Landscape Dynamics

Daniel Friedman and Ralph Abraham

April, 2004

1 Introduction

Advances in mathematics in the 1960s made available a host of new modeling strategies for all the sciences. The framework of global analysis applied to dynamical systems, calculus of variations, partial differential equations of evolution type, game theory, and so on, brought us catastrophe theory, chaos theory, complexity and simplexity, neural network theory, evolutionary game theory, and others. The computer and computer graphic revolutions brought new possibilities of computational modeling, simulation, and scientific visualization. Of all the sciences, those with the greatest potential to benefit from these new methods are the social, behavioral, and economic sciences.

Our proposed research will extend the class of models called *evolutionary games* and open it to new applications in the social sciences. Evolutionary game models analyze strategic interaction over time. Equilibrium emerges, or fails to emerge, as players adjust their strategies in response to the payoffs they earn. Thus far the models have mainly considered situations in which players chose among only a few discrete strategies. Our extensions allow players to choose within a continuous strategy space A.

In this setup, the *current state* is the distribution of all players' choices over A. In any particular application, the current state defines a payoff function on A, whose graph is called the *adaptive landscape*. Players respond to the landscape in continuous time by adjusting their strategies towards higher payoff. Hence the current state (the distribution of chosen strategies) changes, and this in turn alters the landscape. The interplay between the evolving state and the landscape gives rise to nontrivial dynamics. In particular, when players follow the gradient (steepest ascent in the adaptive landscape), the evolving state can be characterized as the solution to a nonlinear partial differential equation, or equivalently, a dynamical system on an infinite-dimensional space.

Our research is intended to advance the arts of mathematical modeling, computer simulation, and scientific visualization of complex dynamical systems encountered in the social, behavioral, and economic sciences. Many of these complex systems involve geometrical spaces of strategies that have not yet been adequately treated.

Our specific motivation for geometric (continuous or lattice) strategy spaces is to model the very common situation in which players actually have an n-dimensional continuum of strategies. The application that we propose to develop in detail is that of financial markets in which traders adjust portfolios consisting of n financial assets. The price and yield of each asset depends on traders' choices, and traders' payoffs depend on the absolute and relative performance of their portfolios. In a later section we sketch how this approach can give insight into financial market bubbles and crashes, while remaining generally consistent with modern finance theory. We will also consider a number of other social science applications, such as political competition for votes, and products positioning by competing firms.

Our primary **objectives** are:

- to advance a mathematical theory, based on the concepts of global analysis, of a new class of evolutionary game models, suitable for massively complex systems, and related to agent-based models;
- to develop a software suite for fitting such a model to data from natural systems, for simulating the model, and for visualizing the natural and simulated data; and
- to fully develop one application for financial markets, and to sketch several other applications in the social sciences.

The **expected significance** is primarily the development and dissemination of new mathematical methods for understanding human social behavior. As explained in the next section, the project draws on insights and techniques from several separate mathematical traditions, unifies and extends the ideas, and creates a mathematical toolkit that can tackle a wide variety of new applications in the economic, behavioral and social sciences.

The **broader impact** will be be twofold. First, current communities of applied mathematicians, economists, finance practitioners, and other social scientists will be exposed to a new set of ideas and techniques that will extend their reach. Second, current and future graduate students, and eventually undergraduate students, will have access to an appealing set of tools and ideas. The proposed toolkit includes computer based models and computer-graphic visualization with tremendous pedagogic value. New methods may attract a new generation of social scientists, expanding interdisciplinary frontiers in new directions.

2 Present State of Knowledge

Nonlinear dynamics and complex behaviours in simple systems and in networks have enjoyed increasing interest in the past two decades. After the pioneering works on cybernetics, catastrophes, chaos, fractals, and neural networks since the Second World War, a wave of new methods for mathematical modeling, computer simulation, and computer-graphic visualization have emerged.

2.1 Complex systems

Complexity, synchronization, self-organization and criticality, analog neural networks, genetic algorithms, games, and agent-based models – all have their special-interest groups, societies, journals, and so on. Our proposed research draws on modeling strategies from this burgeoning interdisciplinary frontier. In particular, our nearest neighbors are games, complex systems, and neural networks.

These models are all special cases of *complex dynamical systems*. Formally, a complex dynamical system comprises a directed graph of simple dynamical systems, adaptively linked with coupling functions. As such, they are actually large-scale dynamical systems with control parameters, so nonlinear dynamical systems theory (chaos theory) is their parent branch of pure mathematics. For this reason, our analytic machinery comes primarily from chaos theory, and its parent in turn, global analysis (Abraham, 1996; www.santafe.edu).

Agent-based models and neural networks are structures of this sort, while games (especially evolutionary games) have a somewhat different personality. Our research methods thus draw on analytical techniques of global analysis – such as bifurcation theory, fractal dimension, attractor reconstruction, Lyapounov exponents, etc – as well as computational methods developed particularly by the dynamical systems community – such as Stella, Vennsim, Madonna, and Phaser, as well as creations of the computer science community – such as Swarm, Starlogo, and Netlogo – and general purpose math tools – such as Maple, Mathematica, and Matlab. (www.mathworks.com)

2.2 Agent-based models

The structures that we will be developing in this project interpolate between evolutionary games and agent-based models, and thus we may make use of software environments created for agent-based modeling, such as Swarm and Netlogo. Netlogo is many advantages for our purpose. It is freeware; it is widely used in universities; it is easily converted to Java applets for distribution via the World Wide Web; and it accommodates experiments with human subjects and with students in the classroom. Thus it combines features for research and for teaching with graphical user interfaces and ease of distribution. (www.econ.iastate.edu/tesfatsi/ace.htm)

While we may do some of our experimental work with our own programs written in C or scripts for Matlab, we would certainly want to prepare applets in Netlogo for our website, conference presentations, and the like. It is interesting to note that the Netlogo home website features some interesting models of financial markets, proving the feasibility of Netlogo for our work. (ccl.sesp.northwestern.edu/netlogo)

2.3 Neural networks

A chief feature of our approach is *learning*. Players of evolutionary games adjust their strategies to improve their payoffs, and in the adaptive cases we will study, the learning algorithms are comparable to those of analog neural networks. The Hopfield network comes to mind as a close relative of our approach. In the more advanced part of our program we would explicate the connections between our modeling strategy and the learning algorithms of neural network theory (Haykin, 1999).

2.4 Games

Game theorists since Von Neumann and Morgenstern (1944) and Nash (1951) have studied strategic interactions among a set of players. Evolutionary game theory – first introduced by Maynard Smith and Price (1973), Taylor and Jonker (1978), Zeeman (1980), and Maynard Smith (1982) – focuses on adjustment dynamics using three basic principles.

- monotone: higher payoff strategies displace strategies with lower payoff ("survival of the fittest");
- inertial: the player population takes real time to change behavior ("evolution not revolution"); and
- game against nature (GAN): players don't try to influence other players' choices ("natural selection").

These principles are consistent with complex adaptive systems theory and with general dynamical systems theory, but are more specific. They mark the boundaries to neighboring branches of game theory. Traditional static game theory assumes no behavioral inertia and assumes complete rationality, a very strict version of monotonicity. Repeated game theory (e.g., Fudenberg and Tirole, 1991, Chapter 5) studies ongoing interactions where (a) your current choice affects other players' choices and hence your future opportunities, and (b) you take this into account in making your current choice. The game against Nature (GAN) principle holds when either (b) fails because you can't reliably assess the indirect future consequences of your current behavior, or (a) fails because no single individual has an appreciable effect on others, as exemplified in the standard price taking assumption in competitive markets.

Evolutionary games were originally aimed at biological applications (e.g., Hofbauer and Sigmund, 1988) but soon began to influence game theorists. Early work includes Binmore (1987), Fudenberg and Kreps (1988) and Friedman (1988, 1991). Evolutionary games enjoyed a vogue among theorists as witnessed by textbooks such as Weibull (1995) and Fudenberg and Levine (1998); monographs such as Cressman (1992), Vega-Redondo (1996), and Samuelson (1997); and special issues of *Games and Economic Behavior* in 1991 and 1993, and *Journal of Economic Theory* in 1992.

The economic and social sciences applications of evolutionary game theory now include isolated articles in international trade (Friedman and Fung, 1996), environmental policy (Dijkstra and De Vries, 2002; McGinty 2002), and several other fields, but the impact on applications clearly lags the impact on theory. Friedman (1998) argues that the evolutionary approach promises distinctive insights and implications in substantive economic applications, but that the theory requires focused development. In particular, the extension of evolutionary game theory from a finite, unordered set of strategies to a geometric space of strategies will empower a host of new applications to economics, and to all the social sciences.

There are already a few papers that treat evolutionary games with continuous strategy spaces. Friedman and Yellin (1997, 2000) lay the groundwork. Bomze (1990, 1991) and Oechssler and Riedl (2001, 2002) extend replicator dynamics to continuous strategy spaces, and Cressman and Hofbauer (2003) further develop the approach and connect it to recent work in theoretical biology. However, as explained in the next section, their approach has a different range of applicability than ours and cannot be interpreted in terms of adaptive landscapes.

3 Plan of Work

We now describe our new modeling strategies in three steps, from well-known to new evolutionary games: replicator dynamics, adaptive landscapes, and adaptive lattices.

3.1 Finite strategy sets and replicator dynamics

A basic evolutionary game concerns a single population of players, each with the same finite set of pure strategies. The instantaneous state of the system consists of the proportion of the population playing each of the pure strategies. The state evolves in continuous time, following some dynamical process that is monotone, inertial and GAN. In biological applications, the players are born, reproduce, and die. Each child has the same strategy as its parent. Meanwhile, players meet at random and play a fixed two-player game whose payoffs represent the players' fitness (number of offspring). An alternative narrative, more attuned to social science applications, is that the players learn and tend to switch to higher payoff strategies.

The mathematical description is as follows. Let k be an integer greater than one, and let $\Delta \subset R^k$ be the unit simplex, whose vertices $e^i = (0, ..., 0, 1, 0, ..., 0), i = 1, \ldots, k$ represent the pure strategies. The simplex will be the state space of our dynamical system. That is, a state $s = (s_1, \ldots, s_k)$ is a tuple of population shares, where each $s_i \geq 0$, and $\sum s_i = 1$. A player choosing strategy i receives payoff $f(e^i, s)$ when the state is $s \in \Delta$. The dynamics are then given by a vectorfield on Δ that depends on the payoff function f. The standard biological example is *replicator dynamics*, given by $\dot{s_i} = (f(e^i, s) - f(s, s))s_i$. That is, the growth rate for each population share $\dot{s_i}/s_i$ is its fitness $f^i(s)$ relative to the population average $f(s, s) = \sum s_j f(e^j, s)$.

More general dynamics are helpful in social science applications. For example, sign-preserving dynamics simply require that the share growth rates have the same sign as the relative fitnesses $f(e^i, s) - f(s, s)$, and are consistent with the principles of monotonicity and inertia.

In most of the literature it is assumed that the fitness function $f(x, s)$ is linear in both arguments. This is consistent with the biological interpretation that players meet at random to play a fixed two-player game, but is too restrictive for some biological applications and the more interesting social science applications. Here we use "playing the field" models with payoff functions f that are non-linear (but smooth) in the state variable s.

It is conceptually straightforward to extend to multiple populations. For example, in biology one might have separate strategy sets for males and females, and in economics one might have separate strategy sets for buyers and sellers. The state s now specifies the strategy shares in each population. Fitness is computed separately for each population but in general depends on the entire state s.

3.2 Adaptive landscapes

Building upon the approach just described, we may now describe the main target of our project, the evolutionary game with a geometric (continuous) space of actions.

The narrative description of a basic adaptive landscape again concerns a population of players, each with the same space of strategies, but the space now is at least one-dimensional and continuous. The instantaneous state of the system consists of a probability measure on the strategy space that describes the population distribution of chosen strategies. A player's fitness (payoff) depends on her choice of strategy, as well as on the current state of the entire system. From the player's perspective, fitness looks like a landscape in which she seeks to go uphill.[1] As play evolves over time, all players continuously adjust their strategies to increase their fitness, so the population distribution changes, and consequently the landscape morphs. A player's fitness changes as the direct result of her own actions, and also indirectly as the state evolves. The dynamics arise from this interplay between state and landscape.

The mathematical description of the basic landscape model is as follows. Let $A = [0, 1]$, the unit interval, be the space of strategies, and let $\mathcal{D}$ be the space of all probability measures (that is, cumulative distributions) on A, with the weak-star topology. $\mathcal{D}$ is the state space for this model, an infinite-dimensional simplex. The *fitness* for any player choosing strategy $x \in A$ when the current state of the game (distribution of all players' strategies) is $D \in \mathcal{D}$, is denoted by $\phi(x, D)$. The application dictates a particular fitness function $\phi : A \times \mathcal{D} \to R$. For any player holding strategy x, the function $A \to R$; $x \mapsto \phi(x, D)$ with D held constant, is the instantaneous *fitness landscape* for that player. It depends on D only, and we think of it as a landscape, or graph in $A \times R$.

The dependence of ϕ on D can take many forms. In some applications it is the expectation

[1] The landscape metaphor goes back to Sewall Wright (e.g., 1949) and has been revived by Stuart Kaufman (e.g., 1993). Wright considered low dimensional continuous landscapes and Kaufman considers high dimensional sequence spaces of discrete-valued traits. Neither considers the dynamically changing (i.e., distribution dependent) landscapes examined here.

$\phi(x, D) = \int g(x, y) dD(y)$ of some two-player game payoff function $g(x, y)$. In some (possibly oversimplified) biological applications the dependence is only via the mean strategy $\mu_D = \int y dD(y)$. In fluid dynamics applications it depends only on the local value $D(x)$ at the chosen strategy x (e.g., Witham, 1974). However, in our featured application, the dependence on the current state D is quite nonlinear and arises from market-clearing prices.

Landscape dynamics are given as a vectorfield (defined almost everywhere) on the infinite-dimensional simplex. A general expression is

$$D_t(x, t) = \Psi(x, D, \phi) \tag{1}$$

where $D(x, t)$ is the cumulative distribution function over $x \in A$ for the state at time t, and D_t denotes the partial derivative of $D(x, t)$ with respect to t.

Consistent with the inertial principle of evolutionary games and with Darwin's dictum *Natura non facit saltum*, we assume that individual adjustment is continuous. That is, a discrete change in a player's strategy x takes a positive amount of time. This restriction might seem innocuous but it is violated by generalized replicator dynamics and other dynamics that don't respect the ordering of the action set $A = [0, 1]$. (The intuition for replicator dynamics is that individuals never adjust, and dynamics arise entirely from differing birth or death rates at different strategies. Apparent jumps occur because new births don't generally appear at the same location as recent deaths.)

Consistent with the monotone principle of evolutionary games, we assume that the direction of adjustment is given by the sign of the gradient, i.e., uphill in the fitness landscape. If also the adjustment speed is proportional to the gradient $\phi_x = \partial\phi/\partial x$, we have a *gradient adjustment system*. In this case dynamics obey the *master equation*,

$$D_t(x, t) = -\phi_x(x, D) D_x(x, t) \tag{2}$$

where D_x denotes the partial derivative of $D(x, t)$ with respect to x. This nonlinear partial differential equation simply states that probability mass is conserved: the rate of change $D_t(x, t)$ in population mass to the left of any point x is equal to the (negative of the rightward) flux past that point. The flux is the product of the density $\rho = D_x$ and the velocity given by the gradient ϕ_x.

Finally, we note that the action space might be n-dimensional, and the cases $n = 2$ and $n = 3$ provide for periodic and chaotic attractors as dynamic equilibria for an adaptive evolutionary game. Connecting with the paradigms of oscillator theory and chaos theory in this way opens up a host of new social sciences applications.

3.3 Adaptive lattices

Without going into detail, the adaptive lattice system is identical to the adaptive landscape system, except that the continuous space of strategies, A, is replaced by a regular lattice of points, K in A. This seems, at first, identical to the basic model above. However, here we may carry over the geometry of A to the finite subset, K. This will always be the case when we wish to simulate an adaptive landscape on a digital computer. These systems were introduced in Friedman and Yellin (1997) for simulation and analysis, under the name *discrete gradient dynamics*.

The category of adaptive lattice games is closely related to other modeling strategies for complex dynamical systems. For example, if the discrete action space is regarded as a physical substrate, the adaptive lattice system becomes a special type of agent-based model, in which the global distribution of agents plays a dynamical role.

In our case, the lattice is a discretized form of the continuous space of actions. Thus, the choice of an action or strategy by a player or agent may be considered as a motion of the player over the lattice. The inertial principle implies that the transition matrix for the adaptive lattice is tri-diagonal. This interpretation puts our discretized adaptive lattice model exactly in the context of agent-based models, permitting direct programming of our models in the environment of Netlogo, for example.

Whether programmed in Matlab or Netlogo, the computer graphic features of these programming environments greatly facilitate the visualization of the model as it runs. In the case of a one-dimensional or two-dimensional action space, both environments have extensive graphics capabilities and are capable of animated representations of the evolutionary game in play. In the three-dimensional case, Matlab has an advantage.

3.4 Specific research agenda

Having laid out the context and its motivation, we may now be more specific about the project. We propose to work upwards from a specific application to financial markets to describe a class of adaptive landscape games sufficiently but not overly general, to reduce these to appropriate adaptive lattice games, to create software tools for their simulation and visualization, and to fit models to historical data from actual markets.

The software tools we envision might consist, for example, of prototype graphical user interfaces and computational scripts written in a common mathematical modeling environment, such as Matlab, Mathematica, Maple, or Vennsim. While built specifically for our exemplary financial market application, these would be easily adaptable for different modeling exercises, and thus reusable by us and by other researchers.

3.5 Consumption Dynamics

Veblen consumption is a good illustrative example because it has already been worked out (Friedman, 2001; Friedman and Yellin, 2000) and because it contains key elements of the financial market application. Thorstein Veblen (1899) popularized the idea that some goods and services (think of Hummers or seldom-used second homes) are consumed largely to gain status, a theme pursued more recently by authors such as Duesenberry (1949), Frank (1985) and Ljundqvist and Uhlig (2000). Such consumption has the desired effect only to the extent that it exceeds the conspicuous consumption of other people, i.e., its utility is rank-dependent.

Consider a single population of consumers with identical incomes. Each consumer chooses a fraction $x \in [0, 1]$ of income to allocate to ordinary consumption, and allocates the remaining fraction $1 - x$ to rank dependent consumption. The state is the cumulative distribution function $D(x)$ of ordinary consumption. Assume standard direct utility $c \ln x$ from ordinary consumption x, where the parameter $c \geq 0$ represents the relative importance of ordinary consumption. Suppose that rank dependent utility arises from envy, i.e., I compare my rank-dependent consumption $1 - x$ to everyone else's and am unhappy to the extent that it falls short. The shortfall is $\min\{0, y - x\}$ when your rank dependent consumption is $1 - y$. After integrating the expected shortfall by parts, one verifies that overall expected utility is

$$\phi(x, D) = c \ln x - \int_0^x D(y)dy, \tag{3}$$

with gradient $\phi_x = c/x - D(x)$.

Figure 1 shows two landscapes with $c = 0.1$ defined by this payoff function for two different initial distributions. In Figure 1A, the distribution D is uniform on $A = [0.0, 1.0]$. In Figure 1B, the distribution D is uniform on $A = [0.9, 1.0]$.

Dynamics are governed by the Master Equation (2). Insert the gradient into (2) to obtain the partial differential equation

$$D_t = D_x[D - (c/x)]. \tag{4}$$

Friedman and Yellin (2000) show that from an arbitrary initial distribution $D(x, 0)$ there is a unique solution $D(x, t)$ to (4), and that as $t \to \infty$ the solution converges to the degenerate (or Dirac delta) distribution at some point $\widetilde{x} \in [0, 1]$ whose value depends on the initial distribution and the parameter c.

For our purposes the transient dynamics are especially interesting. To illustrate, suppose $c = 0$ and the initial distribution is $D(x, 0) = 3x^2 - 2x^3$, i.e., the initial density is the symmetric unimodal (single-peaked) function $\rho(x, 0) = 6x(1 - x)$. It can be shown analytically that $D(x, t)$ has a continuous unimodal density $\rho(x, t)$ for $t \in [0, 2/3)$. The mode (or peak) $x^*(t)$ decreases steadily from $1/2$ at $t = 0$ to $1/6$ at $t = 2/3$, and the height of the mode becomes unbounded as $t \to 2/3$. The intuition is that all consumers decrease ordinary consumption (recall that for $c = 0$ they care only about conspicuous consumption) but gradient dynamics dictate that the modal consumer adjusts more rapidly than consumers with initially lower x and he begins to overtake them at time $t^* = 2/3$ and consumption level $x = 1/6$. Given our assumption of identical underlying preferences and income, this consumer can't actually pass his rivals because his behavior is identical to theirs once he attains the same consumption level. Instead, he clumps together with them, and the clump grows as it overtakes consumers with x just below the mode and is overtaken by consumers with x just above the mode.

Thus, beginning at $t = 2/3$ we get a growing, moving mass of consumers with identical consumption patterns, a homogeneous middle class. To calculate its position $x^*(t)$ and mass $M(t)$ for $t > 2/3$, one uses techniques developed in fluid mechanics to deal with shock waves. The Rankine-Hugoniot conditions (see Smoller, 1994) impose conservation of mass and exploit the weak-star topology to obtain a unique distribution function $D(x, t)$ with a jump discontinuity. It turns out that for $t \in (2/3, 1]$ the position is $x^*(t) = (1-t)/2$ and the jump size is $M(t) = \sqrt{\frac{3}{t^2} - \frac{2}{t^3}}$. Thus the middle class absorbs the entire population by the time it hits the boundary $x = 0$ at time $t = 1$. For $t > 1$, of course, everyone continues to neglect ordinary consumption.

Figure 2 shows the cumulative distribution function D for choices x, shown for time $t = 0$ and subsequent times. At $t^* = 2/3$ the distribution has a vertical tangent at $x^* = 1/6$. At later times the distribution has a jump discontinuity.

Friedman and Yellin (2000) show that the dynamics are qualitatively similar in more complex situations. A shock wave apparently will arise from any local maximum of the initial density far enough above the point $\widetilde{x}$ where the gradient is 0; and for the consumption parameter $c > 0$, the shock velocity decreases as the position approaches $\widetilde{x}$. It is reasonable to include a little behavioral or perceptual noise in the form of a diffusion term in (2), as is assumed in quantal response equilibrium models; see Anderson, Goeree and Holt (1998) for example. In this case the long run equilibrium is slightly smoothed, approximately Gaussian with small variance, and the shock waves are also smoothed. Instead of travelling jump discontinuities in the distribution function, we get travelling steep segments of a continuous distribution, locally approximately Gaussian.

3.6 Financial markets: bubbles and crashes

We propose to use the same mathematical tools to construct new models of financial markets. Consider a single population of portfolio managers whose utility depends on performance. Given performance ratings such as those published four times a year in the Wall Street Journal (and more frequently by Lipper Analytics and Morningstar), it seems clear that relative performance, or rank, matters at least as much as absolute performance. Higher rank brings bonuses and competing job offers, and also indirectly increases managers' compensation by attracting more customers.

The basic model assumes that each manager chooses a single ordered variable, $x \in [0, \infty)$. An interpretation consistent with orthodox financial theory is that x represents the leverage on the composite risky asset, the market portfolio M, with borrowing or lending of a riskless security at a fixed known rate R_o. For a given realized yield R_M on M, the manager realizes a portfolio return $R(x) = R_o + x(R_M - R_o)$. The model can be enriched slightly by assuming that the manager's cost of funds is the risk free rate R_o plus a risk premium $c(x)$, with $c(0) = 0$, $c' > 0$ and $c'' > 0$. The second derivative is positive because the probability of a loss (and its expected size if it occurs) are naturally concave in the exposure defined by x.

The realized yield R_M on M is the expected yield $\bar{R}_M > R_o$ plus the "surprise" yield e_M. The surprise is persistent; one specification is that new surprises are drawn independently at Poisson times from a known distribution with mean 0. (A more sophisticated specification is that e_M is an Ornstein-Uhlenbeck process, i.e., mean-reverting Brownian motion. A complete

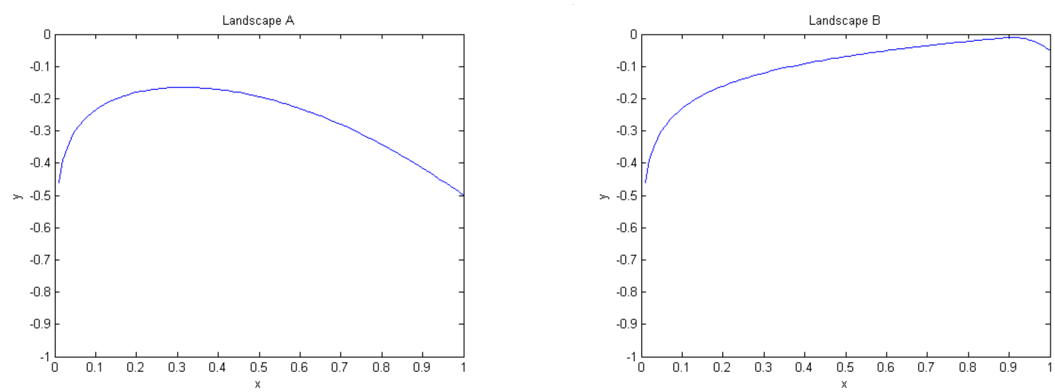

Figure 1: Two landscapes on the unit interval

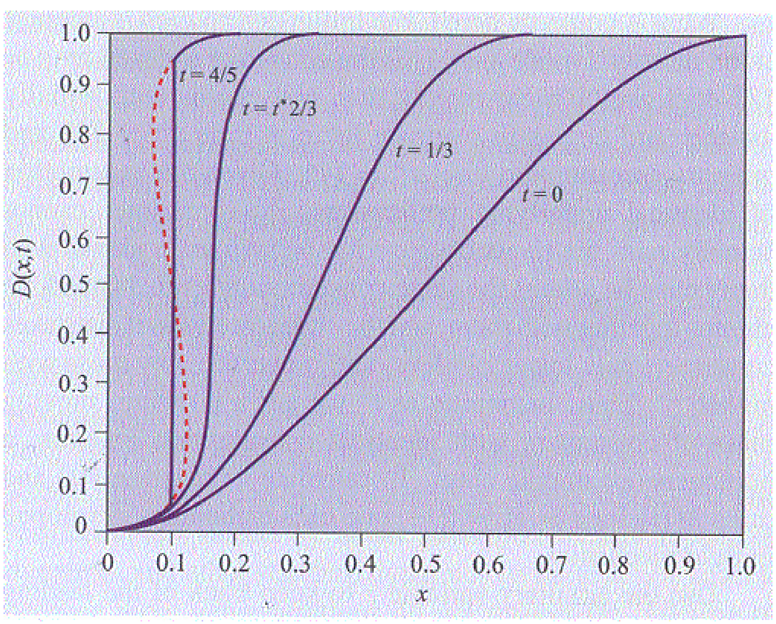

Figure 2: A compressive shock wave

model would derive e_M from shifts in the underlying technology and preferences.) The expected yield $\bar{R}_M$ depends systematically on investors' behavior. Preliminary computations, using only orthodox assumptions in economics and finance, produce an expression of the form $\bar{R}_M = k\bar{x}_D^B + g$, where k, B and g are obtained from structural parameters of the model and $\bar{x}_D$ is the current mean leverage among portfolio managers.

It now is straightforward to derive the portfolio manager's material payoff function

$$\phi(x, D) = z[ax\bar{x}_D^B + bx + e_M x - c(x)] + K, \tag{5}$$

where z is the size of the portfolio, and a, b and K again are obtained from structural parameters of the model. The gradient is $\phi_x(x, D) = z[a\bar{x}_D^B + b + e_M - c'(x)]$, and the spatial rate of change in the gradient is $\phi_{xx}(x, D) = -zc''(x) < 0$.

Gradient dynamics are especially appealing for this application. Short-run adjustment of a trader's position or asset holding x occurs mainly via net selling or buying of the risky asset M. The per-share trading cost increases with the net amount traded in a given short interval of time, mainly because of "price pressure" or limited liquidity. If the increase is linear, then the adjustment cost (net trade times per share trading cost) is exactly quadratic. Proposition 1 of Friedman and Yellin (1997) shows that quadratic adjustment costs are the key condition to obtain precisely gradient dynamics rather than approximate gradient or sign-preserving dynamics.

Even for small weight on rank-dependent utility, gradient dynamics will be quite interesting. Since (5) already has a decreasing gradient, (4) can produce shock waves in the evolving state $D(x, t)$ starting from some initial conditions. For example, suppose that the initial distribution is concentrated at low risk (small x values) and that the current surprise realization e_M is positive so $\phi_x(x, D) > 0$ for all x in the support of D. Then under gradient dynamics all portfolio managers increase their exposure to risk. Of course, $\bar{x}_D$ also is increasing in this case, so the price of the risky asset M increases and its expected yield decreases. This is a bubble, a buying panic, and we can get the same sort of travelling shock wave as in the Veblen consumption model: a travelling (and growing) clump of portfolio managers increase their exposure to risk in lockstep.

A new realization of e_M can turn the gradient negative, resulting in a crash. Even starting from a continuous distribution, we can get a travelling and growing jump discontinuity at $x^*(t)$ in the distribution $D(x, t)$, as managers scramble to unload risky assets. The crash occurs on low trading volume; $x^*(t)$ and portfolio values decrease mainly due to a sharp decline in the price of the risky asset M. Gradient dynamics for probability distributions seem to have the general feature that it is easier to obtain a shock wave that moves downward (in this application, a crash) than one that moves upward (here, a bubble). Indeed, the Veblen consumption model apparently does not permit upward moving shock waves.

4 Conclusion

We are at a special moment for advancing the art of mathematical modeling in the social sciences. Just as mathematical tools, computer simulation, and computer graphic visualization techniques are enabling new methods for the social, behavioral, and economic sciences, our own programs of research have matured and focused on a particularly promising method, the evolutionary game with continuous action. We now propose to accelerate the development of special software and exemplary applications, and to send them out into cyberspace.