UC Berkeley

UC Berkeley Previously Published Works

Title

Applications of dynamical systems to deterministic and stochastic economic models

Permalink

<https://escholarship.org/uc/item/8rp0s1sh>

Author

Hirsch, MW

Publication Date

1999

Peer reviewed

FIELDS INSTITUTE COMMUNICATIONS

THE FIELDS INSTITUTE FOR RESEARCH IN MATHEMATICAL SCIENCES

Topology and Markets

Graciela Chichilnisky
Editor

American Mathematical Society
Providence, Rhode Island

Applications of Dynamical Systems to Deterministic and Stochastic Economic Models

Morris W. Hirsch

Department of Mathematics
University of California at Berkeley
Berkeley, CA 94720-3840, USA
hirsch@math.berkeley.edu

Abstract. Markets and related games are modeled by several types of deterministic and stochastic dynamical systems. In simple cases, long run behavior of all or most trajectories can be determined. Convergence to stable equilibria is emphasized. After reviewing dynamical systems concepts, applications are made to X. Vives' oligopoly games, including Bertrand competition in differentiated markets. Limiting behavior of a class of stochastic processes is shown to be closely related to deterministic dynamical systems. These results are applied to B. Arthur's urn models of stochastic allocation processes, and to repeated games in which players use the long term strategy of fictitious play.

1 Introduction

Dynamical systems theory provides a mathematical way of looking at *processes* of all kinds where long-term behavior is of importance — biological, chemical, physical, etc. Many economic processes have been modeled as mathematical dynamical systems; (see, for example [30], [10]).

There is a peculiarity in economic systems, especially markets, that makes it difficult to model them as dynamical systems: The essence of an economy is *choice*, while the motivating idea of dynamical systems is *determinism*. One way of seeing this problem is to first look at how mathematics deals with somewhat analogous problems in physics and biology. The mathematics of Newtonian mechanics has no difficulty in writing down differential equations of motion for any number of perfectly elastic particles. In principle, then, such equations describe a perfect, ideal gas. But we are then helpless to derive any conclusions from this enormous set of equations. In any case, we are not interested in following the detailed trajectories of these particles. What we want is some useable description of the *ensemble* behavior, in terms of a small number of measurable quantities.

1991 *Mathematics Subject Classification.* 00A69, 90D35, 90Dxx, 57Rxx.
Partially supported by a grant from the National Science Foundation.

This is exactly what statistical mechanics does — in terms of temperature, entropy and pressure. These are properties of the ensemble which are not easily expressible in terms of the variables — position, momentum, acceleration — describing states of individual particles.

The problem in fluid mechanics is much worse — we don't even have language to describe the large-scale behavior of a turbulent fluid, and the relevant partial differential equations are difficult to analyze; even the existence of solutions is problematical.

In these examples the problem lies with the complexity of a system made up of a very large number of well-understood particles. Turning to biology, we find a further problem: each individual "particle" — nerve cell, organism, species — is extremely complex. And while physics has laws, such as conservation of energy and the second law of thermodynamics, which give quantitative restrictions on ensemble behavior, analogous laws in biology have yet to be discovered. In market economics the "particles" are individual economic agents. Here we face both the problem of large scale aggregation of individual activities, as well as the problem of choice, absent in physics.

In modeling large scale ensembles it seems necessary to assume that quantities describing activities of individuals can somehow be lumped or averaged so as to yield a set of numbers sufficient to describe (for modeling purposes) the instantaneous state of ensemble behavior.¹

Once we have settled on such quantities, the next task is to discover the laws describing how the state changes with time. These laws constitute the *dynamic* of the system. Two common modeling methodologies are to postulate (1) deterministic dynamics, or (2) stochastic dynamics.

In the deterministic market models investigated here, we assume that each agent acts in a way which (she thinks) will maximize, or at least increase, her immediate profit, or rate of profit.

In the stochastic game models, we assume agents' actions can be described in probabilistic terms, even though each agent may be acting deterministically in the light of her knowledge. This knowledge, however, may depend on exogenous "random" quantities. This is not the same as postulating that each agent makes a random choice — whatever that may mean. Rather this approach assumes that the distribution of choices over an ensemble of similar agents behaves according to probabilistic laws.

In many writings on economic theory, a third procedure is used for resolving the choice-determinism dilemma: Ignore it, and simply assume that a kind of state defined as an "equilibrium" is reached — without specifying any dynamical system for which it is an equilibrium! While such approach may be useful, it will not be pursued here.

Given the description of the conceivable states of the system, and the dynamic (or the stochastic), the mathematician's task is *to derive laws describing how states of the system evolve in time*. While this problem can be stated as a question in pure mathematics, in reality our intuitions and insights into the original problem are crucial here. We may hope for useful dialog between mathematicians familiar with the relevant branches of dynamical systems, and scientists trying to model a natural system.

¹The celebrated impossibility theorems of Arrow, and Chichilnisky, illustrate the the difficulty of agglomerating individual choices in plausible ways.

While very simple equations can specify extremely complex dynamical systems, one of the most important uses of dynamical theory has been to determine condition enabling us to decide whether a given class of systems is *convergent* — every trajectory tending to an equilibrium (fixed point). This point of view is emphasized throughout, although oscillatory and chaotic systems are considered briefly in Section 11.²

After looking at a very simple dynamical model of duopoly in Section 2 as an introduction to dynamical ideas, we review basic dynamical concepts in the next three sections. These are applied to a more elaborate market model in Section 6. Some recent dynamical theory is presented in the next section, and applied to Bertrand competition in the following section.

In the final three sections we look at a very different class of models; rather than being deterministic, they are stochastic. Yet these particular models — generalized Polya urns and repeated games with uncertainty — are of such a nature that in the long run their behavior is closely connected with deterministic systems of differential equations — a kind of dynamical analog to the law of large numbers. Therefore knowledge of the dynamics of the differential equation leads to knowledge of the limiting behavior of sample paths of the stochastic process. This has long been known for very simple vector fields; recent new results of this kind are presented here.

In Section 10 we look at stochastic allocation processes, which have been used to model competing technologies. There is a precise sense in which the limiting behavior of such a process corresponds to limiting behavior of a vector field associated to the process.

The allocation processes we consider are just the same as generalized Polya urn processes. In Section 11 further dynamic results are applied to these processes. Under certain conditions we conclude that almost surely a sample path tracks a trajectory of the vector field. This leads to the construction of oscillatory and even chaotic urn models.

In Section 12 we turn to another mathematical subject that has been used for economic models, and in fact was inspired by economics, namely Game Theory. Here we treat games played repeatedly by the same players (for simplicity we consider two players), whose payoff matrices are subject to random perturbations, where players adopt long-term policies leading to a description of the evolution of their strategies as stochastic processes very similar to urn models. Again a dynamical system aids in the analysis of long-term behavior.

2 A simple model

We consider a simple example of the oligopoly games treated by X. Vives [29], of the type he calls “Cournot competition in differentiated markets”.

Consider two competing producer-sellers, $\mathcal{A}_1$ and $\mathcal{A}_2$, of the same commodity. Each decides independently what quantity x_i ($i = 1, 2$) to produce; we assume it will all be sold. The market arrives at a unit price³ p , which is a function $p = p(x_1 + x_2)$. Producer $\mathcal{A}_i$ makes profit

$$\Pi_i = x_i(p - C_i)$$

²At present the rigorous mathematical theory of chaotic dynamical systems is very undeveloped. The most useful service it has thus far performed for science is simply to demonstrate beyond doubt that chaotic systems are prevalent — and that most of them are impossible to analyze rigorously, or even describe crudely. We lack not only theorems, but language.

³Vives deals with inverse demand instead of price.

where $C_i = C_i(x_i)$ is her total unit cost. Each agent then revises her production goal.

How should we model this?

The first modeling issue is to describe the *set of conceivable states of the system*. We assume provisionally that the production level x_i can be any nonnegative real number; more realistically we could restrict production levels to some finite interval. We take the state of the system at any time to be a vector $x = (x_1, x_2)$. The *state space* is therefore the space $\mathbf{R}_+^2$ of ordered pairs of nonnegative real numbers.

We also postulate that the price p is a nonnegative function of production.

The next question is: What kind of *time* are we using? In a *discrete time* model all quantities are recalculated at a succession of equal time intervals, usually days or years. In a *continuous time* model, we imagine continuously changing quantities, governed by a differential equation. Both choices have advantages and disadvantages, and in many situations neither is particularly realistic (imagine transactions taking place at random times). But we are not trying to describe reality, only to model it. For a discrete time model, the time variable t takes values in the $\mathbf{Z}$ of integers or in the set $\mathbf{N} = \{0, 1, 2, \dots\}$ of natural numbers. For a continuous time model, t lies in the set $\mathbf{R}$ or real numbers, or in the set $\mathbf{R}_+$ of nonnegative real numbers.

Now the crucial question: how do states evolve as time passes? Suppose we consider a discrete time system, taking $t \in \mathbf{N}$. Given the state $x(t)$ at time t , what can we say about the state $x(t+1)$?

To answer this we must make some assumptions about the producers' behaviors. We assume here that each wants to increase profit. Here the problem of choice versus determinism is evident: There may be many ways to increase profit.

A standard way to proceed is to specify a range of choice for each producer, depending on current production levels and price, and then to analyze all the possible sequences of events that may ensue. The resulting mathematical formalism is not generally thought of as a dynamical system, and the standard results of dynamical system theory do not apply.

To obtain a deterministic dynamical system we need some assumption that completely determines a producer's behavior in terms of the *current* prices, costs and production levels. A simple (to state!), classical assumption is:

Each producer acts to maximize profit assuming that the production level of the other producer is unchanged.

Thus $x_i(t+1) = y_i$, where y_i maximizes

$$y_i(p(y_i + x_j(t)) - C_i(y_i))$$

for $i, j \in \{1, 2\}, i \neq j$. Implicit in this formulation is the unrealistic assumption that *each producer has complete knowledge of the price function p* . More on this later.

The system is not yet deterministic: We need to assume that such y_i are unique. This can be guaranteed by various assumptions, but here we don't need to specify them. Thus we simply assume:

For each i and each $\xi \in \mathbf{R}$, the function

$$\eta \rightarrow \eta(p(\eta + \xi) - C_i(\eta))$$

takes its maximum value at a unique $\eta = f_i(\xi) \in [0, \infty)$.

Under this uniqueness assumption, we see that the production vector $x(t+1)$ depends only on $x(t)$. Thus there is a map

$$F: \mathbf{R}_+^2 \rightarrow \mathbf{R}_+^2 \quad (2.1)$$

such that $x(t+1) = F(x(t))$.

This map is the dynamic of our system. Starting from any initial state $a \in \mathbf{R}_+^2$, we obtain a sequence of production vectors $a = a_0, a_1, \dots$ where $a_{n+1} = F(a_n)$ for $n \in \mathbf{N}$. This sequence is the *trajectory* of the initial state a . We also write this as $a_n = F^n(a)$, where F^n denotes the n 'th iterate of F , that is, the composition $F \circ \dots \circ F$ of F with itself n times (and F^0 denotes the identity map).

Notice that F has the special form

$$F(x_1, x_2) = (f_1(x_2), f_2(x_1)) \quad (2.2)$$

where f_i solves the maximization problem described above.

The main dynamic question is: *What happens in the long run?* To attack this problem we need to know the nature of the map F , which in turn depends on the price function p and the unit cost functions C_i .

It is natural (and almost unavoidable) to postulate that the exogenous function $p: \mathbf{R}_+ \rightarrow \mathbf{R}_+$ is nonincreasing. To simplify matters we assume also that p is continuously differentiable with negative derivative.

What is to prevent our agents from raising production levels past all bounds? Recall that we assumed the market absorbed all production; this is unreasonable if we allow unlimited production. So we need to postulate some feature of our model which prevents unlimited production.

A natural way to do this is to assume that at high production levels, unit costs exceed prices and both profit levels Π_i are negative. Since we are also assuming prices decrease with total production, we postulate:

There exists $z_* > 0$ such that if $z \geq z_*$ then $p(z) < C_i(z)$, $i = 1, 2$.

We of course assume both producers know of such a z_* . This ensures that each producer will always keep her production level below z_* .

With this assumption we can effectively limit the state space to $[0, z_*] \times [0, z_*]$; this space, being compact, has great technical advantages. More precisely, the system is dissipative (Section 3), attracted to a subset of $[0, z_*] \times [0, z_*]$.

It is reasonable to assume that production is feasible, i.e., profitable for each producer, at *some* production levels — otherwise one or both producers would go out of business.

We want to use calculus if possible, so we want all our functions to be continuously differentiable (denoted by C^1). While to assume p and C_i are C^1 seems innocuous, it is not reasonable to assume that F is C^1 , since we have already prescribed a complicated derivation of F from p and the C_i . We had better try to find some plausible additional assumption which implies F is C^1 .

We consider f_1 in detail, the case of f_2 being similar. By definition, $f_1(\xi)$ is that value of η (assumed unique) where the function $\eta \mapsto \eta(p(\eta + \xi) - C_1(\eta))$ takes its maximum value. Define

$$G(\xi, \eta) = \eta(p(\eta + \xi) - C_1(\eta)).$$

Then $f_1(\xi)$ satisfies

$$\frac{\partial G}{\partial \eta}(\xi, f_1(\xi)) = 0.$$

Now $\partial G/\partial \eta$ is the function

$$J(\xi, \eta) = \eta p'(\xi + \eta) + p(\xi + \eta) - C_1'(\eta).$$

Thus $f_1(\xi)$ is locally defined implicitly as the solution to the equation

$$J(\xi, f_1(\xi)) = 0.$$

The implicit function theorem says that f_1 is C^1 provided that

$$\frac{\partial J(\xi, \eta)}{\partial \eta} \neq 0$$

wherever $J(\xi, \eta) = 0$, and in this case

$$f_1'(\xi) = -\frac{J_\xi(\xi, f_1(\xi))}{J_\eta(\xi, f_1(\xi))}$$

where J_η denotes $\partial J/\partial \eta$, etc.

Now

$$\frac{\partial J(\xi, \eta)}{\partial \eta} = \eta p''(\xi + \eta) + 2p'(\xi + \eta) - C_1''(\eta).$$

Some thought shows that as we have already assumed $p' < 0$, the only reasonable way to guarantee this is different from zero (taking $\eta = f_1(\xi)$) is to postulate that each summand is ≤ 0 . There is nothing for it but to make the new assumptions:

$$p'' \leq 0, \quad C_i'' \leq 0.$$

And of course we are now assuming C_i and p are C^2 (continuous second partials).

Are these assumptions plausible? $C_i'' \leq 0$ is easily obtained by the drastic but common assumption of *linear* cost functions: $C_i(x_i) = \alpha x_i + \beta_i$. More generally it can be interpreted as requiring that marginal costs $C_i'(x_i)$ decrease (or don't increase) with increased production. At least this is economically meaningful, and not absurd.

The assumption of a concave price function $p(s)$ (i.e., $p'' \leq 0$) seems less natural to me, but I leave it to economists to justify or discard it.

The assumption of concave pricing has an important mathematical implication, namely that $f_i' < 0$. For by the implicit function theorem we have (suppressing arguments of functions):

$$f_1' = -\frac{J_\xi}{J_\eta};$$

we already know that $J_\eta < 0$, and one easily sees that

$$J_\xi = \eta p'' + p' < 0,$$

proving that $f_1' < 0$; and the proof for f_2 is similar.

This leads to an interesting conclusion about the long-term dynamical behavior of the system:

Theorem 2.1 *Every trajectory of the system approaches an orbit of period 2 (possibly a fixed point).*

This means that for any initial state a there is a state b with the following properties:

1. b has *period 2*, meaning that $F^2(b) := F(F(b)) = b$;
2. the distance from $F^n(a)$ to the orbit of b (the set consisting of b and $F(b)$) goes to zero.

Proof Recalling the formula (2.2) for F , we see that the second iterate F^2 is defined by the formula

$$F^2(x_1, x_2) = (g_1(x_1), g_2(x_2))$$

where

$$g_1 = f_1 \circ f_2, \quad g_2 = f_2 \circ f_1.$$

Then

$$g_1'(z) = f_1'(f_2(z))f_2'(z) > 0$$

and similarly for g_2' . This shows that g_i is strictly increasing; therefore for any initial state (a_1, a_2) , the limit $b_i = \lim_{n \rightarrow \infty} (g_i)^n(a_i)$ exists. The resulting state (b_1, b_2) has the required properties. $\square$

I return to this example in Section 9. Before leaving it, however, consider how technical mathematical considerations, such as differentiability, have led us to postulate economically meaningful properties of markets, such as concavity of cost and price functions. This illustrates an important but little-understood influence of mathematics in science.

3 Dynamical systems

A dynamical system comprises a set X called the *state space* together with a rule called the *dynamic* for how states evolve in time. Here the state space is taken to be Euclidean n -space $\mathbf{R}^n$ comprising all n -tuples real numbers, or more generally a subset such as an orthant or n -cube. Each state vector $x \in \mathbf{R}^n$ is thought of as a numerical coding for a conceivable⁴ state of the economic system of interest. The dynamic is a collection Φ of maps $\Phi_t : X \rightarrow X$, indexed by real numbers t . For convenience we may also write $\Phi_t(x) = \Phi(t, x)$.

For any state $x \in X$ and $t > 0$ we interpret $\Phi_t(x)$ as the state to which x evolves after t units of time have elapsed; and, if negative times are allowed, for $-t < 0$ we think of $\Phi_{-t}(x)$ as the state which evolved into x from t units of time in the past. The map Φ_0 always denotes the identity map of X .⁵

The index t runs either through the set $\mathbf{R}_+$ of nonnegative real numbers — “continuous time” — or the set $\mathbf{N} = \{0, 1, 2, \dots\}$ of natural numbers — “discrete time”. Negative time usually has no practical meaning, but it often makes sense mathematically and can be theoretically useful. When negative times are allowed we always assume the maps Φ_t are one-to-one. Such a system is called *reversible*. In this case the maps Φ_t and Φ_{-t} are inverse to each other.

The dynamic for a continuous time system is usually specified implicitly by a differential equation of the form $dy/dt = F(t, y)$, $y(0) = x$, and defining $\Phi_t(x) = y(t)$.⁶

The theory of dynamical systems is most highly developed for *autonomous systems*, for which the dynamic obeys the *composition law* $\Phi_t \circ \Phi_s = \Phi_{t+s}$. Here the composition $x \mapsto g(f(x))$ of maps g, f is denoted by $g \circ f$. For a reversible

⁴Not necessarily a *possible* state! A dynamical systems model of a real system is nothing more than an elaborate thought-experiment.

⁵To be precise, some of the maps Φ_t may be defined only on a proper subset of X , as for example in the dynamics of the differential equation in one variable $dx/dt = x^2$. In almost all applications it is assumed, or provable without inordinate labor, that trajectories are defined at least for all $t \geq 0$.

⁶Here it is necessary to assume that the function $F(t, x)$ is continuously differentiable, or more generally, continuous in (t, x) and *locally Lipschitz* in x . Such assumptions validate the standard theorems on existence, uniqueness and continuity of solutions to the differential equations.

autonomous system it is easy to see that the maps Φ_t and Φ_{-t} are inverse to each other, so that in this case each Φ_t is a homeomorphism.

The dynamic for a discrete time autonomous system is given by a map g from the state space X to itself; we set $\Phi_k(x) = g^k(x)$ where g^k denotes the composition of g with itself k times (and g^0 denotes the identity map of X).

A continuous-time system is autonomous if it is defined by an autonomous (vector) differential equation, that is, an equation of the form $dx/dt = F(x)$, in which the right-hand side does not explicitly involve t .⁷ For such a system the collection of maps $\{\Phi_t\}$ is called the *solution flow* of the differential equations. We can rarely get an explicit formula for these maps, but many properties of the family of all solutions can be described in terms of them.

All systems are henceforth assumed to be autonomous unless the contrary is stated. The dynamic is assumed to be given by a smooth vector field (for continuous-time systems) or a smooth mapping (for discrete-time systems) unless the contrary is stated.

3.1 Orbits, limit sets, equilibria and cycles. Consider now an autonomous dynamical system (X, Φ) , for either discrete or continuous time. The *trajectory* of a point $x \in X$ is the mapping $t \mapsto \Phi_t(x)$. For a discrete-time systems this is a possibly infinite sequence of points of X ; for a continuous-time system, a parameterized curve in X . For a differential equation the trajectory of a state is the solution that passes through the state at time 0.

The image of the trajectory of state x is a subset of X called the *orbit* of x . The orbit comprises all states to which x can evolve in the future, and (for reversible systems), from which x has evolved in the past. "Orbit" and "trajectory" are loosely used interchangeably.

It often happens that a subset Y of the state space has the property that states in Y can evolve only to other states in Y , so that if y is any state in Y and $z = \Phi_t(y)$ for some positive or negative t , then necessarily $t \in Y$. Such a set Y is called *invariant*. Every orbit is an invariant set, and every invariant set is a union of orbits. The closure of an invariant set is invariant. Restricting attention to states in an invariant set defines another dynamical system, having a smaller state space but the same rule for computing trajectories.

We also speak of a *positively invariant set*, meaning a set Y with the property that if $y \in Y$ and $t \geq 0$, then $\Phi_t(y) \in Y$. An example is the *forward orbit* of x , which means the set of points $\Phi_t(x)$, $t \geq 0$.

A fundamental dynamical concept is that of a *limit point* of a state x : this means a state $y \in X$ such that $y = \lim_{k \rightarrow \infty} \Phi_{t_k}(x)$ for some sequence of times $t_k \rightarrow \infty$. It is important to realize that a state can have more than one limit point, because there may be many such sequences $\Phi_{t_k}(x)$ converging to distinct limits.

The set of all limit points of x is called the *limit set* of x (more properly the *omega limit set*), denoted by $\omega(x)$. We also loosely call it the limit set of the trajectory or orbit of x .

If x has a *unique* limit point p , this means that $\Phi_t(x)$ converges to p as $t \rightarrow \infty$. In this case p is an equilibrium state (defined below), and we say that the trajectory of x stabilizes at p . I emphasize, however, that for an autonomous reversible system the trajectory of x never actually reaches the state p (that is, p is not on the orbit of x), except in the case that it started at p , that is, when $x = p$.

⁷The composition law is proved using the uniqueness theorem for solutions of a differential equation $dx/dt = F(x)$, and the fact that if $x(t)$ is a solution, so is $x(t - b)$ for any constant b .

In the contrary case, a trajectory with more than one limit point wanders around forever near its limit without converging to any definite state.

The following useful facts about limit sets are proved using the composition law and continuity assumptions:

Proposition 3.1 (a) *Every limit set is closed and invariant.*

(b) *A closed, positively invariant set Y contains the limit sets of all states in Y .*

(c) *A state which lies in a compact, positively invariant set K has a limit set which is a nonempty compact subset of K , connected in the continuous-time case.*

A compact nonempty invariant set A is called an *attractor* if it contains the limit sets of all nearby states. More precisely, A is required to have a neighborhood W in the state space such that $\lim_{t \rightarrow \infty} d(\Phi_t(x, A)) = 0$, uniformly in $x \in W$. The set of all points whose omega limits are in A is the *basin* of the attractor.

Most dynamical systems used in modelling have a global attractor, i.e. an attractor whose basin is the whole state space. Such systems are termed *dissipative*.

Many dynamic features can be succinctly defined in terms of orbits and limit sets. For example, a *stationary* state or *equilibrium* is a state p which does not change, that is, $\Phi_t(p) = p$ for all t . Equivalently, $\omega(p) = p$.

For a system defined by a vector field F , a state p is an equilibrium if and only if $F(p) = 0$. For a discrete-time system generated by a map $f : X \rightarrow X$, the equilibria are the *fixed points* of f , i.e., the states p such that $f(p) = p$.

A state a is *periodic* if $\Phi_r(a) = a$ for some time $r > 0$, called a *period* of a . The limit set of a periodic state coincides with the orbit of the state. If a is periodic but not stationary, the smallest such r is called the *minimal period*, and the orbit of a is called a *cycle*. The composition law implies that if a has period r then so does every state on the orbit of a .

Many useful systems have the property that the limit sets of all (or “most”) orbits are equilibria or cycles. It is very important to be able to recognize such systems, since their behavior is reasonably describable; criteria are given in Section 5.1.

But such systems are exceptional! Even very simple-looking systems can have incredibly complicated limit behavior; see the Lorenz system, below. The first such system was discovered a century ago by Poincaré, the inventor of dynamical systems theory, in his study of celestial mechanics [22]. He despaired of describing the dynamic complexity of two planets revolving around the sun under the influence of Newton’s laws.

We can now state more precisely the task of dynamical systems theory: to describe the limit sets of the systems we are interested in. More modestly and realistically, we would like to do this for “most” limit sets of “most” such systems. In the present state of knowledge even this is possible only for very special systems, but nevertheless much practical information can be obtained. Even where complete rigor is not attainable, useful insights may be derived.

4 Stability and persistence

A periodic orbit (possibly an equilibrium state) γ is called *topologically stable* if it is an attractor. This means that if we follow the trajectory of any state that is sufficiently close to γ , eventually it will appear to repeat periodically — although in fact it approaches γ asymptotically without actually repeating any state, unless it was in γ to start with.

In particular, an equilibrium can be topologically stable. This notion makes precise the idea of a state to which the system will always return after any sufficiently small disturbance of the state.

There are well-known conditions which are sufficient (but not necessary) for an equilibrium p to be topologically stable. These assume the dynamics are smooth.

- Theorem 4.1** (a) For a discrete-time system defined by a smooth map $G : \mathbf{R}^n \rightarrow \mathbf{R}^n$ a fixed point p is topologically stable provided every eigenvalue of $DG(p)$ has absolute value less than 1.
- (b) For an autonomous system of differential equations $dx/dt = F(x)$, an equilibrium p is topologically stable provided every eigenvalue of $DF(p)$ has negative real part.

These conditions are not necessary, but they are so useful that when they hold then p is simply called *stable* or more precisely, *linearly stable*. The corresponding weak inequalities on eigenvalues can be shown to be necessary (but not sufficient) for topological stability.

Stability of p (as defined by the eigenvalue conditions) implies that nearby orbits can be proved to converge to p at an exponential rate:

- Theorem 4.2** (a) Let p be a stable equilibrium of a smooth vector field F . Suppose every eigenvalue of $DF(p)$ has real part $< -\lambda < 0$. Then there is a constant $C > 0$ and a neighborhood N of p such that $|x(t) - p| \leq C|x(0) - p|e^{-\lambda t}$ for every trajectory $x(t)$ starting in N .
- (b) Let p be a stable fixed point of a smooth mapping G . Suppose every eigenvalue of $DG(p)$ has absolute value $< \mu < 1$. Then there is a constant $C > 0$ and a neighborhood N of p such that $|G^k(x) - p| \leq C|x(0) - p|\mu^k$ for every $x \in N$.
- (c) There are linear changes of coordinates so that in the new coordinates (a) and (b) are valid with $C = 1$.

A stable equilibrium enjoys another important type of stability, called *persistence*: If the equations defining the system are subjected to any sufficiently small perturbation, then the perturbed system has a unique equilibrium near p , and it is stable for the perturbed system. In other words a stable equilibrium persists under small changes in the dynamic.

An equilibrium can be persistent without being stable: An example is the fixed point of the map x^3 . While unstable equilibria are hard to observe, they exert an important influence on the dynamics because states on nearby orbits change very slowly; this follows from the continuity assumption.

A somewhat weaker kind of stability is the following: Equilibrium p is *orbitally stable* if trajectories starting near p remain near p for all future times. Precisely: for every neighborhood N of p there is a smaller neighborhood W of p which is positively invariant.

A necessary condition for orbital stability of a fixed point p of a smooth mapping F is that every eigenvalue of $DF(p)$ have absolute value less than or equal to 1. For an equilibrium of a smooth vector field H , a necessary condition is that real parts of the eigenvalues of $DH(p)$ be nonpositive.

An equilibrium p of a vector field F is called *simple* if the matrix $DF(p)$ is invertible, or equivalently, 0 is not an eigenvalue. For a map F , an equilibrium (i.e., a fixed point) p is simple if 1 is not an eigenvalue of $DF(p)$. It is generic (see Section 4.2) for all equilibria to be simple. When that is the case, every orbitally stable equilibrium is stable.

A nonequilibrium periodic orbit Γ is *linearly stable* if it is an attractor, and each forward orbit $x(t)$, $t \geq 0$ in its basin is attracted to Γ at an exponential rate:

$$\text{dist}(z(t), \Gamma) \leq Ce^{-\lambda t}$$

for positive constants C, λ (depending on $z(0)$). This is equivalent to the condition that all but one eigenvalue of $D\phi_T(p)$ have absolute value < 1 , where $T > 0$ is the period of the Γ and p is any point of Γ . Unfortunately it is not known how to ascertain this condition purely from equations of the vector field. Nevertheless it has great theoretical importance.

4.1 Long-term behavior of systems. A dynamical system is *convergent* if every trajectory approaches some equilibrium. More generally, a dynamical system is *quasiconvergent* if every orbit approaches the set of equilibrium states, without necessarily converging to a particular one; such orbits are termed quasiconvergent. This is an unstable situation: it can only happen if the equilibrium set is infinite, but under generic assumptions (see Section 4.2) the number of equilibria is finite. Any smooth dissipative system can be approximated by systems whose equilibria are finite in number. In a simulation a quasiconvergent orbit will appear to converge.

It can happen that the trajectories of almost all states are convergent, or quasiconvergent; here “almost all” is used in the sense of measure theory, meaning “except for a set of initial states having Lebesgue measure zero”. Such systems are called *almost convergent*, or *almost quasiconvergent*.

An *oscillatory system* is one for which every limit set is a periodic (perhaps stationary) orbit. The pendulum without friction is an example. The weaker notions of *almost oscillatory*, *quasi-oscillatory*, etc., are defined in the obvious way.

Many systems appear to be oscillatory. It is frequently very difficult to prove they actually are oscillatory; another challenging problem is to identify the mechanism producing the oscillations. Such systems can be produced by coupling individual oscillators.

A system is called *chaotic* if it is very difficult to make reliable long-term predictions for it. Used this way the term is not precise, and in fact there is no generally agreed upon formal definition. Moreover it is exceedingly difficult to verify the formal properties that various writers have proposed for the term. The phrase “sensitive dependence on initial conditions” succinctly conveys the essence of chaotic systems.

A tell-tale sign of chaos is the existence of infinitely many periodic orbits in a bounded region, with unbounded minimal periods. A system that is not almost oscillatory is very likely to be chaotic.

Much of the useful information about the long-term behavior of a system trajectory is contained in the dynamics and geometry of the limit sets. What kinds of limit sets are possible?

We have already discussed periodic limit sets, which are the simplest to understand. Poincaré attributed great importance to periodicity. He conjectured, and C. Pugh and C. Robinson proved [24], that if p is a limit point, the system can be approximated by another system for which p is on a periodic orbit. In practice it is difficult to distinguish between periodic and nonperiodic limit points.

G. D. Birkhoff in [6], one of the founders of dynamical systems theory, wrote:

From the early times the mind of man has persistently endeavored to characterize the motions of the stars by means of periodicities. It seems doubtful whether any other mode of satisfactory description

is possible. The intuitive basis for this is easily stated: any motion of a dynamical system must tend with lapse of time towards a characteristic cyclic mode of behavior.

Birkhoff was thinking about mechanical systems, but the same intuition is relevant to the dynamics of all natural systems. It is interesting that today many scientists do not seem to share his intuition. In any case, taken literally and precisely, his "intuitive basis" (that all dynamical systems are oscillatory) is false (as well he knew); but it is true for many systems. The entire development of dynamics can be viewed as an attempt to find an accurate refinement of Birkhoff's intuition.

Beyond periodic limit sets there are a few well-understood classes of examples, many intriguing but poorly understood simulations, and an unfortunate amount of premature talk about the "new science of chaos". Only recently have we begun to develop concepts and language for describing and dealing with nonperiodic attractors. Not only has very little been rigorously proved, but there is not even a hint of a structure theory for such attractors.

The important thing to realize is that there *are* exotic limit sets and "strange attractors", and they are quite common. Many of them are robust: they are not due to some ingenious choice of coefficients, but on the contrary, they persist under small perturbations of the system. It is usually impossible, in our current state of knowledge, to prove that a particular system is chaotic.

By definition, reliable long-term predictions are not possible in such systems. It is not generally realized that reliable computer simulations are not possible either! The reason is that at every step of computation, errors are introduced by round-off and approximation — it is impossible even to enter the number $1/3$ accurately into digital computers which use base 2 for calculations.⁸ In other words, the simulation of a chaotic system is constantly jumping from one orbit to a nearby orbit. But "sensitive dependence on initial conditions" means that distinct orbits, even though they start out close together, eventually diverge. This phenomenon is easily illustrated with simulations of the *Lorenz system*:

$$\begin{aligned}\frac{dx}{dt} &= -10x + 10y \\ \frac{dy}{dt} &= 28x - y - xz \\ \frac{dz}{dt} &= -\frac{8}{3}z + xy.\end{aligned}$$

4.2 Generic behavior. An important concept for dynamics is that of "generic" phenomena, meaning "typical" or "highly probable"; it is closely allied to notions of stability, persistence and robustness.⁹ In a particular system one speaks of generic states and their trajectories. In dealing with a class of systems one often discusses generic systems in the class.

In discussing generic behavior one must be careful to specify the class of systems under consideration. For example one could consider only frictionless pendulums, for which it is not only generic, but true without exception, that every orbit is periodic. But for pendulums with friction, only two orbits (the equilibria) are periodic — generically, orbits are convergent. The meaning of "generic" depends

⁸Here the Babylonians, who used base 60, had the advantage.

⁹Like most dynamical ideas this goes back a century to Henri Poincaré: In his work on celestial mechanics he excluded "infinitely improbable" situations, such as all the planets moving in a straight line toward the sun.

on the class of systems, and the precise mathematical formulation of "generic". Regrettably there are many such formulations, invented for different situations, and they are not consistent.¹⁰

But all definitions of "generic" share the following characteristics:

- If a property applicable to members of some topological space S is generic, then the members of S which have the property form a dense subset of S . For example, if S is a space of vector fields and P is some generic property, then any vector field in S can be approximated arbitrarily closely by fields having property P .
- The logical product of any finite or countably infinite set of generic properties is generic.

Many useful generic properties are known, but in these notes we will refer only to the following generic property:

Lemma 4.3 *It is C^k generic, for every $k \geq 1$, for every equilibrium to be simple.*

Unfortunately most rigorously proved generic properties apply only to very large classes, such as the class of all C^k systems for some particular k ; and this does not necessarily imply that a particular property will be generic for a smaller class of systems. For example, I don't know if it is generic in some sense for the economic models discussed in this paper to have only simple equilibria (although this is very likely so).

When generic properties are discussed without specifying the class, this usually refers to the class of all C^1 systems.

5 Convergence criteria

The simplest dynamical systems are the convergent ones. It is of prime importance to have at hand convenient criteria for determining whether a given set of equations produces convergent dynamics.

5.1 Liapunov functions. Probably the most useful criterion is the existence of a *Liapunov function*: This means a continuous real-valued function E on the state space which strictly decreases in time along every nonstationary trajectory, and which is bounded from below.¹¹ Examples: the energy function for mechanical systems with friction; the error function for many automatic learning schemes; negative entropy for statistical mechanical systems.

Consider a smooth continuous-time system $dx/dt = F(x)$ and a smooth real-valued function E on the state space. There is a convenient test for E to be Liapunov: By the chain rule

$$\frac{dE(x(t))}{dt} = DE(x)F(x) = \langle \text{grad}E(x), F(x) \rangle, \quad (5.1)$$

where $\langle u, v \rangle = \sum_i u_i v_i$ denotes the inner (or dot) product of vectors $u, v \in \mathbf{R}^n$. This shows that E is *Liapunov precisely if* $\langle \text{grad}E(x), F(x) \rangle < 0$ at all nonstationary states x . Notice that this algebraic condition can be checked without computing

¹⁰A common definition of a generic property in some complete metric space S is that it defines a subset which contains the intersection of a countable family of open dense sets. Another definition, applicable to measure spaces, is that the property holds for a set whose complement has measure zero. In $\mathbf{R}^n$ the second definition implies the first.

¹¹For a dissipative system it is not necessary to assume Liapunov functions are bounded below, as this is automatically satisfied.

solutions of the differential equation. It follows from this condition that a state p is an equilibrium precisely when it is a critical point for the Liapunov function E , that is, when $\text{grad } E(p) = 0$.

The most natural class of systems admitting Liapunov functions is the class of *gradient systems*. Such a system is generated by a system of differential equations of the form:

$$\frac{dx_i}{dt} = -\frac{\partial E}{\partial x_i}, \quad i = 1, \dots, n \quad (5.2)$$

or in vector form $dx/dt = -\text{grad } E$. Here it is usually assumed that E is differentiable of class C^2 so that the vector field $-\text{grad } E$ will be C^1 , allowing application of the standard theory of differential equations.

It is readily computed that along a solution curve $x(t)$ we have

$$\frac{d}{dt} E(x(t)) = -|\text{grad } E|^2.$$

It follows that E is a Liapunov function.

There is a simple test for determining whether a smooth vector field F on $\mathbf{R}^n$ is the gradient of some function: F is a gradient if and only if all the Jacobian matrices DF are symmetric, i.e., $\partial F_i / \partial x_j = \partial F_j / \partial x_i$.

The main facts about Liapunov functions are as follows:

Theorem 5.1 *Let $E : X \rightarrow \mathbf{R}$ be a Liapunov function for a dissipative dynamical system (X, Φ) . Then:*

- (a) *The system is quasiconvergent, so there are no periodic orbits other than equilibria.*
- (b) *Every local minimum point for E is an orbitally stable equilibrium.*
- (c) *The topologically stable equilibria are exactly the strict local minimum points of E .*
- (d) *The system is convergent in each of the following two cases:*
  - (i) *$\text{grad } E = 0$ at only a finite number of states;*
  - (ii) *the system is the gradient of a real-analytic function.*

Unfortunately there are no general methods for deciding if a system admits a Liapunov function, or for constructing one if it exists.

Sometimes a Liapunov function for a system can be discovered by noting that the system is closely related to a gradient system:

Theorem 5.2 *Let E be a smooth function on $\mathbf{R}^n$. Let H be a vector field of the form*

$$H_i(x) = a_i(x) \frac{\partial E}{\partial x_i} \quad (5.3)$$

where each $a_i : \mathbf{R}^n \rightarrow \mathbf{R}$ is a smooth nonnegative function which is positive wherever $H \neq 0$. Then E is a Liapunov function for the system

$$\frac{dx_i}{dt} = -H_i(x). \quad (5.4)$$

Proof First a geometric proof. The vector $-\text{grad } E(x)$, if it is nonzero, is orthogonal to the level surface S of E through x , pointing in the direction of decreasing E . Each component of $-H_i(x)$ is a positive multiple of the corresponding component $-\text{grad } E(x)$. Therefore $-H(x)$ points towards the same side of S . It

follows that E decreases on the trajectory of $-H$ through x , since that curve is tangent to $-H(x)$ at x .

Here is a shorter but (to me) less illuminating analytic proof: By the chain rule,

$$\begin{aligned}\frac{d}{dt}E(x(t)) &= -\sum_i \frac{\partial E}{\partial x_i} H_i \\ &= \sum -a_i(x) \left(\frac{\partial E}{\partial x_i} \right)^2\end{aligned}$$

which is negative if $H(x) \neq 0$. □

6 Another economic model

We consider an example similar to that in Section 2, but with a large number n of producers $\mathcal{A}_i$ making the same product. Production is now an n -vector $x = (x_i)$. The price of a unit of product is given by an exogenous function $p(s)$, $s = x_1 + \dots + x_n$. The cost of production to $\mathcal{A}_i$ is $C_i(x_i)$. We assume the dynamic is describable by differential equations governing the production rates dx_i/dt .

Each producer wants to increase her profit $\Pi_i = p(x)x_i - C_i(x_i)$, but rather than maximize it, she merely tries to increment it. Thus we assume that x_i is adjusted so as to increase Π_i . Since

$$\frac{d\Pi_i(x(t))}{dt} = \frac{\partial \Pi_i}{\partial x_i} \frac{dx_i}{dt},$$

we postulate:

$$\text{Sign } \frac{dx_i}{dt} = \text{Sign } \frac{\partial \Pi_i}{\partial x_i}, \quad (6.1)$$

where for any real number s we define $\text{Sign } s$ to be $s/|s|$ if $s \neq 0$, and to be 0 if $s = 0$. Therefore we assume there is a positive function $a_i(x)$ such that

$$\frac{dx_i}{dt} = a_i(x) \frac{\partial \Pi_i}{\partial x_i}.$$

Computing $\partial \Pi_i / \partial x_i$ gives us the system of differential equations determining the dynamic:

$$\frac{dx_i}{dt} = a_i(x)[p(x) + p'(x)x_i - C_i'(x_i)]; \quad i = 1, \dots, n. \quad (6.2)$$

Unfortunately this system seems difficult to analyze. We need a simplifying assumption.

There are two assumptions that have mathematically similar effects, but different economic interpretations. One is to assume zero (or constant) costs, so that $C_i' = 0$; we leave this analysis as an exercise. Instead we make the psychological assumption that *producers ignore the effect of their production changes on the price*. In other words they compute $\partial \Pi_i / \partial x_i$ holding p constant. This is not unreasonable if there are very many producers, or if production changes are very small (and here they are infinitesimal!). So instead of system (6.2), we now postulate:

$$\frac{dx_i}{dt} = a_i(x)[p(x_1 + \dots + x_n) - C_i'(x_i)]; \quad i = 1, \dots, n. \quad (6.3)$$

Notice that an equilibrium of (6.3) is a state x at which all marginal production costs equal the unit price.

System (6.3) is almost a gradient system:

Lemma 6.1 *There is a function $V : \mathbf{R}^n \rightarrow \mathbf{R}$ such that*

$$\frac{\partial V}{\partial x_i} = p(x_1 + \dots + x_n) - C_i'(x_i),$$

and this is > 0 if x is not a constant solution.

Proof Define

$$V(x_1, \dots, x_n) = \int_0^s p(r) dr - \sum_i C_i(x_i); \quad s = \sum_i x_i.$$

□

We conclude that trajectories of (6.3) tend to critical points of V .

At this stage we impose a typical "generic" hypothesis (except that we don't know that it is actually generic!): We assume that the function V has only a finite number of critical points (where its gradient vanishes) and that equilibria of (6.3) (which are the critical points of V) are simple. Finally we assume that system (6.3) is dissipative, which means there is a number $R > 0$ such that every solution eventually enters and stays in the ball of radius R about the origin. This is a consequence of plausible assumptions on prices and costs.

It then follows that every solution $x(t)$ is defined for all $t \geq 0$, and approaches an equilibrium as $t \rightarrow \infty$. Moreover the assumption of simple equilibria implies that except for a closed set of measure zero, the trajectory of every state converges to a stable equilibrium (which is the same as a local minimum point of $-V$).

This is a powerful conclusion: Every trajectory converges to an equilibrium, and most converge to stable equilibria.

We can learn more by computing the Hessian of $-V$ at any state x : This is the symmetric matrix

$$-DV(x) = -p'E + \text{diag} \{C_1'', \dots, C_n''\}$$

where E is the $n \times n$ matrix with every entry equal to 1, and "diag" denotes a diagonal matrix.

We assume for simplicity that all production rates x_i are positive at equilibria.¹²

If e is an equilibrium, then it is stable if and only if all the eigenvalues of $-DV(e)$ are negative. A standard result in matrix theory, Gerschgorin's Theorem¹³ [19], implies:

Lemma 6.2 *The eigenvalues of $-DV(e)$ are negative provided that*

$$\max_i C_i''(e_i) < np' \left(\sum e_i \right).$$

If this condition holds at every equilibrium, then there is only one equilibrium.¹⁴

We can interpret Lemma 6.2 as follows. Since C_i'' is the rate of change of marginal demand with respect to supply, and p' is the rate of change with respect to price, we conclude:

¹²This assumption can be interpreted as the assumption of a mature economy, in which only profit-making firms have survived. But in fact the phenomenon of firms leaving and entering the economy is a difficult one to treat dynamically, as it changes the dimension of the state space.

¹³Gerschgorin's Theorem for a real or complex square matrix A states that for each i there is an eigenvalue in the closed disc centered at A_{ii} with radius $\sum_{j, j \neq i} |A_{ij}|$.

¹⁴The basins of distinct stable equilibria are disjoint open sets. But the state space is connected—it cannot be the union of disjoint nonempty open sets. Therefore if every equilibrium is stable there can be only one equilibrium.

If increasing production causes marginal costs to decrease sufficiently more rapidly than demand, then System (6.3) is globally convergent.

Assuming that a globally convergent market system is desirable (which is by no means clear), how could it be established? One possibility is for the government to tax revenues at a sufficiently high rate r , $0 < r < 1$. In this case profits are recomputed as

$$\Pi_i = (1 - r)px_i - C_i,$$

so holding p constants gives

$$d\Pi_i/dx_i = (1 - r)p - C_i'.$$

Instead of (6.3) we have:

$$\frac{dx_i}{dt} = a_i(x)[(1 - r)p(x_1 + \cdots + x_n) - C_i'(x_i)]; \quad i = 1, \dots, n. \quad (6.4)$$

As before there is a Liapunov function.

The Jacobian matrix at an equilibrium of the vector field defined by (6.4) is:

$$\text{diag}\{a_i\}[(1 - r)p'E - \text{diag}\{C_1'', \dots, C_n''\}].$$

Gershgorin's Theorem implies stability of equilibria, and hence global convergence to a unique equilibrium, in case

$$\max_i C_i'' < (1 - r)np' \quad (6.5)$$

at every equilibrium. Evidently this can be achieved by a tax rate r which is sufficiently confiscatory (close to 1).

By allowing the positive function a_i in (6.3) to depend on time we sacrifice autonomy of the system of ODEs, but we allow the interpretation that producers can make different production choices at different times, even under identical market conditions. This leads to the slightly nonautonomous system

$$\frac{dx_i}{dt} = a_i(x, t)[p(x_1 + \cdots + x_n) - C_i'(x_i)]; \quad i = 1, \dots, n. \quad (6.6)$$

The nonautonomy is not serious because one can show that the same function V is Liapunov for (6.4), in the sense that it decreases along nonconstant solutions. Under the further assumption that every point x has a neighborhood U such that

$$\inf\{a_i(x, t) : x \in U, t \geq 0\} > 0,$$

one can prove that every solution tends toward a critical point of V .

7 Monotone dynamics

The *vector order* in $\mathbf{R}^n$ is the partial order relation defined by:

$$x \leq y \Leftrightarrow x_i \leq y_i, \quad i = 1, \dots, n.$$

A dynamical system is called *monotone* if it preserves this order for positive time. If the dynamic is Φ , this means that

$$t \geq 0, x \leq y \Rightarrow \Phi_t(x) \leq \Phi_t(y).$$

Such systems have very special dynamical properties, yet they arise in many situations. Before giving examples in the following sections, we state some results about them.

A C^1 vector field F in $\mathbf{R}^n$ is *cooperative* (or *quasimonotone*) if $\partial F_i / \partial x_j \geq 0$ whenever $i \neq j$.

Theorem 7.1 (MÜLLER-KAMKE) *The solution flow of a cooperative vector field is monotone.*

Proof See the book by Coppel [8]. □

An $n \times n$ matrix M is *irreducible* if for every partition of the set $\{1, \dots, n\}$ into two disjoint nonempty subsets S, T , for any $i \in S$ there exists $j \in T$ with $M_{ij} \neq 0$.

A vector field is called *irreducible* if its Jacobian matrices are irreducible.

The following result shows that cooperative dynamical systems have very special dynamics, namely, generic orbits are quasiconvergent:

Theorem 7.2 *Consider the dynamics of a cooperative vector field F in $\mathbf{R}^n$.*

- (a) *Every attractor contains an orbitally stable equilibrium. In particular there are no attracting periodic orbits.*
- (b) *Suppose F is irreducible. Then there is a dense invariant set X whose complement has measure zero, such that every bounded forward trajectory in X is convergent. If the set of equilibria is countable (or finite), then every bounded trajectory in X is convergent.*
- (c) *Suppose $n = 2$. Then every bounded forward or backward trajectory converges.*

Proof Part (a) can be derived from Hirsch in [13]. Part (c) is proved by Hadeler and Glas in [11]; see also Hirsch [14]. Versions of part (b) have been proved by several authors: Smith and Thieme [27], Poláčik [23], Takáč [20]. See also Hirsch [15]. □

There are examples satisfying the hypotheses of Theorem 7.2 that have arbitrarily complex dynamics in invariant subsets — but these sets have measure zero, and cannot be attractors. Such examples show that we cannot expect general cooperative systems to have Liapunov functions.

A vector field F in $\mathbf{R}^n$ is called *competitive* if it is C^1 and the off-diagonal terms of Jacobian matrices are ≤ 0 ; this is equivalent to $-F$ being cooperative. S. Smale in [26] showed that an arbitrary system of ODEs in $\mathbf{R}^{n-1}$ can be embedded as an attractor in a competitive system in $\mathbf{R}^n$, so this is not a very powerful assumption from a general dynamical point of view. For two dimensions, however, we have the following result:

Theorem 7.3 *Let $x(t)$, $0 < t < \infty$ be a bounded trajectory of a competitive vector field in $\mathbf{R}^2$. Then there is an equilibrium e such that $x(t) \rightarrow e$, with each component converging monotonically (nonincreasing or nondecreasing).*

In fact Theorem 7.3 and Theorem 7.2(c) are valid for slightly nonautonomous systems of the form:

$$\frac{dx_i}{dt} = a_i(x, t)G_i(x), \quad i = 1, 2 \quad (7.1)$$

where the functions a_i are positive, and the vector field G is competitive (for Theorem 7.3) or cooperative (for Theorem 7.2(c)).

We turn now to monotone discrete time systems. It is not hard to show that a C^1 map F from a convex subset of $\mathbf{R}^n$ to $\mathbf{R}^n$ is monotone provided all its partial derivatives are nonnegative. The following result is an analog of Theorem 7.2(a); it can be derived from Hirsch [12], [13]:

Theorem 7.4 *Let $X \subset \mathbf{R}^n$ denote either $\mathbf{R}^n$, an orthant of $\mathbf{R}^n$, or an n -cube. Let $F : X \rightarrow X$ be a continuous monotone map which defines a dissipative dynamical system. Then:*

- (a) F has an orbitally stable fixed point.
- (b) Every attractor contains an orbitally stable periodic orbit.

This shows that an attractor A in such a system cannot be very chaotic: If A is not a finite set, no orbit is dense in A , nor can a union of periodic orbits be dense in A — properties often used to define chaos.

8 A cooperative model of competition

Here we consider another of X. Vives' oligopoly games, this time "Bertrand competition in differentiated markets".

Let there be n producers $\mathcal{A}_i$ producing products which are gross substitutes. $\mathcal{A}_i$ continually adjusts the price p_i of her product. The demand for the i 'th product is a positive function $h_i(p)$ of the price vector $p = (p_i)$. $\mathcal{A}_i$ produces exactly enough to meet the demand. The production cost of the i 'th good is $C_i(h_i(p))$. The profit to $\mathcal{A}_i$ is $\Pi_i = p_i h_i - C_i$.

Each agent acts so as to increase profits. We interpret this to mean that there are positive functions $a_i(p_i)$ such that

$$dp_i/dt = a_i(p_i)\partial\Pi_i/\partial p_i.$$

This time we assume that our oligopolists understand perfectly the effect of price changes on demands and costs. Computing $\partial\Pi_i/\partial p_i$ yields the following dynamical system, whose state space is the space of price vectors p :

$$\frac{dp_i}{dt} = a_i(p_i)[\{p_i - C_i'(h_i(p))\}\frac{\partial h_i(p)}{\partial p_i} + h_i(p)] := F_i(p). \quad (8.1)$$

We make the natural assumptions $\partial h_i/\partial p_i < 0$ and $C_i' > 0$.

The following additional mathematical assumptions are made to guarantee that vector field F is cooperative; we attempt to interpret them economically. First of all, as the products are gross substitutes we have:

- (i) $\frac{\partial h_i}{\partial p_j} > 0$ for $i \neq j$.

We postulate what has been called "submodularity" in a related context ([28], [29]):

- (ii) Suppose all prices are held fixed except for p_i . Let $j \neq i$. At a higher level of p_j , the marginal demand for good i with respect to p_i greater than at a lower level of p_j :

$$\frac{\partial}{\partial p_j}\left(\frac{\partial h_i}{\partial p_i}\right) \geq 0.$$

We also need to assume:

- (iii)

$$C_i''\frac{\partial h_i}{\partial p_i} < 1.$$

To interpret (iii), suppose all prices, demands and costs are fixed except those for good i . Ignoring the other goods and recalling that $\partial h_i/\partial p_i < 0$, we see that (iii) is equivalent to

$$\frac{dC_i'}{dh_i} > \frac{dp_i}{dh_i}.$$

If we identify h_i with the production and hence the supply of good i , this says that as supply increases, marginal cost does not fall as quickly as price. Note that (iii) holds in the special case of linear costs.

Finally we make the drastic hypothesis that producers choose prices which at least equal marginal costs:

$$(iv) \quad p_i - C_i' \geq 0.$$

This holds, for example, if there are only fixed costs ($C_i' = 0$).

Lemma 8.1 *Under assumptions (i)–(iv), System (8.1) is cooperative and irreducible.*

Proof A computation yields, for $i \neq j$:

$$\frac{\partial F_i}{\partial p_j} = a_i \left[(p_i - C_i') \frac{\partial^2 h_i}{\partial h_j \partial h_i} + \left(1 - \frac{\partial h_i}{\partial p_i} C_i'' \right) \frac{\partial h_i}{\partial p_j} \right]. \quad (8.2)$$

Coefficient a_i is positive, the first summand is nonnegative by (iv) and (ii), and the second is positive by (iii) and (i). Therefore Jacobian matrices of F have positive off-diagonal terms. $\square$

Applying Theorem 7.2 we obtain:

Theorem 8.2 *Under assumptions (i)–(iv), there is a dense invariant set X having full measure, such that every bounded forward trajectory approaches the set of equilibria.*

9 The simple model revisited

We consider again the mapping in Section 2

$$F : \mathbf{R}_+^2 \rightarrow \mathbf{R}_+^2,$$

having the form

$$F(x_1, x_2) = (f_1(x_2), f_2(x_1))$$

with

$$f_i' < 0, \quad i = 1, 2.$$

This system is dissipative, but not monotone. It can be made monotone, however, by the simple coordinate change

$$\begin{aligned} y_1 &= x_1, \\ y_2 &= -x_2. \end{aligned}$$

In these coordinates the dynamic is given by the map

$$G : X \rightarrow X, \quad X = \{y \in \mathbf{R}^2 : y_1 \geq 0, y_2 \leq 0\}$$

where

$$\begin{aligned} G(y_1, y_2) &= (g_1(y_2), g_2(y_1)), \\ g_1(y_2) &= f_1(-y_2), \\ g_2(y_1) &= -f_2(y_1). \end{aligned}$$

Then $g_i' > 0$, implying that G is monotone.

Now we apply Theorem 7.4 to G , to conclude that G has an orbitally stable fixed point. Reverting to the original system, we have:

Theorem 9.1 *The system F has an orbitally stable fixed point.*

Under the (hopefully) generic assumption of simple fixed points, there is a stable fixed point.

10 Stochastic allocation processes and deterministic dynamics

In this section and the next we consider discrete-time systems where the current state of the system determines not the next state, but rather the *probability distribution* of the next state. We suppose further that as time goes on, the probability of any particular state change goes to zero.

Example 10.1 (Allocation Processes) Paraphrasing B. Arthur [1]: Suppose a unit addition or allocation is made to one of d categories at each time, with probabilities that are a function of the proportion of units currently in the categories. Time here is event time, not clock time. In practice we might be considering the build-up of market shares by observing the "allocation" of adopters, one at a time, to d technologies; or consumers to d product brands; or, in regional economics, firms to d locations. Thus, the next unit is added to category i with probability $g_i(x)$ where $x = (x_1, \dots, x_d)$ is the vector of current proportions or market shares. Here g is a deterministic, exogenously given map of the $d - 1$ simplex Δ^{d-1} into itself:

$$g : \Delta^{d-1} \rightarrow \Delta^{d-1},$$

$$\Delta^{d-1} = \left\{ x = (x_1, \dots, x_d) \in \mathbf{R}^d : 0 \leq x_i \leq 1, \sum_{i=1}^d x_i = 1. \right\}$$

We start the process at time $k = 0$ with an initial vector $y = (y_1, \dots, y_d)$ of allocations, and initial total allocation $w = \sum y_i > 0$. The initial vector $x_0 \in \Delta^{d-1}$ has components $x_0(i) = y_i/w$. At time $k \in \mathbf{N}$ we have allocated $w + (k - 1)$ units. If $Y_k \in \mathbf{R}^d$ denotes the vector of allocations, then the vector of proportions is $x_k = Y_k/(w + k - 1)$. We select an integer $j_{k+1} \in \{1, \dots, d\}$ at random, the probability that $j_{k+1} = i$ being $g_i(x(k))$. We then increase the allocation to category $j(k)$ by one unit. Let e_{k+1} denote the random unit vector whose i 'th coordinate is 1 if $i = j_{k+1}$ and 0 otherwise. Then

$$x_{k+1} = x_k + \frac{1}{w + k + 1}(e_{k+1} - x_k) \quad (10.1)$$

We rewrite this as

$$x_{k+1} = x_k + \frac{1}{w + k + 1}(g(x_k) - x_k) + \frac{1}{w + k + 1}(e_{k+1} - x_k). \quad (10.2)$$

Now it is easy to see that the conditional expectation of the random variable $U_{k+1} = e_{k+1} - x_k$, given the current state x_k , is 0. We emphasize this by rewriting Equation (10.2) as

$$x_{k+1} = x_k + \frac{1}{w + k + 1}(g(x_k) - x_k + U_{k+1}), \quad (10.3)$$

with

$$\mathbf{E}(U_{k+1}|x_k) = 0. \quad (10.4)$$

From this it is easy to calculate the expected value of x_{k+1} given the current state:

$$\mathbf{E}(x_{k+1}|x_k) = \frac{1}{w + k + 1}(g(x_k) - x_k). \quad (10.5)$$

There is a strong resemblance between Equation (10.3) and the Cauchy-Euler approximation to the solution of the differential equation¹⁵

$$\frac{dx}{dt} = g(x) - x; \quad (10.6)$$

with variable stepsize $\frac{1}{w+k+1}$:

$$x_{k+1} = x_k + \frac{1}{w+k+1}(g(x_k) - x_k). \quad (10.7)$$

Thus Equation (10.3) can be viewed as the result of a noisy numerical solution to Equation (10.6).

It is thus reasonable to expect a close relationship between the limit $L\{x_k\}$ of a sample path $\{x_k\}$ of Equation (10.3) (i.e., the set of limit points of subsequences) and limit sets of solutions of Equation (10.6). Such a relationship has long been known in case the dynamics of Equation (10.3) are very simple, for example of gradient type, or globally convergent (see the books by Benveniste et al. [5] and Kushner and Clark [18]).

Recently a very general result has been proved by M. Benaïm [2]. To state this result we need the following dynamical concept. A set $L \subset \Delta^{d-1}$ is called *internally chain transitive* for Equation (10.3) if it has the following properties:

- (i) L is nonempty, compact and connected,
- (ii) L is invariant: it contains the complete solution curve through every point of L ,
- (iii) $\forall \epsilon > 0, T > 0$ we can pass from any point of L to any other point by following orbits in L and allowing, at times separated by at least T , a finite number of jumps of size $< \epsilon$ to other orbits in L .

An equivalent definition is that the flow defined in L defined by following solution curves to Equation (10.3) has no proper attracting or repelling sets.

It can be shown that if F has a Liapunov function, and only finitely many equilibria, then the only internally chain transitive sets are equilibria. Likewise for a vector field in the plane with negative divergence almost everywhere. In general, however, internally chain transitive sets can be very complicated.

The following is a special case of the main result of Benaïm [2]:

Theorem 10.2 *Assume the map g is Lipschitz (e.g., continuously differentiable). Then with probability one, the limit set $L\{x_k\}$ of a sample path $\{x_k\}$ of process Equation (10.3) is internally chain transitive for Equation (10.6).*

In favorable circumstances we can therefore derive information about the allocation process (10.3) from the dynamics of the vector field F in Δ^{d-1} representing the right-hand side of Equation (10.6). This idea was exploited by Benaïm and Hirsch in [3]. A simple example:

¹⁵Here we regard the map $x \mapsto g(x) - x$ as a vector field $f : \Delta^{d-1} \rightarrow T\Delta^{d-1}$ on Δ^{d-1} with values in the *tangent space* to Δ^{d-1} , namely the linear subspace of $\mathbf{R}^d$ defined by

$$T\Delta^{d-1} = \left\{ x \in \mathbf{R}^d : \sum x_i = 0 \right\}.$$

If we identify Δ^{d-1} with a simplex $\sigma^{d-1} \subset \mathbf{R}^{d-1}$ by a convenient affine homeomorphism, then f is a vector field on σ^{d-1} .

Corollary 10.3 *Assume that system (10.6) has a Liapunov function, and that every trajectory converges to one of a finite set of equilibria. Then with probability one, a sample path of the allocation process (10.3) converges to one of these equilibria.*

11 Urn processes

Stochastic processes such as Equation (10.3) are examples of *generalized Polya urn processes*. In this interpretation, an urn initially contains w balls of $d \geq 2$ different colors. At stage k we add one new ball whose color is chosen as follows. Let $x_k \in \Delta^{d-1}$ be the vector whose i 'th coordinate is the proportion of balls of color i currently in the urn. Then the probability that we add a ball of color i is $g_i(x_k)$. It is not hard to see that Equations (10.3), (10.4) and (10.5) describe the evolution of the process $\{x_k\}$. Thus Theorem 10.2 applies to these urn processes.

Much more can be said about the relation between sample paths and dynamics. One can ask, for example, whether a given equilibrium is likely to be the limit set of a sample path, or more generally, whether a given internally chain transitive set has positive probability of containing the limit set. The following theorem states some results proved in [21] and [3]:

Theorem 11.1 *Assume that g in Equation (10.6) is C^2 .*

(a) *If w is a linearly stable equilibrium then*

$$P\{\lim z_k = w\} > 0;$$

(b) *If w is a linearly unstable equilibrium then*

$$P\{\lim z_k = w\} = 0.$$

(c) *If Γ is a linearly unstable periodic orbit then*

$$P\{L\{x_k\} \cap \Gamma \neq \emptyset\} = 0.$$

(d) *Assume g maps Δ^{d-1} into its interior (g need only be Lipschitz now.) If Λ is an attractor then*

$$P\{L\{x_k \subset \Lambda\} > 0.$$

Suppose there is no reason to expect orbits of Equation (10.6) to converge; the system may be thought to be oscillatory or chaotic, for example. We may then inquire whether sample paths are likely to closely track trajectories. The following result, a consequence of [2], gives conditions under which this is true.

In order to state this result we introduce an numerical invariant called the *logarithmic expansion rate* $l_{\exp}(F)$ of the vector field F on Δ^{d-1} , which is hard to compute but often easy to estimate. Let $\{\Phi_t\}_{t \in \mathbf{R}}$ denote the flow generated by Equation (10.3). Then

$$l_{\exp}(F) = \liminf_{t \rightarrow \infty} \min_{x \in \Delta^{d-1}} \frac{1}{t} \|(D\Phi_t(x))^{-1}\|^{-1}.$$

If $l_{\exp}(F) > -c$ for some $c \geq 0$ then it can be shown that for all sufficiently large $t > 0$, the map Φ_t has the following property. For sufficiently small $r > 0$,

$$\Phi_t B(x, r) \subset B(\Phi_t x, e^{-ct} r)$$

where $B(x, r)$ denotes the set of points in Δ^{d-1} at distance $\leq r$ from $x \in \Delta^{d-1}$.

In [16] I gave several estimates for $l_{\exp}(F)$, including the following:

(a)

$$l_{\exp}(F) \geq \inf_x \beta(x), \quad (11.1)$$

where $\beta(x)$ denotes the smallest eigenvalue of the matrix

$$\frac{1}{2}(DF(x) + DF(x))^T.$$

(b) If F is Morse-Smale and $\gamma \in \mathbf{R}$ is the smallest characteristic exponent of periodic and stationary orbits, then

$$l_{\exp}(F) \geq \gamma. \quad (11.2)$$

In the following consequence of a theorem of Benaïm, we assume given an urn process (10.3) and the corresponding dynamical system (10.6) with vector field $F(x) = g(x) - x$. Set

$$\tau_n = \tau(n) = \sum_1^n \frac{1}{w+n}.$$

It is easy to see that there exists $\beta > 0$ and a sequence $\epsilon_n \rightarrow 0$ such that

$$\tau_n = \log(w+n) + \beta + \epsilon_n. \quad (11.3)$$

Theorem 11.2 *Assume*

$$l_{\exp}(F) = -c > -\frac{1}{2}.$$

Then with probability one there is a random variable y in Δ^{d-1} with the following property:

$$\lim_{n \rightarrow \infty} \|x_n - \Phi(\tau_n, y)\| = 0. \quad (11.4)$$

This means that there is solution $u(t)$ to Equation (10.6) such that the sequence $\{u_{\tau(n)}\}$ is asymptotic with the sample path $\{x_n\}$.

The initial value $u(0)$ is the random variable y , and cannot be computed in most cases. Nevertheless this result tells us that the sample path closely follows *some* trajectory of Equation (10.6) — albeit with a slowing down of time (the difference between the sequences $\{n\}$ and $\{\tau_n\}$).

One can construct generalized Polya urns — or allocation processes — with novel long run behavior by exploiting Theorem 11.2. Suppose for example that H is a C^2 vector field on Δ^2 which points transversely inward on the boundary, has a linearly unstable repelling equilibrium b in the interior, and a unique nonstationary periodic orbit Γ which is linearly stable and which attracts all forward orbits except b . For sufficiently $\epsilon > 0$ the map $g : x \mapsto x + \epsilon H(x)$ sends Δ^2 into its interior. The urn process (10.3) with $g = g_\epsilon$ corresponds to the vector field $\epsilon h(x)$ via Equation (10.6). By (11.1) we choose $\epsilon > 0$ sufficiently small that

$$l_{\exp}(\epsilon H) > -\frac{1}{2}.$$

Theorem 11.3 *Let g_ϵ be as above and consider sample paths $\{x_n\}$ of the 3-color urn process (10.3) with $g = g_\epsilon$. Then with probability one there is a unique periodic trajectory $z : \mathbf{R} \rightarrow \Gamma$ of h_ϵ such that*

$$\lim_{n \rightarrow \infty} \|x_n - z(\log(w+n))\| = 0. \quad (11.5)$$

Proof It is well known, however, that any trajectory in the basin of a linearly stable periodic orbit Γ is asymptotic with a trajectory $z(t)$ in Γ (Hirsch and Smale [17]; see Hirsch [16] for generalizations). Therefore from Theorem 11.2 we conclude that almost surely there exists $y \in \Gamma$ satisfying Equation (11.4) for the process under consideration. The theorem now follows from Equation (11.3) by choosing the initial value of $z(t)$ to be $\Phi_\beta y$. $\square$

In a similar construction one can start from a vector field H in the 4-simplex for which all trajectories except for a linearly unstable equilibrium are attracted to a chaotic trajectory, and thereby construct an urn model for which sample paths almost surely track chaotic trajectories.

12 Repeated randomized games and deterministic dynamics

Consider a game played repeatedly by two players, each having 2 pure strategies, called here *actions*. At play k she chooses her pure strategy to maximize her expected payoff, under the natural but mistaken assumption that her opponent plays the mixed strategy corresponding to her past frequencies of pure strategies. To formalize this, let N_k be the number of times that the second player has played strategy 1 in the first k games. The *empirical frequency vector* of player 2 after play k is:

$$q_k = (y_k, 1 - y_k) \in \Delta^1, \quad y_k = \frac{1}{k} \sum_{j=1}^k N_j.$$

The empirical frequency vectors $p_k = (x_k, 1 - x_k)$ of player 1 are defined similarly. Let $q_0 = (y_0, 1 - y_0), p_0 = (x_0, 1 - x_0)$ be arbitrary vectors in Δ^1 .

Let player 1 have the 2×2 payoff matrix A , so that A_{lm} is her payoff if she and her opponent play $l, m \in \{1, 2\}$ respectively.

Player 1 assumes at play $k+1$ that her opponent will play the mixed strategy q_k , so she plays the action which maximizes her expected payoff under this assumption: At play $k+1$ she plays strategy $l \in \{1, 2\}$, where

$$l = \operatorname{argmax}_{m \in \{1, 2\}} (Aq_k)_m, \quad (12.1)$$

with any convenient rule for the cases where the argmax is not unique. Player 2, with payoff matrix B , behaves similarly. These behavior rules constitute what is called *fictitious play* (Brown [7], Robinson [25]).

Thus far we have a deterministic system. We now assume, however, that the payoff matrices change in a random way. The first player's payoff matrix for game k has the form $A_k = A + E_k$ where $\{E_k\}_{k \in \mathbb{N}}$ is a sequence of independent, identically distributed, matrix-valued random variables with common mean zero. We assume each entry $E_{k,lm}$ lies in a finite interval, and has a probability distribution given by a smooth (continuously differentiable) density function supported on that interval; and that different entries are independent random variables. The second player's payoff is given by a similar sequence of random matrices $B_k = B + F_k$.

We assume for each k : *at the beginning of play k each player knows her own payoff matrix, and her opponent's empirical frequency vector*. Again each player chooses the action which maximizes her expected payoff under the assumption that her opponent plays the mixed strategy given by her empirical frequency vector. Thus player 1 plays action $l_k \in \{1, 2\}$ given by

$$l_k = \operatorname{argmax}_{m \in \{1, 2\}} (A_k q_k)_m, \quad (12.2)$$

while player 2 plays strategy

$$m_k = \operatorname{argmax}_{m \in \{1,2\}} (B_k p_k)_m. \quad (12.3)$$

It can be shown that the assumptions on the E_k imply that, with probability one, (12.2) and (12.3) specify unique strategies.

The *state* of the game immediately after play $k \geq 1$ is defined to be the vector $z_k = (x_k, y_k) \in \Delta^1 \times \Delta^1$ specifying the empirical frequencies with which each player has chosen strategy 1; the initial state is the arbitrary vector $(x_0, y_0) \in \Delta^1 \times \Delta^1$. Because of randomness in the payoff matrices, the sequence $\{z_k\}_{k \in \mathbf{N}}$ is a stochastic process.

We now calculate the conditional expected value of z_{k+1} given z_k . We do this for the first component x_{k+1} , the second component being similar. Fix k and set $E_{k,ij} = E_{ij}$. Let P_k denote conditional probability given z_k . Then from the definition of x_k we have

$$(k+1)x_{k+1} = kx_k + X_{k+1} \quad (12.4)$$

where $X_{k+1} \in \{0, 1\}$ is a random variable such that:

$$\begin{aligned} P_k\{X_{k+1} = 1\} &= P_k\{(A_k q_k)_1 > (A_k q_k)_2\} \\ &= P_k\{E_{11}y_k + E_{12}(1 - y_k) > E_{21}y_k + E_{22}(1 - y_k)\} \\ &= P_k\{y_k(E_{11} - E_{12} - E_{21} + E_{22}) > E_{22} - E_{12}\}, \\ P_k\{X_{k+1} = 1\} &= g(y_k), \end{aligned} \quad (12.5)$$

where g is a smooth bounded function which can be computed from the probability densities of the E_{ij} . Therefore from Equation (12.4) we derive:

$$E(x_{k+1}|z_k) = x_k + \frac{1}{k+1}(-x_k + g(y_k)). \quad (12.6)$$

Now define a random variable V_{k+1} by the equation

$$\frac{V_{k+1}}{k+1} = x_{k+1} - E(x_{k+1}|z_k); \quad (12.7)$$

then V_{k+1} is bounded (independently of k), and has mean zero conditional on z_k . From (12.6) and (12.7) we obtain

$$x_{k+1} = x_k + \frac{1}{k+1}(-x_k + V_{k+1});$$

an analogous formula relates y_{k+1} and y_k . Finally we obtain the key recursion:

$$z_{k+1} = z_k + \frac{1}{k+1}(-z_k + h(z_k) + U_{k+1}) \quad (12.8)$$

where the smooth map $h : [0, 1] \times [0, 1] \rightarrow \mathbf{R} \times \mathbf{R}$ has the form

$$h(x, y) = (g_1(y), g_2(x)), \quad (12.9)$$

and $\{U_{k+1}\}_{k \in \mathbf{N}}$ is a bounded sequence of mean-zero random variables with values in $\mathbf{R}^2$, such that

$$E(U_{k+1}|z_k) = 0;$$

in technical jargon they are martingale differences.

The theorem of Benaïm is applicable here: *almost surely the limit set $L\{z_k\}$ of a sample path $\{z_k\}$ is an internally chain transitive set for the flow of the differential*

equation:

$$\frac{dx}{dt} = -x + g_1(y) \quad (12.10)$$

$$\frac{dy}{dt} = -y + g_2(x). \quad (12.11)$$

What information can we get from this? There are two steps: First, investigate internally chain transitive sets for this system; second, interpret them in terms of the stochastic process.

The first thing to observe is that the vector field

$$F(x, y) = (-x + g_1(y), -y + g_2(x)) \quad (12.12)$$

representing the right-hand side of Equation (12.12) is negative. This implies, by standard theory, that the solution flow (defined by following solution curves) decreases areas. This leads to the conclusion that *every internally chain transitive set is a nonempty compact connected set of equilibria*. An equilibrium means a fixed point of the flow, or equivalently, a vector (x, y) such that $F(x, y) = (0, 0)$. From the formula for $F(x, y)$ we see that the set of equilibria is just the intersection of the graphs $x = g_1(y)$ and $y = g_2(x)$. The only compact connected nonempty subsets of the graph of a smooth function are points and smooth arcs. Thus we conclude: *with probability one, sample paths limit at either a point or a smooth arc*.

Since generically, in some sense, two smooth graphs intersect at isolated points, we are strongly tempted to conclude: "generically, almost surely the sequence of empirical frequencies of each player converges". To justify this rigorously, however, would require showing that this conclusion holds for a "generic" set of density functions for the payoff matrices. This is an interesting question which is probably not too difficult to decide one way or the other. Part of the problem is to specify the precise meaning of "generic".

Regardless of genericity, we can conclude:

Theorem 12.1 *If the vector field F has only isolated zeroes, then almost surely the sequences of empirical frequencies converge.*

By a finer analysis (Benaïm and Hirsch [4]) one can prove:

Theorem 12.2 *Assume that the functions g_1 and g_2 are strictly positive and have continuous second derivatives, and that at every equilibrium the Jacobian determinant of F is nonzero. Then almost surely the following hold:*

- (a) *The sequence $\{z_k\}$ of empirical frequency vectors converge to an equilibrium of F ;*
- (b) *if w is a linearly stable equilibrium then*

$$P\{\lim z_k = w\} > 0;$$

- (c) *if w is a linearly unstable equilibrium then*

$$P\{\lim z_k = w\} = 0.$$

Suppose every equilibrium of F is either linearly stable or linearly unstable — this is generic for vector fields, and probably generic for the probability distributions under consideration. Then, as the set $\mathcal{S}$ of linearly stable equilibria is finite, it follows that with probability one, sample paths converge to an element of the finite set $\mathcal{S}$.

The following result, left as an exercise, is also of interest:

Proposition 12.3 Assume (in addition to the earlier assumptions):

- (a) the probability densities of the entries in the random matrices E_k and F_k are real analytic in the interiors of their support intervals,
- (b) There exists $0 < s < 1$ such that $f(g(s)) \neq s$.

Then the equilibrium set $\mathcal{E}$ of F is finite, and with probability one, sample paths converge to elements of $\mathcal{E}$.

We can give a game-theoretic interpretation to an equilibrium (u_1, u_2) of the game vector field. Denote by U_i the random matrix variable which has the same distribution as the payoff matrix for player $i \in \{1, 2\}$ for any game $k \in \mathbb{N}$ (these have identical distributions for all k). Following [9] we define a *Nash distribution equilibrium* to be a vector $u = (u_1, u_2) \in \Delta^1 \times \Delta^1$ with the following property for both choices of $i, j \in \{1, 2\}$, $i \neq j$: If $l_j = l_j(U_j, u_i) \in \{1, 2\}$ is the action which maximizes the expected payoff to player j , assuming that player i plays mixed strategy $(u_i, 1 - u_i)$, then it is required that u_j be the expected value of the random variable l_j .

From the definition (12.12), (12.5) of the game vector field F there follows:

Theorem 12.4 $u \in \Delta^1 \times \Delta^1$ is an equilibrium of the game vector field if and only if u is a Nash distribution equilibrium.

Applications to Prisoners Dilemma and other games are given in Benaïm and Hirsch [4].

References

- [1] Arthur, B. [1988], *Self-reinforcing mechanisms in economics*, in *The Economy as an Evolving Complex System* (P. W. Anderson, K. J. Arrow, and D. Pines eds.), Addison-Wesley, pp. 10–31.
- [2] Benaïm, M. [1996], *A dynamical system approach to stochastic approximation*, *Siam Journal of Optimisation and Control* **34**, 437–472.
- [3] Benaïm, M. and Hirsch, M. W. [1993], *Dynamics of morse-smale urn processes*, Mathematics Department, U. California at Berkeley, Technical Report. *Ergodic Theory and Dynamical Systems* **15**, 1005–1030.
- [4] Benaïm, M. and Hirsch, M. W. [1994], *Learning processes, mixed equilibria and dynamical systems arising from repeated games*, Mathematics Department, University of California at Berkeley, preprint. Submitted.
- [5] Benveniste, A., Métivier, M. and Priouret, P. [1990], *Stochastic Approximations and Adaptive Algorithms*, Springer, Berlin, Heidelberg, New York. Translated from: *Algorithmes adaptatifs et approximations stochastiques*, Masson, Paris 1987.
- [6] Birkhoff, G. D. [1920], *Recent advances in dynamics*, *Science* **51**, 51–55.
- [7] Brown, G. W. [1951], *Iterative solutions of games by fictitious play*, in *Activity Analysis of Production and Allocation* (T. C. Koopmans, ed.), Wiley, New York.
- [8] Coppel, W. A. [1965], *Stability and asymptotic behavior of differential equations*, Heath, Boston.
- [9] Fudenberg, D. and Kreps, K. [1993], *Learning mixed equilibria*, *Games and Economic Behavior* **5**, 320–367.
- [10] Gori, F., Geronazzo, L. and Galeotti, M. [1993], (editors), *Nonlinear Dynamics in Economics and Social Sciences*, Springer, Berlin, Heidelberg, New York. Second Informal Workshop at Certosa di Pontignano, Siena, Italy, May 27–30, 1991.
- [11] Hadeler, K. P. and Glas, D. [1983], *Quasimonotone systems and convergence to equilibrium in a population genetic model*, *J. Math. Anal. Appl.* **95**, 297–303.
- [12] Hirsch, M. W. [1995], *Fixed points of monotone maps*, *J. Diff. Equations*, 171–179.
- [13] Hirsch, M. W. [1985], *Attractors for discrete-time monotone dynamical systems in strongly ordered spaces*, in *Geometry and Topology* (J. Alexander and J. Harer, eds.), *Lecture Notes in Mathematics* **1167**, pages 141–153, New York. Proceedings of Special Year at University of Maryland, Springer-Verlag.

- [14] Hirsch, M. W. [1985], *Systems of differential equations that are competitive or cooperative II: Convergence almost everywhere*, SIAM J. Math. Anal. **16**, 432–439.
- [15] Hirsch, M. W. [1988], *Stability and convergence in strongly monotone dynamical systems*, J. reine angewandte Math. **383**, 1–53.
- [16] Hirsch, M. W. [1993], *Asymptotic phase, shadowing and reaction diffusion systems*, in Control Theory, Dynamical Systems and Geometry of Dynamics (K. D. Elworthy, W. N. Everitt and E. B. Lee, eds.), Chapter: Asymptotic phase, shadowing and reaction diffusion systems, pages 87–99. Marcel Dekker, New York. Proceedings of a festschrift in honor of Lawrence Markus.
- [17] Hirsch, M. W. and Smale, S. [1974], *Differential Equations, Dynamical Systems, and Linear Algebra*, Academic Press, New York.
- [18] Kushner, H. J. and Clark, D. S. [1978], *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer, Berlin, Heidelberg, New York.
- [19] Noble, B. and Daniel, J. [1988], *Applied Linear Algebra*, Prentice Hall, Englewood Cliffs, NJ.
- [20] Takáč, P. [1991], *Domains of attraction of generic ω -limit sets for strongly monotone semiflows*, Zeits. Analysis **3**, 275–317.
- [21] Pemantle, R. [1990], *Nonconvergence to unstable points in urn models and stochastic approximations*, Annals of Probability **18**, 698–712.
- [22] Poincaré, H. [1892, 1893, 1899], *Les Méthodes Nouvelles de la Mécanique Céleste*, volume 1,2,3. Gauthiers-Villars, Paris.
- [23] Poláčik, P. [1989], *Convergence in smooth strongly monotone flows defined by semilinear parabolic equations*, J. Diff. Equations **79**, 89–110.
- [24] Pugh, C. C. and Robinson, C. [1983], *The C^1 closing lemma*, Ergodic Th. Dyn. Sys. **3**, 261–313.
- [25] Robinson, J. [1951], *An iterative method for solving a game*, Annals of Mathematics **54**, 296–301.
- [26] Smale, S. [1976], *On the differential equations of species in competition*, J. Math. Biology **3**, 5–7.
- [27] Smith, H. L. and Thieme, H. [1991], *Convergence for strongly order-preserving semiflows*, SIAM J. Math. Anal. **22**, 1081–1101.
- [28] Topkis, D. M. [1979], *Equilibrium points in non-zero sum n -person submodular games*, SIAM J. Control and Optimization **17**(6), 773–787.
- [29] Vives, X. [1990], *Nash equilibrium with strategic complementarities*, J. Math. Econ. **19**, 305–321.
- [30] Zhang, W.-B. [1991], *Synergetic Economics*, Springer-Verlag, Berlin and Heidelberg.