Abstract

Evolutionary game models analyze strategic interaction over time; equilibrium emerges (or fails to emerge) as players/traders adjust their actions in response to the payoffs they earn. This paper sketches some early and some recent evolutionary game models that contain ideas useful in modeling financial markets. It spotlights recent work on adaptive landscapes. In an extended example, the distribution of player/trader behavior obeys a variant of Burgers' partial differential equation, and solutions involve travelling shock waves. It is conjectured that financial market crashes might insightfully be modeled in a similar fashion.

Towards Evolutionary Game Models of Financial Markets

Daniel Friedman*

16 August 2000

1. Introduction

Mainstream economics and finance theory assumes equilibrium, i.e., that all market participants choose optimally given the choices of other participants. Equilibrium definitions have increased in sophistication over the last 50 years to encompass interdependent markets (general competitive equilibrium, e.g., Arrow and Hahn, 1971), participants with market power (Nash equilibrium, e.g., Shubik, 1982), and participants with private information (Bayesian Nash equilibrium, e.g., Hirshleifer and Riley, 1994).

Physicists and market participants can learn much from equilibrium models, but in the end are likely to be dissatisfied. Equilibrium models do not tell us when and how equilibrium might be achieved, in terms of how traders respond in actual time to cues generated by other traders. The adaptation process allows traders to make big profits and losses, and it may lead to equilibrium quickly or slowly or not at all. It deserves careful study.

Unfortunately there is no canonical model of the adaptation process in financial markets or other economic settings. Numerous models of economic adaptation have been proposed in the last decade or so. They increasingly rely, implicitly or explicitly, on evolutionary game theory. As defined in Friedman (1998), evolutionary games have three primary characteristics that embody the mottoes "survival of the fittest," "evolution not revolution," and "natural selection" respectively:

- monotone: higher payoff strategies displace strategies with lower payoff;

*My thanks to the Santa Fe Institute workshop participants for useful comments and to Steve Anderson for reserach assistance.

- inertial: the player population takes real time to change behavior; and
- price-taking: players regard the present and future play of others as exogenous as in a game against Nature; they don't try to influence other players' future choices.

On this definition, evolutionary game theory covers a vast intellectual territory beyond the borders of equilibrium theory (which assumes no behavioral inertia and assumes extreme rationality rather than monotonicity) and repeated game theory (where players try to influence others' future behavior). Much of the territory remains uninhabited, but parts of it support intense recent activity, witnessed in more than a dozen books and hundreds if not thousands of articles. I will not attempt a systematic survey here, but more of a helicopter tour. Like any tour guide with limited time, I'll disproportionately cover my favorite parts and those I know best. I'll pay special attention to work that seems promising for models of financial markets. Generally I cite good textbook treatments, not always the original source or the latest work.

The next two sections lay out the main ingredients of an evolutionary game model, and show how the classic work of the 1980s used the ingredients. Section 4 lists the underlying evolutionary processes that might be important in modeling financial markets, and section 5 describes two illustrative models from the recent literature. Section 6 points out some of the areas of greatest recent activity, particularly models of learning in games. Section 7 presents unpublished work on evolutionary games in which actions are chosen from an ordered continuum (rather than a discrete unordered set), as in many financial market choices. Here the payoff function can be visualized as a landscape that morphs as players adjust their actions trying to increase payoff.

Section 8 summarizes an application of such gradient dynamics to the choice of conspicuous vs ordinary consumption. It shows that, depending on specification details, the long run equilibrium can either be clumped (e.g., all consumers choose the same bundle) or dispersed. Perhaps of more interest, it explains how transient dynamics can produce a moving, growing clump (a jump discontinuity in the cumulative distribution function, interpreted here as an emerging "middle class"). As in ocean surf, the discontinuity occurs because the velocity of the upper part of the distribution exceeds the velocity of the lower part, and a piece of the distribution function becomes vertical. The last section conjectures how the various ideas might be combined into an evolutionary model of financial markets that can insightfully describe bubbles and crashes.

2. Evgame Model Ingredients

Evolutionary game models are constructed from the following ingredients:

- 1. $k \geq 1$ populations, each with its own action set. The simplest example is a single population playing a 2x2 symmetric bimatrix game, so the action set is just the two alternative strategies. In a market game, we might have a population of traders each with a price q above which a unit is sold and below which a unit is purchased. We might include a second population of market makers, each of which has a bid price b and an ask price a. The current state s is the distribution of actions chosen in each population.

- 2. A payoff (or fitness)[1] function. The simplest example is a matrix whose (i, j)-th element is the payoff to a player choosing the i-th action when the opponent chooses the j-th action, e.g., $A = \begin{pmatrix} 1 & -1 \\ 2 & 0 \end{pmatrix}$ for a version of the prisoner's dilemma. In a market game, the payoff typically would be the trading profit given the trader's buy or sell order(s) and the price(s) determined by all traders' actions. In general, the payoff for a player in population k is a real-valued function $f_k(x, s)$ that is linear in own (mixed) action x and possibly nonlinear in the state variable s.

- 3. Dynamics describing how behavior in each population responds to realized payoffs. Standard examples such as replicator dynamics and smoothed fictitious play will be given below.

- 4. One or more definitions of equilibrium. We will discuss leading examples below, both static and dynamic.

3. Basic Evolutionary Game Theory and Applications

Maynard Smith and Price (1973; see also Maynard Smith, 1982) founded evolutionary game theory on a static equilibrium concept called Evolutionarily Stable Strategy (ESS). For a single population game the definition can be expressed in terms of the current state s and a perturbed version $s' = \varepsilon x + (1 - \varepsilon)s$. The perturbed version is referred to as a (ε-) small invasion of (x-) mutants. For a single population the formal definition is as follows. The current state s is an ESS if $f(s, s') > f(x, s')$ for all $x \neq s$ and for all $\varepsilon > 0$ sufficiently small.

[1] In this paper fitness and payoff are synonymous. In some other papers, fitness is constructed from a given payoff and other factors such as transmission rates.

The idea is that mutants in any small invasion have lower fitness and therefore no invasion can find a toehold to displace the current state. Standard game theory would refer to ESS as a proper Nash equilibrium of a symmetric game. For linear payoff functions (e.g., in matrix games) the inequality is independent of invasion size and the definition can be re-expressed in its original form: either $f(s,s) > f(x,s)$, or $f(s,s) = f(x,s)$ but $f(s,x) > f(x,x)$. The generalization to multipopulations is conceptually straightforward; see Cressman (1995) for technical details.

As an equilibrium concept, ESS has two shortcomings. First, it is not hard to find games with multiple ESS or with no ESS. Second, although the idea is dynamic (mutant invasions fail), the definition is static. One would like to predict behavior beginning at an arbitrary initial state, and note when we have convergence to specific equilibrium points. Such predictions, of course, depend on the dynamics.

Replicator dynamics (Taylor and Jonker, 1978) are the first example of evolutionary game dynamics and are still quite prominent. In biology the relevant payoff is fitness, defined as the growth rate. Abstracting from genetic complications, it is natural in biology to assume that the growth rate of any action (or trait) relative to alternatives is its payoff relative to alternative payoffs. The standard definition of replicator dynamics assumes a single population with a finite number n of alternative actions, so the state $s = (s_1, ..., s_n)$ is a point in the simplex $\sum s_i = 1$, $s_i \geq 0$. The i^{th} vertex of the simplex, denoted $e^i = (0, ...1, ..., 0)$, represents an individual choosing pure strategy i or (as the second argument in the payoff function) the state where everyone plays i. It is easy to check that the population-weighted average payoff at state s is $f(s,s)$. Replicator dynamics then are defined by the condition growth rate = relative payoff,[2] $(ds_i/dt)/s_i = f(e^i, s) - f(s,s)$ for each population fraction i, or

$$ds_i/dt = s_i(f(e^i, s) - f(s,s)), i = 1, ..., n.$$

One nice result is that dynamically stable equilibria (i.e., the locally asymptotically stable steady states) of replicator dynamics include all ESS, as well as a few other points such as the vertices. Multipopulation generalizations are straightforward. Discrete time versions are conceptually straightforward; see Weibull (1995) for technical details.

In economic applications it is perhaps too strong to assume the growth rate (or, alternatively, the rate of increase) in the fraction of the population employing action i is directly proportional to the relative payoff. It seems more reasonable

[2] Another way to describe replicator dynamics is as the radial projection onto the simplex of Lotka-Volterra dynamics (Friedman, 1991).

to assume only that over time higher payoff actions will crowd out lower payoff actions. This more general case, called monotone dynamics, preserves many of the nice results from replicator dynamics (Friedman, 1991; Samuelson and Zhang, 1992; Weibull, 1995).

I am aware of two financial market applications of early evolutionary game theory. Conlisk (1980), apparently working independently of the biological literature, finds conditions under which agents that optimize at some cost can coexist in dynamic equilibrium with cheap imitators. Cornell and Roll (1981) use a variant on biologists' 2x2 "hawk-dove" matrix to address a same theme in financial markets: when will costly analysis of fundamental value coexist with cheap rules of thumb in choosing stock purchases? They find an interior (mixed) ESS of costly informed traders and uninformed traders, characterized by a marginal information cost = marginal benefit condition and an endogenous market share condition with sensible comparative statics.

The Cornell and Roll model, like other early evolutionary game models, assumes pairwise random interactions described by payoff matrices and therefore the payoff function is linear in the state variable s. This assumption is difficult to reconcile with the more complex forms of market interaction that determine asset price and asset returns. For example, in the simple competitive model described in section 5 below, the the price is set by certain order statistics of s (the marginal bids or asks) and the payoff function is highly nonlinear. Asymmetries or stochastic components play a role in other market models.

4. Sources of Dynamics

The distinctive features of financial markets demand that evolutionary models be constructed from scratch, not simply borrowed from biological or other applications. The dynamics are of special interest. From the outset we must recognize that there is not just one all-purpose adaptation process that is appropriate for all markets and all time scales. Rather, there are several quite distinct processes:

- Entry and exit (and mergers and acquisitions). The exit of bankrupt producers and the entry of producers with new technology is perhaps the most economically important example, e.g., Nelson and Winter (1982). Another example, central to biologists but usually unimportant on economists' time scales, is birth and death leading to genetic evolution of agent populations via natural selection. In financial markets we have new securities and new investors entering occasionally and others exiting.

- Endogenous market shares. Even when the trader population is constant and individual traders do not change behavior, we can have market-level adjustment as traders with less profitable behavior lose wealth and market share to traders with more profitable behavior, as in Blume and Easley (1991). A fine grained variant considers how an individual trader employing a portfolio (or mixed) strategy increases weights on more successful components.[3]

- Adaptive learning. Arguably the most important process over the most relevant time scales (minutes to months) is that traders systematically change their actions in response to personal experience. This will be the focus of the next two sections.

- Learning variants include: observational learning in response to other traders' experience, direct imitation of more successful traders' actions, and active learning by trying actions more for their informativeness than for their direct profitability.

- Institutional evolution. The market rules themselves change in response to competitive pressures, e.g., the NYSE no longer forbids after hours trading in listed stocks. One could imagine formalizing some of the insights of North (1990).

5. Learning in Markets: Two Examples

Two recent examples may help fix ideas. Many physicists regard the minority game (Challet and Zhang, 1997) and the El Farol problem (Arthur, 1994) as models that capture the congestion aspect of financial markets. The El Farol problem can be described as a single population evolutionary game with two choices for each of 100 players: stay home, or go to the bar El Farol, whose seating capacity is limited. The payoff to going to the bar is -1 if 60 or more players choose to go and is +1 if fewer than 60 do so. The payoff to staying home is always 0. Bell et al (1999) take the probability p with which 1 is chosen as the strategy, so the state s is a point in $[0,1]^{100}$. There is a unique symmetric Nash equilibrium in mixed strategies where everyone chooses $p \approx 0.6$ and receives expected payoff 0. By contrast, there are lots ($\binom{100}{60}$) of asymmetric pure strategy

[3]For example, Peter Muller of Morgan Stanley said in the May 2000 SFI workshop, "We feed the strategies that work." Other practitioners echoed this theme and some mentioned specific adjustment dynamics such as the Kelly rule, intended to maximize the growth rate of portfolio value. From a population perspective, such changes are the same as adaptive learning.

equilibria with exactly 60 players choosing to go, and these equilibria are efficient with average payoff 0.6.

Bell et al propose a discrete time learning algorithm in which each player who goes to the bar increments the attendance probability p by an amount proportional to (60 - current attendance), and players who stay home don't update p. The new values of p are truncated at below at 0 and above at 1. Using techniques related to stochastic approximation (but with weaker assumptions), they show that the stability of the discrete time learning dynamic on $[0,1]^{100}$ is the same as that of the ordinary differential equation $ds/dt = G(s)$, where each component of G is the expected value of the increment (60 - current attendance) for a given player. Then they show that the ordinary differential equation (and therefore the original learning dynamic) has a stable fixed point at each of the efficient asymmetric Nash equilibria, but not at the inefficient symmetric mixed Nash equilibrium. Convergence is exponential. The interpretation is that players who learn from personal experience will rapidly and efficiently sort themselves out into regular attendees and non attendees.

Let us now consider learning in the simplest true market institution, the call market. A call market can be modeled as a two-population evolutionary game of buyers (each with a privately known value) and sellers (each with a privately known cost). Each period each buyer i submits a bid (the highest acceptable purchase price for a single indivisible unit) b_i and each seller j submits an ask (the lowest acceptable sale price) a_j. The demand revealed in $\{b_i\}$ and the supply revealed in $\{a_j\}$ then are cleared at a uniform price p^*.[4] The payoffs are profits, $v_i - p^*$ for a buyer with value v_i and $p^* - c_j$ for a seller with cost c_j.

When buyers' values and sellers' costs are drawn randomly every period, the relevant static equilibrium concept is Bayesian Nash Equilibrium (BNE).[5] Here each buyer optimally reduces her bid below value (and each seller increases his ask above cost) to the point that (i) the marginal loss from the reduced probability of transacting just matches (ii) the marginal gain conditional on transacting. Rustichini, Satterthwaite and Williams (1994) compute BNE bid and ask functions. The question is whether human traders learn (to behave as if they use) these functions. For example, Figure 1 shows buyers' values and actual bids in one experimental session. Do these bids tend over time to come closer to the

[4] That is, p^* maximizes transaction volume subject to the constraint that transacted bids are are at least p^* and transacted asks are at most p^*. Existence of such a competitive equilibrium price is guaranteed, but with indivisible units there is often a nondegenerate interval of such prices, as explained for example in Cason and Friedman (1999).

[5] When values and costs are constant over time, then the relevant static equilibrium concept is competitive equilibrium. See Friedman and Ostroy (1996) for an analysis of call market convergence in this setting.

graph of the BNE bid function?

Cason and Friedman (1999) reduce the buyer's learning problem to adjustment of a single parameter, the slope of the bid function, i.e., the mark-down ratio of bid relative to private value. Similarly, sellers try to learn the best mark-up ratio of asks relative to private cost. The learning model assumes partial adjustment of these ratios towards the ex post optimum. It can be written in terms of the logarithm r_t of the mark-up (or -down) ratio in period t and the value r_t^o that would have been most profitable in hindsight:

$$r_{t+1} - r_t = a + b(d, e, o)(r_t^o - r_t).$$

The adjustment coefficient b may depend on variables (d, e, o) mentioned below.

The learning model is estimated with pooled data from all traders (4 buyers and 4 sellers) in a series of laboratory market sessions, each with about 100 trading periods.[6] The parameter estimates indicate negligible autonomous trend ($a \approx 0$) and positive adjustment towards the ex post optimum ($b > 0$). There is a strong recency effect: the learning rate b is rather insensitive to a discount factor d that represents accumulated experience. Perhaps the most striking finding is a strongly asymmetric response to different kinds of ex post error (e): traders respond strongly when they miss a trade ($e = m$) because they marked up (or down) too aggressively, but hardly respond at all when they trade at a less favorable price by not being aggressive enough. The asymmetry is only slightly attenuated in observational learning from other traders' ex post errors (o).

Simulations show that the model can account for the main systematic deviations from BNE predictions observed in the laboratory data. Indeed, convergence in the simulations and in the laboratory is often closer in relevant cases to full revelation ($r = 0$, or $b = v$ and $a = c$, consistent with competitive equilibrium) than to the BNE predictions, $r \approx \ln 0.8$ or $\ln 0.9$, depending on details of how the clearing price is selected when it is not unique. The simulations track the median trader's markup (or markdown) rather well, but greatly understate variability across individual traders.

6. Recent Developments

I now pause the tour to list some important areas we shall skip, and to point out useful surveys and references. The standard textbook reference on evolutionary game theory is Weibull (1995). Classic books include Maynard Smith (1982),

[6] Separate estimates are also made by type of market session (e.g., whether traders have previous experience) and for buyers and sellers, but the data are insufficient to allow separate estimation of individual traders.

and Hofbauer and Sigmund (1988). Important single author collections include Cressman (1992), Samuelson (1997), and Vega-Redondo (1996). Selten (1991) is an amusing introduction to some modeling issues. Friedman (1998), a brief "user's manual," includes a guide to the literature.

The text by Fudenberg and Levine (1998) briefly covers the classic material but focuses on learning dynamics in bimatrix games (with two chapters on extensive form games). It is a good introduction to recent stochastic models. One chapter emphasizes smoothed fictitious play, and connects it to stochastic approximation (introduced to microeconomists in the early 80s by Arthur, Ermolev and Kaniovski, and to macroeconomists in the late 1980s by Marcet and Sargent; (see also Benaim and Hirsch, 1999), to logit estimation, and to probability matching behavior and reinforcement learning. It also presents stochastic adjustment models in the style of Kandori, Mailath and Rob (1993) and Young (1993), and notes that such models give sharp predictions on equilibrium selection but neglect important issues regarding time scale and depth of basins of attraction. The book also touches on important recent developments, such as local interaction (e.g., Ellison, 1993), machine learning (e.g., Kivanen and Warmuth (1995), and rule-based learning (e.g., Stahl, 1999).

Camerer (1999) summarizes recent laboratory evidence on learning in games. His own work (most of it with Ho) hybridizes two of the main types of learning models, reinforcement and belief learning. The logit choice model just mentioned plays the key role in the quantal response equilibrium (QRE) model of McKelvey and Palfrey (1995). QRE fits a variety of laboratory data pretty well, especially when it allows the error amplitude to decline over time (e.g., Anderson, Goeree and Holt, 1998). Chen, Friedman and Thisse (1997) independently develop QRE and extensions to true learning models.

7. Landscape Learning

Some matrix game experiments allow estimation of individual player data (e.g., Cheung and Friedman, 1997), and the data can be interpreted using well established theoretical models of heterogeneous populations, e.g., ESS or replicator dynamics. There is no practical obstacle to running market experiments that allow individual trader estimates, but there are no well established theoretical models to help interpret the data. The problem is that traders' action sets are ordered continua (e.g., a range of prices or a range of markup ratios) rather than the small discrete unordered sets assumed in standard evolutionary games.

In this section I describe some recent work on evolutionary game models for continuous ordered action sets first introduced in Friedman and Yellin (1997, denoted FY97 below). See also Basov (1999). For concreteness, take a large single population of players (e.g., traders), each with action set $A = [0,\ 1]$, the unit interval. As in the finite action case, it is convenient to normalize the population mass to 1, i.e., to express all variables in per capita terms. Now the state is a point in an infinite dimensional simplex, the set of probability measures on $[0, 1]$ endowed with the usual weak-star topology. It is more convenient to represent the state at time t by a probability density $\rho(x, t)$ on A (including improper densities known as Dirac delta functions), or equivalently by a cumulative distribution function $D(x, t) = \int_0^x \rho(y, t)dy$, for $x \in [0, 1]$.

In this setting, a player's fitness $\phi(x, D)$ depends on her chosen action $x \in [0, 1]$ and the current state, the entire distribution D of other players' actions.[7] The graph of the fitness function, holding constant the current state D, defines a fitness landscape.[8] The player adjusts the choice x to reach a higher point on the landscape. However, as all players adjust, the distribution D changes and the landscape morphs. Players then respond to the new landscape, further changing the distribution. This interplay between fitness landscape and distribution can lead to nontrivial dynamics.

The rest of this section assumes that players' adjustment speed is proportional to the fitness gradient, $\phi_x = \partial\phi/\partial x$. Proposition 1 of FY97 shows that such adjustment results from quadratic adjustment costs. The other propositions in FY97 show that the main qualitative results require only that adjustment is in the same direction as the gradient, i.e., uphill in the fitness landscape. The main substantive restriction is that adjustment doesn't involve a player jumping instantaneously from current action x to a distant action y. This restriction reflects Darwin's dictum *Natura non facit saltum* (also favored by the economist Alfred Marshall). The restriction is violated by direct generalizations of replicator dynamics and other standard dynamics on finite action spaces, because these generalizations don't respect the ordering of the action set $A = [0, 1]$.

Dynamics in a gradient adjustment system can be characterized by a population mass conservation law: the rate of change 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, so we have the equation

[7] The dependence on the distribution D can take many forms. In some applications ϕ depends on a statistic of D such as the mean or variance or an order statistic. Below I describe an economic application where ϕ depends on the local value, $D(x)$.

[8] 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 dynamically changing (i.e., distribution dependent) landscapes as we do here.

$$D_t(x,t) = -\phi_x(x,D)D_x(x,t), \quad (7.1)$$

which holds at all points $x \in (0,1)$ where the current distribution is continuous. To preserve population mass in the interval [0,1], we impose the condition $\phi_x(1,t) \leq 0$ at $x = 1$, and $\phi_x(0,t) \geq 0$ at $x = 0$. Discontinuities in the distribution, or "clumps," are especially interesting because they represent a positive fraction of the population choosing precisely the same action. The next section illustrates how such discontinuities arise and how to work with them.

Gradient dynamics on distributions have potential applications in biology (e.g., evolution of continuous traits such as beak size, as in Eshel, 1983), political science (evolution of party or candidate positions in issue space, as in Kolman et al, 2000) and economics (e.g., location of firms in characteristic space, as in Sonnenschein, 1982) as well as in physics (e.g., Lam, 1997). In such applications, the modeler constructs a fitness function (and an initial distribution) from the data, and asks whether the landscape and distribution settle down to a steady state. If so, one asks whether the steady state is a continuous distribution (so heterogeneous choices persist) or a degenerate single point distribution (everyone makes the same choice in long run equilibrium) or some other distribution.[9] Transient dynamics are sometimes of even greater interest.

8. Consumption Dynamics

It's time for an illustrative example of a landscape model. A Veblen consumption model is a good choice for three reasons. It cleanly illustrates a nontrivial fitness function that leads to interesting transitory states and simple but varied asymptotic states. Second, it features rank dependent fitness, which may be an important aspect of financial markets. Third, it has been worked out in recent work by Friedman and Yellin (2000, denoted FY00 below).

Perhaps taking a cue from Rae (1834), Veblen (1899) popularized the idea that some goods and services (think of yachts, or even cars and 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 con-

[9] For example, the Bell et al version of the El Farol game features a noisy version of gradient dynamics. The long run steady state is a unique degenerate two-point distribution with population mass 0.6 on the point $x = 1$ and mass 0.4 on $x = 0$.

spicuous consumption of other people, i.e., its utility is rank-dependent.[10] Frank emphasizes the point that conspicuous consumption has a nonpecuniary negative externality: increasing my own rank decreases the rank of the people I pass.

The evolutionary model here is single population, with action $x \in [0,1]$ representing the fraction of income an individual allocates to ordinary consumption, so $1-x$ is allocated to rank dependent consumption. We assume that an individual receives direct utility $c \ln x$ from ordinary consumption x,where the parameter $c \geq 0$ represents the importance of ordinary consumption relative to conspicuous. We specify rank dependent preferences driven by envy, so an individual choosing the fraction of rank dependent consumption $1-x$ when everyone else chooses $1-y$ receives disutility proportional to the amount by which $1-y$ exceeds $1-x$, viz.

$$r_E(x,y) = \min\{0, (1-x) - (1-y)\} = \min\{0, y-x\}. \tag{8.1}$$

If others' choices of ordinary consumption y are distributed according to the cumulative distribution function $D(y)$, then from (8.1) the rank dependent utility component is

$$\int_0^1 r_E(x,y)dD(y) = \int_0^x (y-x)dD(y) = -\int_0^x D(y)dy, \tag{8.2}$$

and the total payoff (i.e., the fitness function) is

$$\phi^E(x,D) = c\ln x - \int_0^x D(y)dy. \tag{8.3}$$

Dynamics are governed by the gradient $\phi^E_x = c/x - D(x)$.

Alternatively, one could specify rank dependent preferences driven by pride, the degree to which own conspicuous consumption exceeds that of others. Here the pairwise kernel is

$$r_P(x,y) = \max\{0, (1-x) - (1-y)\} = \max\{0, y-x\}, \tag{8.4}$$

the rank dependent component given choice distribution $D(y)$ is

[10] Why might people have a rank dependent preference component? High rank signals prosperity, which may be profitable to some people (think of lawyers and brokers). Perhaps the preference is innate and part of our evolved psychology. E.g., in Barkow, Cosmides and Tooby (1992) one finds the argument that status seeking was adaptive for our ancestors because higher status brought better access to valuable resources. We won't pursue such arguments here but will simply take as given a taste for conspicuous consumption, and explore the dynamic consequences.

$$\int_0^1 r_P(x,y)dD(y) = \int_x^1 (y-x)dD(y) = \langle x\rangle - \int_0^x [1-D(y)]dy, \quad (8.5)$$

and the overall payoff function is

$$\phi^P(x,D) = c\ln x + \langle x\rangle - \int_0^x [1-D(y)]dy, \quad (8.6)$$

where $\langle x\rangle = \int_0^1 x dD(x)$ is the population-mean choice. The gradient of the pride payoff is the same as the envy payoff gradient from with the distribution $D(x)$ replaced by the survival function $1-D(x)$. If both pride and envy are present, then straightforward calculations disclose that the mixed gradient is a renormalization of the stronger component, e.g., of ϕ_x^E if pride has weight a and envy has weight $b \in [0,a)$.

Inserting the Envy gradient $\phi_x^E(x,D) = c/x - D(x)$ in the Master Equation (7.1), we obtain the partial differential equation

$$D_t = D_x[D - (c/x)]. \quad (8.7)$$

Theorem 1 of FY00 shows that given any initial distribution $D(x,0) = F(x)$, the distribution converges monotonically under (8.7) to a degenerate steady state distribution with everyone choosing the same action $\widetilde{x}$, the solution to $\widetilde{x}F(\widetilde{x}) = c$.[11] The intuition is captured in the following example. Assume a uniform initial distribution $F(x) = D(x,0) = x$. Then $\widetilde{x} = \sqrt{c}$ solves the equation $xF(x) = c$ (or $\phi_x^E(x,F) = 0$). At $t = 0$ in equation (8.7), we see that $D_t \leq 0$ at points $x < \widetilde{x}$, and $D_t \geq 0$ at points $x > \widetilde{x}$. The slope D_x increases at $x = \widetilde{x}$, but the value $D(\widetilde{x},t)$ remains unchanged because no mass passes through $\widetilde{x}$. Hence the zero of the gradient remains at $x = \widetilde{x}$, and this point is an attractor toward which mass flows from both directions. Eventually the entire population clusters at $x = \widetilde{x}$.

In contrast to Envy, under the Pride specification there is a specific dispersed long run equilibrium that is independent of initial condition. Note that the gradient $\phi_x^P = c/x - 1 + D(x,t) > 0$ for $x < c$, so in long run equilibrium we have $D^*(x) = 0$ for $x \in [0,c)$. The boundary sign restriction allows a mass point at $x = 1$. All remaining mass must reside at points $x \in [c,1)$ such that $\phi_x^P(x,D^*) = 0$, i.e., $D^*(x) = 1 - c/x$. Thus for $0 < c < 1$, the asymptotic distribution $D^*(x)$ follows the hyperbolic arc $1 - c/x$ inside the unit square $0 \leq x, D^* \leq 1$. The asymptotic distribution follows the lower edge $D^*(x) = 0$

[11] If the initial distribution has a mass point and no such solution exists, then one uses the more general characterization $\widetilde{x} = \sup\{x \in [0,1] : xF(x) < c\}$.

(upper edge $D^*(x) = 1$) when the arc is below (above) the square. Thus all consumers end up on the same indifference curve given by the hyperbolic arc $D^*(x) = 1 - c/x$. On this indifference curve the marginal gain $1 - D^*(x)$ from increasing rank dependent consumption is exactly offset by the corresponding marginal cost $-c/x$ of reducing ordinary consumption. The payoff function is constant and maximal along the arc and at the upper endpoint $x = 1$, but takes lower values $c \log x - x$ in the unpopulated zone $[0, c)$.

What can be said about transient dynamics? The non-linear Envy equation (8.7) can be regarded as an integrable family of ordinary differential equations (ODEs) linked by the initial condition $D(x, 0) = F(x)$. To characterize solutions, it is useful to define an auxiliary variable $z = z(x, t)$ implicitly given by $D(x, t) = F(z)$. The Appendix in FY00 shows that (8.7) can be integrated to obtain the general implicit solution

$$t = \frac{z - x}{F(z)} + \frac{c}{F^2(z)} \ln \left(\frac{|c - zF(z)|}{|c - xF(z)|} \right). \tag{8.8}$$

For initial conditions F and parameter values c that allow (8.8) to be rewritten uniquely in the form $z = z^*(x, t)$ for each $x \in [0, 1]$ and $t \geq 0$, the PDE has a unique and continuous solution $D(x, t) = F(z)$. On the other hand, if $z^*(x, t)$ is multiple valued for some (x, t) then the solution incorporates a shock wave as explained below.

To fix ideas, we find the explicit solutions for the pure conspicuous consumption case $c = 0$. Given a uniform initial distribution $D(x, 0) = x$, (8.8) gives $z = \frac{x}{1-t}$. Hence, using the Heaviside notation $\Theta(x) = 0$ for $x < 0$ and $= 1$ for $x \geq 0$, the solution is

$$D(x, t) = \frac{x}{1 - t} \Theta(1 - t - x), \;\; 0 < t < 1, \tag{8.9}$$

with associated density

$$\rho(x, t) = \frac{1}{1 - t} \Theta(1 - t - x), \;\; 0 < t < 1. \tag{8.10}$$

The solution says that during the time interval $0 < t < 1$ households decrease their consumption of ordinary goods and increase rank dependent consumption so as to maintain a uniform distribution in x on the (shrinking) interval $0 < x \leq 1 - t$. At $t = 1$ the time paths for all households cross, and all consumers cluster in a delta-function singularity at $x = 0$. Conservation of probability mass in the unit interval and gradient adjustment imply there is no further change in the distribution for $t \geq 1$; every household continues to devote all resources to rank dependent consumption, not surprisingly since ordinary consumption has no value when $c = 0$.

Things become interesting even in the $c = 0$ case when the initial distribution has an interior mode. For example, take $F(x) = 3x^2 - 2x^3$ so the initial density $f(x) = 6x(1-x)$ is unimodal and symmetric about $x = 1/2$. The general solution (8.8) yields the cubic expression $-2tz^3 + 3tz^2 - z + x = 0$. The solution becomes multiple valued when z_x becomes infinite, i.e., at a time t^* such the characteristics (indexed by z) pile up at some point x^*. Taking the derivative of the cubic expression with respect to x we find $z_x(6tz - 6tz^2 - 1) + 1 = 0$. Of course at $t = 0$ we have $z_x = 1$ and for a short period of time z_x remains finite because the expression in parenthesis remains positive. But at $t^* = 2/3$ the expression in parenthesis has the real root $z^* = 1/2$, and hence the solution has a singularity. Plugging z^* and t^* into the cubic expression and solving for x we see that the singularity emerges at choice $x^* = 1/6$.

Some intuitive remarks may be in order before proceeding. The initial distribution is steepest at the density's mode $x = 1/2$; the consumer indexed by $z = 1/2$ decreases x more rapidly than households with initially lower x and begins to overtake them at time $t^* = 2/3$ and ordinary consumption level $x^* = 1/6$. Given our assumption of identical underlying preferences, 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 lower z indexes and is overtaken by those with higher z. Thus, beginning at t^* we get a growing, moving mass of consumers with identical consumption patterns, a homogeneous middle class. To calculate its position $s(t)$ and mass $M(t)$ for $t > t^* = 2/3$, one uses shock wave techniques developed in fluid mechanics.

Use the auxiliary variable z to keep track of the leading and trailing edges, $z_L(t)$ and $z_R(t)$,of the clump or middle class, whose mass thus is $M(t) = F(z_R) - F(z_L)$. Three equations (called the Rankine-Hugoniot conditions) in the three variables z_L, z_R and s characterize the desired quantities. The first two equations apply the general solution (8.8), here simply the cubic expression above) to z_L and z_R, at position $x = s(t)$. The third equation integrates the conservation of mass equation (8.7) across the shock and takes the limit as the upper and lower limits of integration converge to $s(t)$. The result is $-M(t)ds/dt = (F^2(z_R) - F^2(z_L))/2$, which simplifies to $ds/dt = -\frac{1}{2}(F(z_R) + F(z_L)) = -\frac{1}{2}$. The last equality follows from $F(z_R) + F(z_L) = 1$, using the symmetry of the initial distribution around $z^* = 1/2$. Recall that $s(2/3) = 1/6$, so for $t^* \leq t \leq 1$ the shock position must be $s(t) = (1-t)/2$.

The first equation (8.8) can be written in the form $F(z_R) = (z_R - x)/t$. Write $z_R = \frac{1}{2} + a$ to get $F(\frac{1}{2} + a) = [\frac{1}{2} + a - (1-t)/2]/t = \frac{1}{2} + \frac{a}{t}$. Again using the symmetry of F, write $z_L = \frac{1}{2} - a$ and obtain $F(\frac{1}{2} - a) = \frac{1}{2} - \frac{a}{t}$. Subtracting the second equation from the first, obtain $M(t) = F(\frac{1}{2} + a) - F(\frac{1}{2} - a) = \frac{2a}{t}$. Substituting $F(z) = 3z^2 - 2z^3$, straightforward algebra reveals that the shock mass is $M(t) = \sqrt{\frac{3}{t^2} - \frac{2}{t^3}}$, valid for $t^* \leq t \leq 1$. Thus the middle class absorbs the entire population by the time it hits the boundary $x = 0$ at time $t = 1$. Again, everyone continues to neglect ordinary consumption after $t = 1$. See Figure 2.

FY00 shows that dynamics are qualitatively similar in more complex situations. When the consumption parameter $c > 0$, the shock velocity decreases as the position approaches $\widetilde{x}$, and nice closed form solutions are unavailable when the initial distribution is asymmetric. However, theorem 2 in FY00 shows that a moving interior shock will appear in finite time for any smooth initial density function if $c >$ [redacted] [illegible]

9. Discussion

Evolutionary games provide a broad framework for models of interaction over time. Equilibrium emerges (or fails to emerge) in real time as players/traders adjust their actions in response to the payoffs they earn. In a financial market model, the payoffs depend on the prices (or returns) resulting from all traders' actions as well as on own action choice. The adjustment dynamics can produce standard homogeneous equilibria in which all traders of the same type choose the same optimal action, or dispersed heterogeneous equilibria. Even in the homogeneous case, the dynamics can be interesting, involving travelling clumps corresponding to shock waves in fluid mechanics.

New models of asset price bubbles might be constructed using these tools. Consider a single population of portfolio managers whose compensation depends on performance rank. Each manager chooses a single ordered variable, the portfolio risk measured (say) by beta, the normalized covariance of own return with the overall market return. One would write out how asset returns depend on the distribution of choices, and how managers' payoff depends on their realized returns (given own portfolio choice and the vector of asset returns) relative to the returns realized by other managers. I conjecture that such a model will have some phases where equilibrium (e.g., the Capital Asset Pricing Model) is quite descriptive and other phases (e.g., a crash phase where managers scramble to lower portfolio risk) where the dynamics are better described by shock waves.

A parting thought. Equilibrium theorists rightly point out that good theoretical models should conform to basic theoretical principles as well as to the empirical facts. Some theoretical discipline is needed to replace the standard assumptions of optimality and equilibrium. Evolutionary models include just such a discipline: evolvability. In a financial market model one asks, e.g., can statistical arbitrageurs invade? With this discipline in mind I am optimistic that evolutionary models have a bright future in finance.

10. Bibliography

Anderson, S.P., Goeree, J.K. and C.A. Holt (1998) Rent seeking with bounded rationality: an analysis of the all-pay auction. *Journal of Political Economy.* 106(4): 828-853.

Arthur, W.B. (1989) A Nash-Discovering Classifier System for Finite-Action Games. *manuscript.* Santa Fe Institute.

Arthur, W.B. (1994) Inductive Reasoning and Bounded Rationality: The "El Farol" Problem. *American Economic Review: American Economic Association Papers and Proceedings.* 84 (May): 406-411.

Arthur, B., Ermol'ev, Y. and Y. Kanioskii (1983) A Generalized Urn Problem and Applications. *Cybernetica.* 19: 61-71.

Arrow, K. and F. Hahn (1971) *General Competitive Analysis.* Holden-Day, San Francisco.

Barkow, J., Cosmides, L. and J. Tooby (1992) *The Adapted Mind: Evolutionary Psychology and the Generation of Culture.* Oxford University Press, New York.

Basov, S. (1999) Social Learning and Stochastic Decision Making. *preliminary draft manuscript.* Boston University.

Bell, A.M., Sethares, W.A. and J.A. Bucklew (1999) Coordination Failure as a Source of Congestion in Information Networks. *manuscript NASA Ames Research Center.*

Benaim, M. and M. Hirsch (1999) Learning Processes, Mixed Equilibria and Dynamical Systems Arising from Fictitious Play in Perturbed Games. *Games and Economic Behavior.* 29: 36-72.

Blume, L. and D. Easley (1992) Wealth Dynamics and the Market Selection Hypthesis, *Journal of Economic Theory*, 58(1): 9-41.

Camerer, C. (1999) Behavioral Game Theory. *Draft Book Manuscript.* Caltech Division of HSS.

Cason, T.N. and D. Friedman (1999) Learning in a Laboratory Market with Random Supply and Demand. *Experimental Economics.* 2: 77-98.

Challet, D., and Y.-C. Zhang (1997) On the Minority Game: Analytical and Numerical Studies. *Physica A*. 256: 514-532.

Chen, H.C., Friedman, J.W. and J.F. Thisse (1997) Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach. *Games and Economic Behavior*. 18(1): 32-54.

Cheung Y.W. and D. Friedman (1997) Individual Learning in Games: Some Laboratory Results. *Games and Economic Behavior.* (April) 19(1): 46-76.

Conlisk, J. (1980) Costly Optimizers versus Cheap Imitators. *Journal of Economic Behavior and Organization*. 1: 275-293.

Cornell, B. and R. Roll (1981) Strategies for Pairwise Competitions in Markets and Organizations. *Bell Journal*. 12: 201-213.

Cressman, R. (1992) *The Stability Concept of EvolutionaryGame Theory : A Dynamic Approach*. Springer-Verlag, Berlin; New York.

Cressman, R. (1995) Evolutionary Game Theory with Two Groups of Individuals. *Games and Economic Behavior.* 11: 237-253.

Duesenberry, J. (1949). *Income, Saving, and the Theory of Consumer Behavior*. Harvard University Press, Cambridge.

Ellison, G. (1993) Learning, Local Interaction and Coordination, *Econometrica* 61(5): 1047-1071.

Eshel, I. (1983) Evolutionary and Continuous Stability. *Journal of Theoretical Biology*. 103: 99-111.

Frank, Robert H. (1985) *Choosing the Right Pond*. Oxford University Press, New York.

Friedman, D. (1991) Evolutionary Games in Economics. *Econometrica*. 59: 637-666.

Friedman, D (1998) On Economic Applications of Evolutionary Game Theory. *Journal of Evolutionary Economics.* 8(1): 15-43.

Friedman, D. and J. Ostroy (1995) Competitivity in Auction Markets: An Experimental and Theoretical Investigation. *The Economic Journal*. 105 (January): 22-53.

Friedman, D. and J. Yellin (1997) Evolving Landscapes for Population Games. *draft manuscript*. University of California, Santa Cruz.

Friedman, D. and J. Yellin (2000) Dynamics of Conspicuous Consumption. *draft manuscript*. University of California, Santa Cruz.

Fudenberg, D. and D.K. Levine (1998) *The Theory of Learning in Games*. MIT Press, Cambridge MA

Hirshleifer J. and J.G. Riley (1994) *The Analytics of Uncertainty and Information*. Cambridge University Press, Cambridge.

Hofbauer, J. and K. Sigmund (1988) *The Theory of Evolution and Dynamical Systems*. Cambridge University Press, Cambridge.

Kandori, M., Mailath, G.J. and R. Rob (1993) Learning, Mutation, and Long Run Equilibria in Games, *Econometrica*. 61(1), 29-56.

Kaufman, Stuart A.(1993) *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York.

Kivinen, J. and M. Warmuth (1995) Additive versus Exponential Gradient Updates for Linear Prediction. *Proceedings of the 27th Annual ACM Sumposium on Theory of Computing*. ACM Press, New York; 209-218.

Kollman K., Miller, J. and S.E. Page (2000) Political Institutions and Sorting in a Tiebout Model. *American Economic Review* (forthcoming)

Lam, Lui (1997) *Introduction to Nonlinear Physics* (esp Chapter 15). Springer, New York.

Ljungqvist, L. and H. Uhlig (2000) Tax Policy and Aggregate Demand Management Under Catching Up with the Joneses. *American Economic Review*. 90(3) (June): 356-366.

Marcet, A. and T. Sargent (1989) Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory* 48: 337-368.

Maynard Smith, J. (1982) *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

Maynard Smith, J. and G. Price (1973). The Logic of Animal Conflict. *Nature*. 246: 15-18.

McKelvey, R.D. and T.R. Palfrey (1995) Quantal Response Equilibria for Normal Form Games. Games and Economic Behavior.10(1): 6-38.

Nelson, R. and S. Winter (1982) *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge MA.

North, D.C. (1990) Institutions, Institutional Change, and Economic Performance. Cambridge University Press, Cambridge ; New York.

Rae, John, 1834. *Statement of Some New Principles on the Subject of Political Economy*. Hilliard Gray and Company, Boston.

Rustichini, A., Satterthwaite, M. and S.Williams (1994) Convergence to Efficiency in a Simple Market with Incomplete Information.

Econometrica. 62: 1041-1063.

Samuelson, L. (1997) *Evolutionary Games and Equilibrium Selection.* MIT Press. Cambridge, Mass.; London.

Samuelson, L. and J. Zhang (1992) Evolutionary Stability in Asymmetric Games. *Journal of Economic Theory*. 57(2): 363-391.

Selten, R. (1991) Evolution, Learning and Economic Behavior. *Games and Economic Behavior*. 3: 3-24.

Shubik, Martin (1982) *Game Theory in the Social Sciences: Concepts and Solutions*. MIT Press. Cambridge, Mass.; London.

Sonnenschein, H. (1982) Price Dynamics Based on the Adjustment of Firms. *American Economic Review*. 72(5): 1088-1096

Stahl, D.O. (1999) Evidence-Based Rule Learning in Symmetric Normal Form Games. *International Journal of Game Theory*. 28: 111-130

Taylor, P.D. and L.B. Jonker (1978) Evolutionary Stable Strategies and Game Dynamics. *Mathematical Biosciences*. 40: 145-156.

Veblen, Thorstein (1899) *The Theory of the Leisure Class*. MacMillan Co., London.

Vega-Redondo, F. (1996) *Evolution, Games and Economic Behavior*. Oxford University Press, New York.

Weibull, J. (1995) *Evolutionary Game Theory*, MIT Press, Cambridge MA.

Wright, Sewall (1949) Adaptation and Selection. (L. Jepsen, G.G. Simpson, and E. Mayr eds.) *Genetics, Paleontology, and Evolu-*

tion. Princeton University Press. Princeton, N.J.

Young, P. (1993) The Evolution of Conventions. *Econometrica.* 61(1): 57-84.

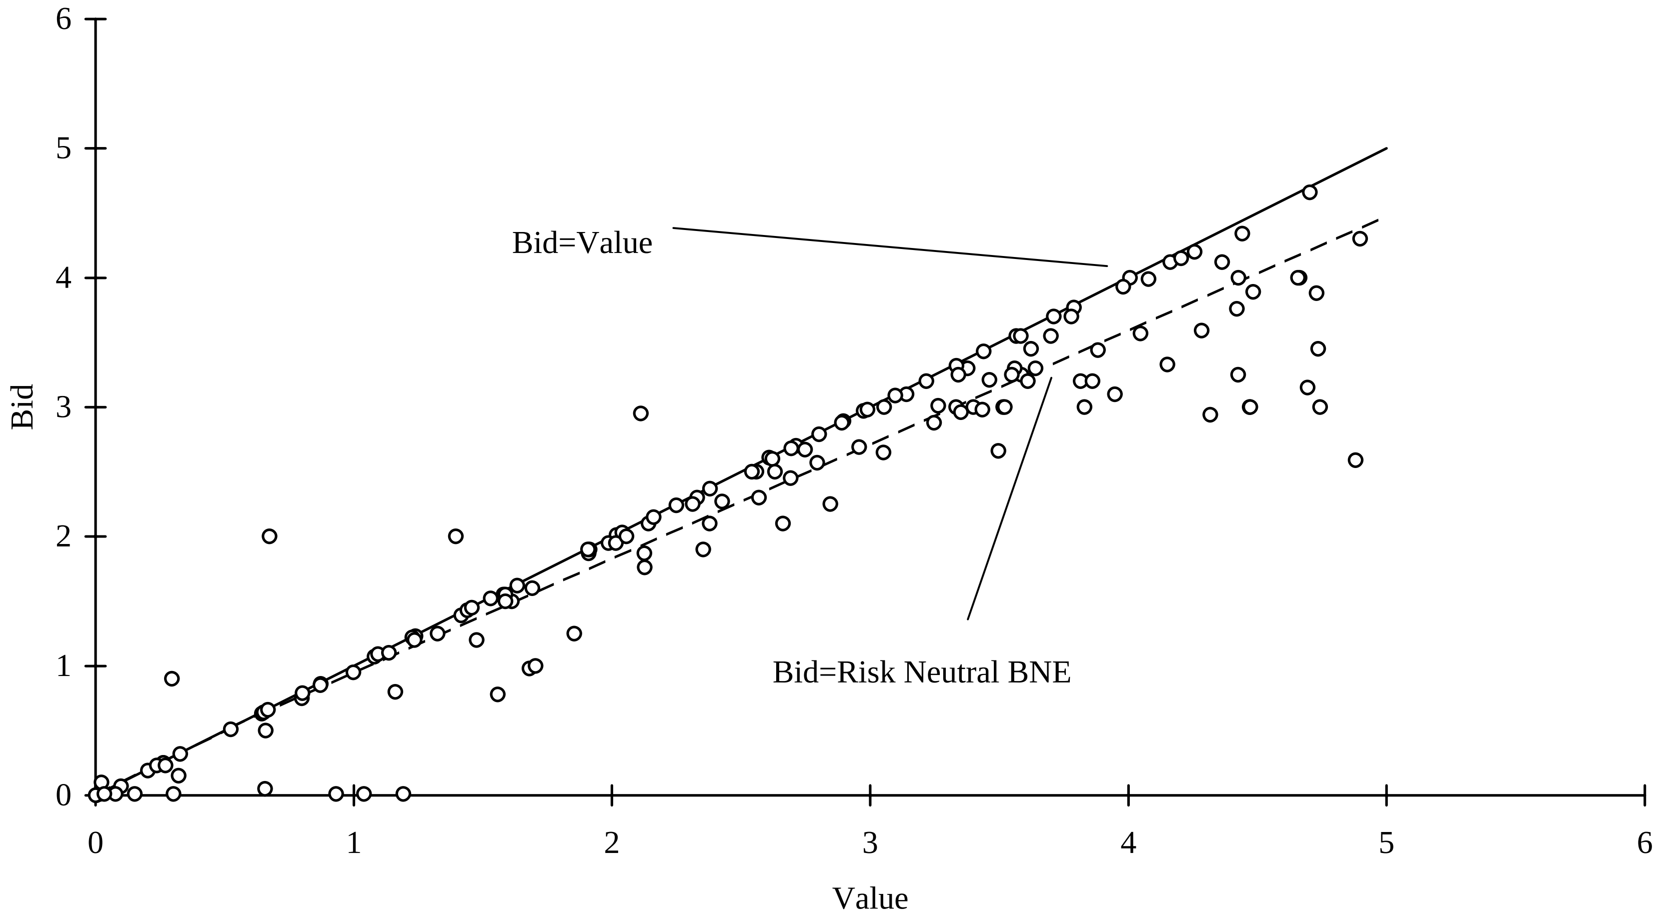

Bid
6
5
4
3
2
1
0
0
1
2
3
4
5
6
Value
Bid=Value
Bid=Risk Neutral BNE

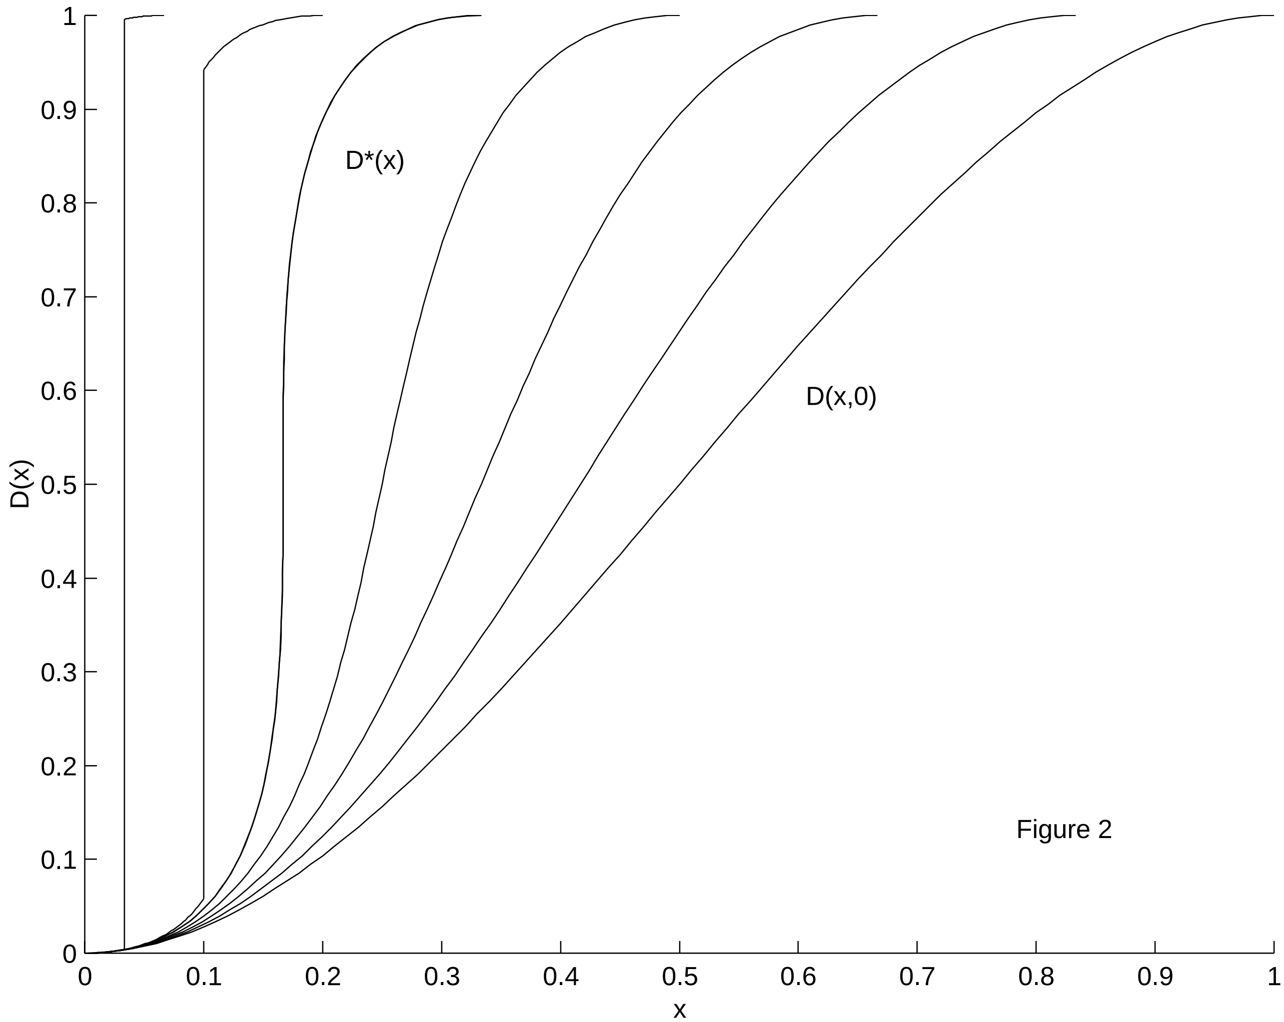

Figure 2