then letting it go with no shove causes it to swing indefinitely. Remember, there is no friction in the hinge and no air in the way. The representation of this motion as a trajectory in Newton's model is shown here in four steps.

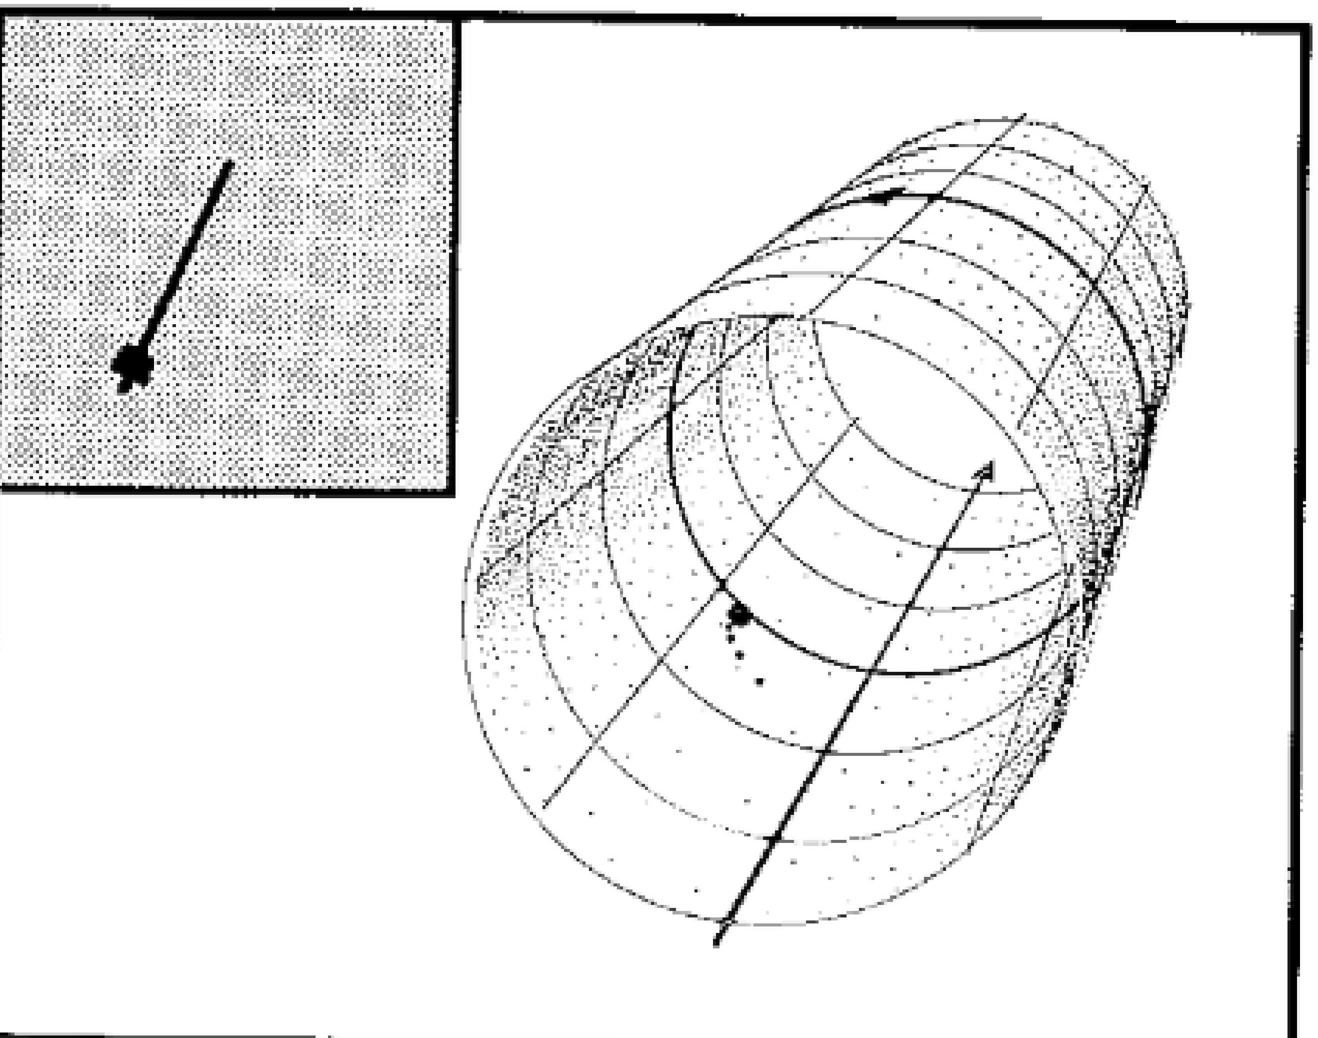

2.1.6. Step 1. Immediately after the pendulum is released, the representative point is on the circle of $R = 0$ to the left of the origin, moving away from us as the rate R increases.

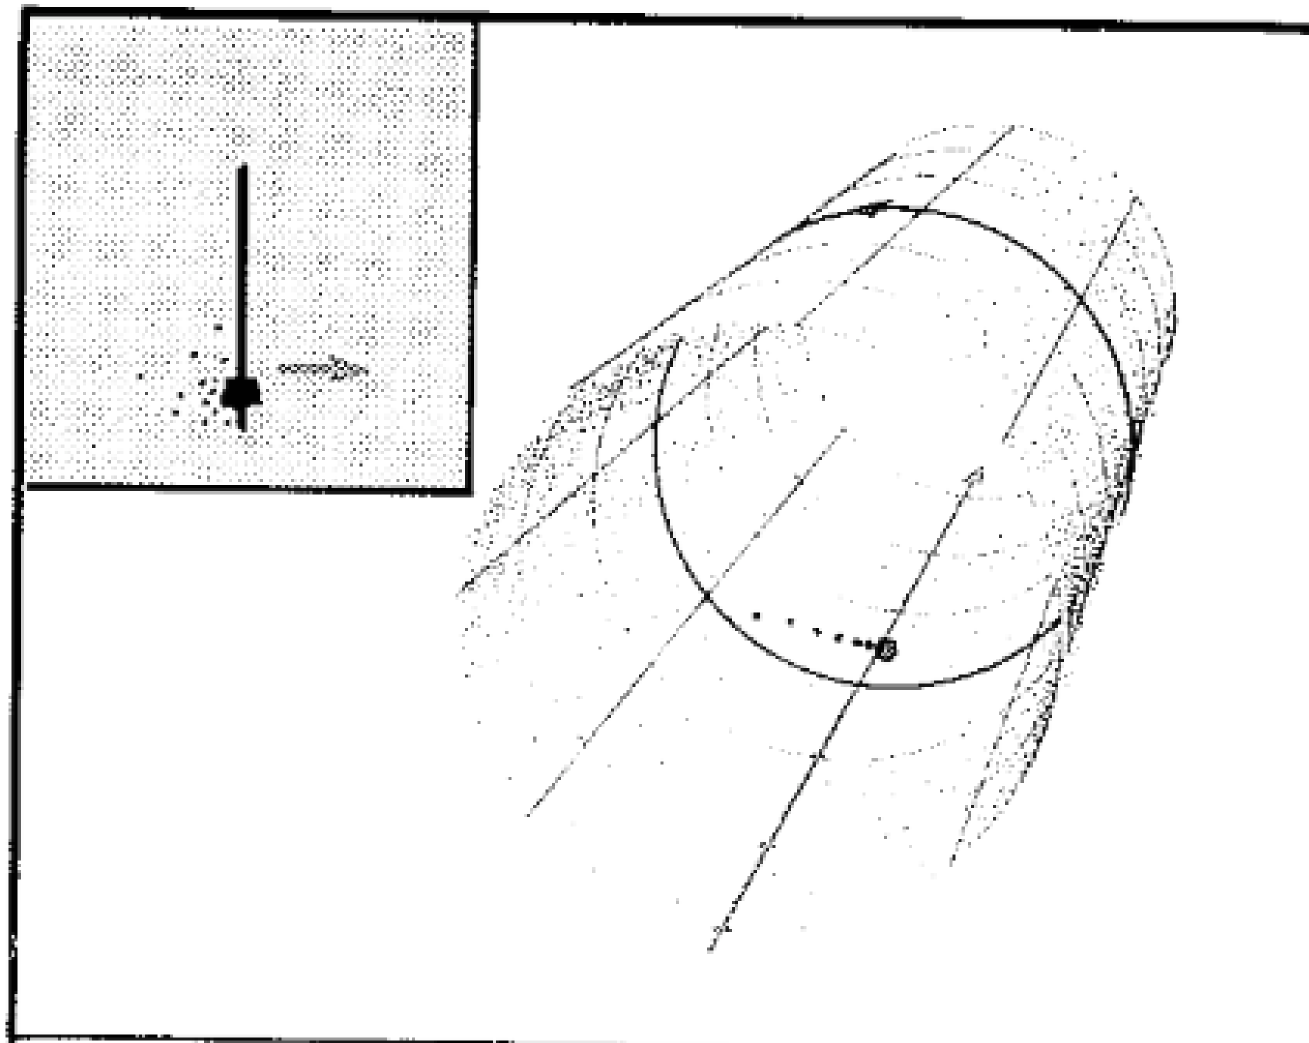

2.1.7. Step 2. It also moves to the right as the inclination increases. Here it has just reached the axis, $A = 0$, as the pendulum goes by its bottom point.

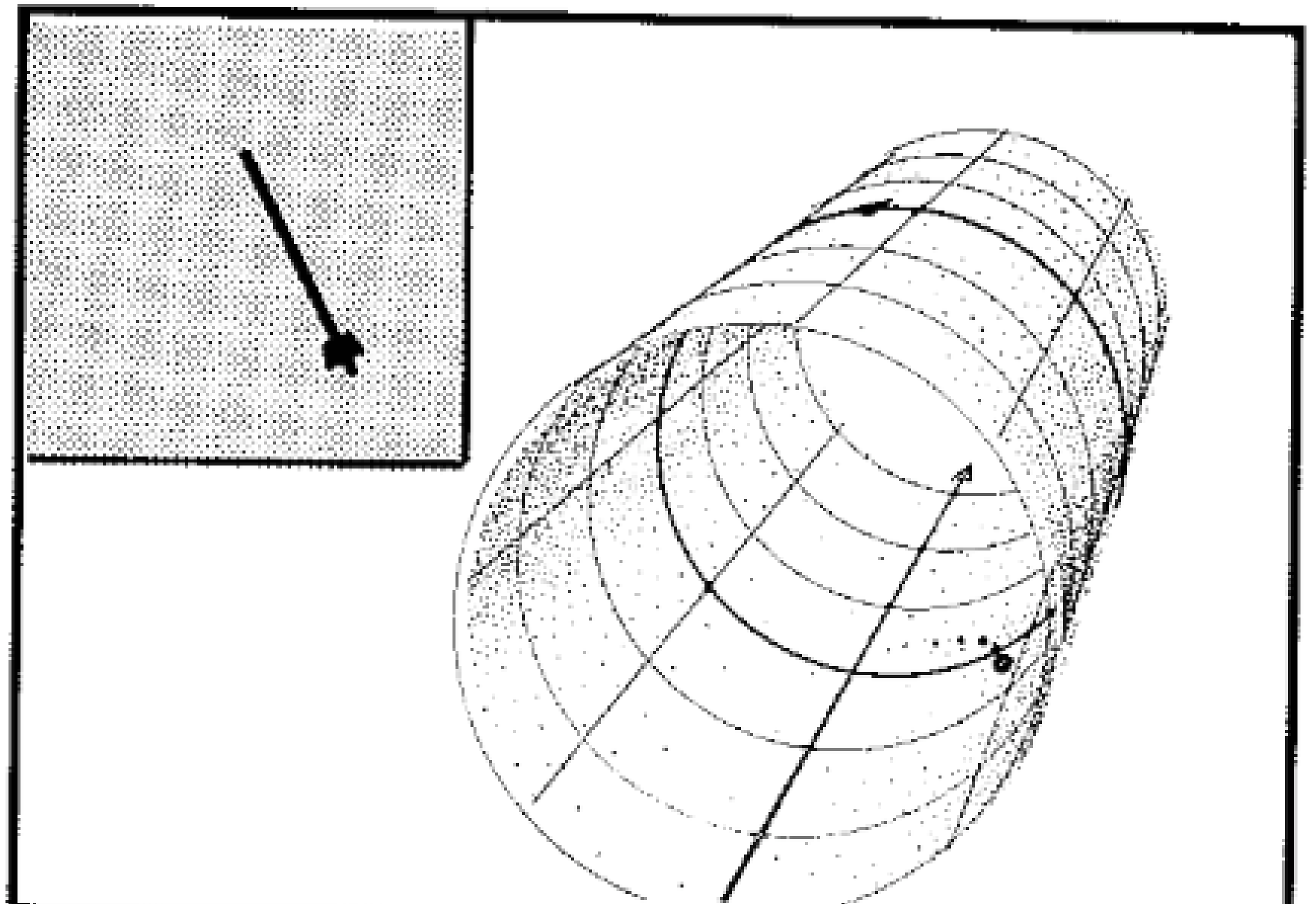

2.1.8. Step 3. It continues to move to the right, moving towards us rather than away, as R decreases. It reaches the circle of $R = 0$, when the pendulum attains its maximum swing to the right, and turns to fall again toward its bottom, $A = 0$.

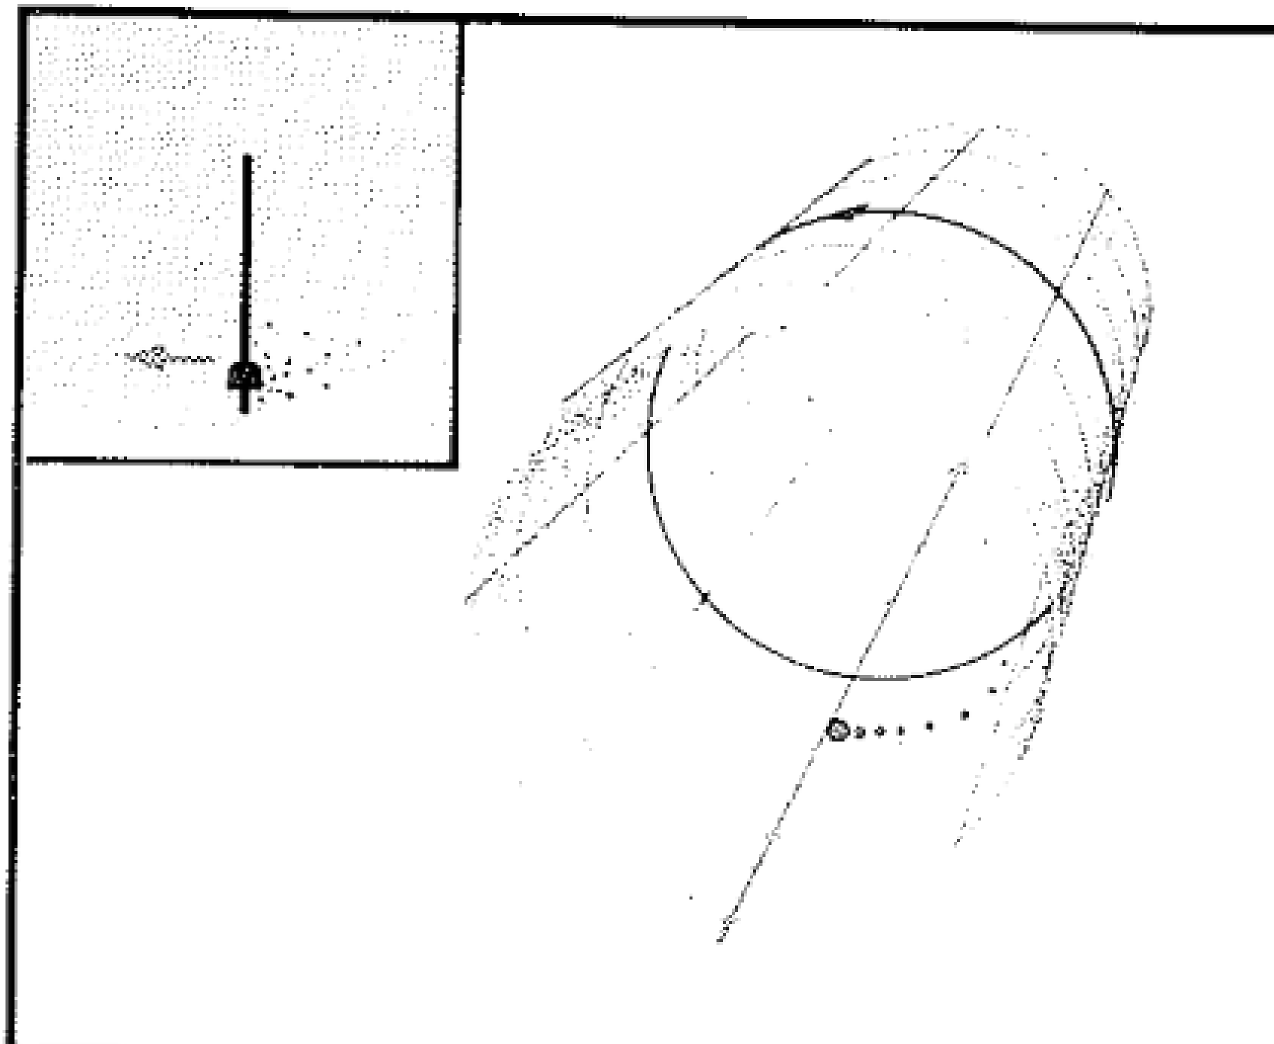

2.1.9. Step 4. It approaches us and moves to the left, as the pendulum falls. It crosses the axis, $A = 0$, as the pendulum swings through bottom.

Then the cycle begins again. The full trajectory in the state space, corresponding to this oscillating motion of the pendulum, is a cycle, or closed loop.

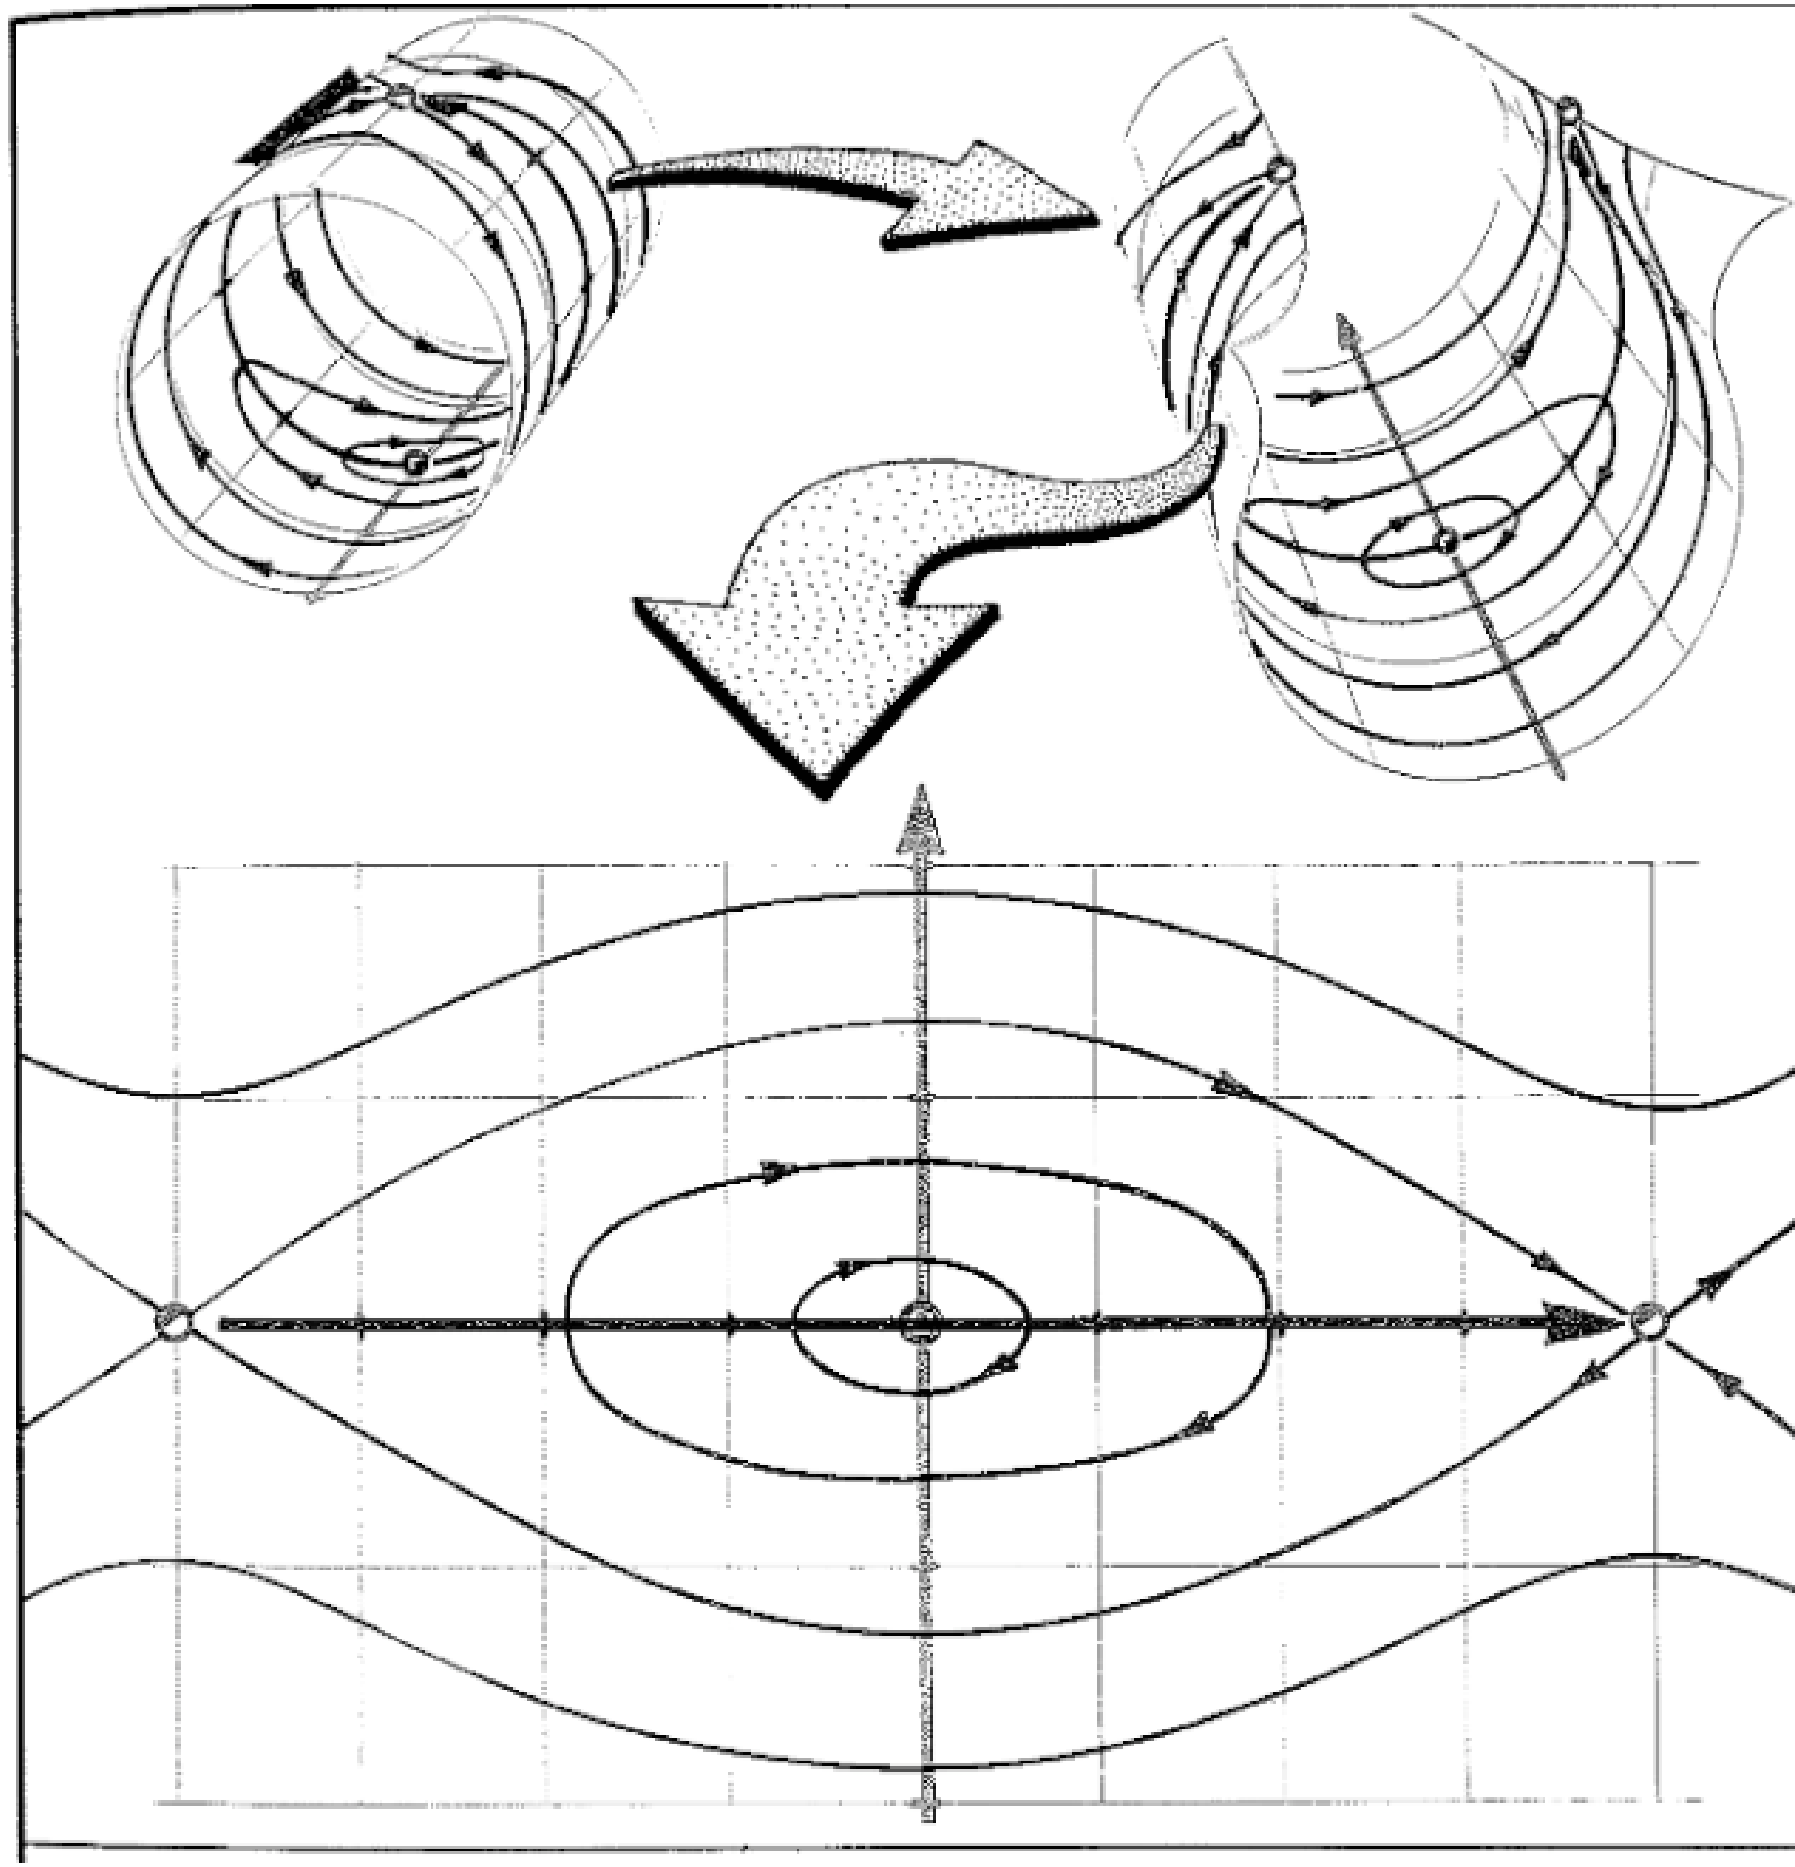

Adding a little friction

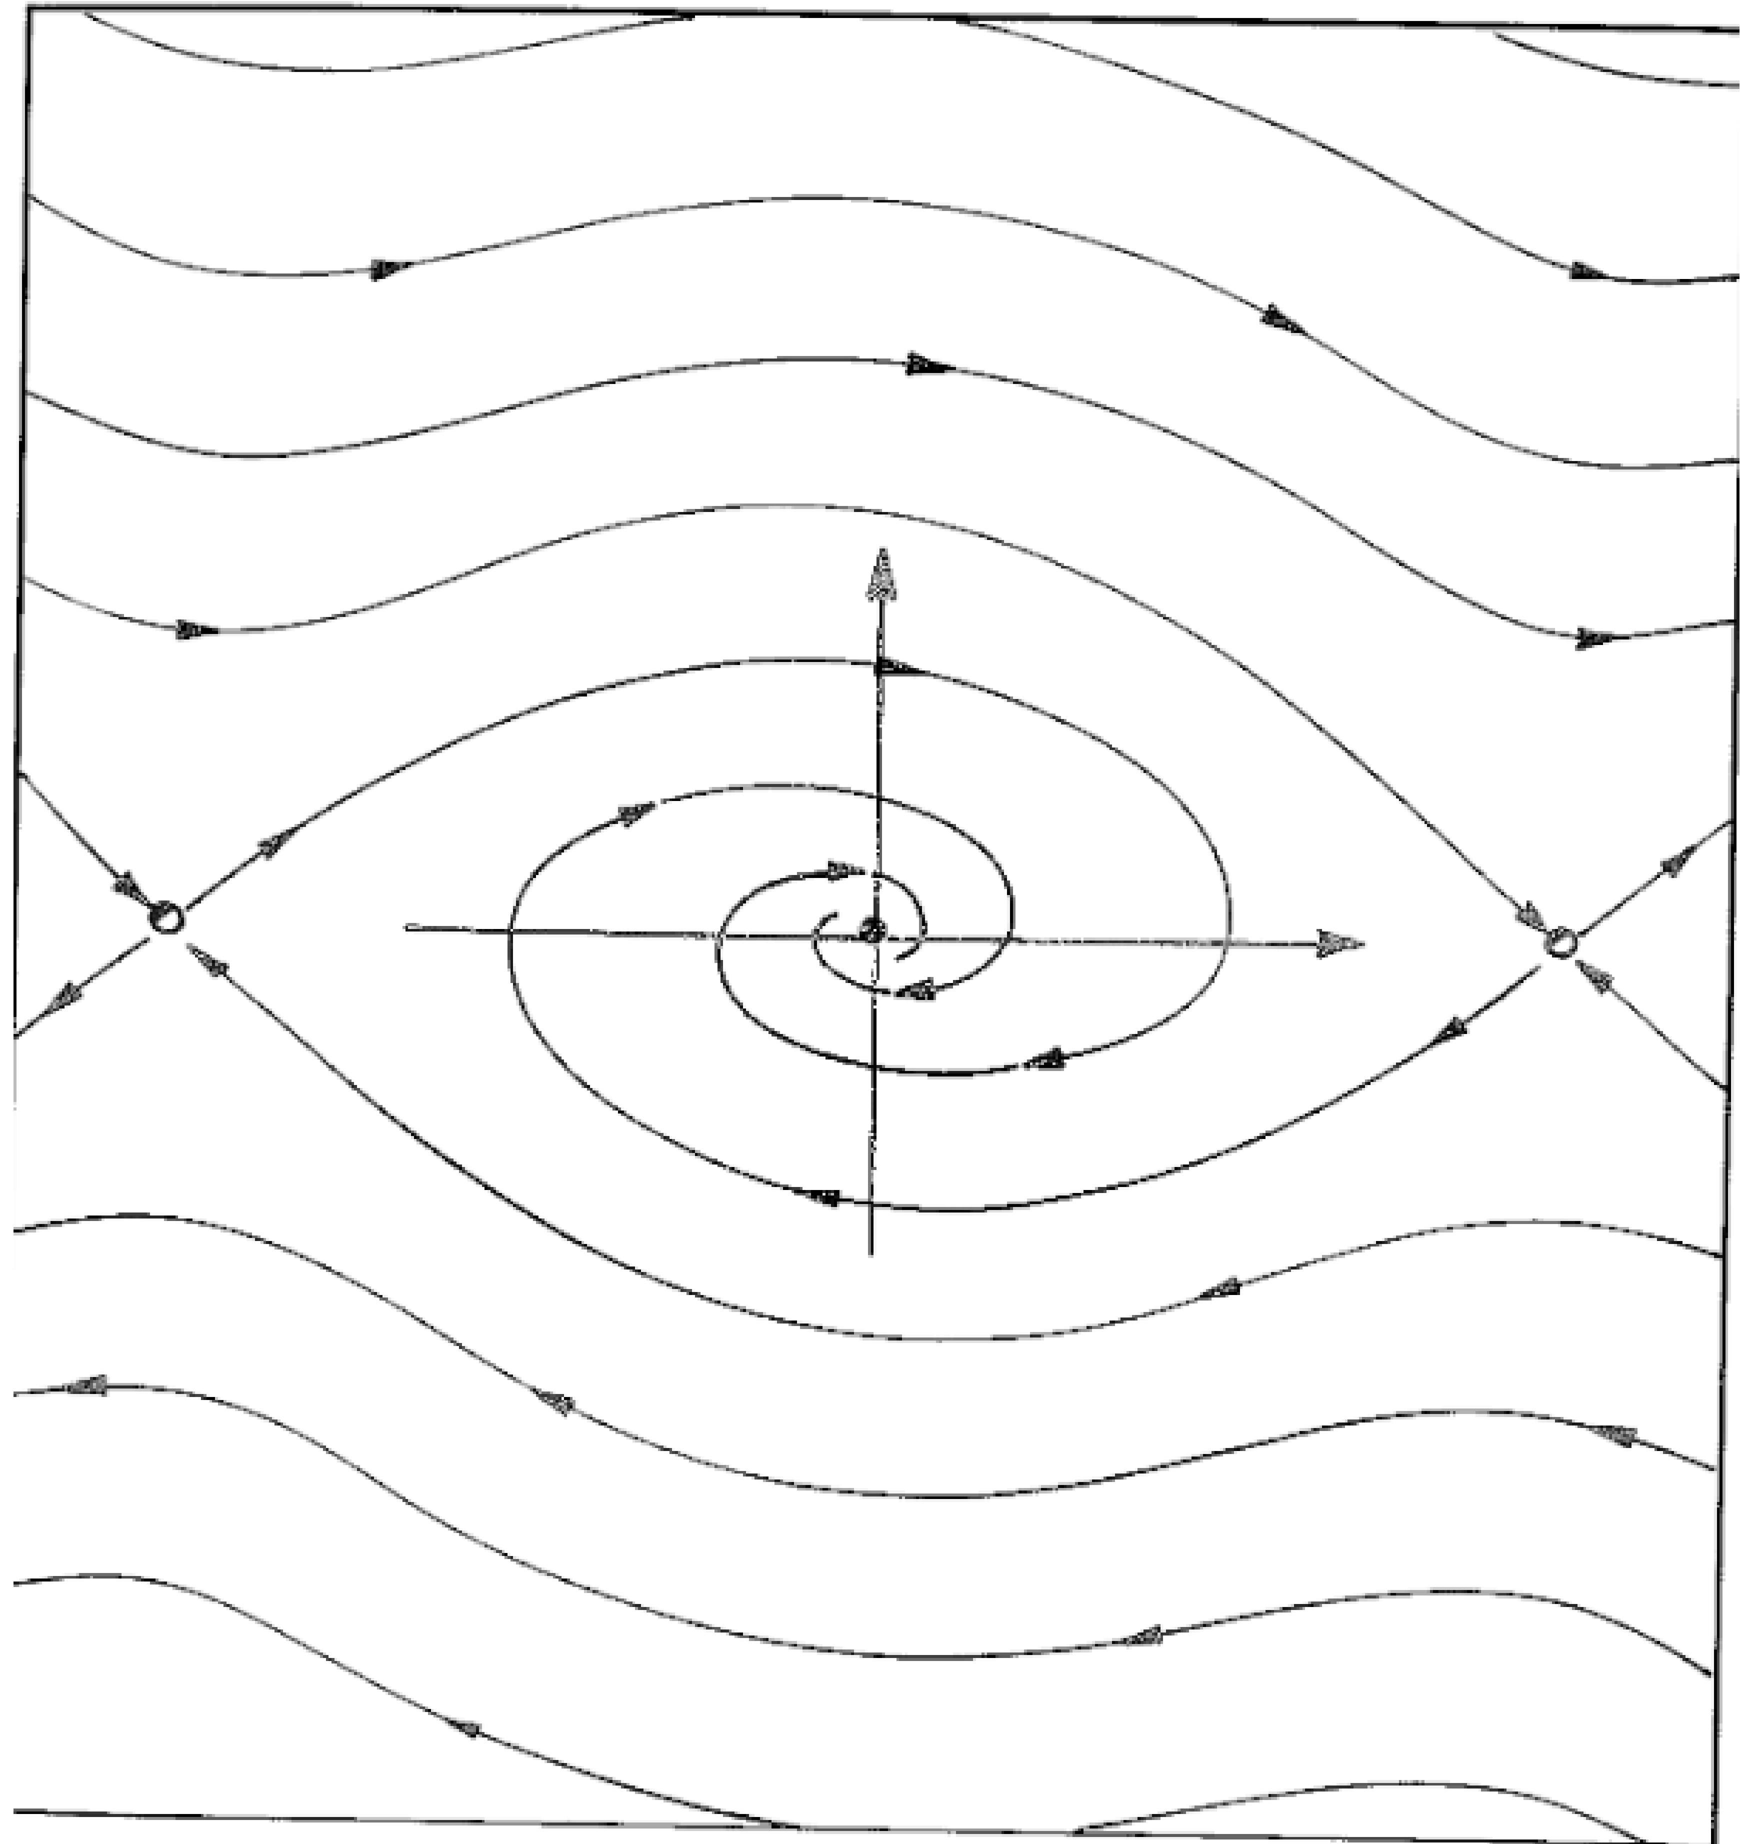

As we gradually alter the parameters of a system, we may cause a sudden shift in the asymptotic behavior:

Reducing gravity, would result in a slowed rate of fall, until gravity became negative........

A system which tends towards a fixed point, may develop oscillations

Oscillations may become irregular, or chaotic...

Or a complex system may collapse to a single fixed state........

DST is the lingua franca of

Physics (Newtonian Mechanics)

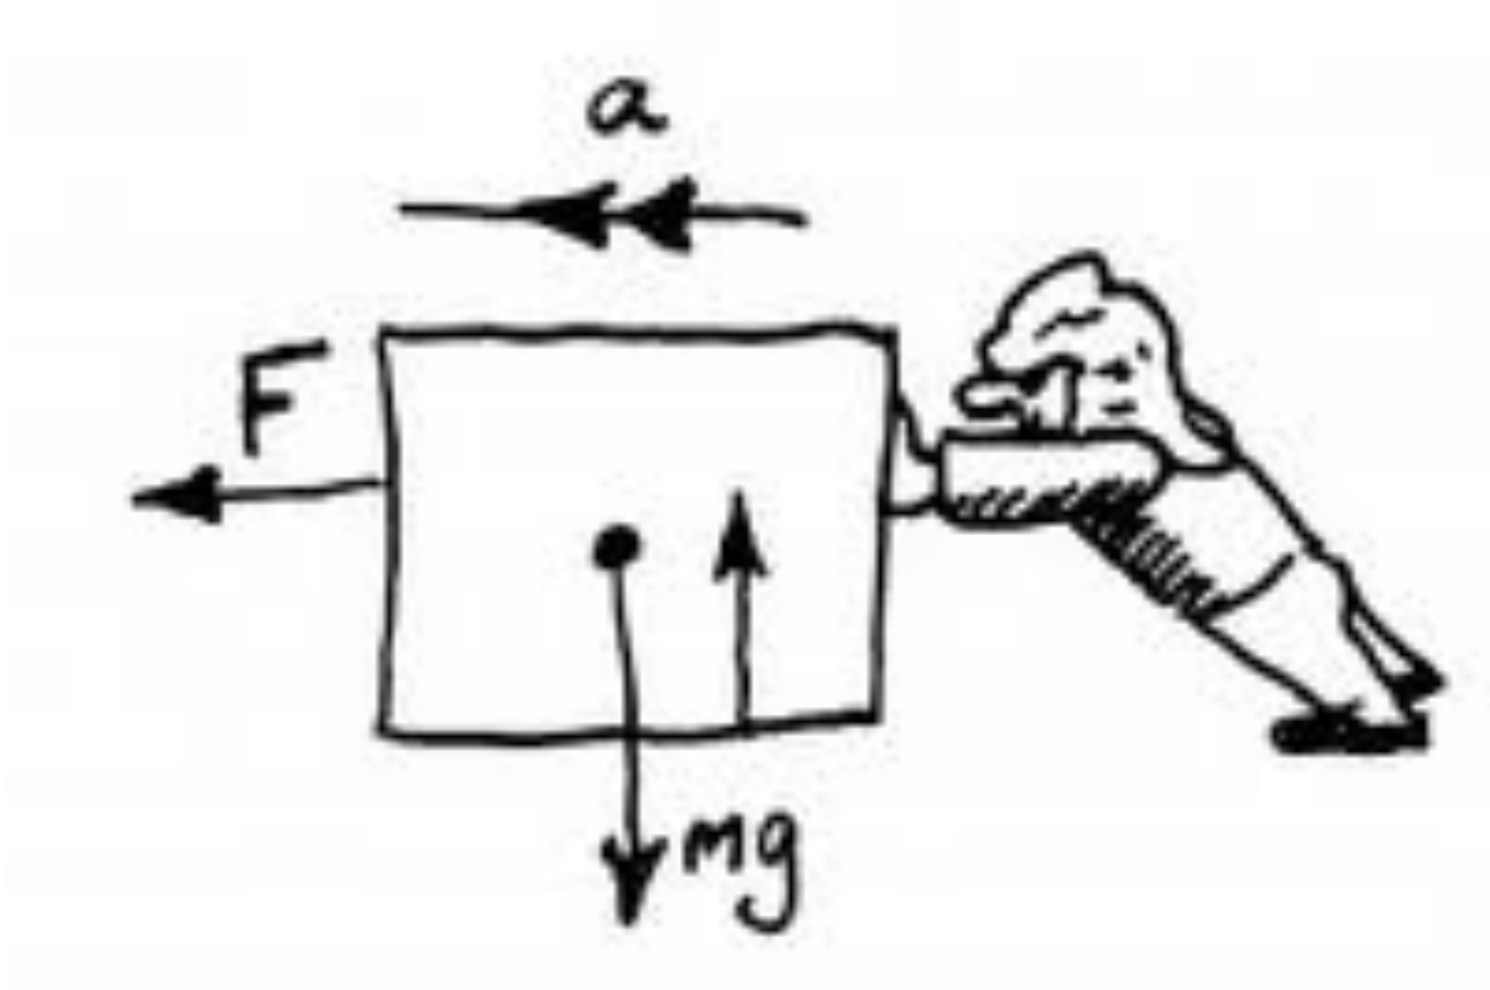

DST is the lingua franca of

Physics (Modern)

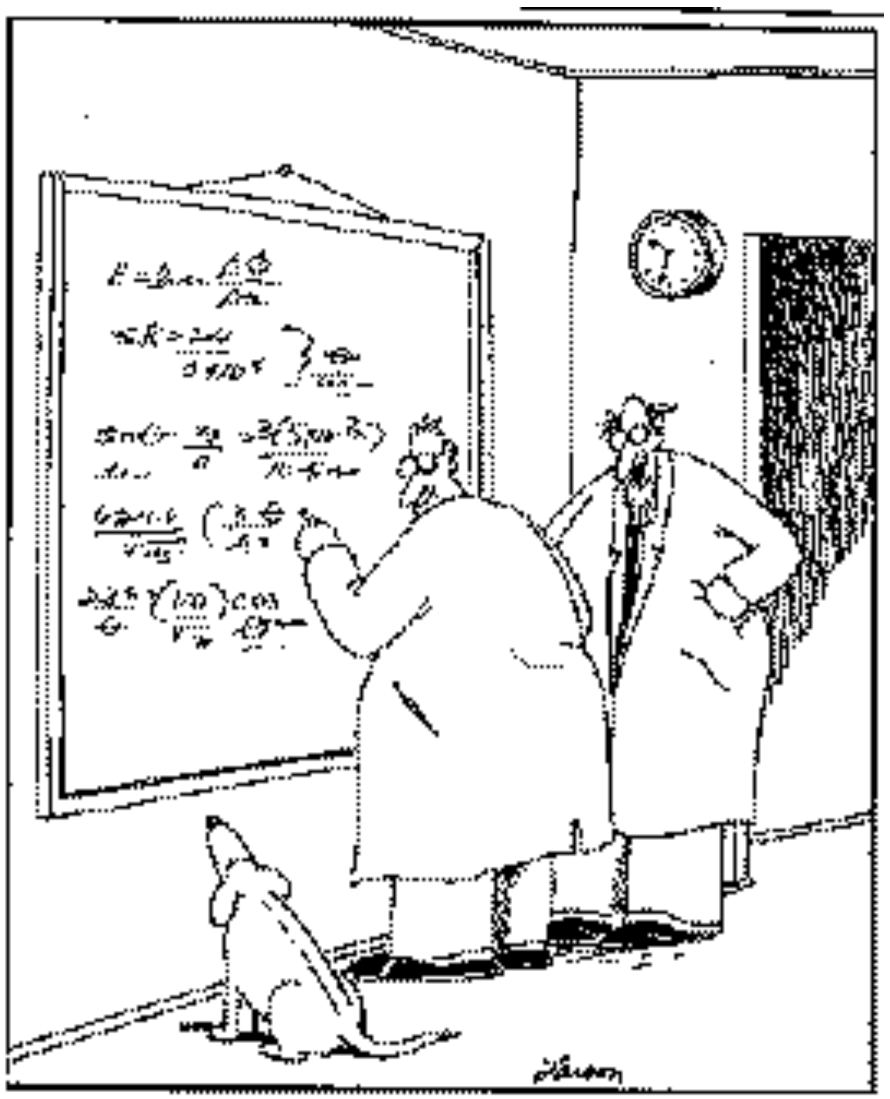

They are so cute when they try to understand quantum mechanics

DST is the lingua franca of

Biology

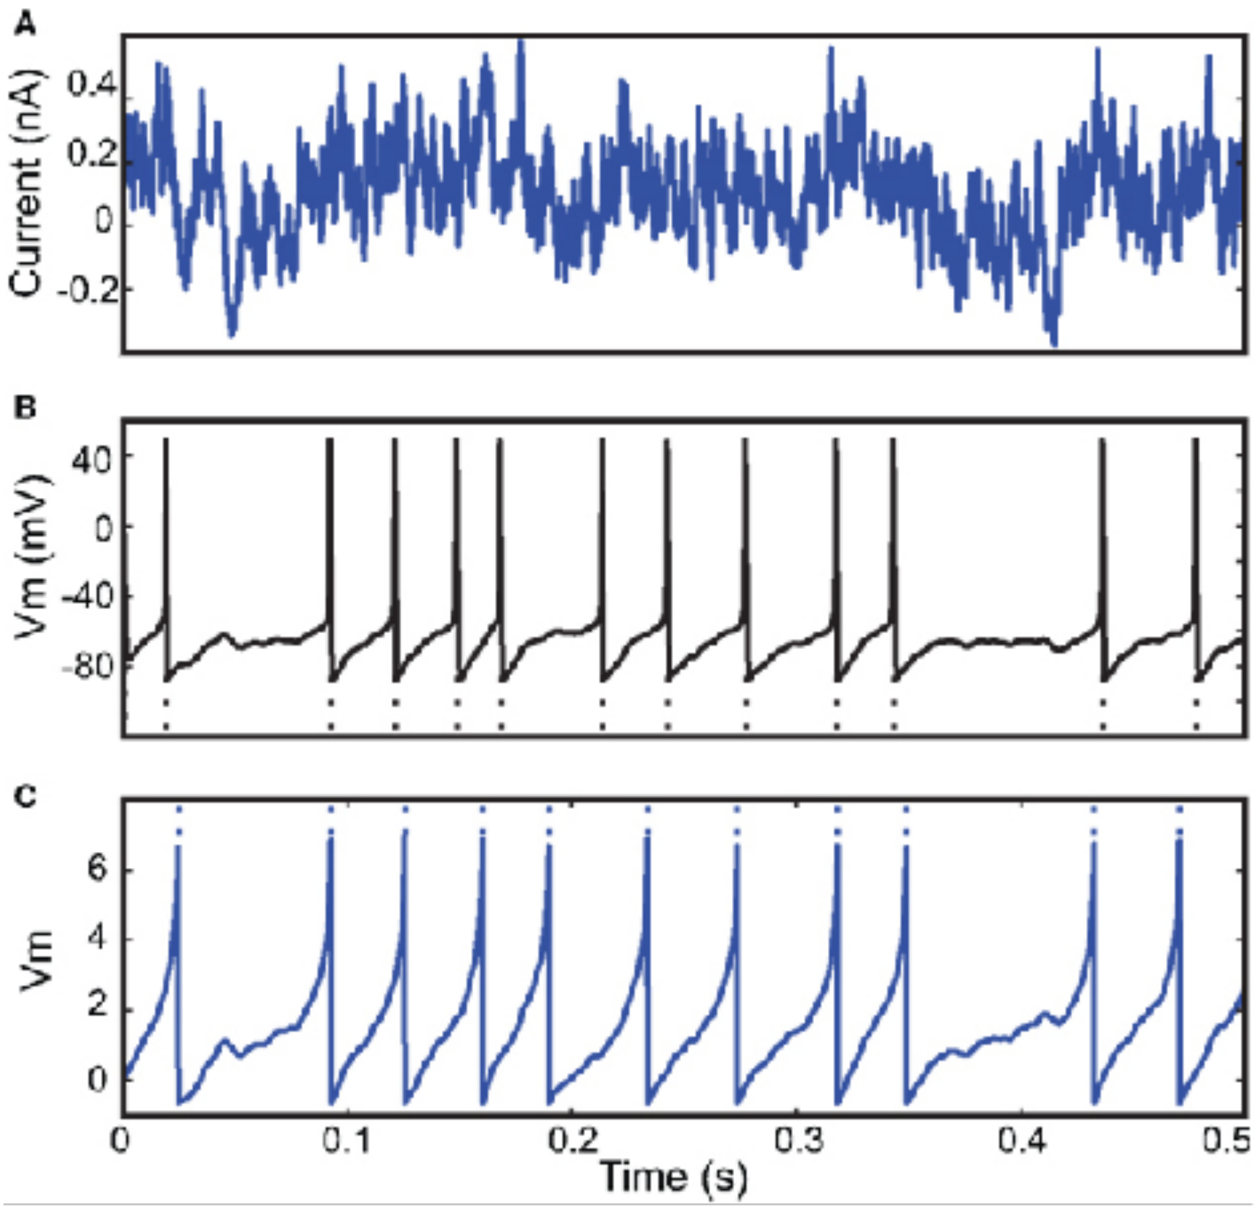

DST is the lingua franca of

Economics (??!!!)

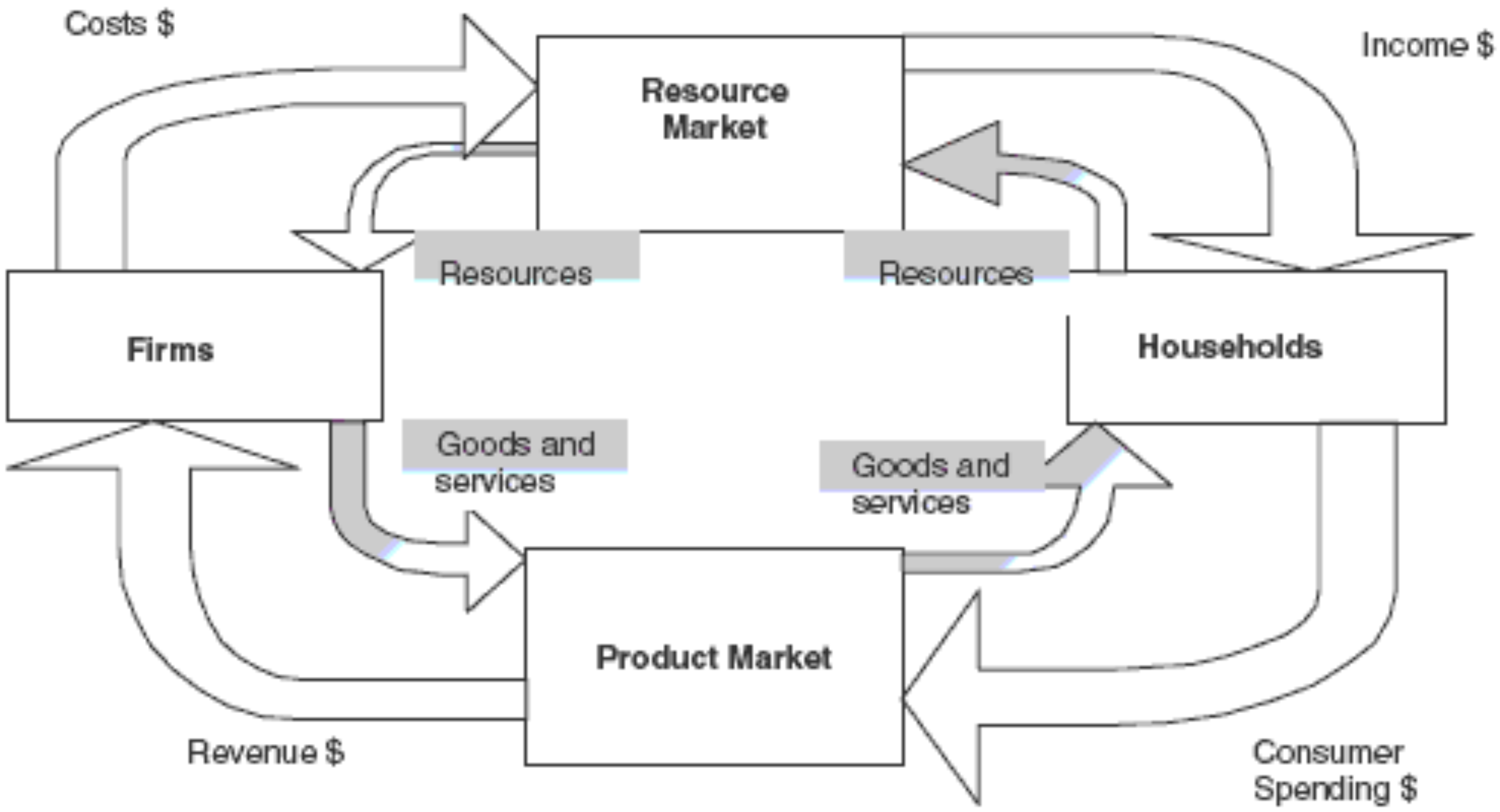

Figure 12.1

DST does not give us a high road to the truth.

It *may* allow us to understand phenomena in a unique and insightful way.

The seductive appeal of mechanics

Sir Isaac Newton

Newton's Second Law of Motion

$$F = ma$$

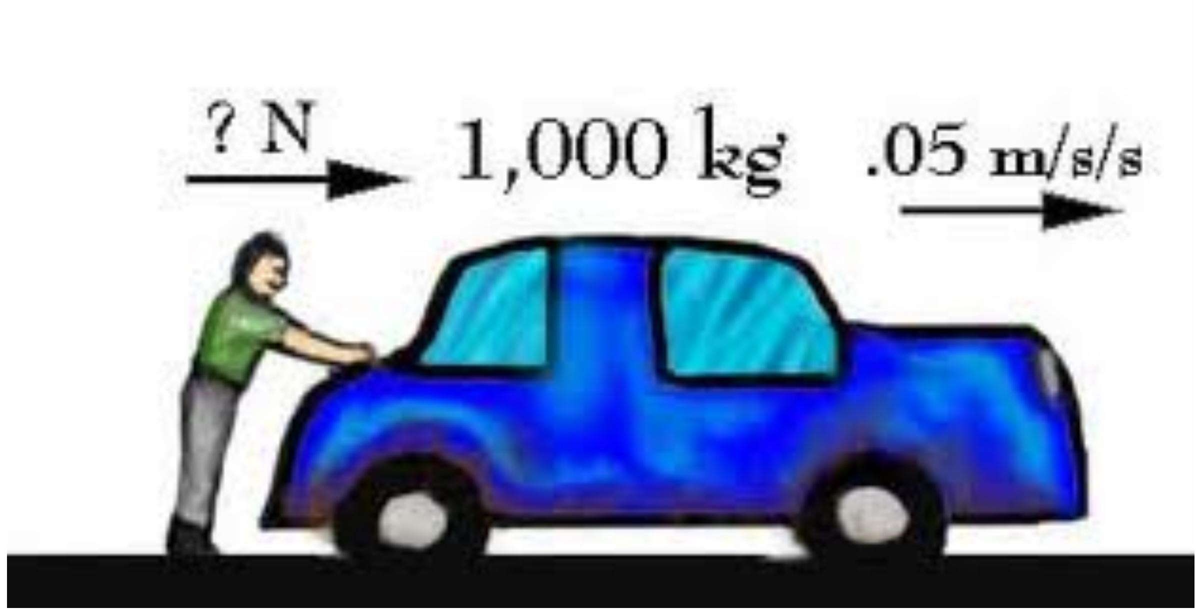

Seeks to provide a deterministic account of the relation between *motions* of *massive bodies* and the *forces* acting upon them

DST approaches go far far beyond mechanics!

The *relevant features* of components are described by two numbers per component

$$\left.\begin{array}{ll}\text{Position} & x \\[2em] \text{Velocity} & \dot{x}\end{array}\right\}$$

STATE

$$\ddot{x} = f(x, \dot{x})$$

DYNAMIC

In the domain of *mechanics*, where we seek to account for the motion of massive bodies, we know what variables to choose. Position and Velocity of point masses. Always.

The description of the lawful change of state over time extends far beyond mechanics.

In most domains, and for most systems, we will not have an explicit dynamic available to us.

But the language of dynamics continues to be of use in describing systems and their interactions.

This works for discrete dynamical systems too:

State describes the entire system at x_t

A dynamic is a rule expressing $x_{t+1} = f(x_t, \ldots)$

Most of what follows describes continuous dynamical systems, but we will use discrete systems where it can provide us with useful examples.

DEFINITION:

A Dynamical System comprises

[1] A *state* description

[2] A rule governing state change over time
(= a *dynamic*)

We might consider the "system" to be just a few pieces of the body

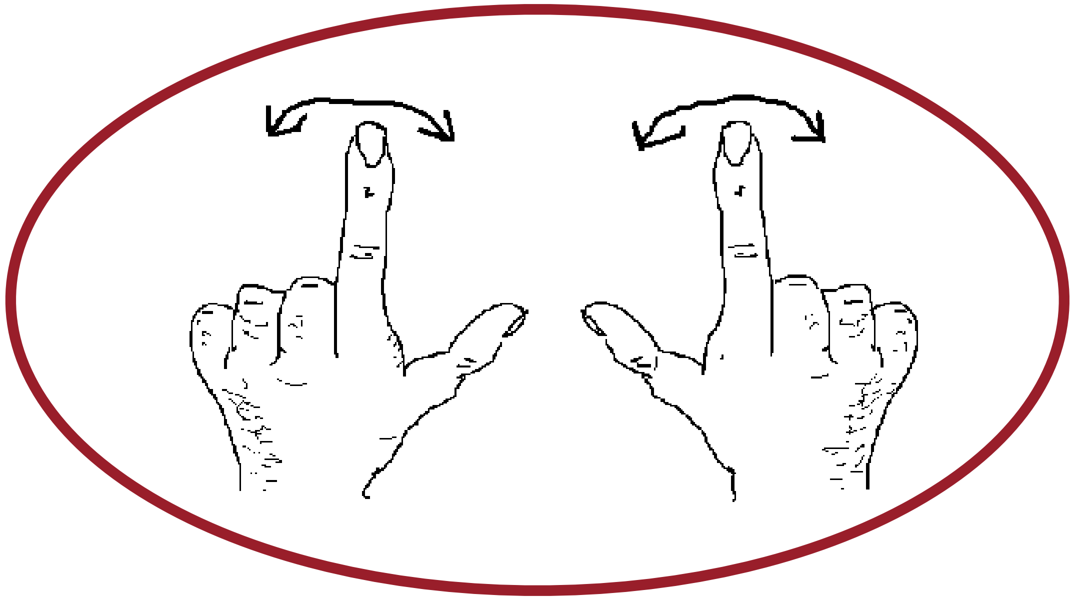

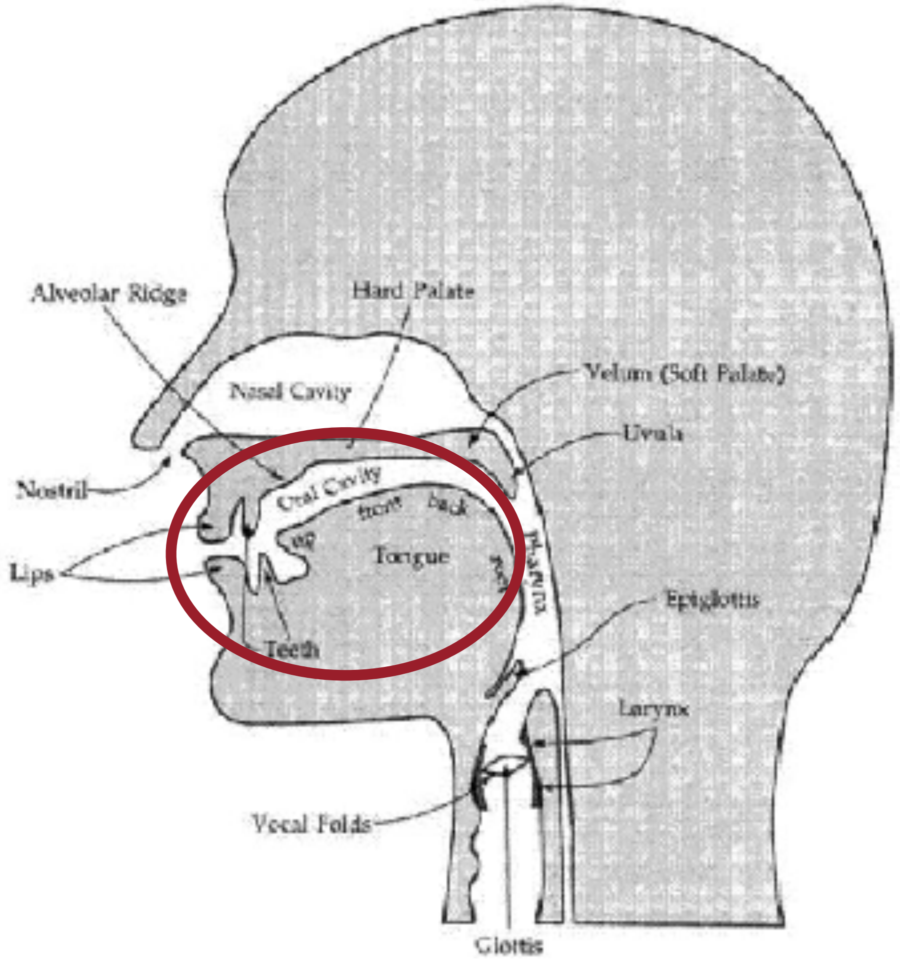

Or the system might be the whole body:

Or body + tool

Or multiple bodies . . .

Important point:

We do not know in advance what the "system" is that we are dealing with. That will depend on the phenomenon or behavior being studied.

Related point:

There is no point in arguing that something *is* or *is not* a dynamical system. DST may provide a useful way of regarding a system, or it may not.

Thought for the day

Heisenberg: Science alters and refashions the object of investigation. In other words, method and object can no longer be separated.

Suppose we have some dynamic system which we are fortunate enough to be able to express as a differential equation(s):

$$\frac{d\mathbf{x}}{dt} = f(\mathbf{x}, \mathbf{P})$$

We might want to know how a specific instance of this system behaves over time.

For example, we have an equation describing a *simple harmonic oscillator*, such as an ideal mass-spring system:

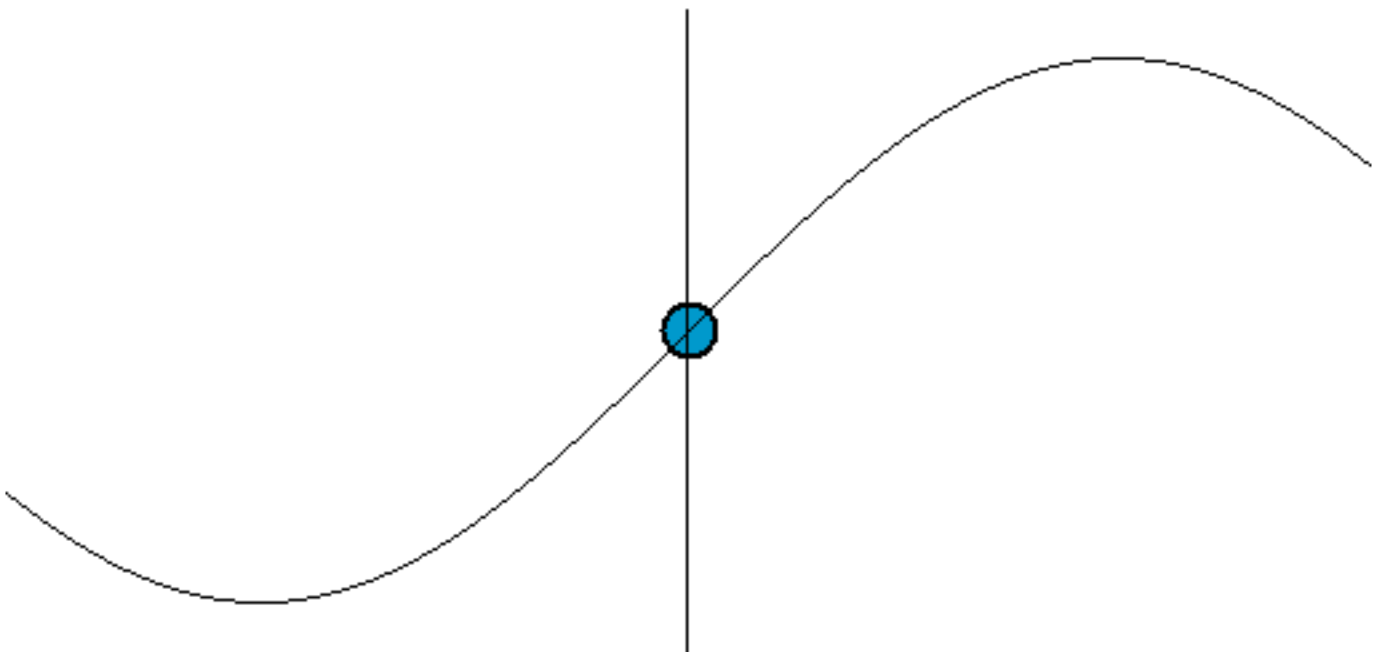

$$\frac{dx_1}{dt} = x_2$$

(x_1 is position; x_2 is velocity = rate of change of position)

$$\frac{dx_2}{dt} = -kx_1$$

But we would like to have an equation for how x_1 changes over time

In this simple case there are means (*waves hands*) to go from the differential equation(s):

$$\frac{dx_1}{dt} = x_2$$

$$\frac{dx_2}{dt} = -kx_1$$

To an explicit statement of how the position, *x1,* varies as a function of time:

$$x_1(t) = x_1(0)cos(\sqrt(k)t) + x_2(0)sin(\sqrt(k)t)$$

which describes a particular system, starting at ($x_1(0)$,$x_2(0)$)

This is known as a *solution* to the original equation(s)

But for any reasonably complicated (and hence interesting) system, we do not have such means.

This means we cannot usually use a dynamical analysis to predict the future state of a system as a function of time.

The absence of *this kind of solution* does not stop us learning many other kinds of information about the system we are studying.

If we can not solve a system analytically, what can we say about it?

Qualitative analysis:

Transient response to perturbation
Asymptotic behaviors - Attractors
Stability properties
Structural stability
Dimensionality
.....

Attractors, Transients, Initial Conditions

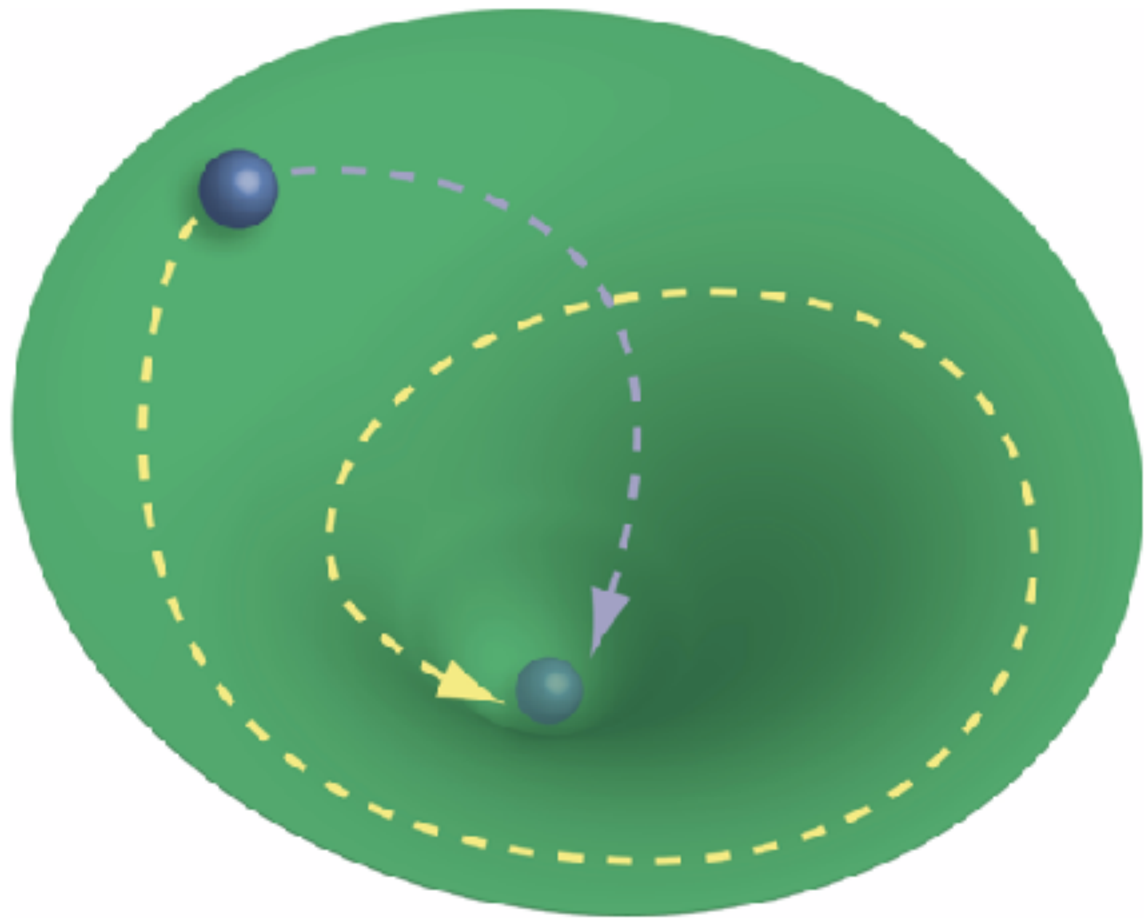

To what kinds of behaviour does a given system tend over the long-term?

The short-term?

For a given set of parameters, and a set of initial conditions, we can watch a system and see what its long-term (asymptotic) behaviour is.

A rolling ball comes to rest at the bottom of the hill

A frictionless ideal mass-spring system oscillates regularly

Weather patterns display chaotic (non-random, but non-predictable, behaviour)

Attractors

If, over the long-term, the system occupies only a limited part of its state space, and if when perturbed to another part, it returns to the former, that part is called an *attractor*.

That volume of state space

There are exactly *three* kinds of attractors in a deterministic dynamical system.

The Point Attractor

Trajectories converge to a single point in state space.

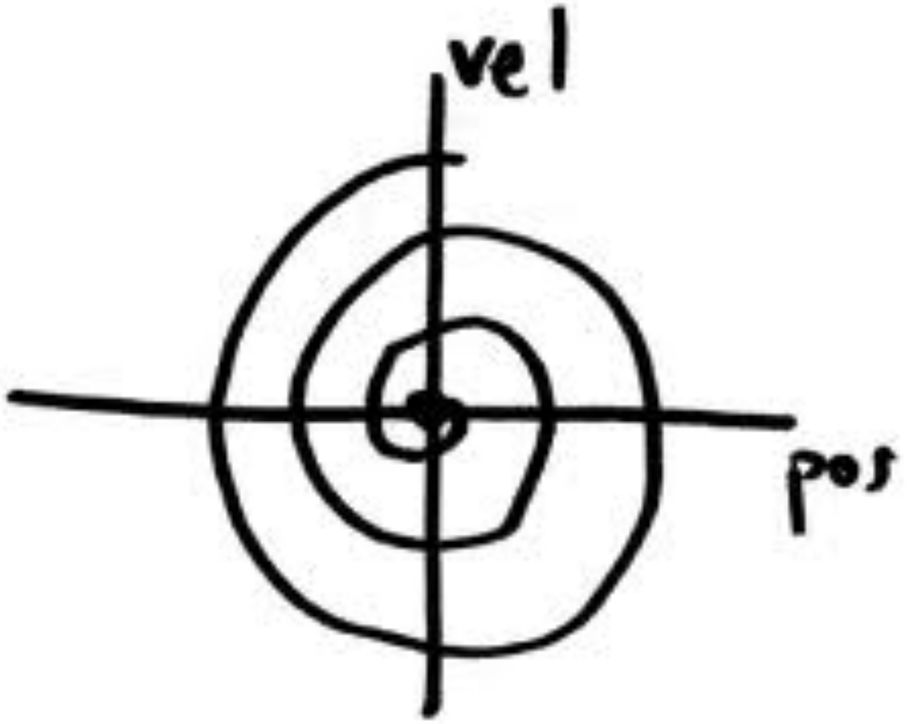

Here, a pendulum (state space = position x velocity) runs down and stops. Always.

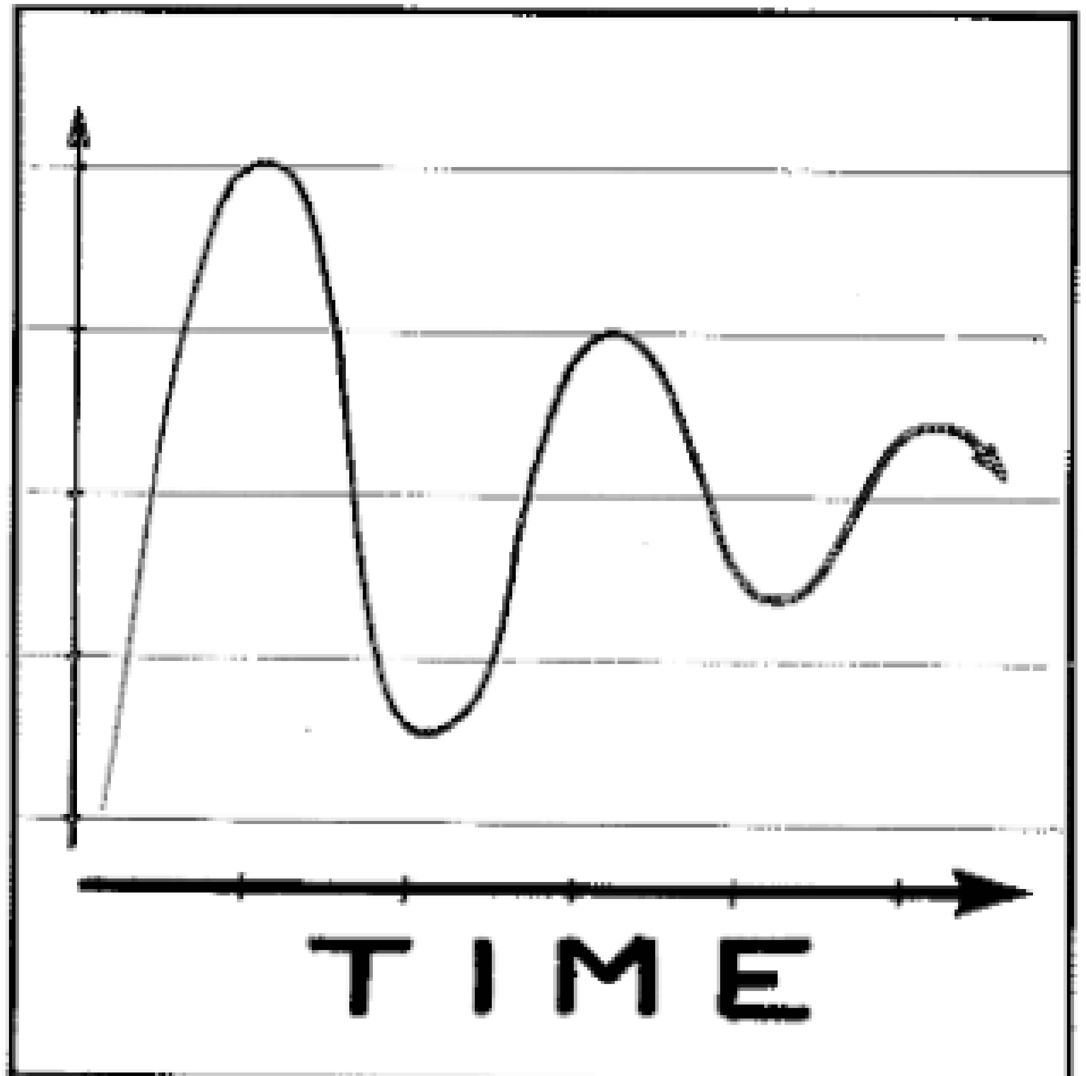

The time series of a plucked string decays.

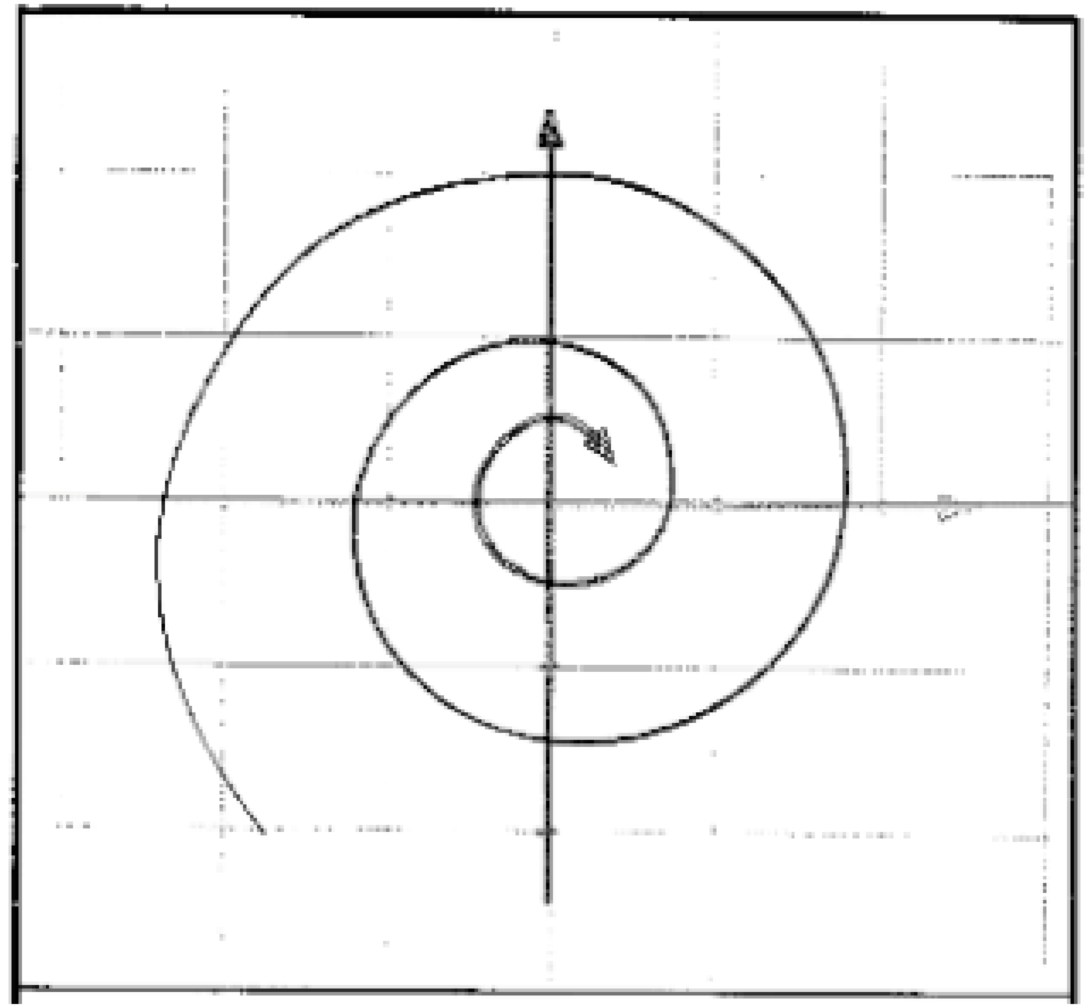

Here is the trajectory in state space (position & velocity are the variables)

The Limit Cycle

An invariant trajectory is
repeatedly traced out
in state space.

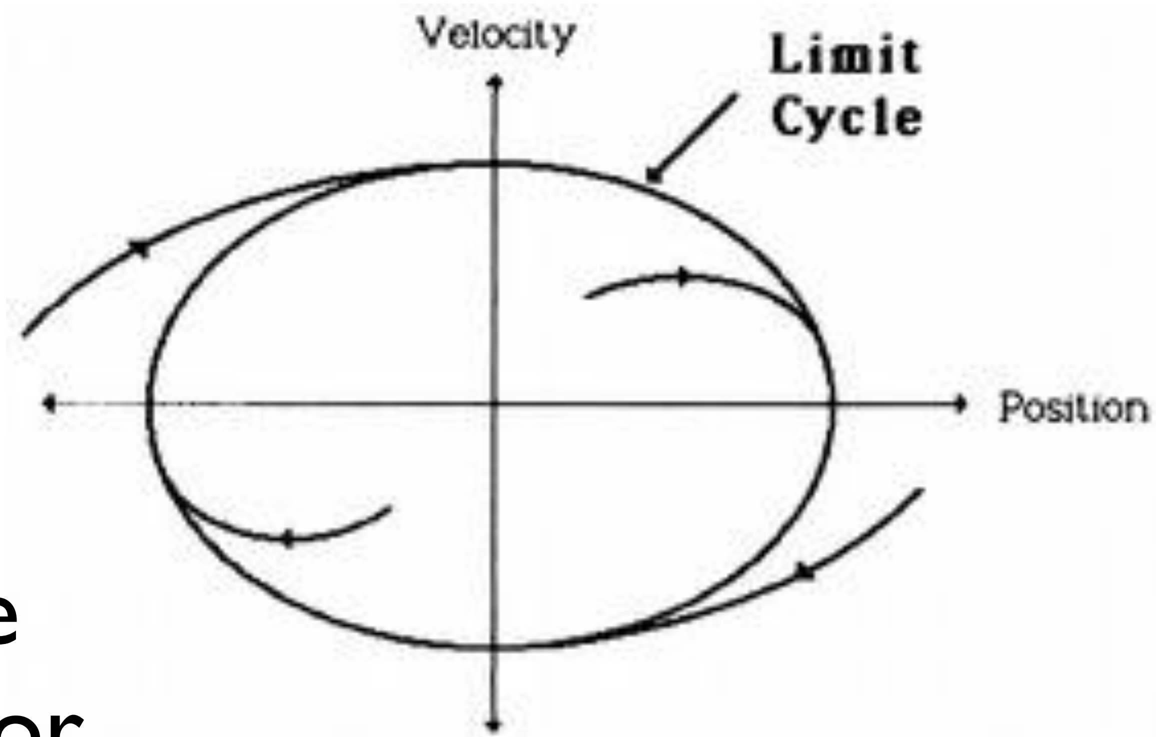

Any system displaying a limit cycle
behavior may be called an oscillator

Phase is straightforward to define for a limit cycle.

Point Attractor

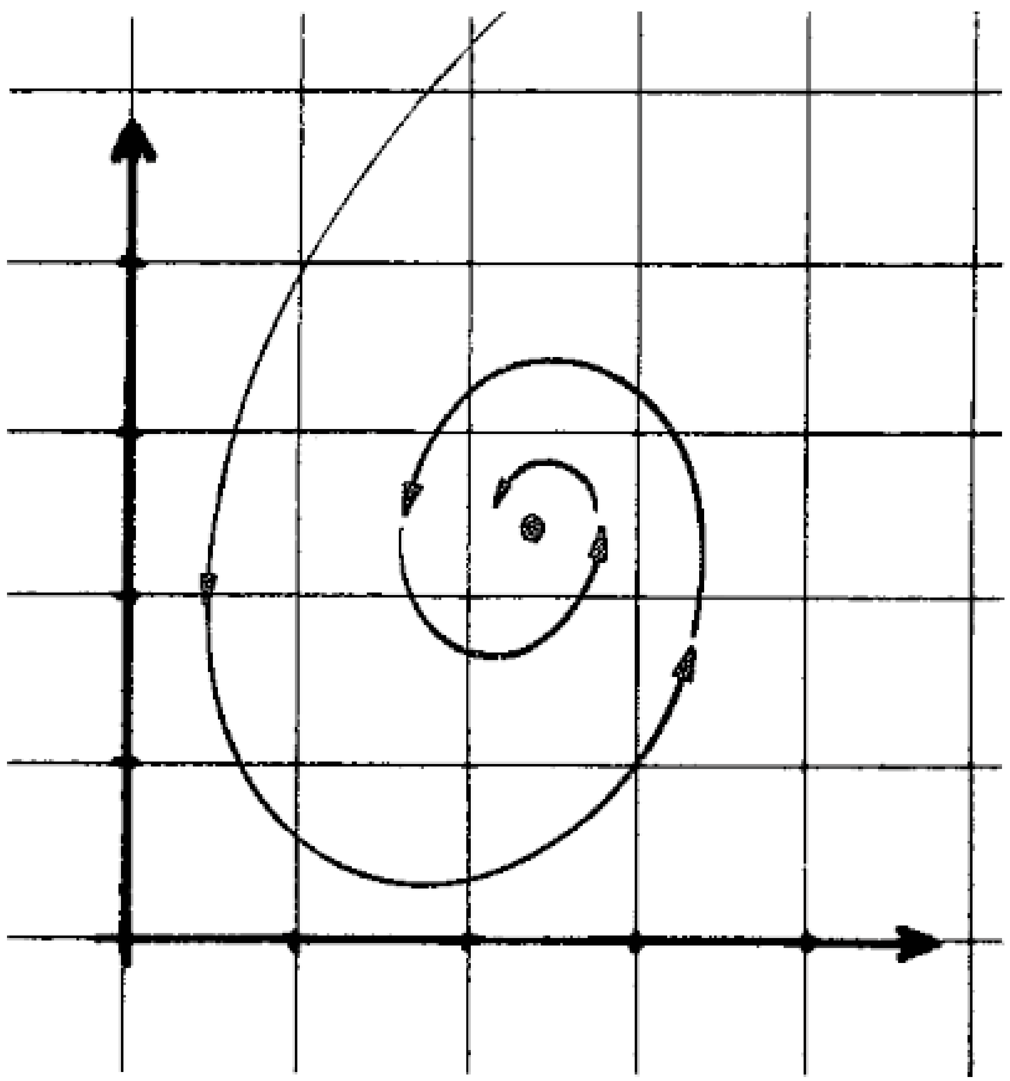

Limit Cycle Attractor

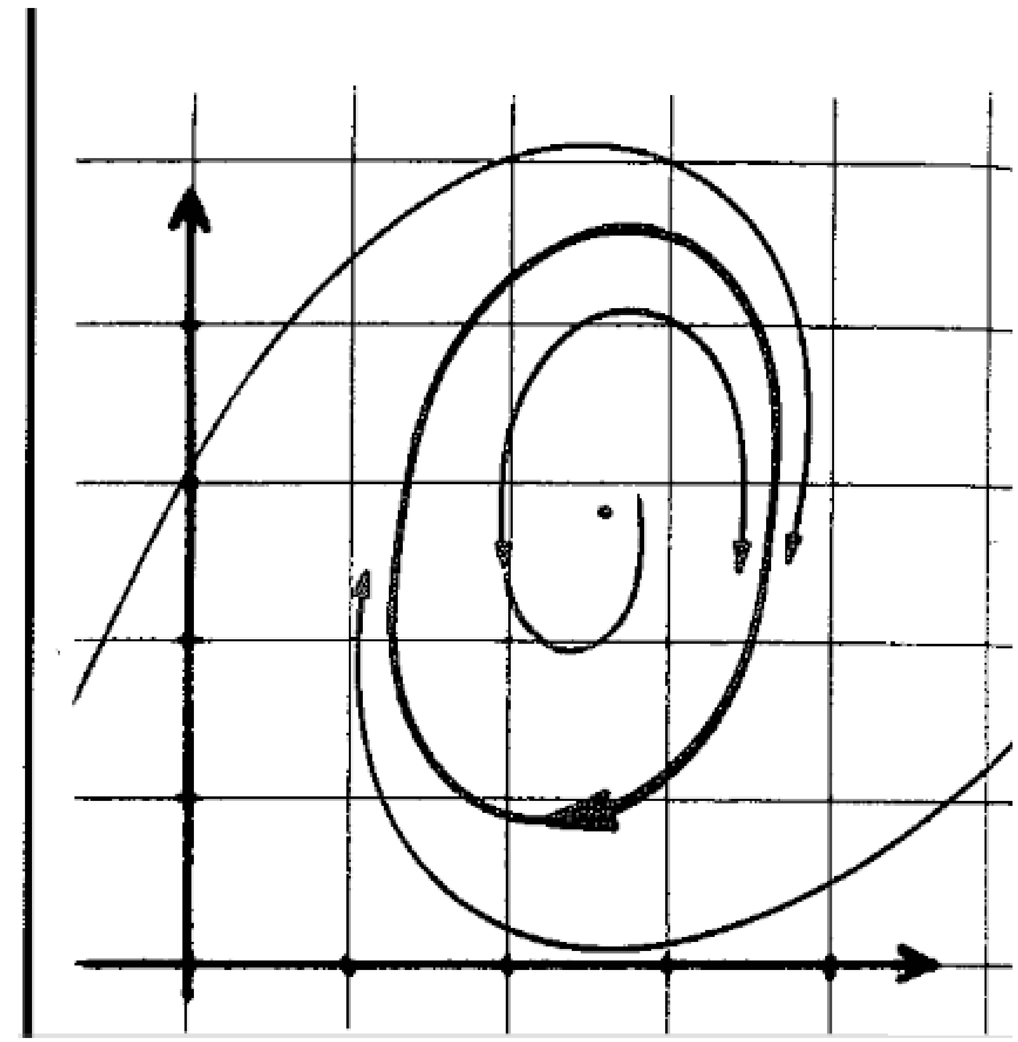

The Chaotic Attractor

The system occupies a finite subset of state space, but without repetition.

Two trajectories that start very close will diverge rapidly, losing all mutual predictability

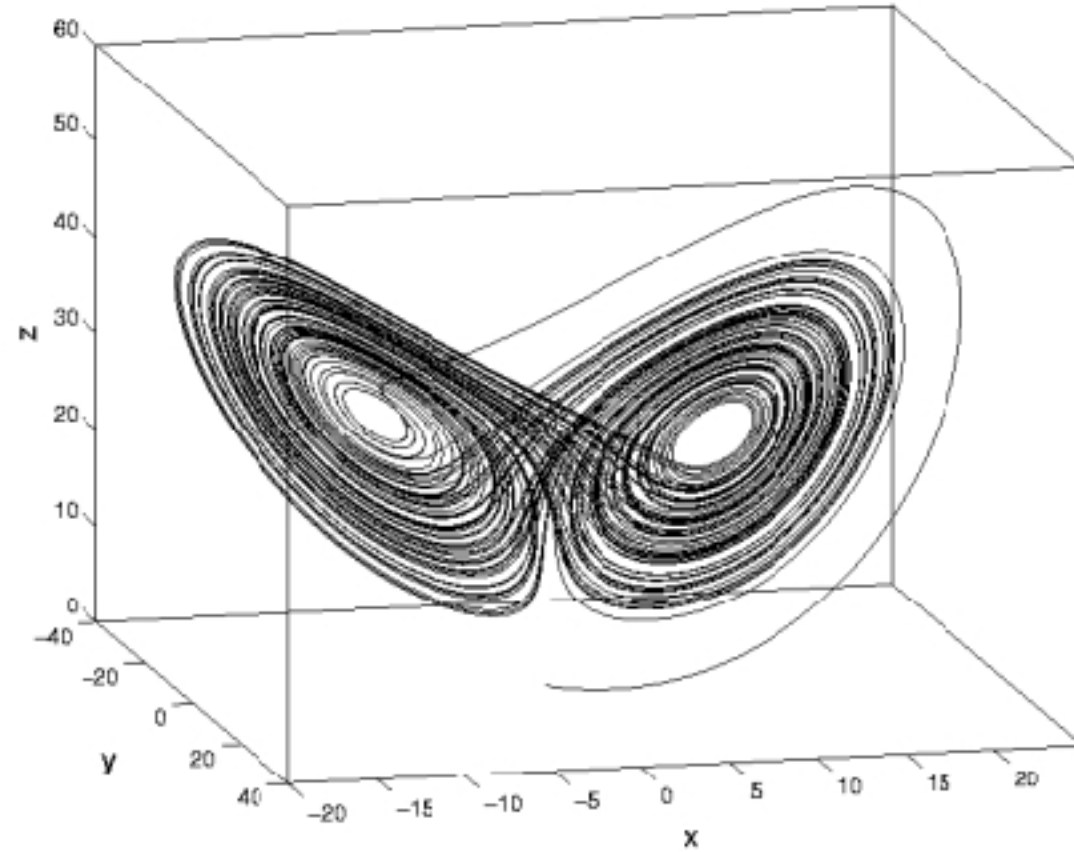

A chaotic attractor is a small part of the state space in which the system is to be found, but without any simple repeating structure.

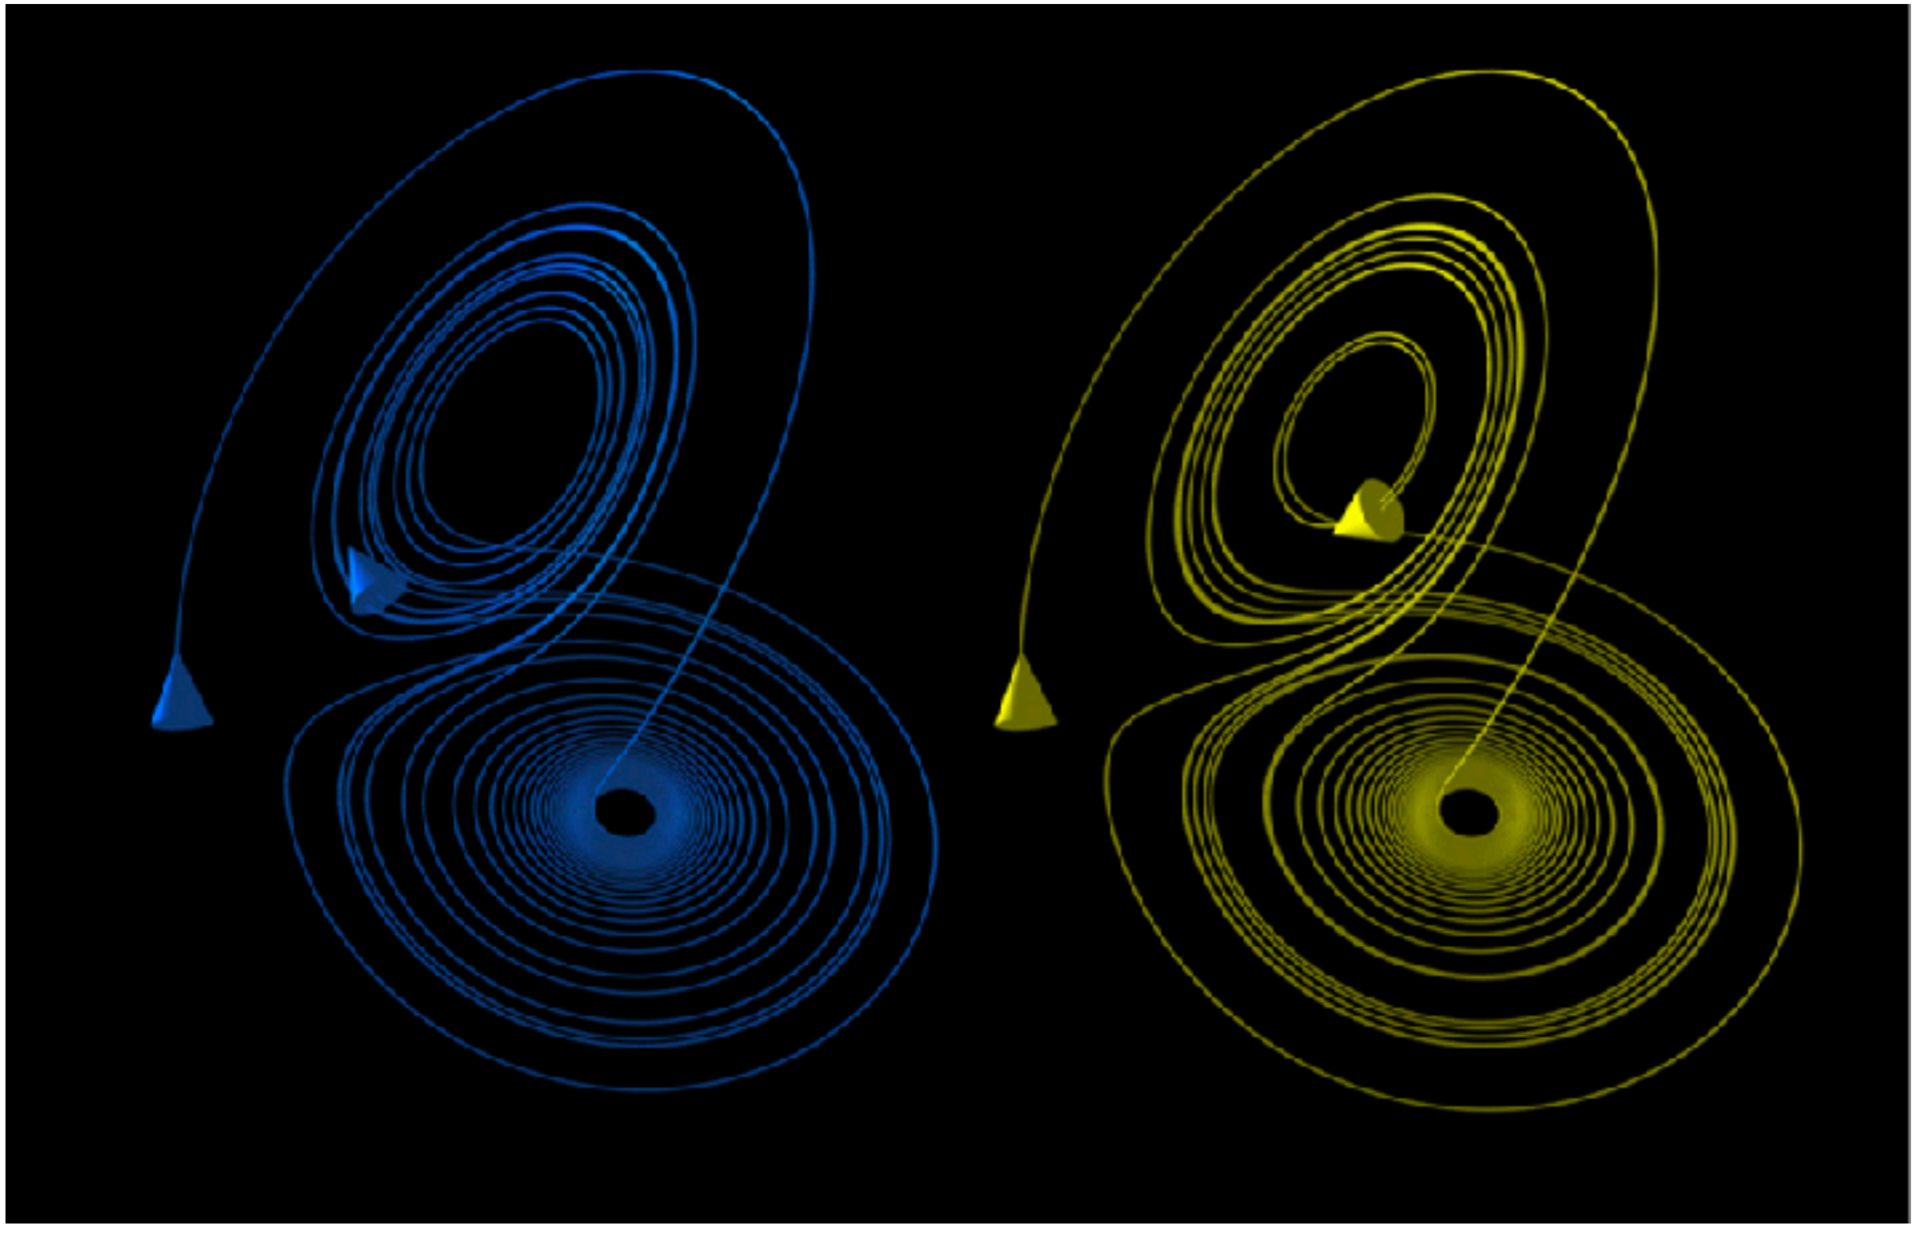

Lorenz attractor, two very similar starting points

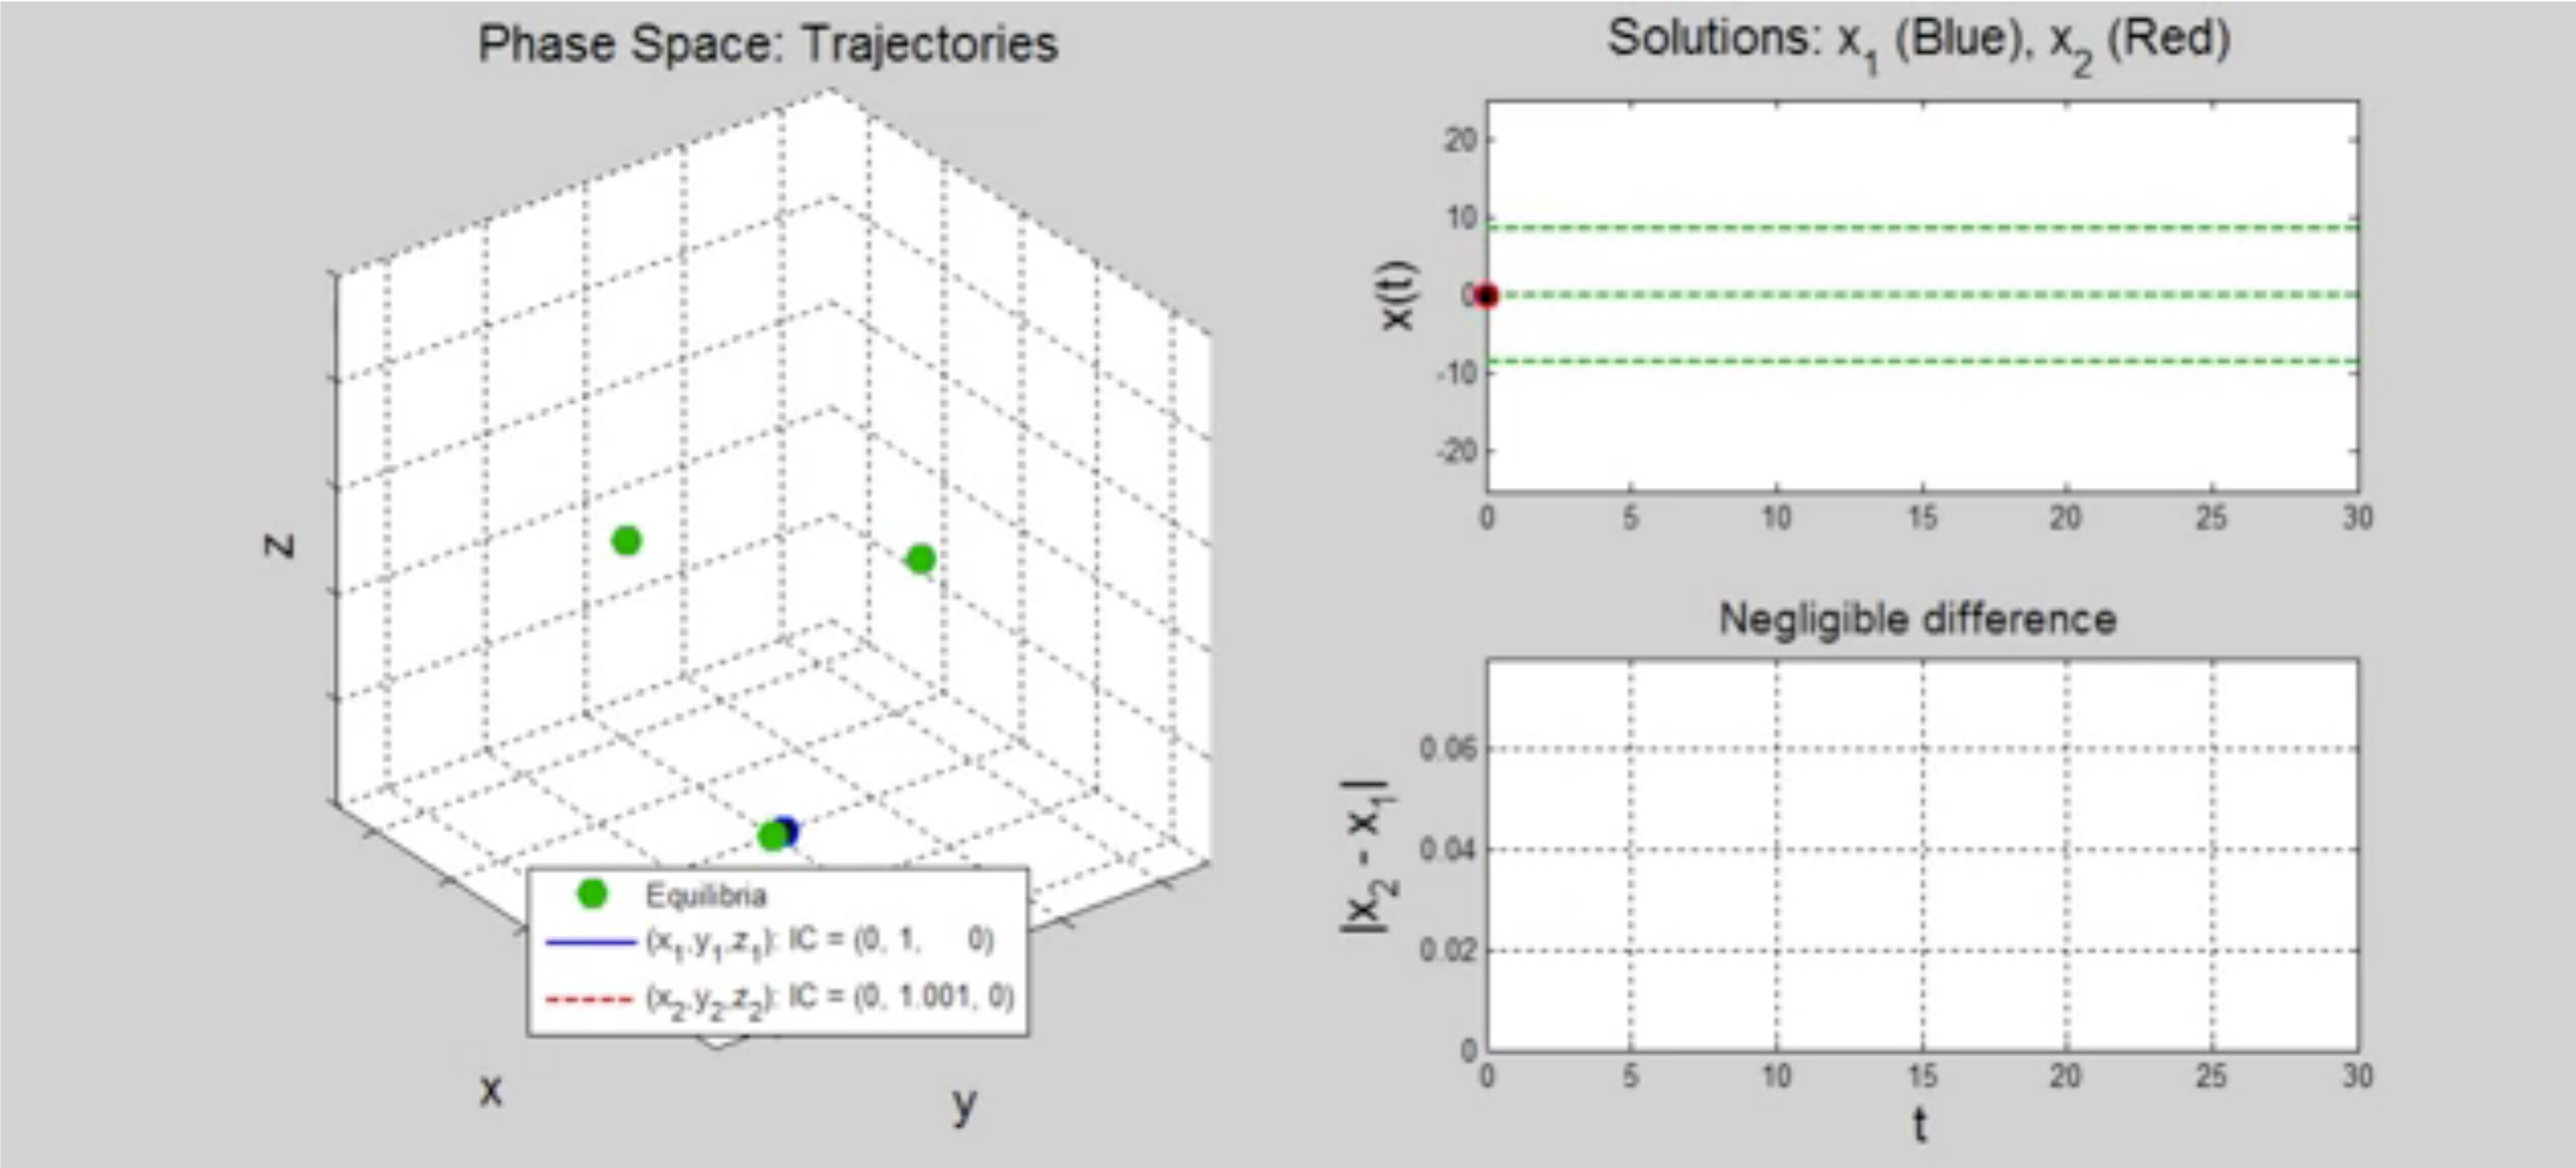

Two trajectories that evolve from *almost* the same initial conditions (x1 = (0,1,0); x2 = (0, 1.001, 0))

Induction of a limit cycle attractor (to be done in lab)

Input: 1 1 0 1 0

Output: 1 0 0 1 1

Test: 1 1 1 1.....1 1 1

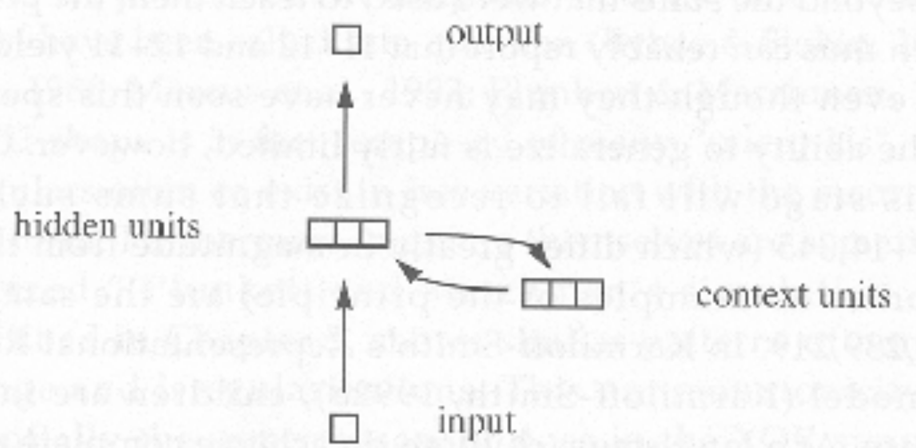

FIGURE 4.27 Simple recurrent network architecture used to learn the odd/even task. Context units store the state of the hidden units from the prior time step.

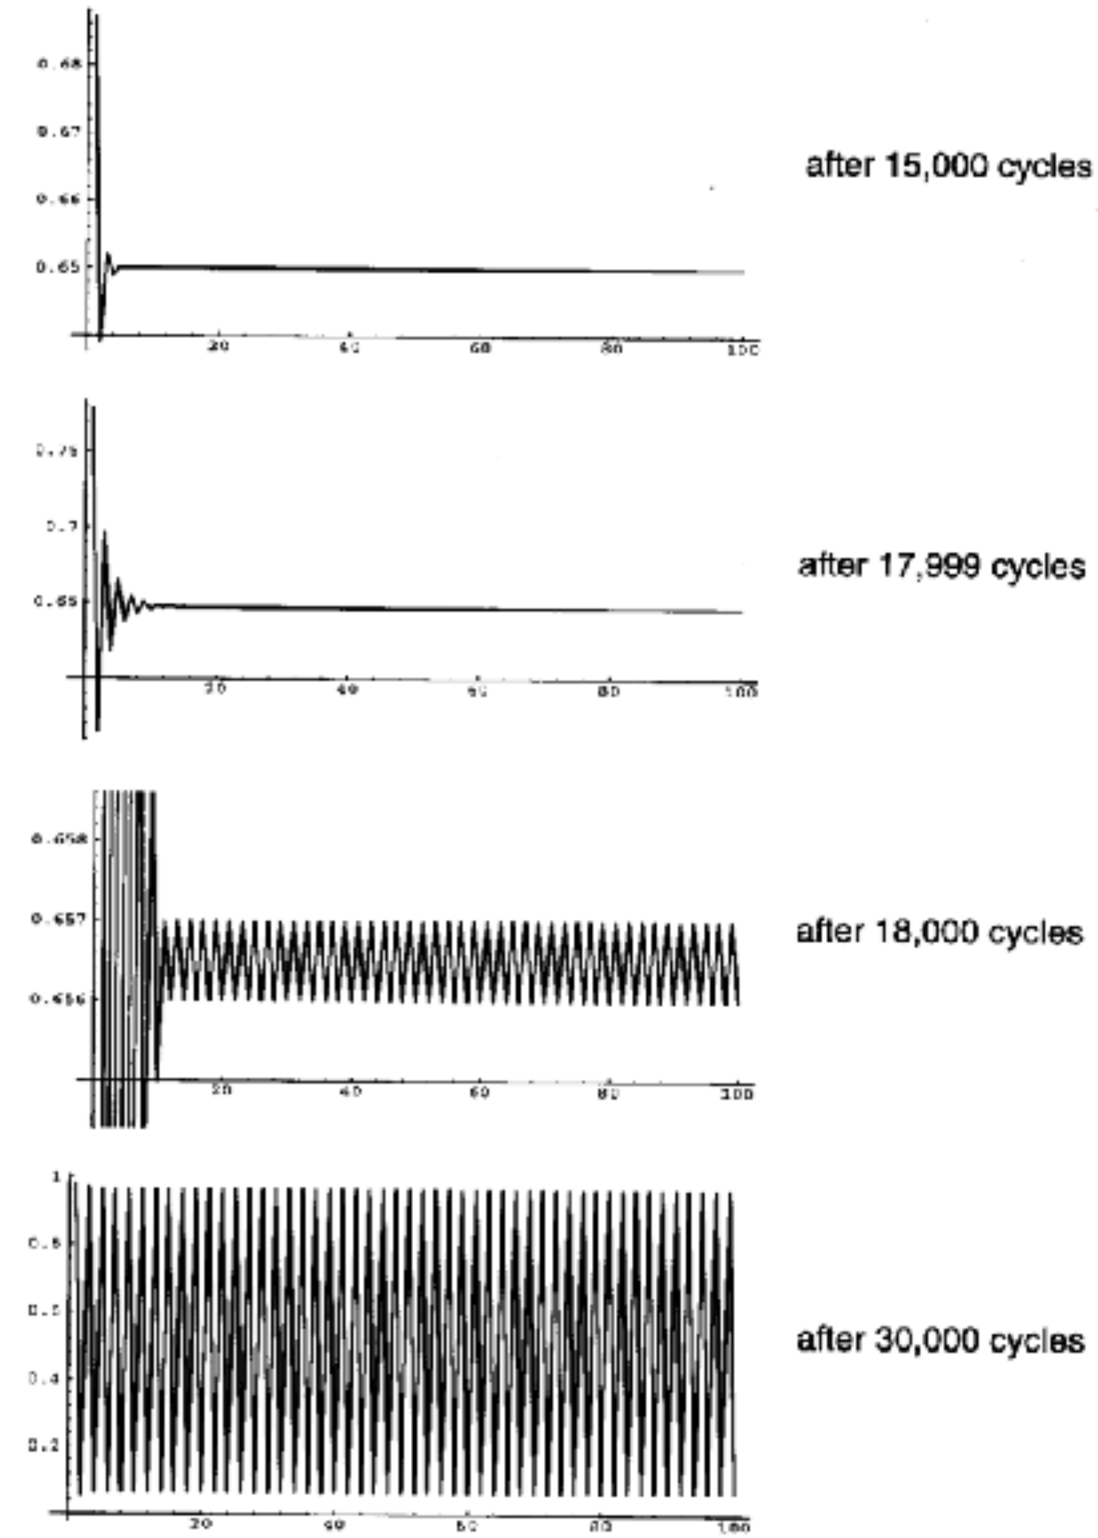

FIGURE 4.28 Performance of the simple recurrent network while learning the odd/even task at several points in time (measured in cycles). Generalization occurs at 17,999 cycles but is limited. One learning cycle later, however, the network is able to extend the generalization indefinitely.

In lab 9:

You will make smooth gradual changes to the network parameters (weights).

At some point, the network behavior changes qualitatively. That is, it went from doing one *kind* of thing, to doing a *different kind* of thing.

It went from a point attractor to a limit cycle attractor.

A qualitative change like this is known as a *bifurcation*.

Worked example: the logistic map

Example: the logistic map

$$x_{t+1} = r x_t (1 - x_t)$$

http://tuvalu.santafe.edu/~joshua/LogisticTools.html
or
http://cogsci.ucd.ie/Connectionism/Labs/logistic/logistic.html

This simple dynamical system illustrates point attractors, limit cycle attractors, bifurcations, chaos, sensitivity to initial conditions, etc.

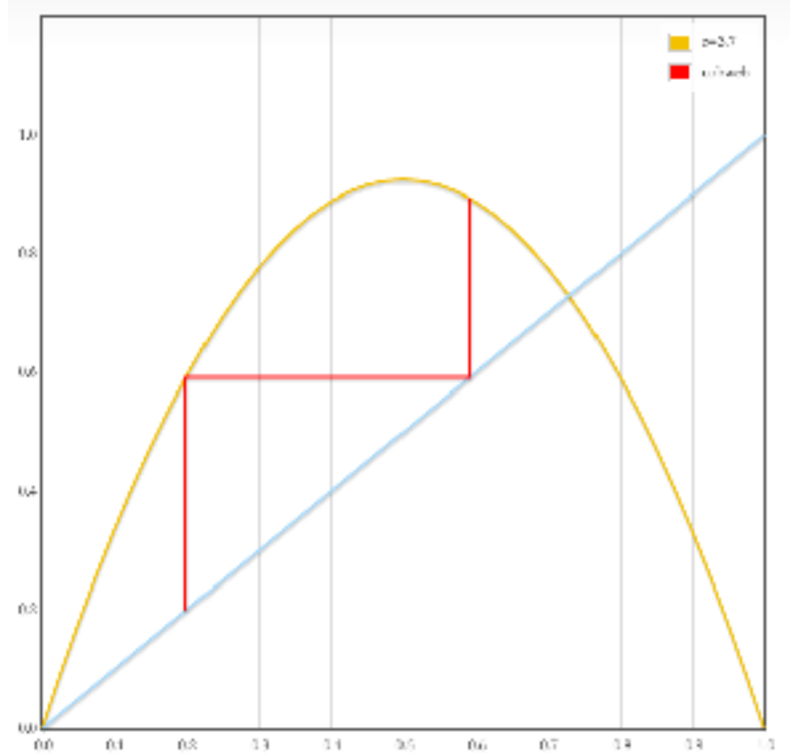

How the logistic app works.

For a given value of r, we plot x_{t+1} as a function of x_t.
That is the brown line.

We then pick a starting value, x_0.

We use x_0 to work out $x_{1,2,3 \ldots}$ The quick way to do this is to draw a line up to the brown curve, then over to the blue diagonal (the line y=x), back up to the curve, back to the blue diagonal, etc. (Make sure you understand why this works).

After many iterations, the system will have settled down to an attractor. This may be a single value (e.g. r=2) or a periodic attractor with period = 2, 4, 6, etc or it may be a chaotic attractor.

These asymptotic values are illustrated in this plot:

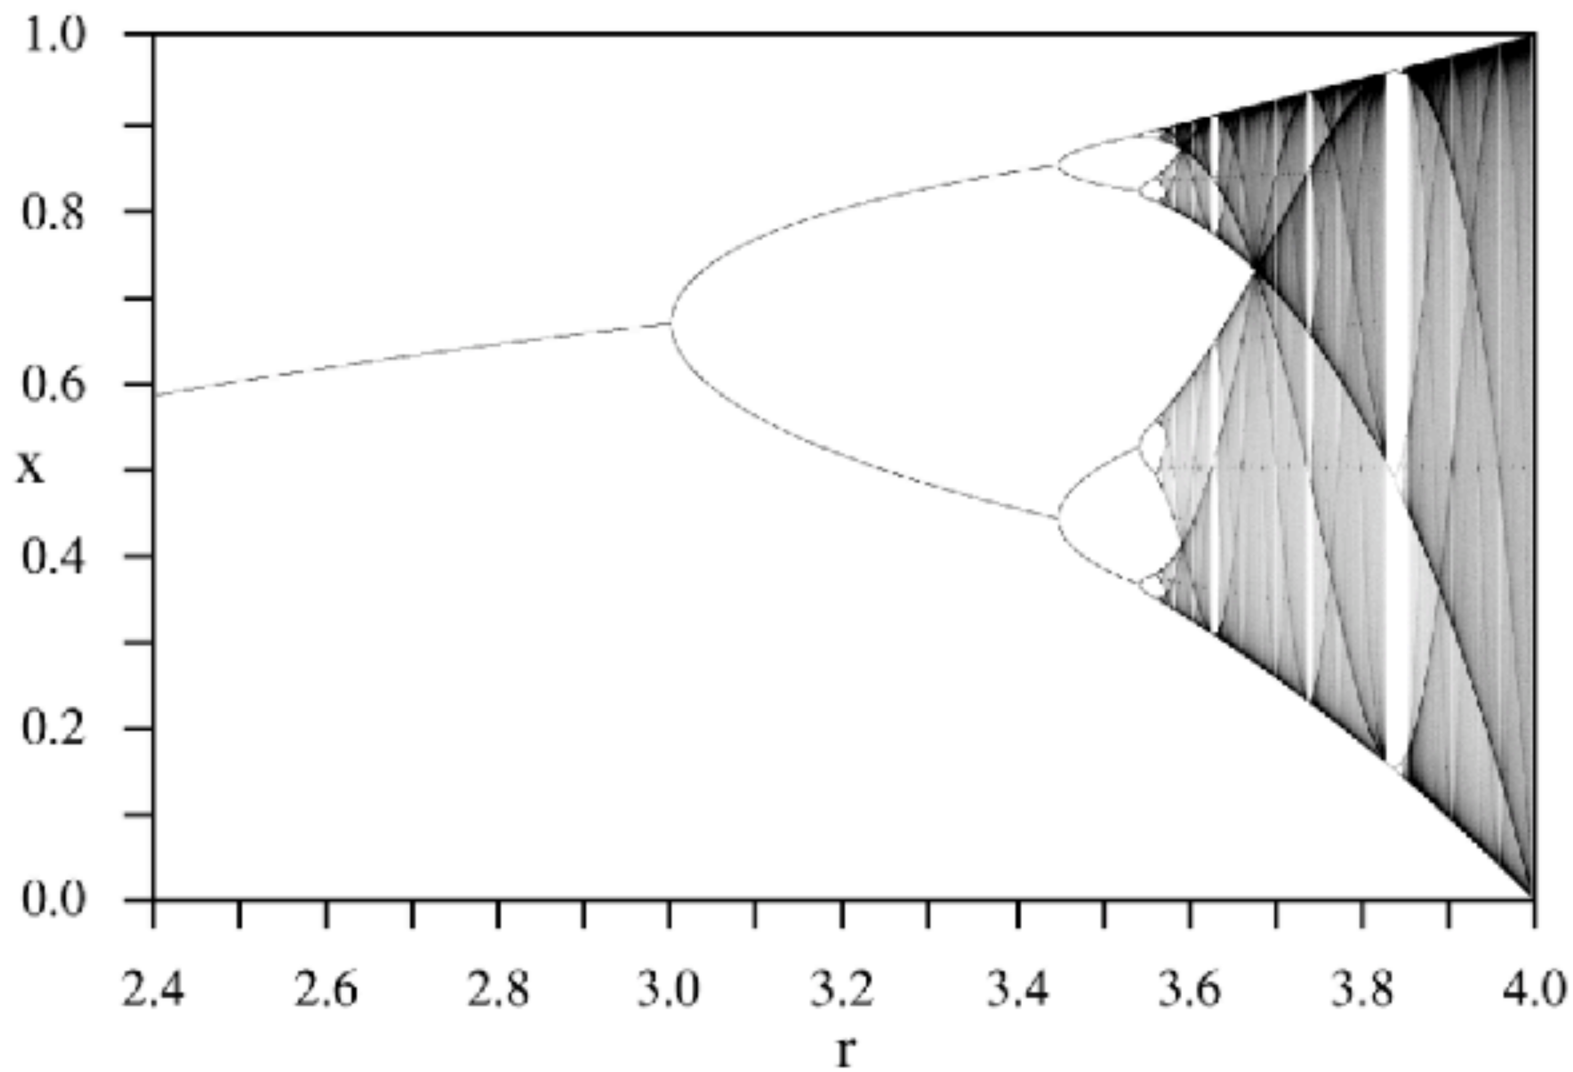

r=2.6 Point attractor

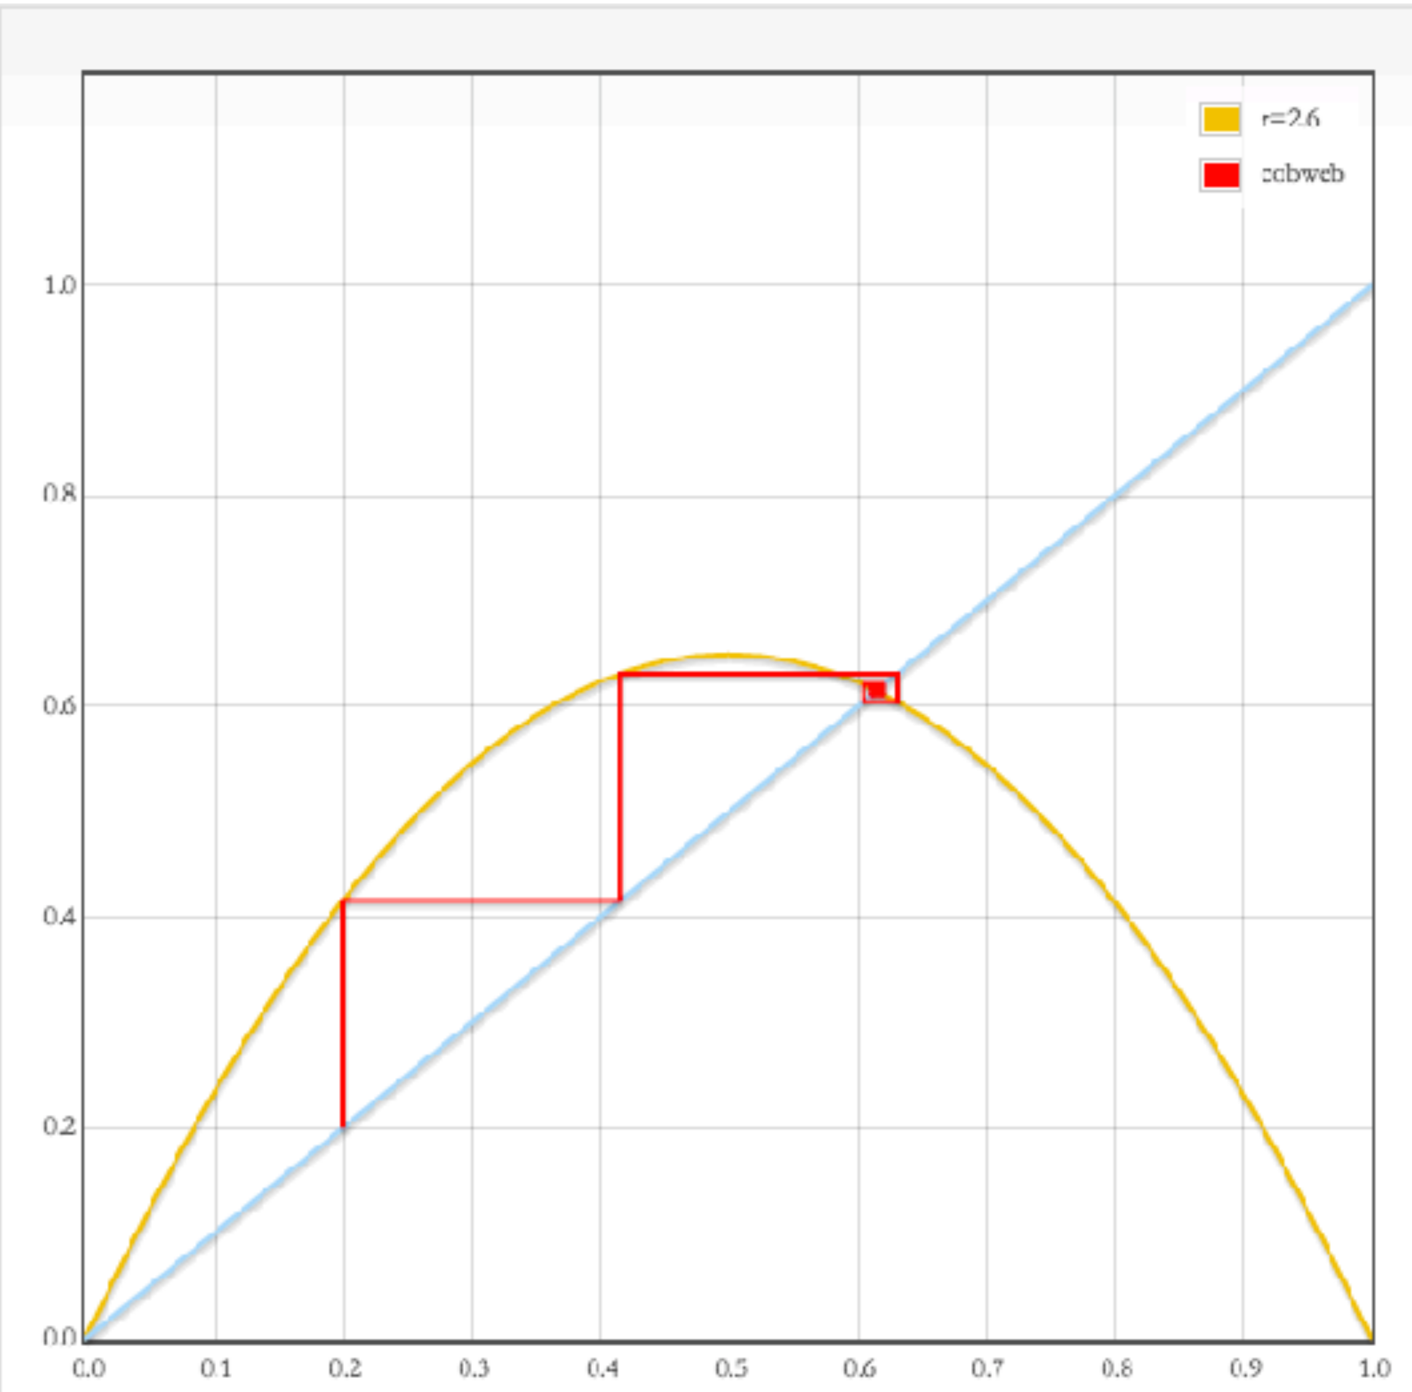

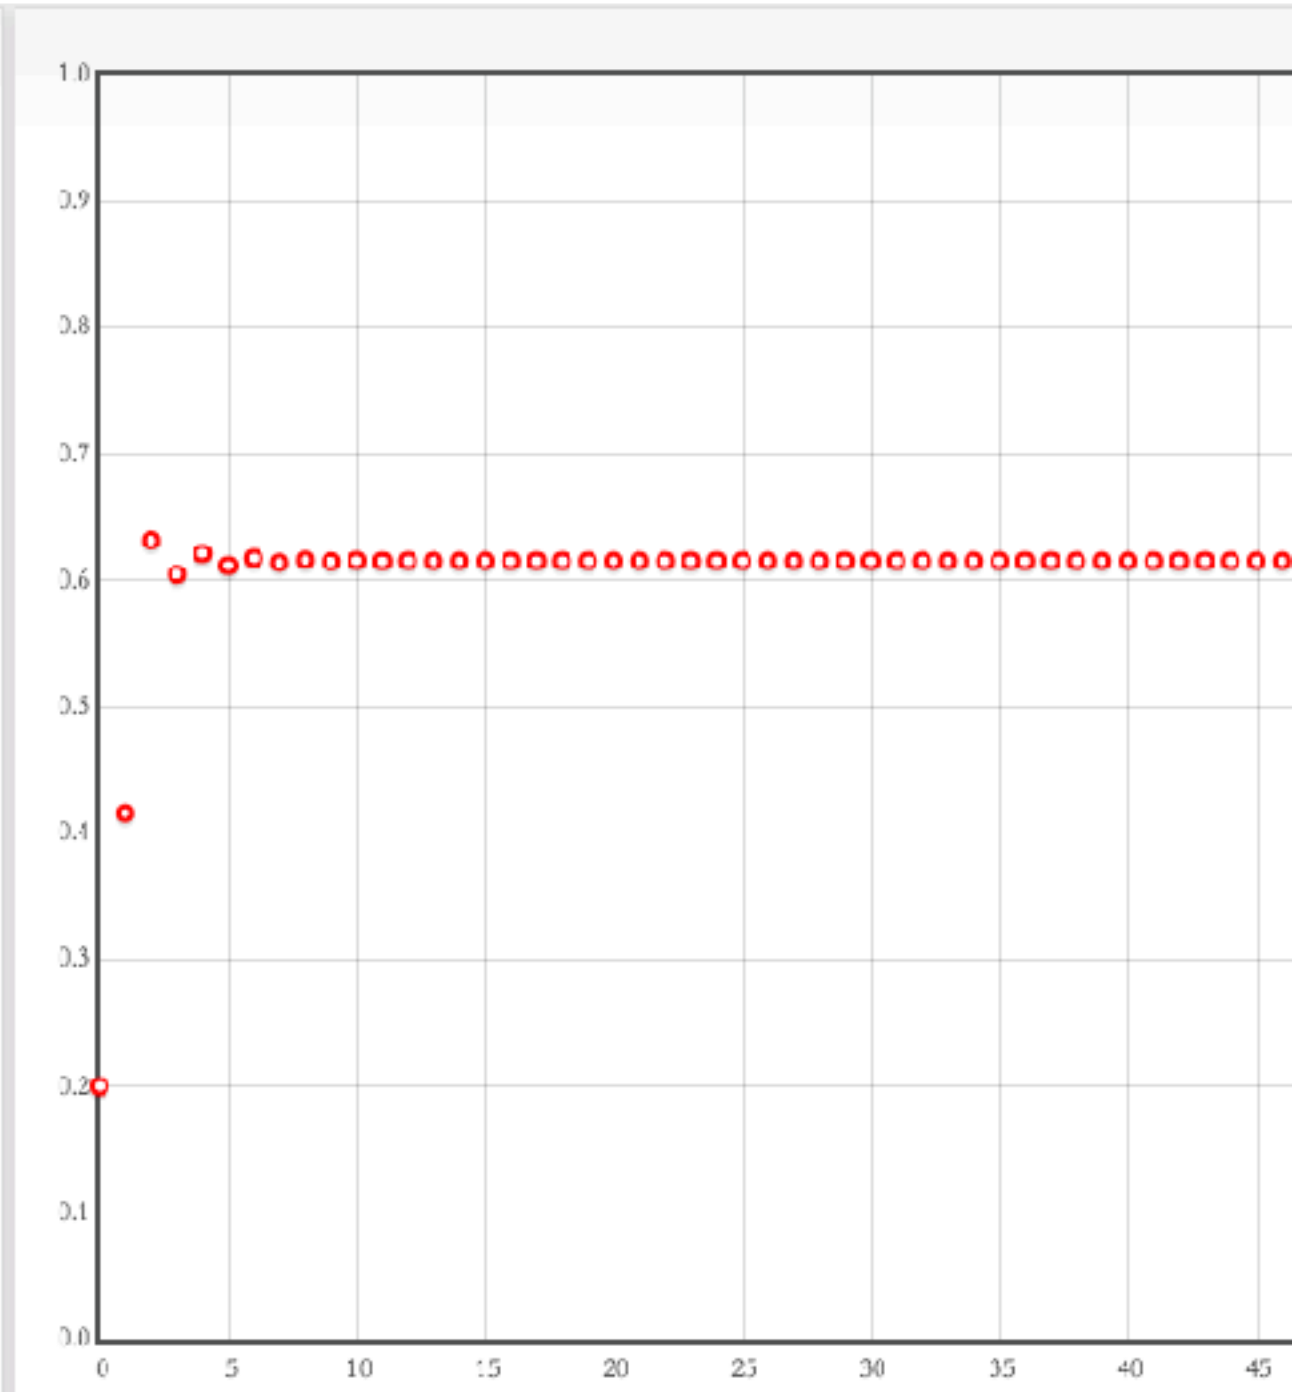

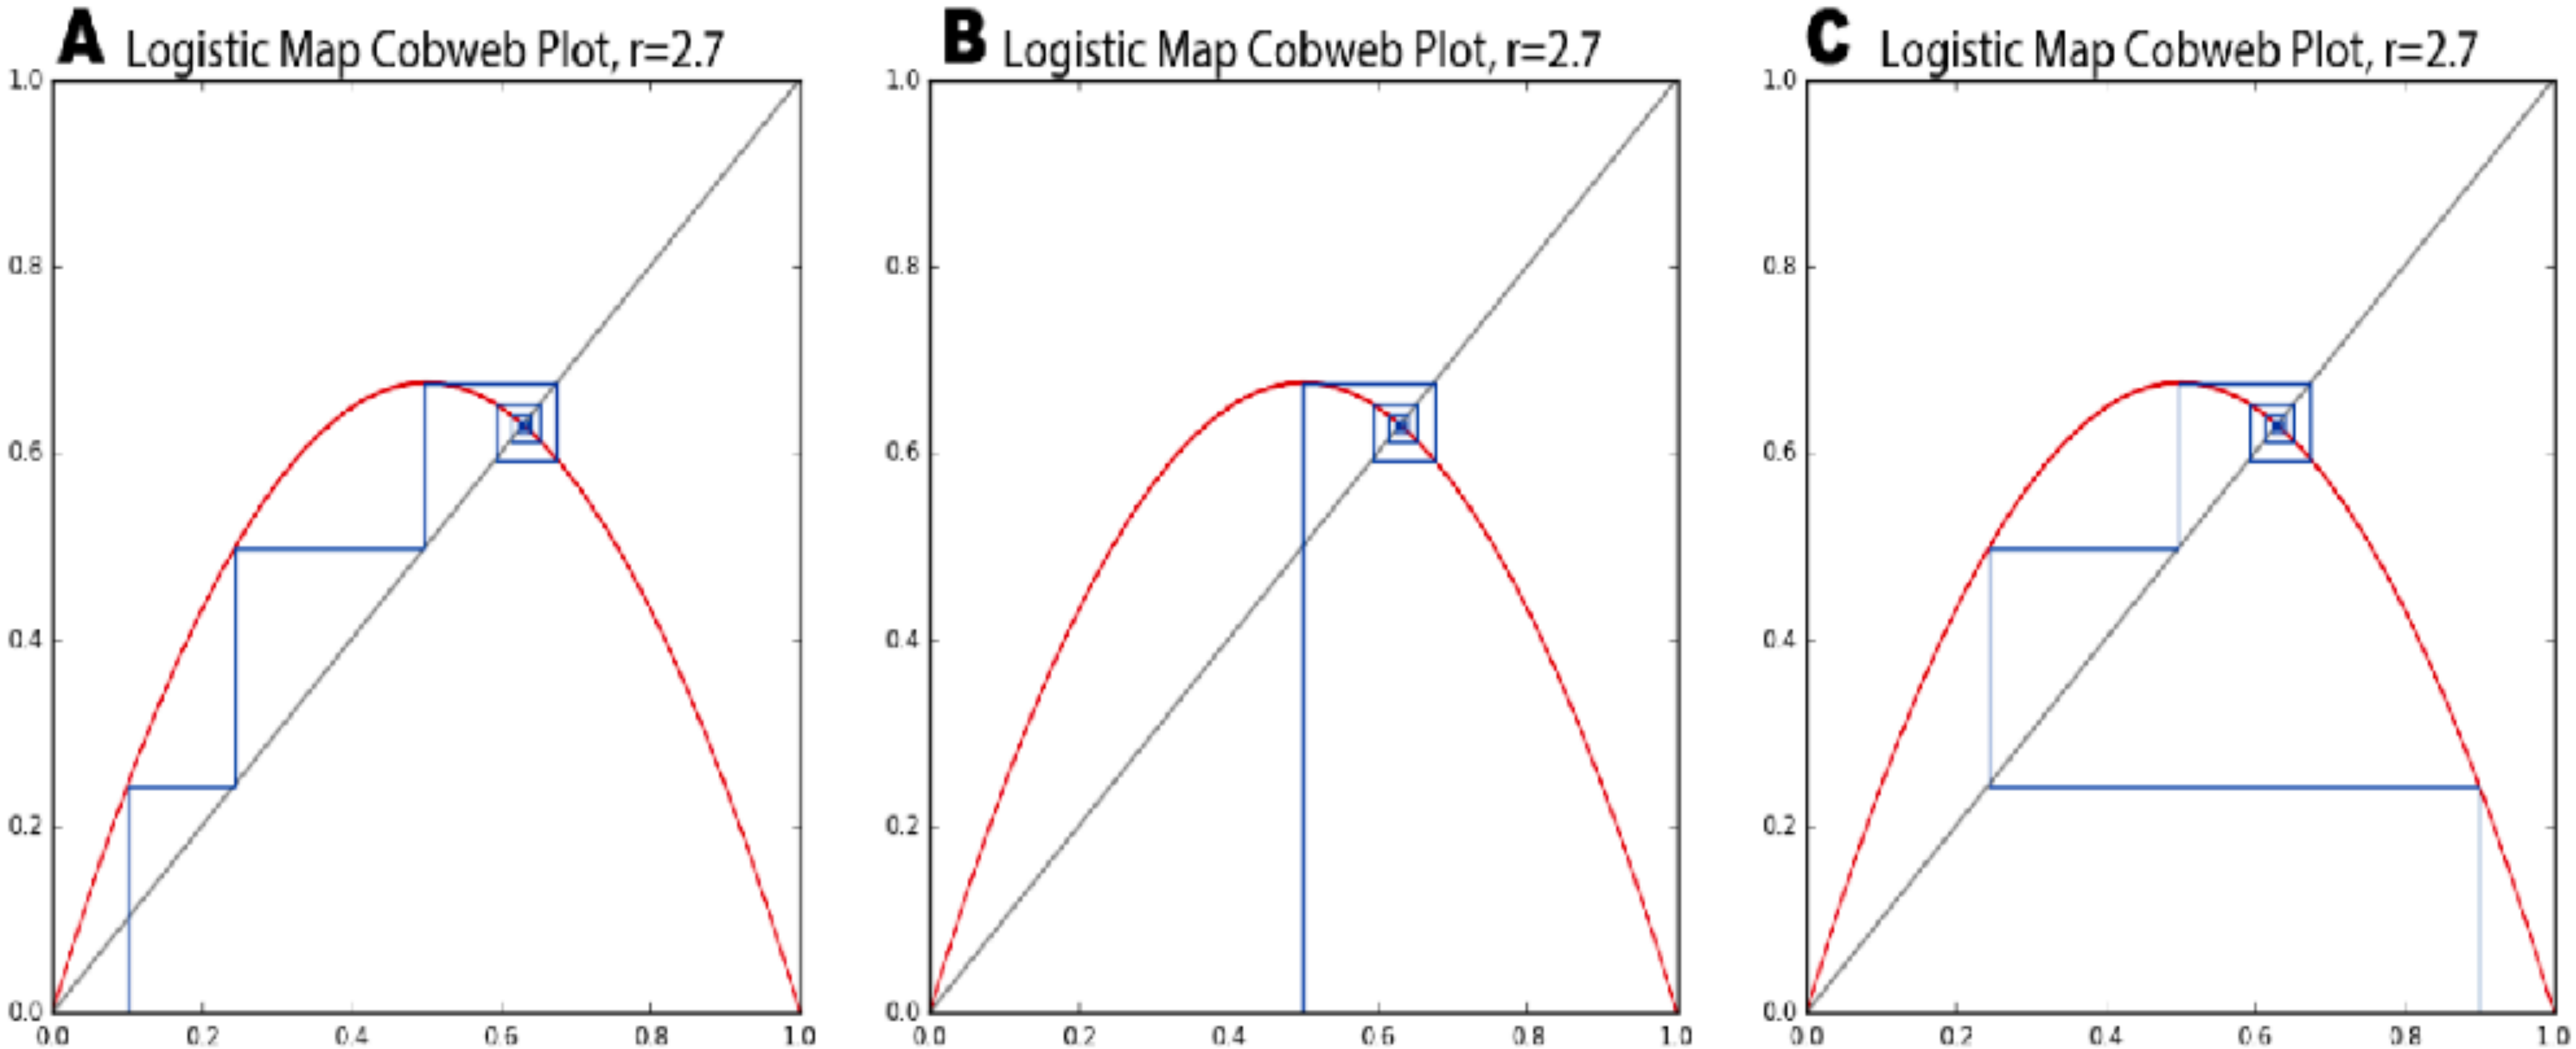

Figure 12. Cobweb plots of the logistic map pulling initial population values of 0.1 (**A**), 0.5 (**B**) and 0.9 (**C**) into the same fixed-point attractor over time. At this growth rate parameter value of 2.7, the Lyapunov exponent is negative.

r=3.2

Period 2 attractor

The system has undergone a bifurcation!

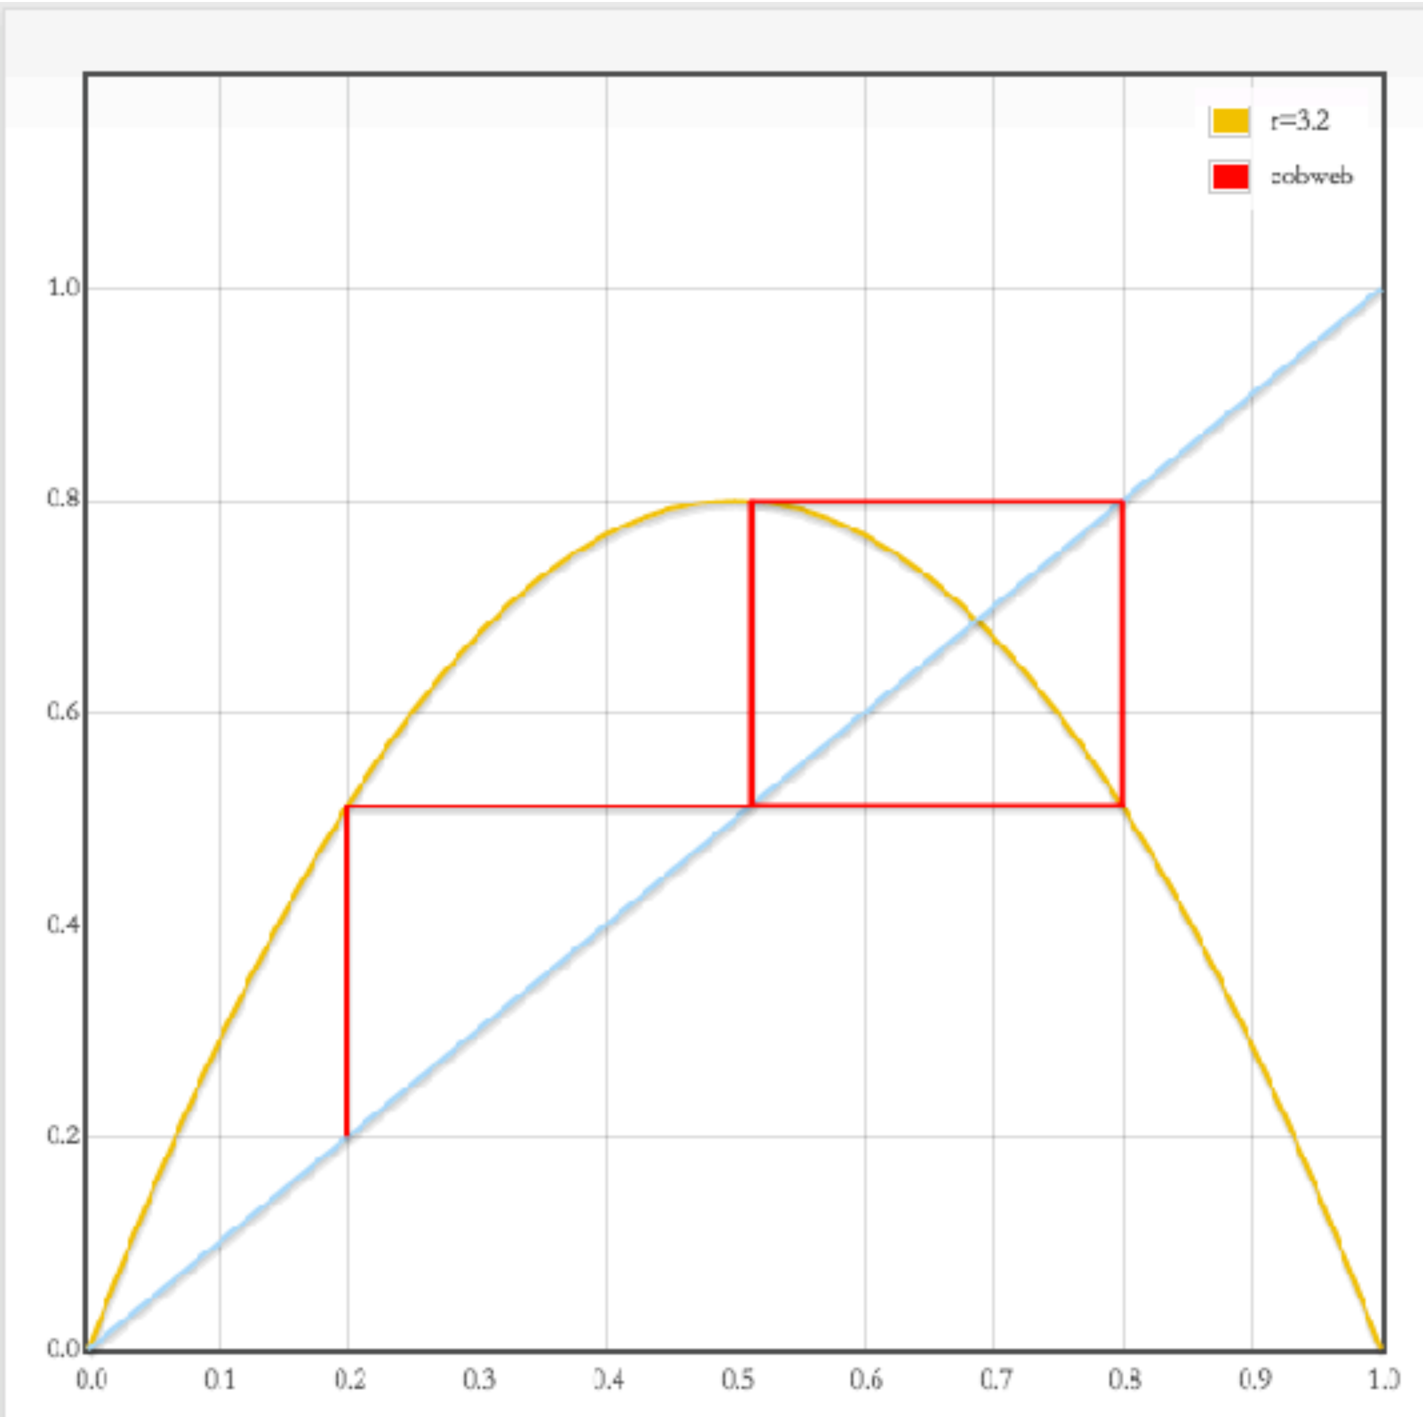

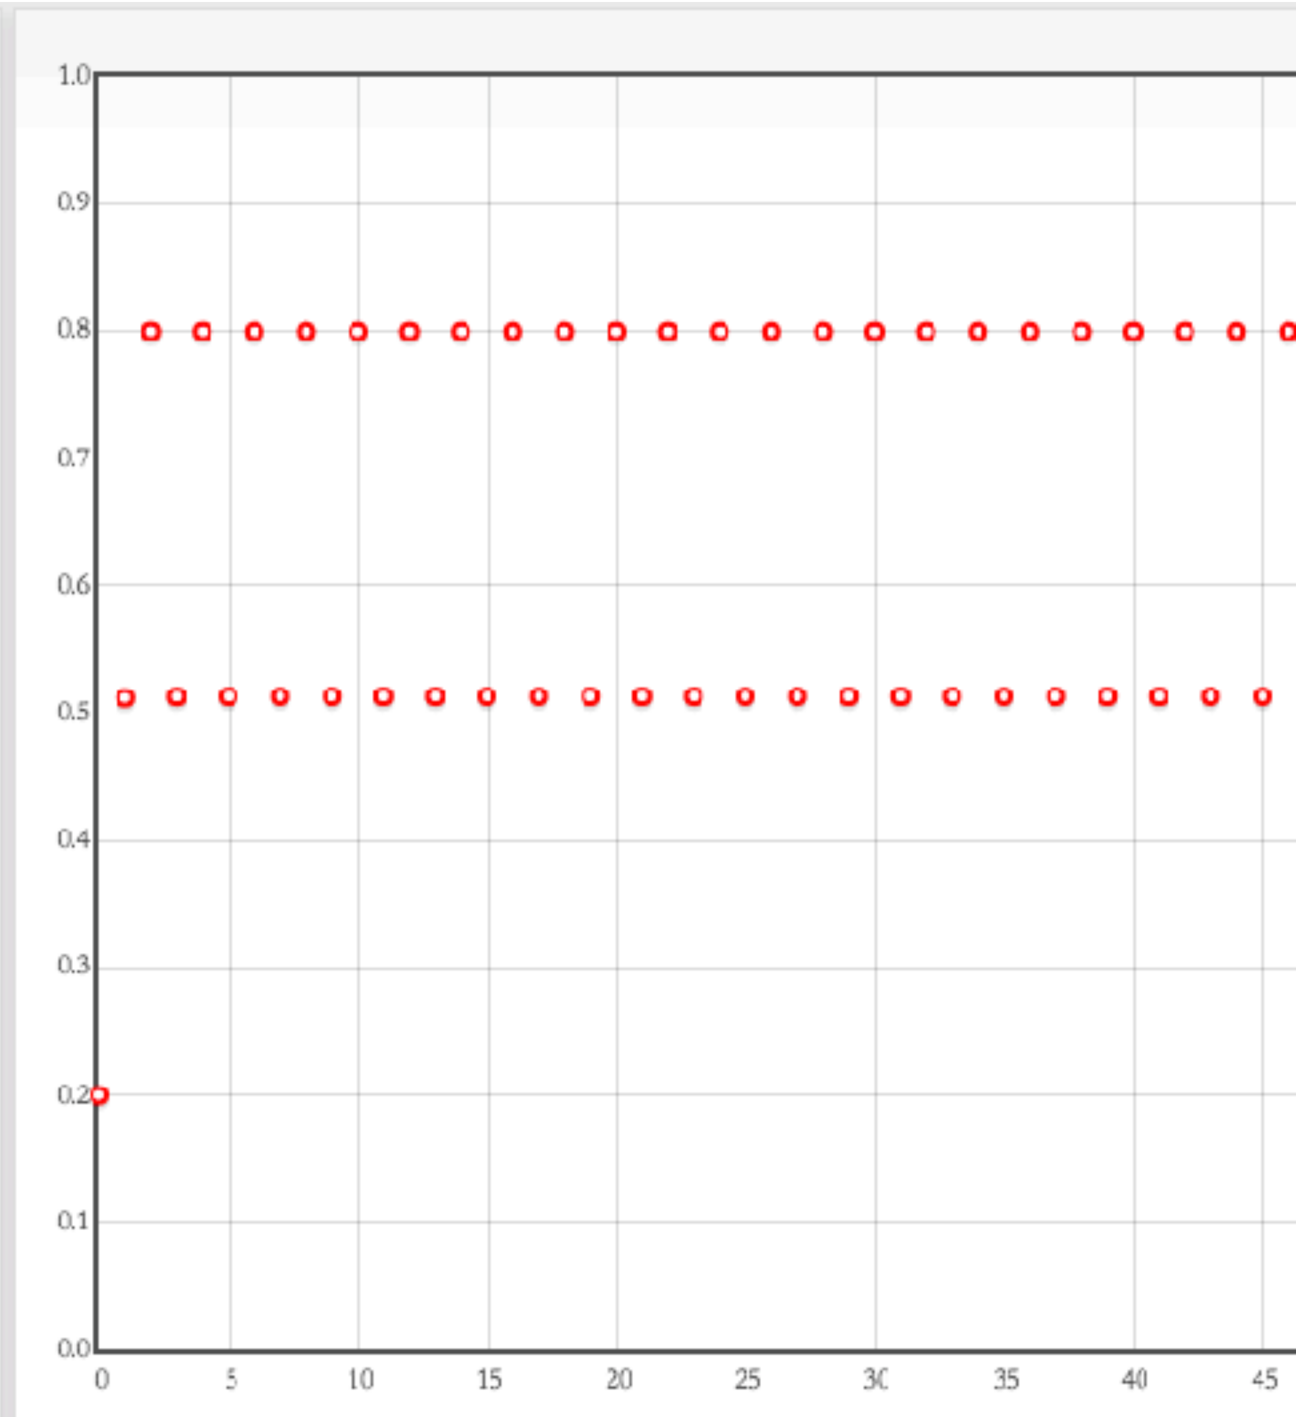

r=3.5 Period 4 limit cycle

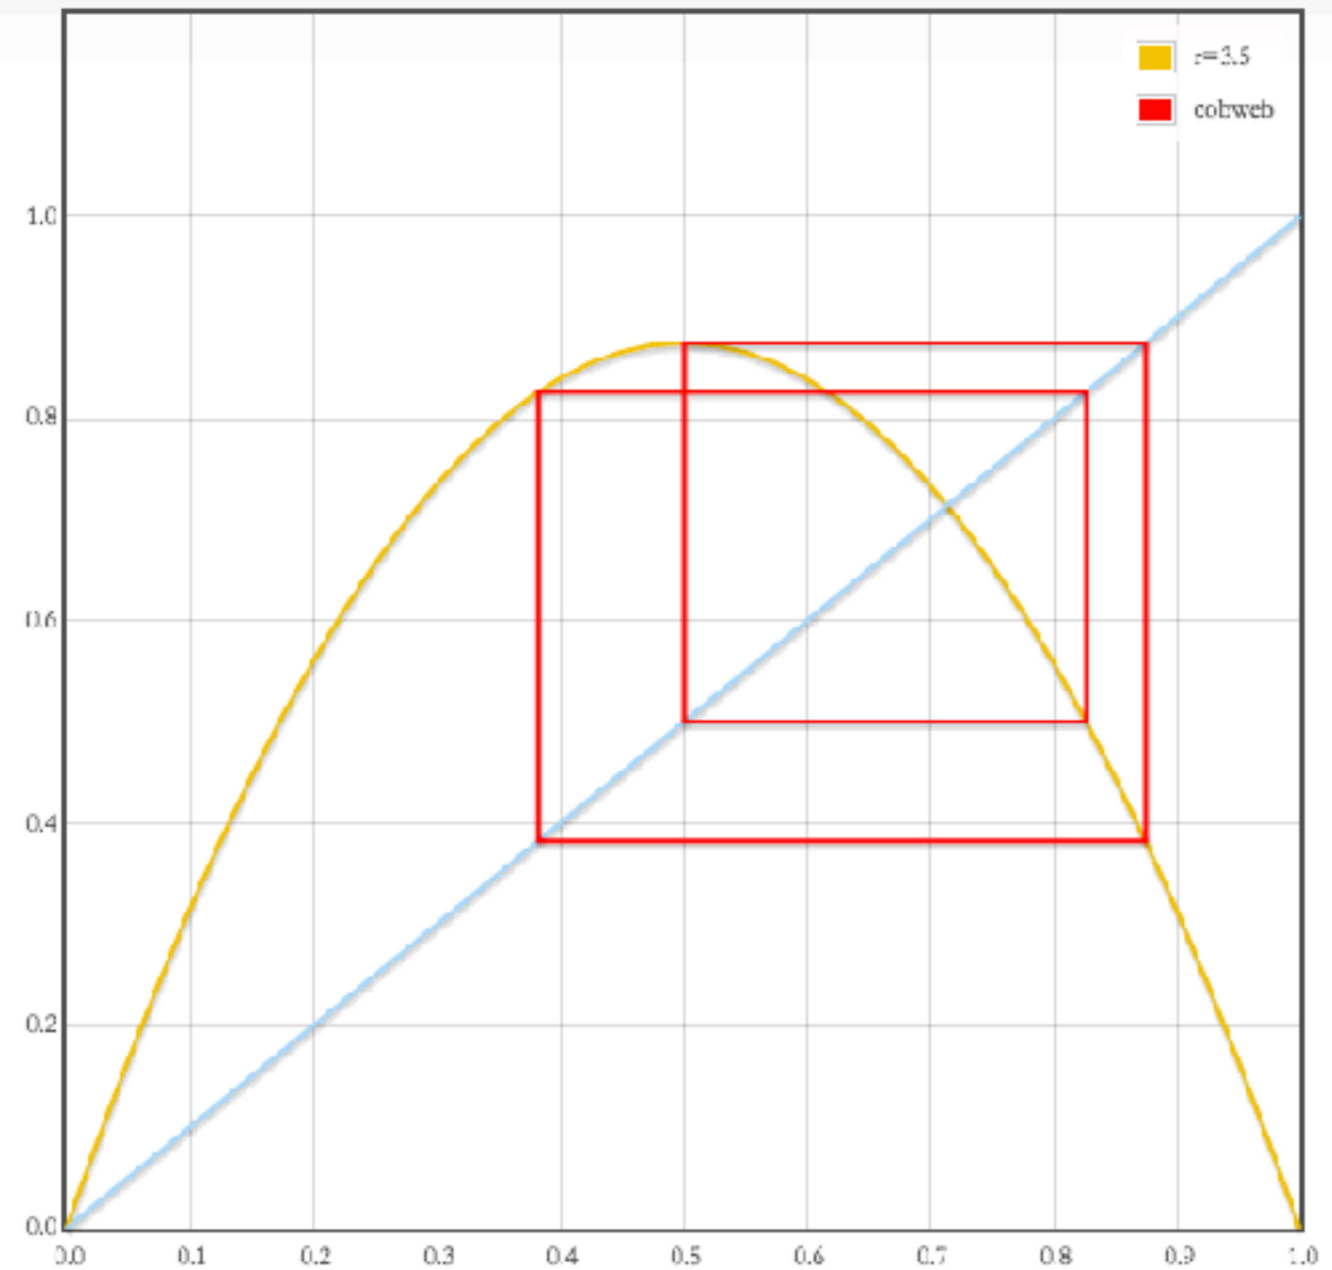

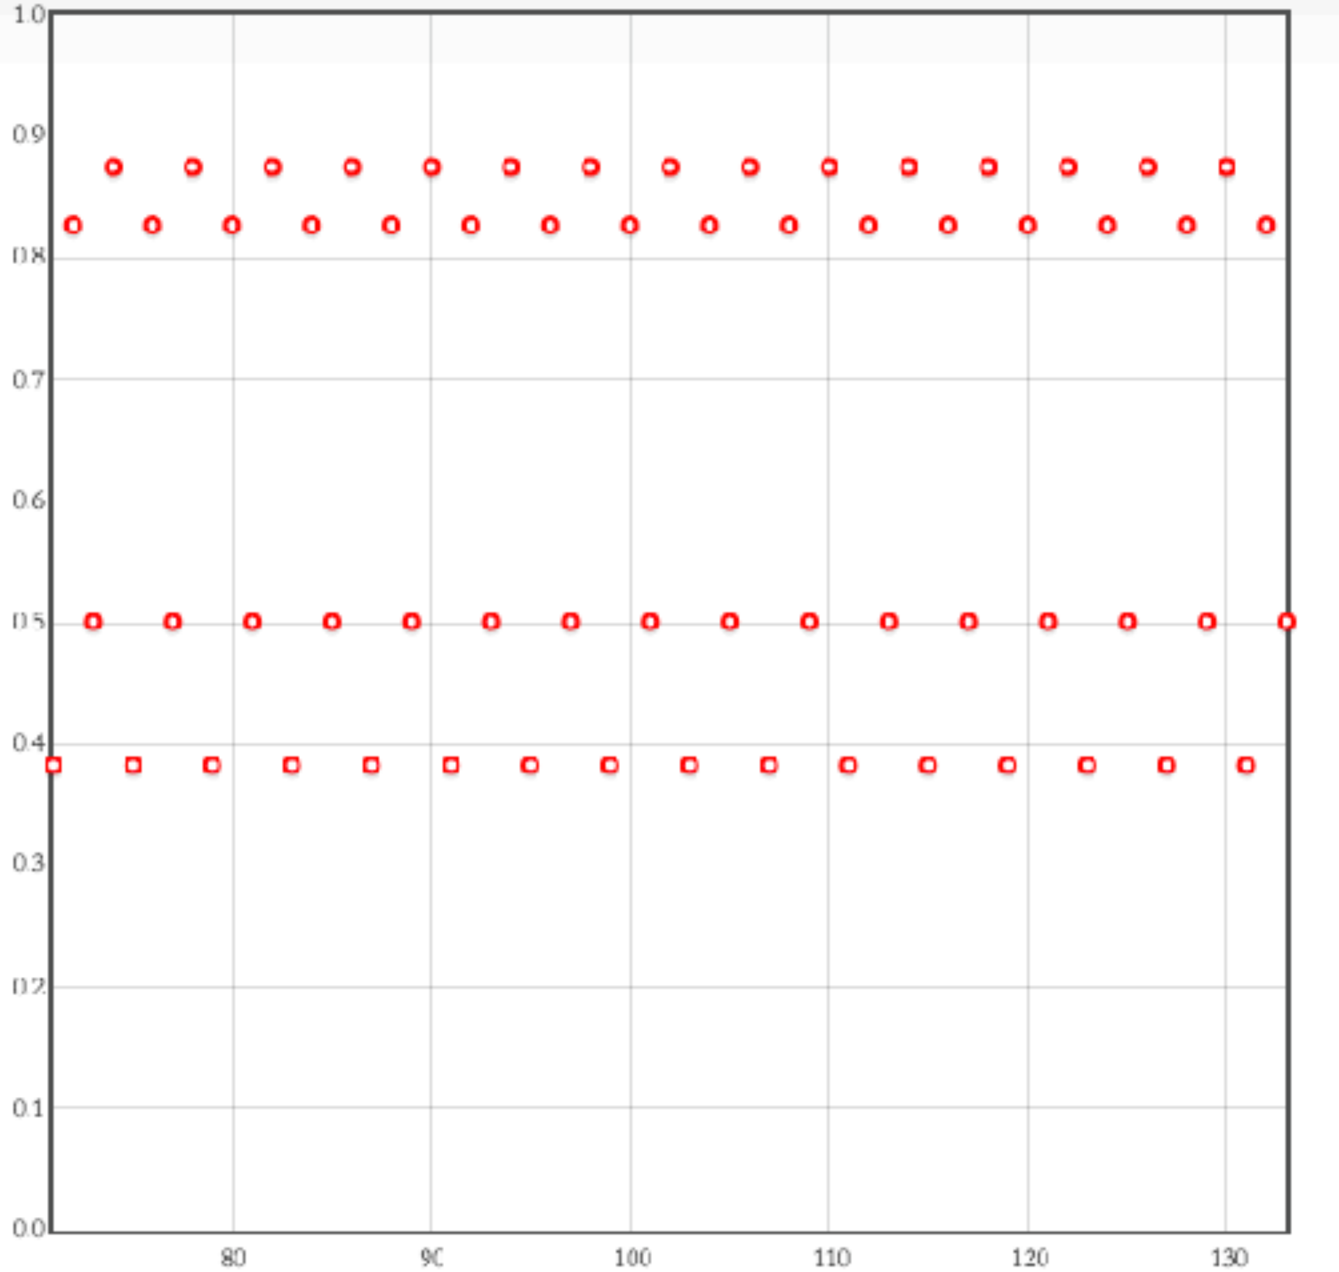

r=3.7 Chaos!

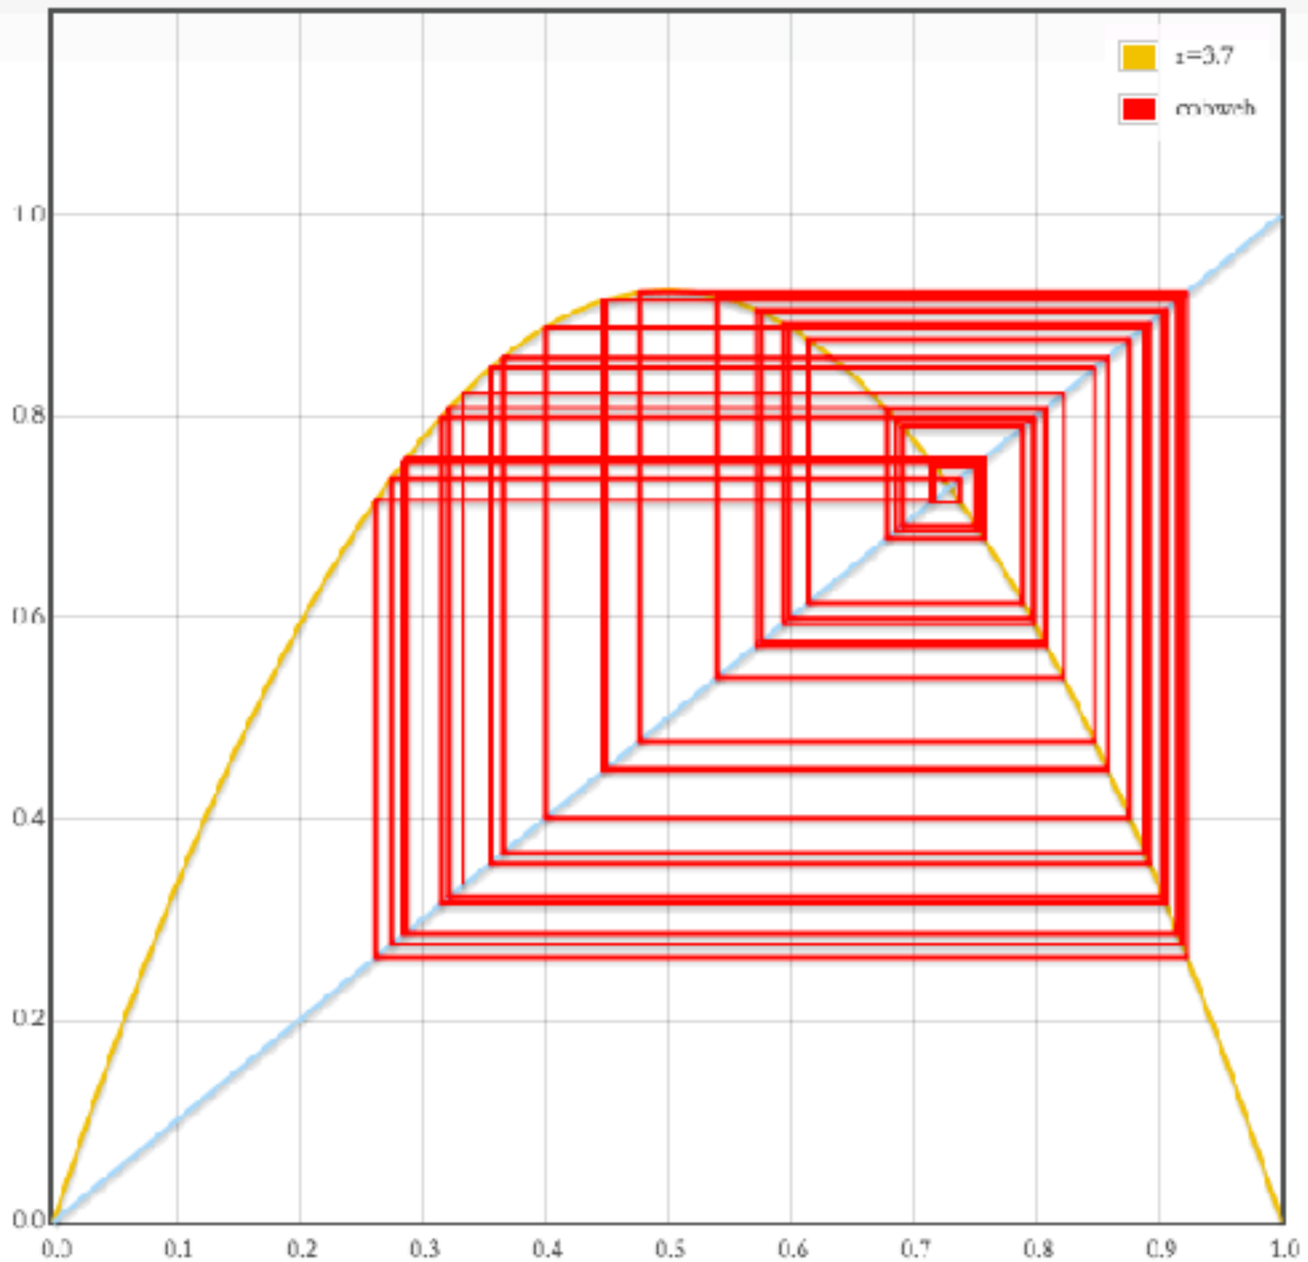

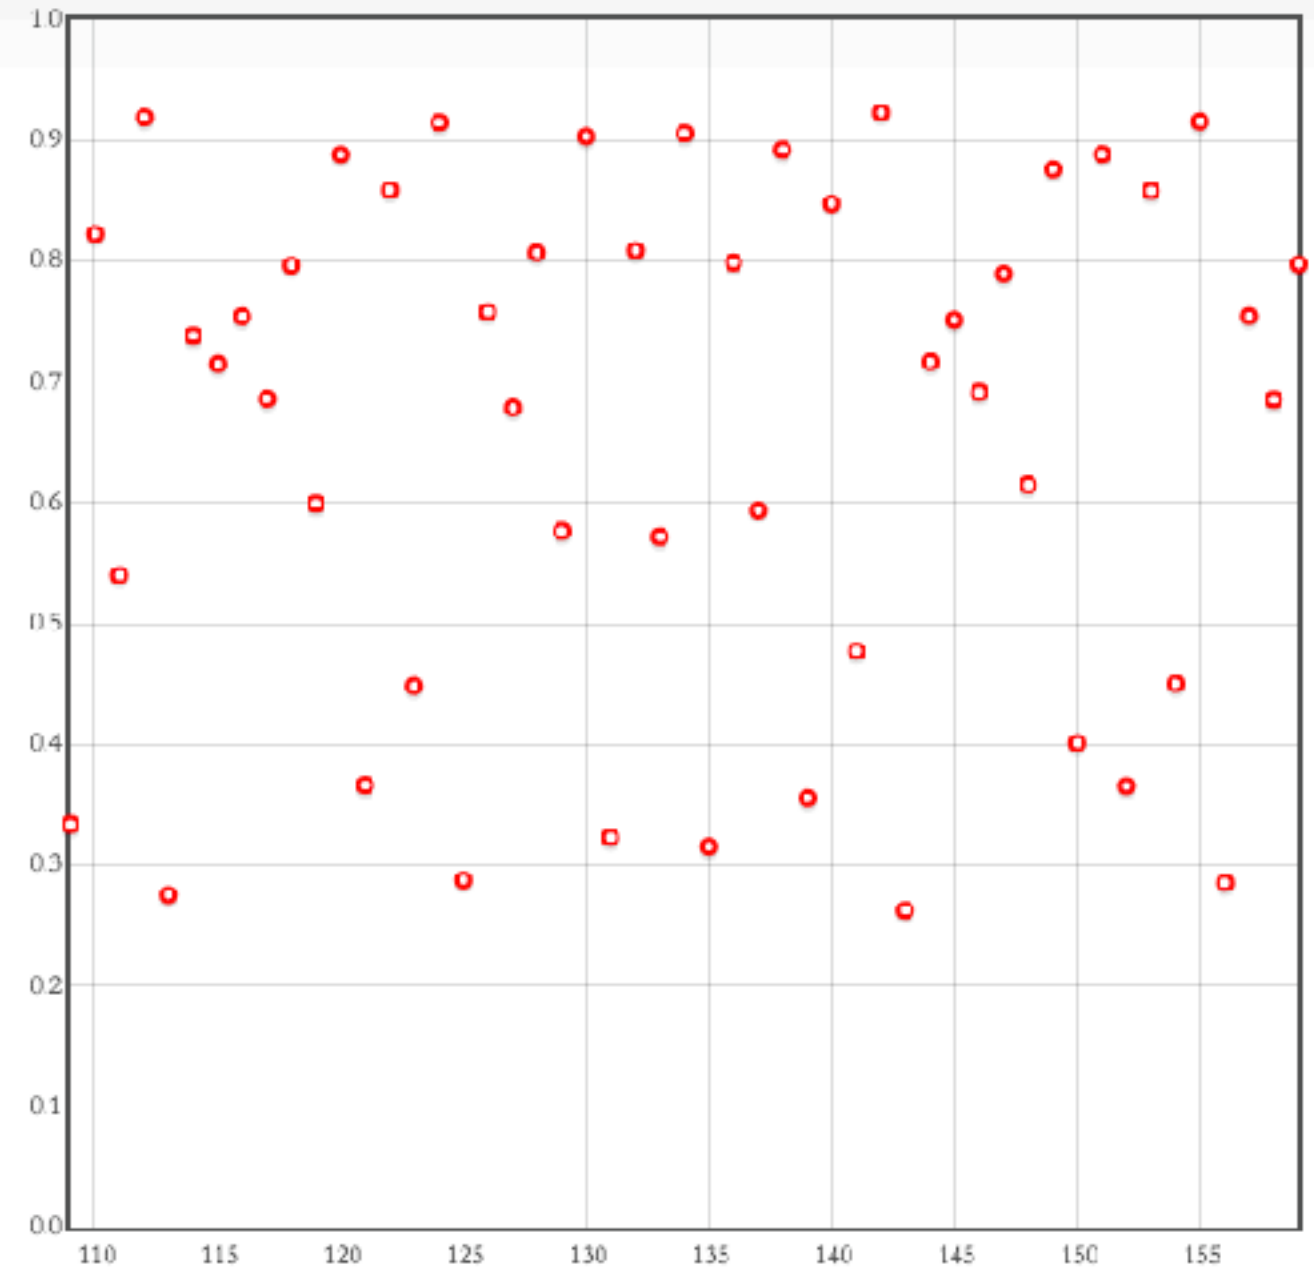

Chaos does not mean random!

It does entail unpredictability though.

A chaotic system exhibits *sensitivity to initial conditions*

Two trajectories that start arbitrarily close together will, over time, diverge, so that one contains no information about the other

A double-rod pendulum is a very simple system that illustrates a chaotic evolution in time.
Starting from a very slightly different position will result in an entirely different trajectory.

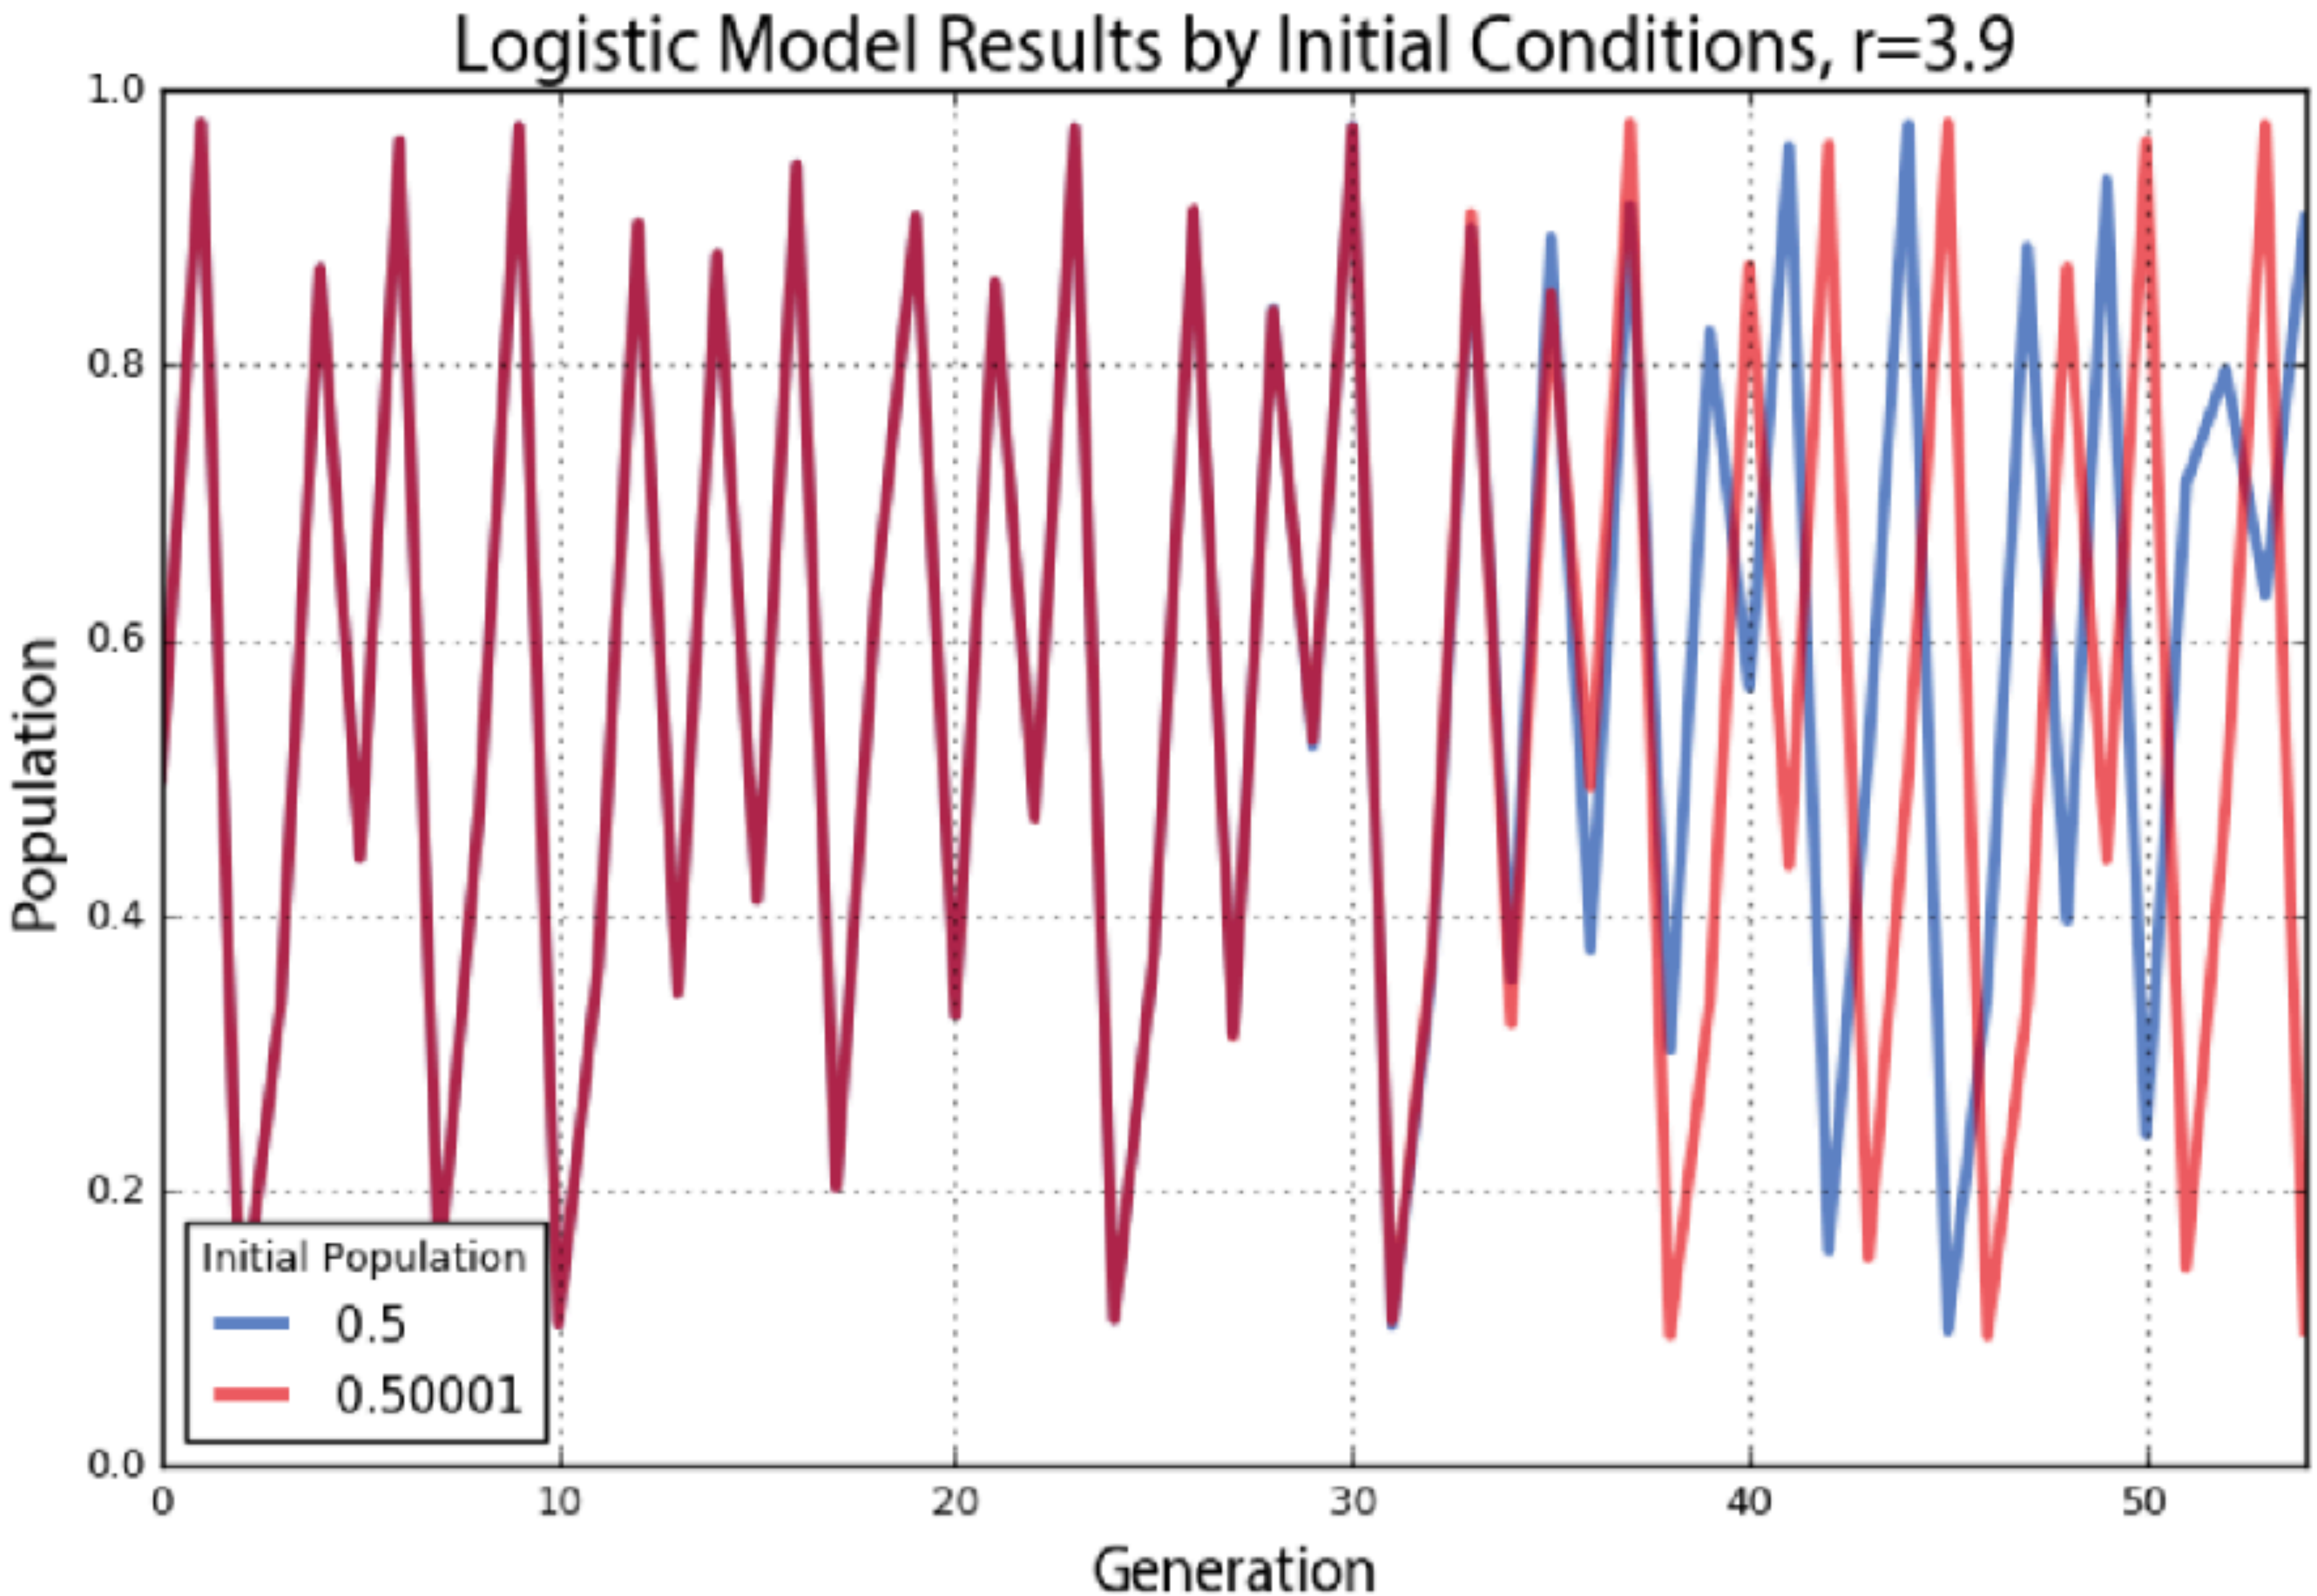

Figure 13. Plot of two time series with identical dynamics, one starting at an initial populati value of 0.5 (**blue**) and the other starting at 0.50001 (**red**). At this growth rate parameter value 3.9, the Lyapunov exponent is positive; thus, the system is chaotic, and we can see the nearby poir diverge over time.

When the qualitative form of the attractor layout changes as we vary a parameter (here, *r*), we say the system has undergone a bifurcation.

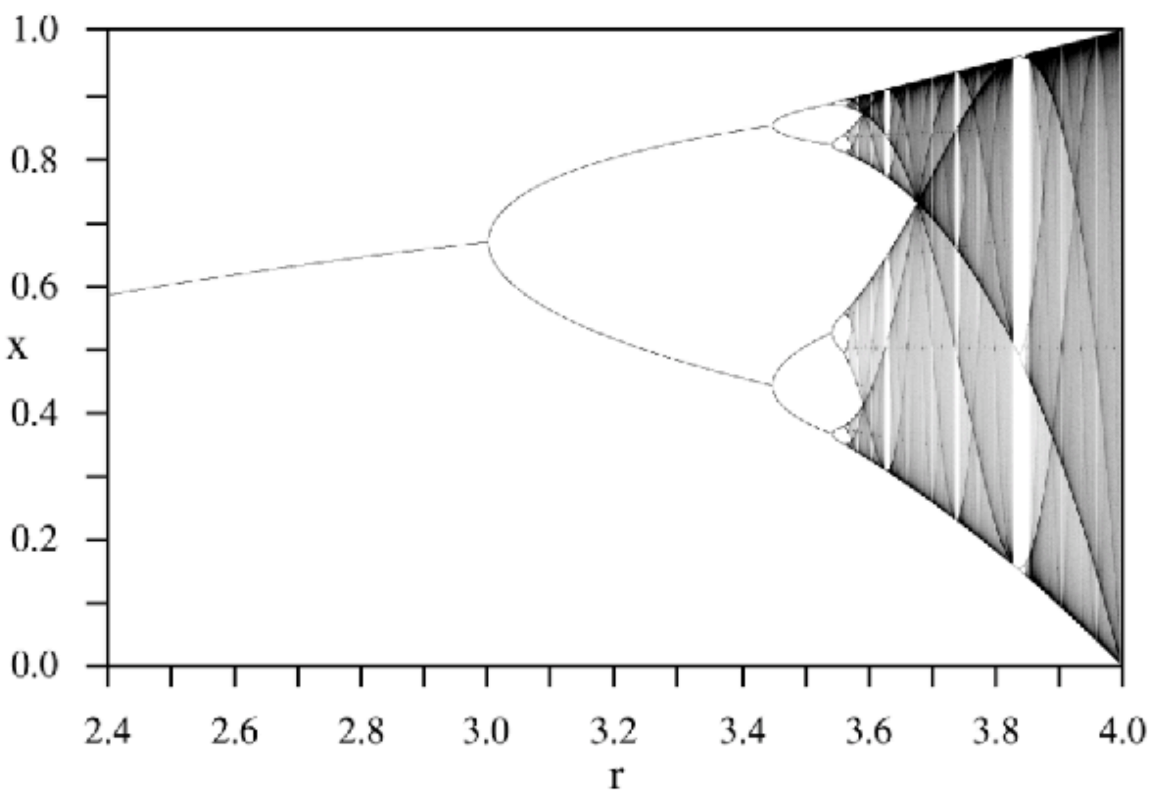

Bifurcation diagram for the logistic map

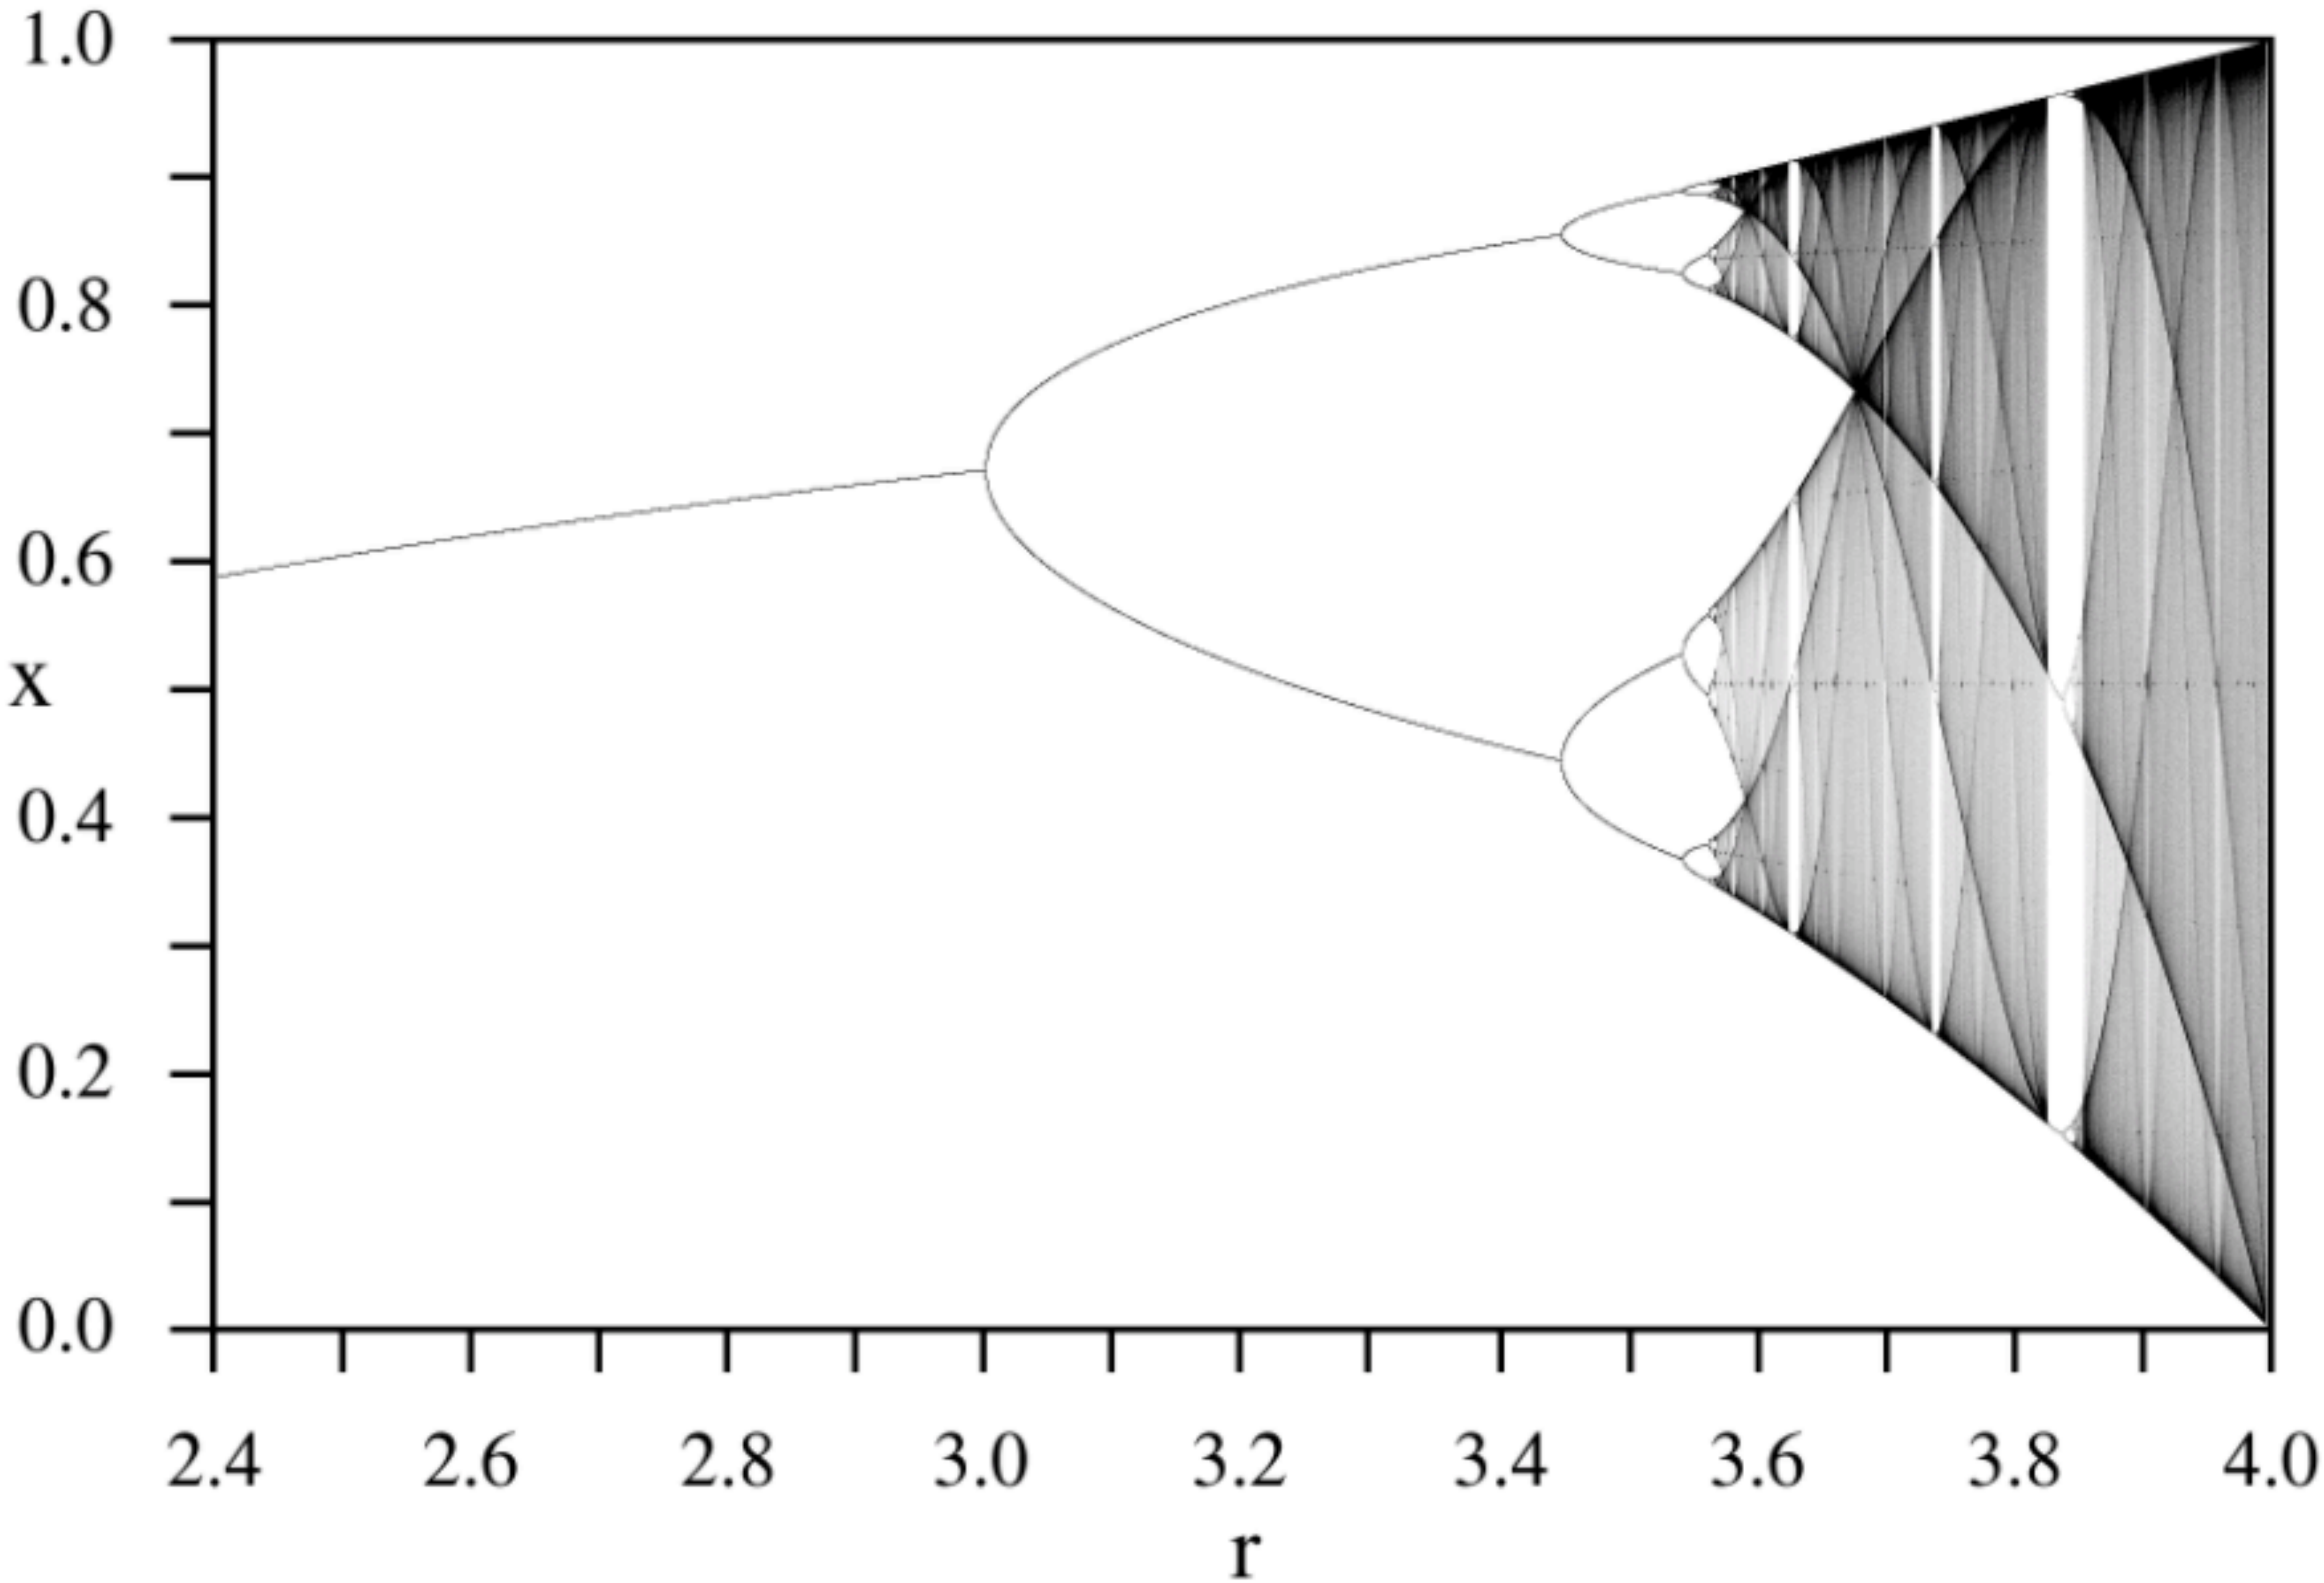

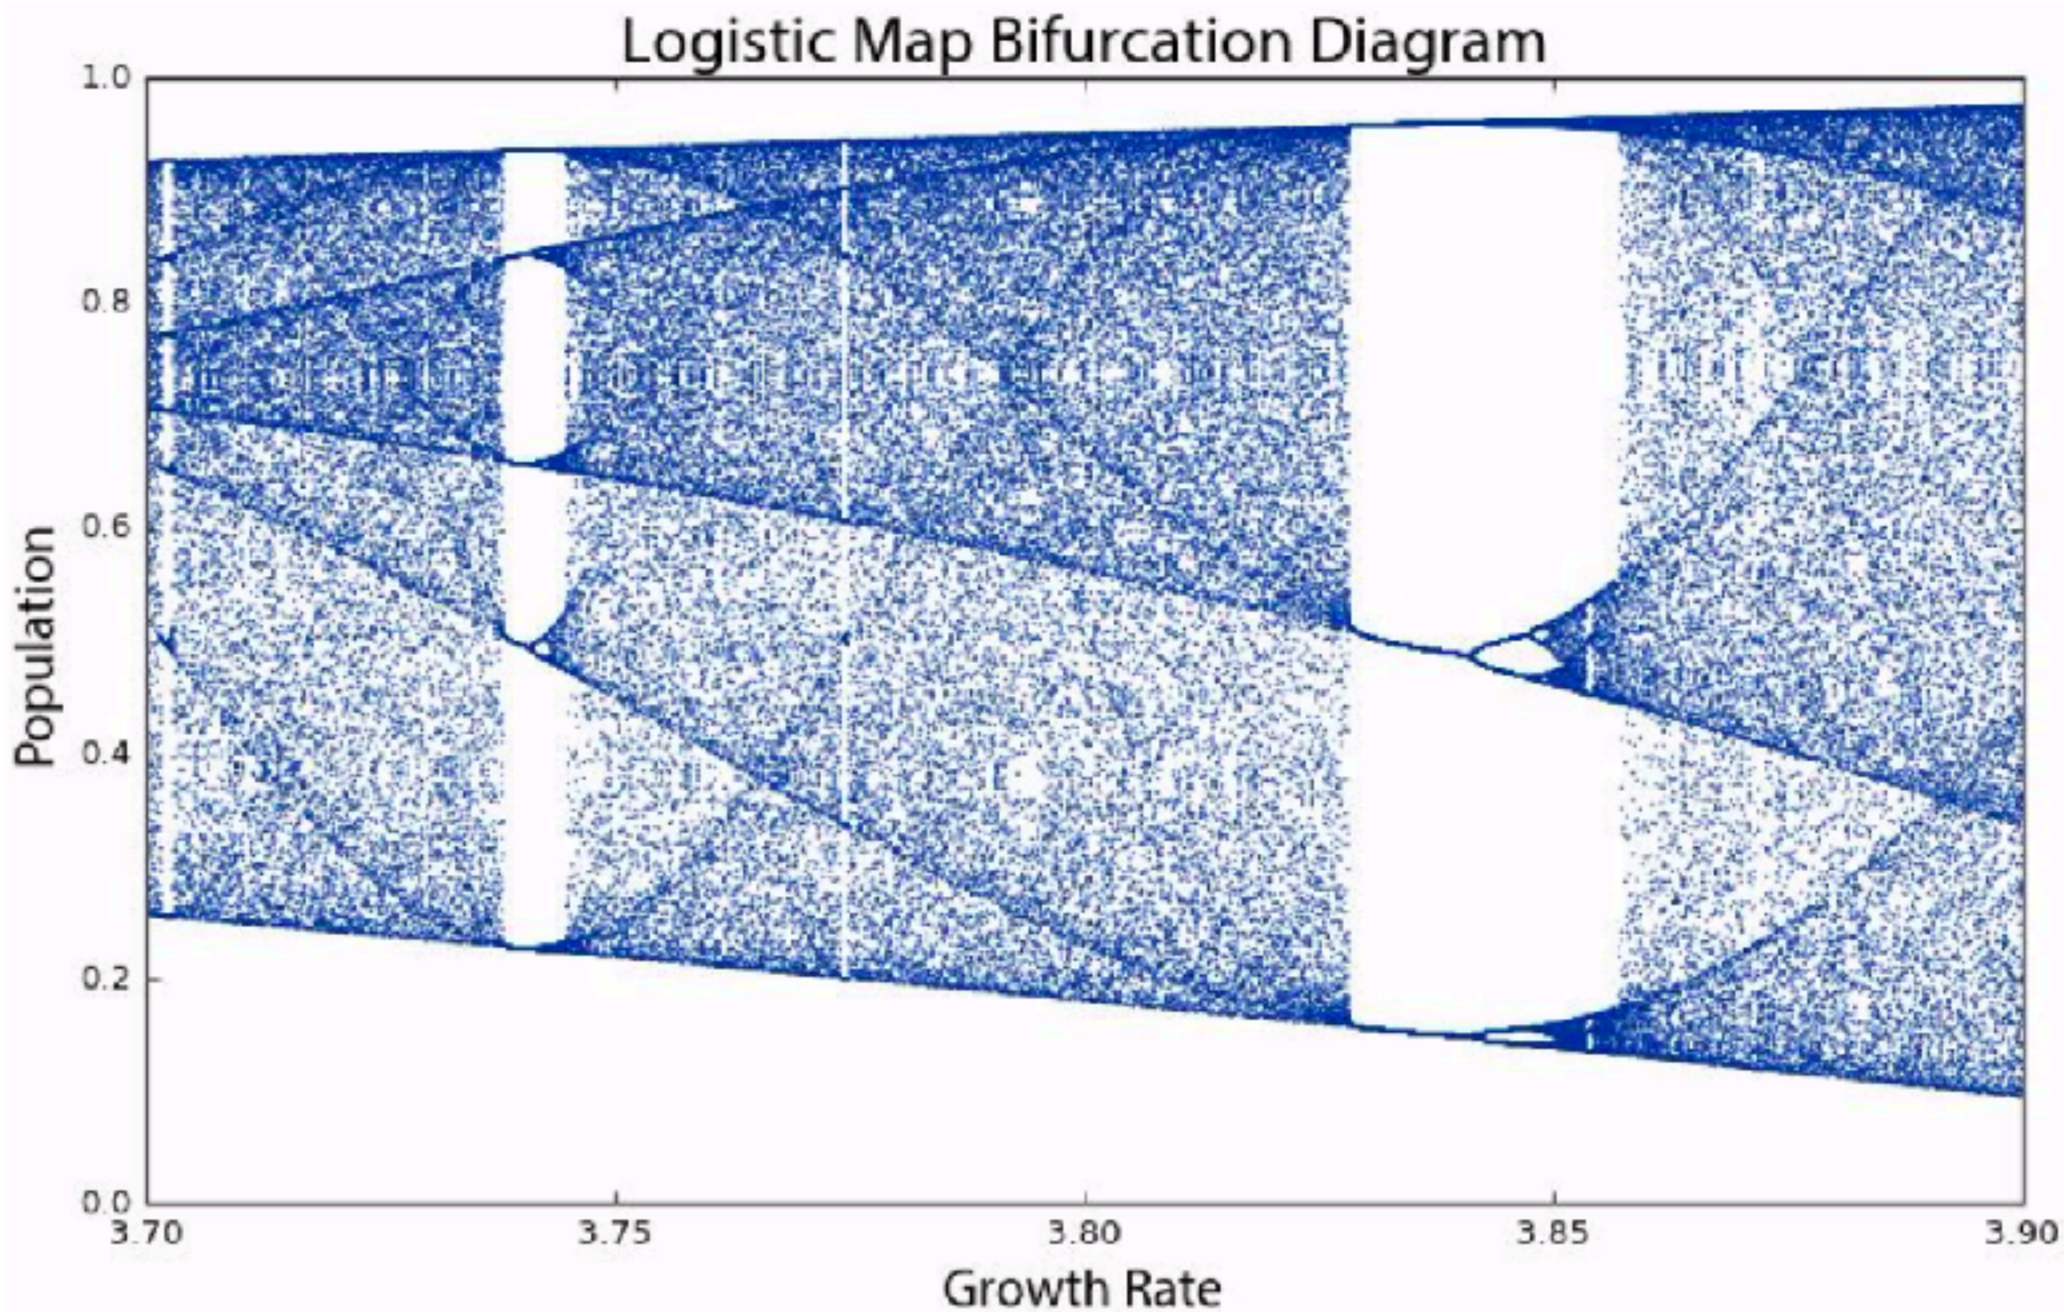

Figure 4. Bifurcation diagram of 100 generations of the logistic map for 1000 growth rate parameter values between 3.7 and 3.9. The system moves from order to chaos and back again as the growth rate is adjusted.

The bifurcation diagram for the logistic map has a *fractal* structure: self-similarity at all scales

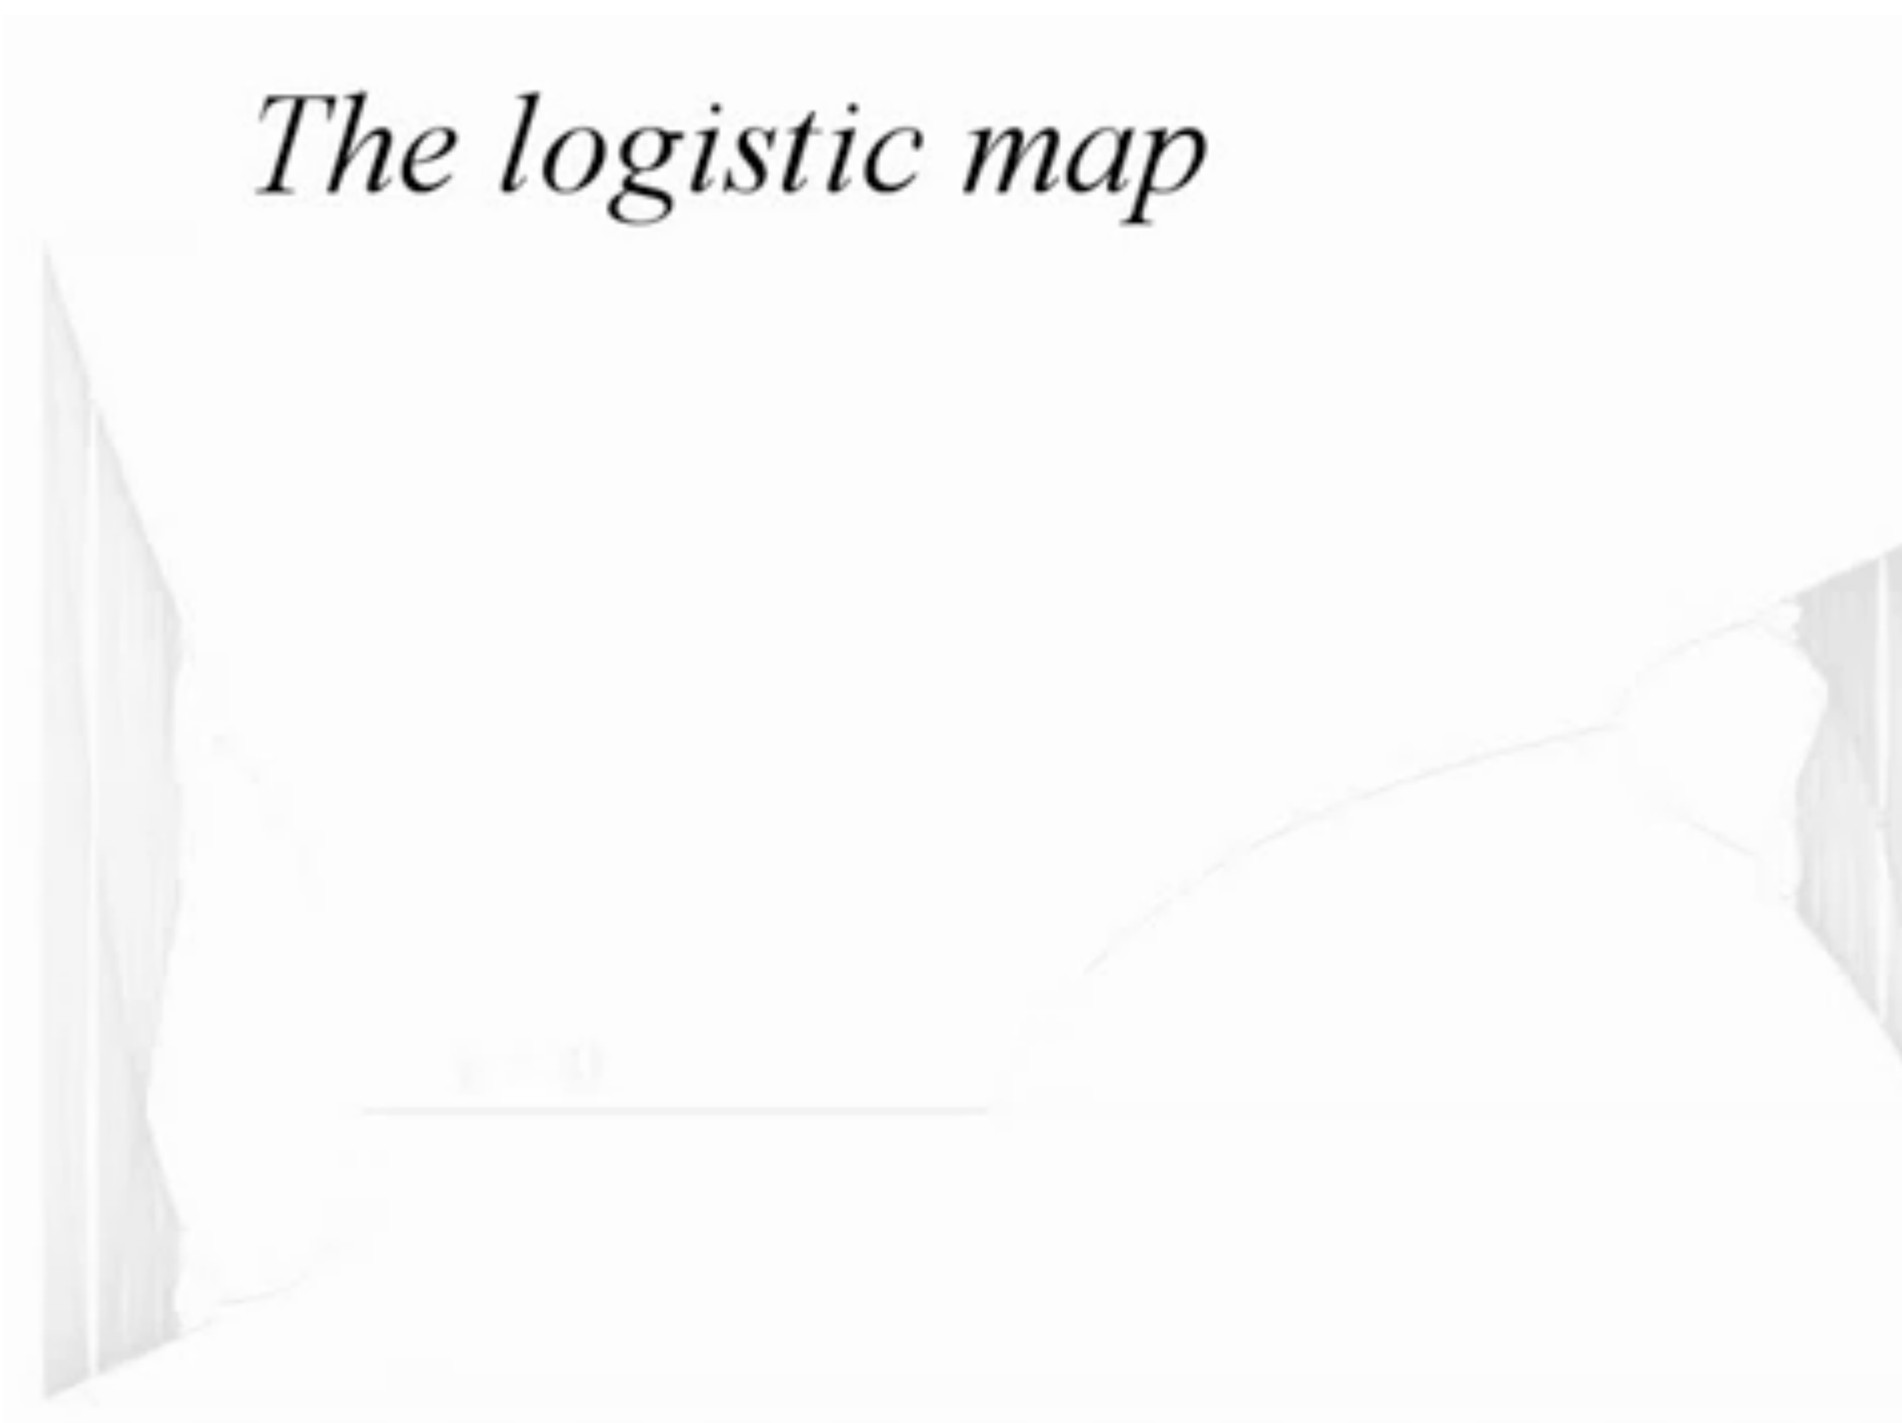

Transients

If attractors represent the ultimate fate of a system, we might also be interested in the short term behaviour:

These are *transients*

Transients: as the network settles into the solution state, error for individual patterns changes, and not always monotonically.

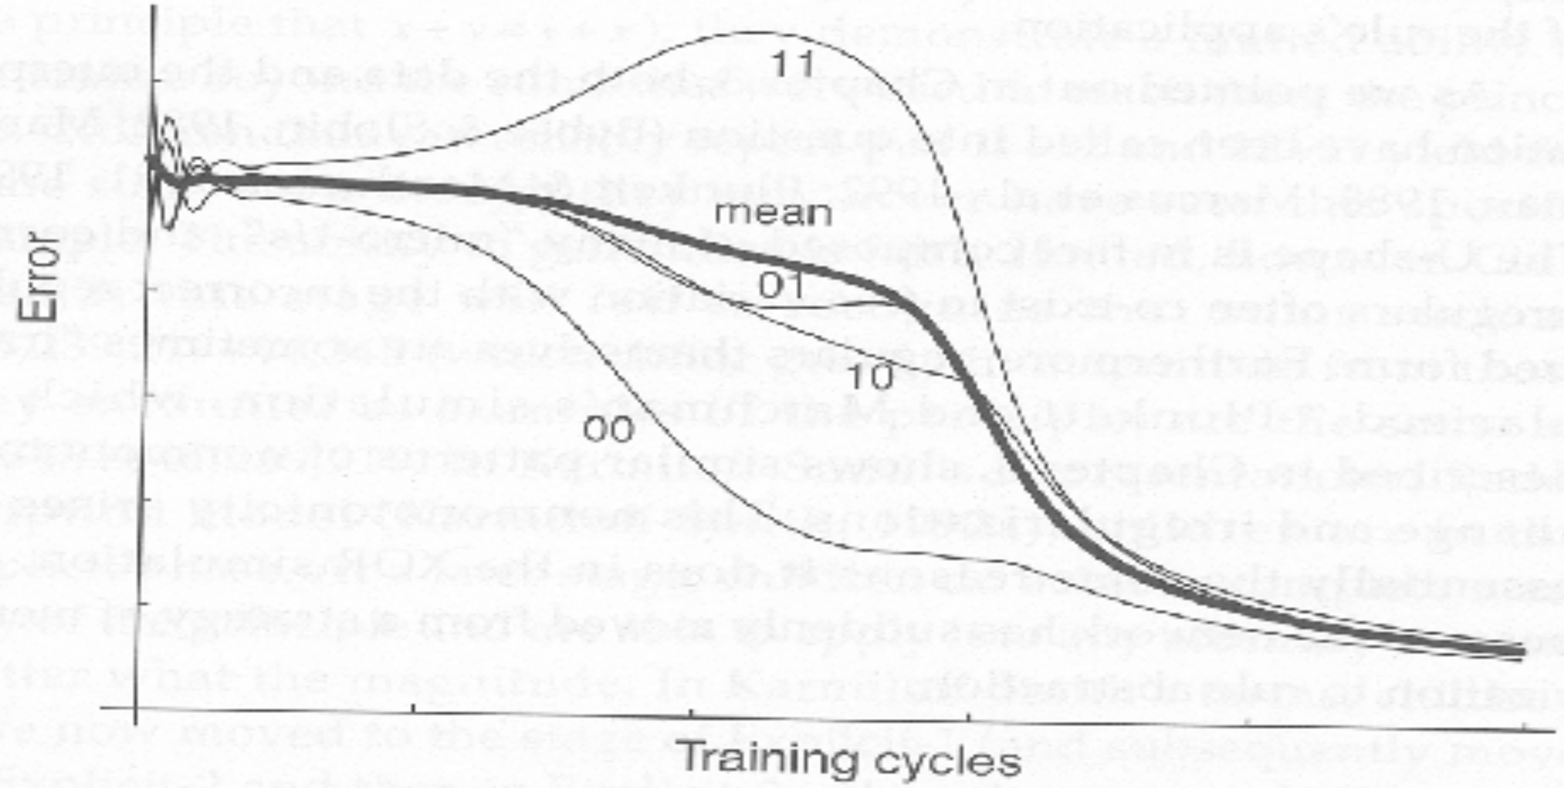

FIGURE 4.26 Error plot during learning the XOR function. The heavy line shows the mean error calculated over successive batches of 100 training cycles. The lighter lines show the error produced by each of the four patterns which comprise XOR. Two of the patterns (00 and 10) have monotonically decreasing error; the other two (01 and 11) temporarily increase error before a solution is found which accommodates all four patterns.

When you pluck a guitar string, the vibration that occurs is a transient.

The equilibrium state is a state of rest (silence).

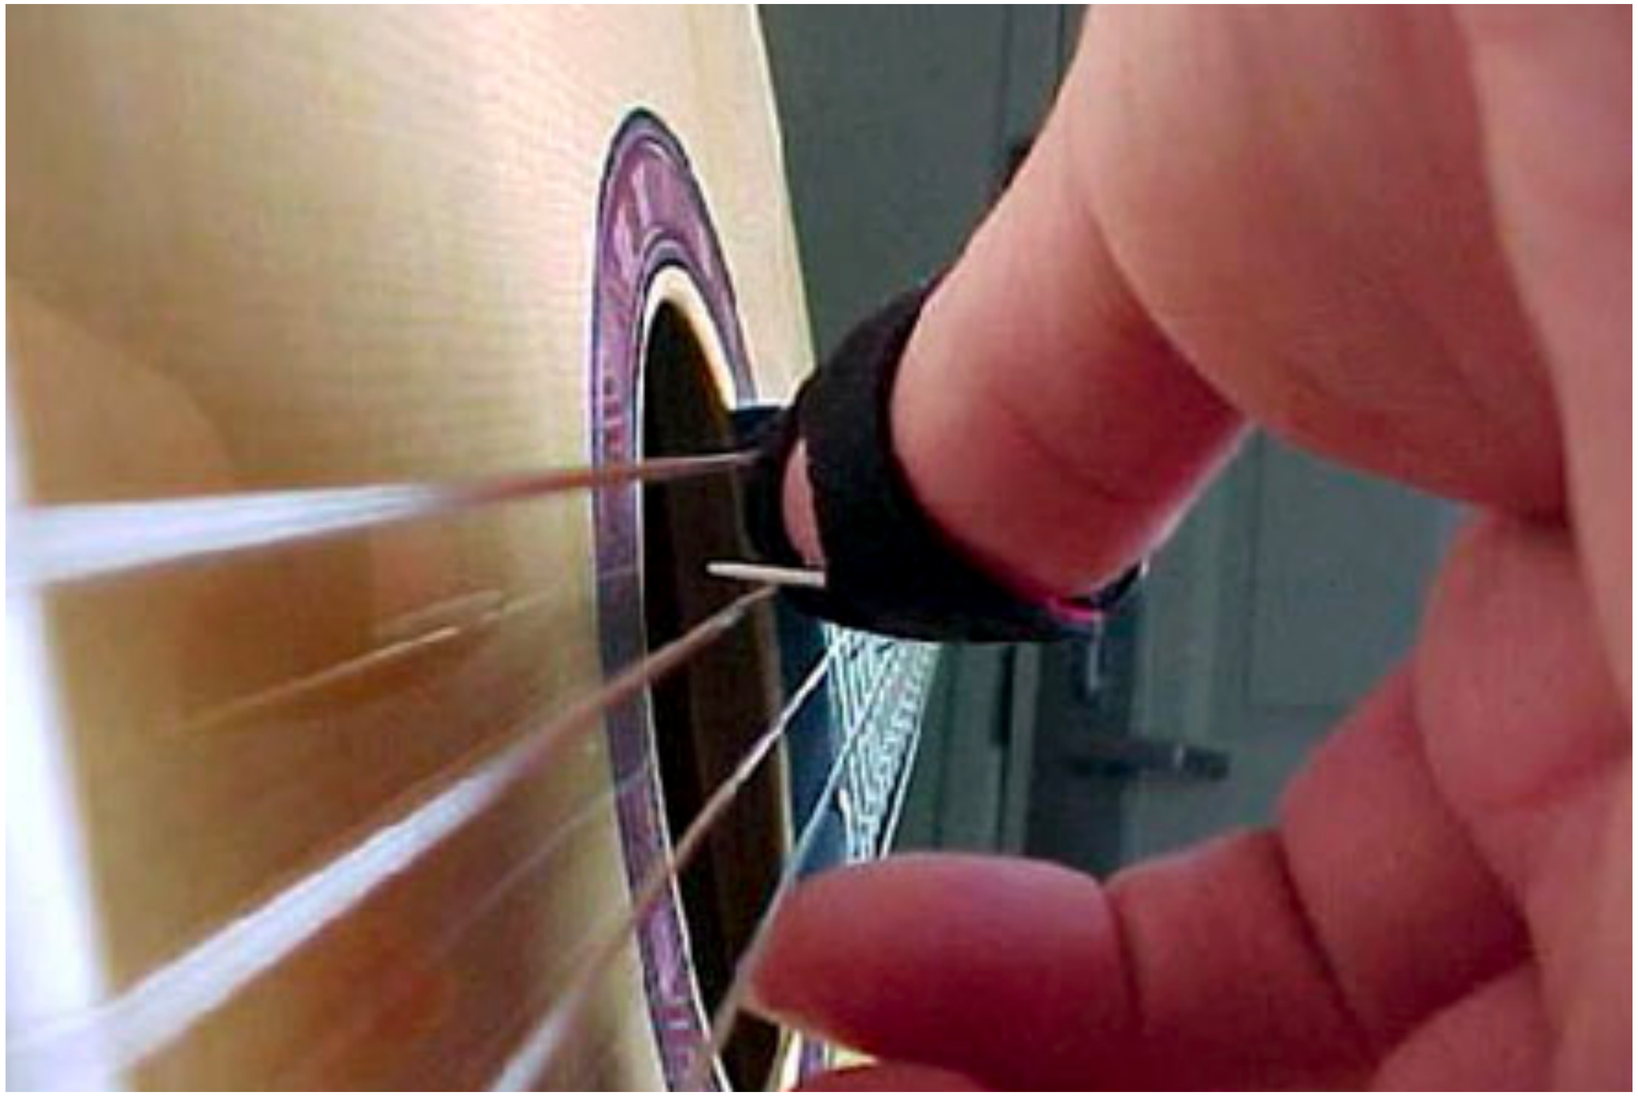

Another key concept that finds expression in the vocabulary of dynamic systems: **Autonomy**

Autonomy in everyday language: being a law unto yourself.

Consider
* the autonomy of the body,
* or of the region of Catalonia,
* or of the Judicial branch of government.

There is a mathematical analogue

(That is all mathsmatics ever provides: a model or analogue of things in the world)

If we have identified a system with state $\boldsymbol{X(t)}$

Then we distinguish between an autonomous system:

$$\boldsymbol{dX/dt = f(x, parameters)}$$

and a driven system:

$$\boldsymbol{dX/dt = f(x, parameters, t)}$$

A small difference in the maths. A huge conceptual difference.

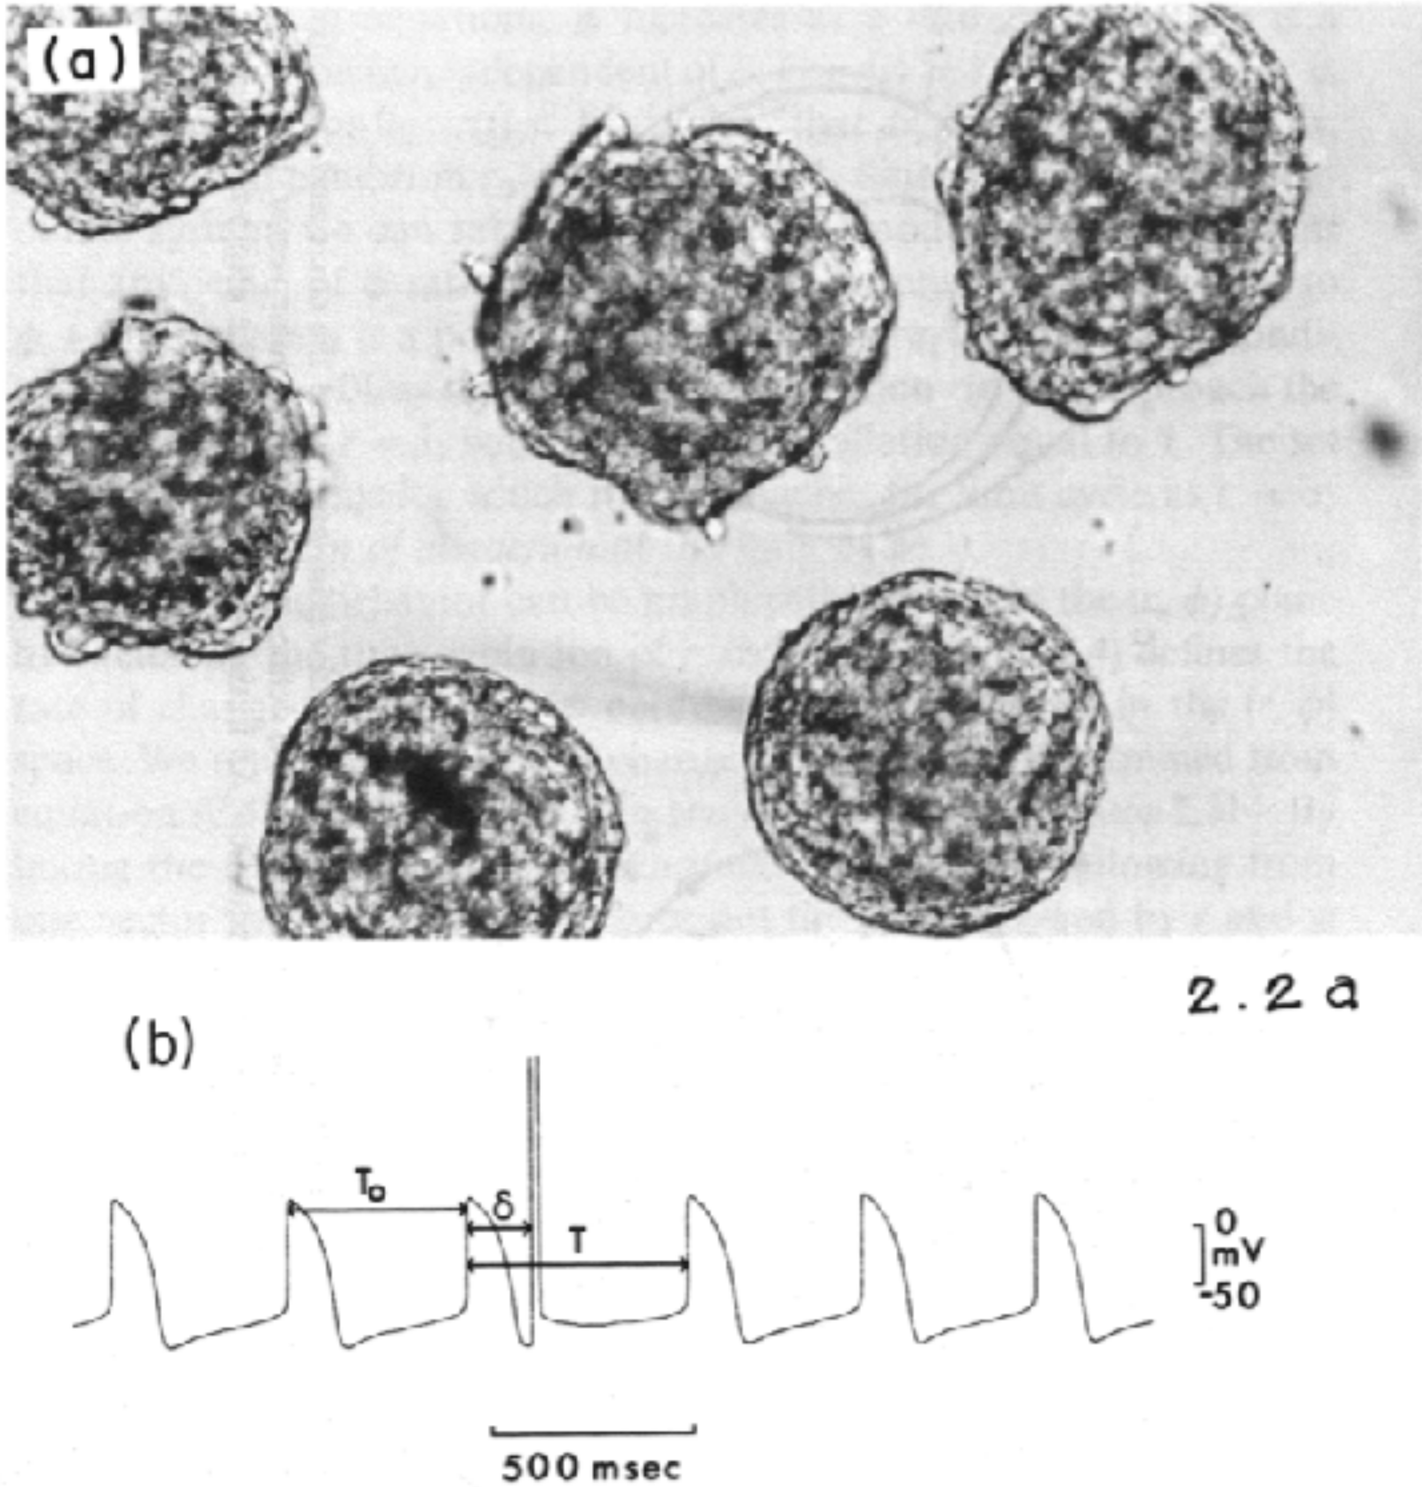

2.2. (a) Aggregates (about 100 mμ in diameter) of spontaneously beating heart cells derived from the ventricles of 7-day-old embryonic chicks. All cells in a single aggregate are electrically coupled and beat with the same intrinsic frequency. Photograph provided by A. Shrier. (b) Transmembrane potential from an aggregate showing spontaneous electrical activity and the effect of a 20-msec, 9-nA depolarizing pulse delivered through an intracellular microelectrode. The control cycle length is T_0 and the perturbed cycle length T. From Glass et al. (1984).

Chick Heart Cells display their own periodic behaviour.

A perturbation disturbs the activity, but it soon re-establishes itself.

The cells are exhibiting autonomy.

In general, if a system is autonomous, it may be perturbed or influenced, but its reaction to a perturbation is governed by its own dynamic.

Autonomy is the hall mark of the living.

Autonomous systems may enter into many kinds of coupling, but they remain in principle decomposable.

Let’s look at a counter-example: predator-prey relations speak of mutual dependence that does not decompose

Interactive Exploration of a Dynamical System

Bret Victor / May 17, 2011
http://worrydream.com

Predator + Prey form a *single* system, not two coupled systems

If we take away one, the other goes away as well. Neither has its own, intrinsic, behavior.