

Evolutionary economics goes mainstream: A review of the theory of learning in games*

Daniel Friedman

Economics Department, University of California at Santa Cruz, 212 Social Sciences I,
Santa Cruz, CA 95064, USA (e-mail:dan@cash.ucsc.edu)

Abstract. The *Theory of Learning in Games* by Fudenberg and Levine surveys a key branch of evolutionary economics from a mainstream perspective. Its publication provides an opportunity to reassess the prospects and goals for evolutionary economics.

Key words: Learning in games – Evolutionary games – Evolutionary economics

JEL-classification: A12; B41; C73

1 Introduction

Evolutionary economics in recent decades has defined itself as a radical alternative to mainstream economics. The mainstream studies simple interactions of unboundedly clever agents, and assumes that the agents instantaneously achieve mutual consistency, as in competitive equilibrium or Nash equilibrium. The intent is to illuminate the fundamental role of tastes and technology in determining economic outcomes, and to reveal unsuspected consequences of policies that alter private opportunities and incentives (e.g., Mas-Colell, Whinston and Green, 1995).

By contrast, evolutionary economics emphasizes the adaptation processes of simple people interacting through specific economic institutions. It

* My thanks to JEE Editor Steven Klepper for inviting an article reviewing Fudenberg and Levine's new book in "a larger context of work in learning, evolution, and games." Some of the material below draws on "Adaptive Process Economics," work in progress by Axel Leijonhufvud, myself, and Peter Howitt. I am also indebted to Donald Wittman for useful comments.

does not take mutual consistency for granted; equilibrium may (or may not) be achieved over time. The intent is to illuminate the role of time and circumstance and to reveal unsuspected dynamical consequences of economic institutions (e.g., Boulding, 1991; Dosi, 1991). A continental divide thus seems to separate evolutionary economics from the mainstream.

This intellectual landscape is permanently altered with the publication of *The Theory of Learning in Games*, by Drew Fudenberg and David K. Levine (MIT Press, June 1998, xvi + 276pp). Both authors are 1981 MIT PhDs, avidly courted by top US economics departments to steer their graduate programs. Levine is currently at UCLA and recently was associate editor and coeditor of *Economic Theory* and the *Journal of Economic Theory*, and Fudenberg is currently at Harvard and editor of *Econometrica*.

Despite their extraordinary mainstream credentials, the authors deal with the intellectual core of evolutionary economics. In evolutionary economics, the dynamical processes of primary interest are not genetic, but rather changes in individual behavior in response to experience, i.e., learning processes. And the setting of primary interest is not isolated individual choice, but rather strategic interaction, i.e., games.

The Theory of Learning in Games merges evolutionary economics with mainstream theory. It is an elegant mix of formal theory and commentary in a relatively slim volume, about half the heft of Fudenberg and Tirole's standard game theory text. The target audience appears to be mainstream theorists and applied theorists, including advanced graduate students.

After sketching the book's substance, I will list some issues the book leaves unresolved. I will close with some speculations on the future of evolutionary economics in the new landscape.

2 The book

The introduction quickly solves the puzzle of why mainstream theorists are so interested in evolutionary economics. The authors are not opposed to equilibrium theory; to the contrary, they believe it deserves a firmer foundation than that provided by the now-traditional common knowledge rationality explanation.¹ They see evolutionary economics, specifically learning dynamics, as raw material for a solid foundation for equilibrium theory.

Accordingly, the authors “downplay [an] antiequilibrium argument... focus primarily on the long-run properties of the models...[and] emphasize the convergence results.” [p. 2–3] Thus the book does not treat everything in evolution and learning, but rather just the results that bear most closely on equilibrium theory.

¹ This explanation supposes that everyone is rational, knows that everyone else is rational and that they know that they know it, etc., and argues that only Nash equilibrium behavior (or some close variant of Nash) can then arise. Fudenberg and Levine cite three problems for this sort of explanation: (a) multiple equilibria, (b) the fragility of the conclusion to slight relaxation of the common knowledge assumption, and (c) unsupportive laboratory data.

Narrow though their focus might seem to some evolutionary economists, it still demands that the authors assemble a lot of recent (and some not so recent) theoretical literature. They succeed admirably in weaving together diverse strands of theory and making them accessible to well prepared graduate students and applied theorists. The mathematical tools are also quite diverse, ranging from the classic differential geometry of Liouville's theorem and asymptotic stability in dynamical systems, to Markov chains and multi-armed bandit problems. The book's seven mathematical appendices briefly but expertly summarize the tools and point to the most useful primary and secondary sources.

The first three chapters treat the most venerable models – Cournot, fictitious play and replicator dynamics – and their interconnections. Chapter 1 sets notation and briskly covers the Cournot (or best response) model. Even in this introduction most readers will find one or two unfamiliar results. For instance, I hadn't previously noticed that every 2×2 bimatrix game with a unique NE in mixed strategies is best response equivalent to a zero sum game. Chapter 2 treats fictitious play (i.e., players best respond to the empirical distribution of play observed so far) and briskly summarizes the surprisingly large literature, mainly from the 1950s and the 1990s. Again, most readers will find several new perspectives and results. For instance, under a logarithmic time transformation, discrete time fictitious play looks almost like an inertial version of the Cournot dynamics in continuous time, but there are some asymptotic discrepancies.

Chapter 3 is the longest in the book. It surveys the vast (and now relatively mature) evolutionary games literature spawned in the late 1970s by Taylor and Jonker's replicator dynamic and, a few years earlier, by the static equilibrium concept of Maynard Smith and Price called Evolutionarily Stable Strategy (ESS). Compared to Weibull (1995), Fudenberg and Levine's treatment is a bit cursory, but it offers some useful new perspectives. The chapter also has some new applications to extensive form games and some recent results connecting replicator-like dynamics to individual learning processes. I agree with the authors' conclusion that the replicator dynamic is too special to take literally but is a convenient first cut for a large class of empirically sensible specifications.

Chapters 4 and 5 treat several recent stochastic techniques. The presentation begins with smoothed fictitious play, and then connects it to stochastic approximation (introduced to microeconomists in the early 1980s by Arthur, Ermolev and Kaniovski, and to macroeconomists in the late 1980s by Marcat and Sargent), to logit estimation (p. 119, via the classical entropy function!), and to probability matching behavior and reinforcement learning. Chapter 5 centers on the stochastic adjustment models in the style of Kandori, Mailath and Rob (1993). Fudenberg and Levine offer a cookbook on the construction of such models and appraise their strengths – e.g., they often give sharp and persuasive predictions on equilibrium selection – and weaknesses – e.g., the relevant transition times make glacial look short run and the results depend on the width of the basins of attraction (radius and coradius) but not on the depth.

Chapters 6 and 7 treat learning in extensive form games. The analysis is delicate because players' experience is typically concentrated on a small

subset of the terminal nodes; players may not achieve Nash equilibrium because they never learn to share beliefs about what happens at rare terminal nodes. Even by the book's high standards, the exposition of this rather difficult material is especially clean, perhaps because most of the seminal papers were written by Fudenberg, Kreps and Levine.

Chapter 8 considers several distinct ways of modelling more sophisticated learning processes. The first is strict Bayesian. With appropriate conditions on prior beliefs, play converges to a Nash equilibrium. Unfortunately the conditions often are impossible to satisfy. The authors recommend "...learning procedures that are robust in the sense that, unlike Bayesian learning, they continue to perform well even if none of the alternatives viewed as possible turn out to be true." [p232] To that end, the authors turn to the computer science literature on evaluating "experts" or strategy choice algorithms. They use the concepts of universal consistency (a worst-case performance criterion) and calibration (roughly, statistically consistent beliefs) to extend earlier results to history-contingent strategies, and to prove asymptotic convergence to correlated (but not necessarily Nash) equilibrium.

3 Remaining controversies

Fudenberg and Levine's clear exposition will end some lingering confusions and controversies, but it will only sharpen other points of contention. I have in mind three ongoing debates.

Level of Rationality

In numerous articles, Al Roth and various coauthors have argued that, until thoroughly falsified, learning models should assume only simple reinforcement learning, the lowest possible level of rationality (e.g., Erev and Roth, 1998). Several other economists, myself included, have argued for explicit representation of beliefs in models of human learning (e.g., Cheung and Friedman, 1997). Most mainstream economists are comfortable with learning models only if they are strictly Bayesian. See Selten (1991) for a hilarious early description of the debate.

Fudenberg and Levine clearly favor models that have a high but bounded level of rationality, but for them apparently it is simply a matter of taste. I believe that proper resolution of the debate is mainly an empirical matter: which models can best capture regularities in human interactions over a wide variety of environments? Evidence accumulates but is not yet decisive.²

² Camerer (1999) summarizes recent laboratory evidence. He distinguishes reinforcement from belief learning models by a single parameter that measures the impact of hypothetical payoffs to strategies not chosen relative to the actual payoff from the chosen strategy. Estimates of this parameter vary considerably across environments but usually fall between 0 (strict reinforcement learning) and 1.0 (strict belief learning). Friedman and Aoki (1992, p. 265ff.) show that asset price bubbles can arise from such discounting of hypotheticals.

Sources of errors

Empirical learning models need some sort of error structure in order to fit data. The error term can simply be tacked on, as in estimating deterministic replicator dynamics (e.g., Cheung and Friedman, 1998). Preferably errors (i.e., deviations from apparent best response) are an integral part of the model, as in the belief learning model in the same paper and as in Fudenberg and Levine's chapters 4, 5 and 7. Are errors best modelled as blind mutations (or trembles or noise), unintended and unanticipated? Or as "experimentation," informational investment intended to find out what happens off the beaten track? Is the error distribution uniform or responsive to expected payoff differences? Do players correctly anticipate their own error distribution and that of other players? The book touches on these questions but doesn't try to answer them.

Empirical researchers need answers. For example, McKelvey and Palfrey's (1995) quantal response equilibrium (QRE) assumes an exogenous payoff-responsive error distribution that players correctly anticipate. 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). But QRE is a noisy rational expectations model, not a learning model per se, and an exogenous declining error amplitude at best is a first approximation to a true learning process (Chen, Friedman and Thisse, 1997). Empirical researchers have not yet formed a consensus on (a) whether the good fits are attributable entirely to the nice functional forms for the errors (usually logit, sometimes probit) or whether the equilibrium assumption actually helps; and (b) whether individual differences in error processes (or in learning rates) need to be taken into account.

Empirical answers have theoretical implications. For example, Friedman (1996) finds that the risk dominant Nash equilibrium (RDNE) in a simple laboratory coordination game attracts far less attention than predicted in the Kandori-Mailath-Rob theoretical model. That model assumes uniform error amplitudes, but in the experiment the error amplitude is much higher at RDNE than elsewhere. This finding seems to support the alternative theory of Bergin and Lipman (1995), which chapter 5 of the book lists in the bibliography but I can't find cited in the text. Clearly what we need, but do not yet have, is an empirically grounded theoretical account of human errors.

Objects of learning

Most of the book examines models of how people might learn to choose specific, uncontingent actions from a finite list of alternatives. Chapters 6 and 7 expand the objects of learning to extensive form strategies, i.e., information set-contingent actions. The last chapter considers more sophisticated choices among behavior rules, i.e., history-contingent actions. This progression towards rule learning is mirrored in the empirical literature (e.g., Stahl, 1998). Unfortunately, the set of possible behavior rules is uncountably infinite so the modeller is forced to choose arbitrarily a short list

of a priori plausible rules. How long a list should he choose? What range of simplicity to sophistication should he consider? The dilemma is acute because the success of a belief learning model depends sensitively on the arbitrarily chosen short list. The dilemma is embarrassing to those of us who have given Bayesians a hard time for the analagous difficulties they face in specifying priors.

Two other issues regarding the objects of learning (or evolution) merit a brief mention. Most recent empirical learning models use up many of their free parameters to fit first period choices, with consequences extending over all periods. As noted below, initial choice is a separate problem from how people learn from experience. Therefore it seems to me more reasonable to use fewer parameters and not to fit first period choices. The other issue is that individuals may well adapt to (or learn from) behavior in their own population as well as to (or from) behavior in another population with which they interact, but standard asymmetric evolutionary game specifications assume away such own-population effects. Fudenberg and Levine's generous cite on p7 misses the real point of my earlier work: the elegant Liouville argument against stable interior equilibria becomes irrelevant given own-population effects.

4 What next?

Fudenberg and Levine's book presents evolutionary economists with the opportunity to reassess our own goals. Some of us will endorse their goal of constructing evolutionary foundations for mainstream economics. There is a nice precedent from physics. Clausius, Kelvin and other mid-nineteenth century scientists were able to organize a large body of data on the basis of general optimization and equilibrium principles called the two (later three) Laws of Thermodynamics. Much later, Gibbs, Boltzmann, and other early 20th century scientists explained the austere Laws in terms of underlying molecular-level processes. The process theory, statistical mechanics, does not replace the classical theory. Rather, it provides rigorous but intuitive foundations and establishes the range of applicability for the classical theory.

Contemporary mainstream economic theory is similar to classical thermodynamics one hundred years ago. The optimization principle of rational choice together with mutual consistency principles (such as market-clearing or Nash equilibrium) provide austere but empirically useful explanations of large bodies of economic data. However, anomalies remain and the standard theoretical models are indeterminate in some cases (e.g., multiple Nash equilibria). Note well that mainstream equilibrium theory involves postulates as mysterious as (classical) increasing entropy – how do people make calculations that are impossible for existing computers or even for Turing machines? How do people achieve common knowledge of each others' unobservable beliefs and plans?

The foundations goal, then, is to replace these dubious postulates with simple adaptive processes of evolution and learning. Suppose for the moment that we have empirically valid models of initial behavior and of how

behavior adapts over time. Suppose also that the adaptive processes converge rapidly and reliably to equilibrium. Then the goal is achieved. The evolutionary foundation is appealing because it relies on tangible processes rather than teleological postulates of optimization and equilibrium. It connects economics to cognitive and social psychology and sister social science disciplines, because these disciplines investigate behavior and its adjustment and have always been process-oriented. Best of all, it uses a conceptually clear and empirically accessible criterion, process convergence, to establish the range of applicability for equilibrium theory.

But Fudenberg and Levine's book falls far short of reaching that goal. We do not yet have empirically solid models of initial behavior nor of how behavior adapts over time. Once we do, analysis of convergence may require results beyond those gathered by Fudenberg and Levine. Their work is but a first step.

Of course, not all evolutionary economists will endorse the foundations goal. Some will resist joining evolutionary economics with the mainstream, regarding it as a disguised hostile takeover bid. Perhaps evolutionary economics would lose its distinctive identity if it were harnessed to serve mainstream needs.

I personally endorse seeking evolutionary foundations for mainstream economics, but believe that should not be our final goal. Evolutionary economics (game theoretic and otherwise) should vigorously pursue its own agenda. As I see it, the agenda includes several distinctive items.

Institutions

Evolutionary economists recognize that adaptation (and initial behavior) unfolds within an institutional context, market or nonmarket. For example, wholesale grain markets are continuous auctions. The auction grinds out a unified price that responds almost instantly to new information about crop prospects, and that price coordinates production, storage and milling decisions at national or even global scale. The labor market has very different institutions and market performance. It is personalized "customer" market; price (wage) adjustment tends to be slow, and even in the same city and the same job classification there simultaneously can be unemployment and unfilled vacancies.

In the last decade or so, mainstream theorists have begun to study the equilibrium properties of institutions, but (as Fudenberg and Levine's book shows) they have not yet thought seriously about how institutions affect learning processes. There may be general principles underlying all human learning processes, but it seems apparent that empirical learning models must be institution-specific. The evolution of the institutions themselves is an equally important area, so far equally neglected by mainstream theorists.

Behavioral economics

Once we drop unbounded rationality as a maintained assumption, we need empirically valid and institution-sensitive models of how people make

initial choices and how they adjust those choices over time. A nascent theory of initial choice can be seen in some of the work following Kahneman and Tversky (1979). Recently this interdisciplinary effort has referred to itself as the field of behavioral economics. See Camerer (1997) and Rabin (1998) for recent progress reports.

Evolutionary economists may help the field mature by emphasizing the roles of institutions and adaptation processes. We should definitely monitor its progress, because behavioral economics is the most promising source for the theory of initial behavior that we need to complete our evolutionary models.

Adaptive processes

We must first recognize that there is not just one all purpose process (smooth fictitious play, say) that will be appropriate for all institutions and all time scales. Rather, there are several quite distinct processes, most of which Fudenberg and Levine at least mention in passing.

a. 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.

b. Endogenous market shares. Even when the agent population is constant and individual agents do not change behavior, we can have market-level adjustment as agents with less profitable behavior lose market share to agents with more profitable behavior, as in Blume and Easley (1991).

c. Adaptive learning. Systematic changes in individual behavior in response to personal experience are featured in most of Fudenberg and Levine's book.

d. Active learning or "experimentation," as featured in chapter 7. The economically most important case is R&D expenditure. See Klepper (1996) for an empirically oriented theoretical model that combines two kinds of R&D with entry and exit to explain the product life cycle.

d. Imitation and social learning. Humans are social creatures and need not wait for (or invest in) personal experience when they can observe the behavior of others. Parts of Fudenberg and Levine's chapters 3 and 5 touch on such processes.

f. Institutional evolution. Demographic and technological innovations alter the fitness of existing institutions; e.g., central districts lose retail business to suburban shopping malls, which may eventually lose share to electronic commerce.

Serious applications

The publication of *The Theory of Learning in Games* certifies that learning and evolution are now fashionable among mainstream economic theorists.

But for how long? Theorists' fads come and go with alarming speed. To secure a permanent and central place in economics, evolutionary economics has to deliver the empirical goods. That is, models incorporating evolution and learning must deal convincingly with important real world issues that elude applied mainstream theory. I have in mind several possibilities: short-run price and volume dynamics in financial asset markets; the dynamics of bank crises and currency crises; the evolution of monetary institutions and their robustness to high inflation and deflation; and evolution of markets, e.g., making sense of the varied experiences in the transition from socialism. Evolutionary economists' experience in thinking about institutions and adaptive processes gives them an important advantage over mainstream economists in dealing with such issues.

Successful pursuit of this agenda will allow evolutionary economics to construct a foundation for the mainstream while preserving its own distinctive identity. But it will do more. An empirically successful theory of adaptive processes and institutions that reduces (where appropriate) to standard equilibrium theory would amount to a friendly takeover of the mainstream. As I read it, Fudenberg and Levine's book signals that the time is ripe to prepare that takeover bid.

References

- Anderson SP, Goeree JK, Holt CA (1998) Rent seeking with bounded rationality: an analysis of the all-pay auction. *Journal of Political Economy* 106(4): 828–853
- Bergin J, Lipman BL (1995) Evolution with state-dependent mutations. *Econometrica* 64: 943–956
- Blume L, Easley D (1996) Evolution and market behavior. *Journal of Economic Theory* 58: 9–40
- Boulding KE (1991) Some thoughts on the promises, challenges and dangers of an "evolutionary perspective" in economics. *Journal of Evolutionary Economics* 1: 9–18
- Camerer C (1999) Behavioral game theory. Draft book manuscript, Caltech Division of HSS
- Camerer C (1997) Progress in behavioral game theory. *Journal of Economic Perspectives* 11: 167–188
- Chen, HC, Friedman, JW, Thisse, JF (1997) Boundedly rational Nash equilibrium: A probabilistic choice approach. *Games and Economic Behavior* 18: 32–54
- Cheung YW, Friedman D (1998) A comparison of learning and replicator dynamics using experimental data. *Journal of Economic Behavior and Organization* 35: 263–280
- Cheung YW and Friedman D (1997) Individual learning in games: Some laboratory results. *Games and Economic Behavior* 19: 46–76
- Dosi G (1991) What is evolutionary economics? *Journal of Evolutionary Economics* 1(1): 5–8
- Erev I, Roth AE (1998) On the role of reinforcement learning in experimental games: The cognitive game theory approach. In: Budesu D, Erev I, Zwick R (eds) *Games and human behavior: Essays in honor of Amnon Rapoport*. Lawrence Erlbaum Associates, Mahwah, NJ
- Friedman D (1998) On economic applications of evolutionary game theory. *Journal of Evolutionary Economics* 8: 15–43
- Friedman D (1996) Equilibrium in evolutionary games: Some experimental results. *Economic Journal* 106: 1–25

- Friedman D, Aoki M (1992) Inefficient information aggregation as a source of asset price bubbles. *Bulletin of Economic Research* 44: 251–279
- Fudenberg D, Levine DK (1998) *The theory of learning in games*. MIT Press, Cambridge, MA
- Kandori M, Mailath GJ, Rob R (1993) Learning, mutation, and long run equilibria in games. *Econometrica* 61: 29–56
- Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47: 263–291
- Klepper S (1996) Entry, exit, growth and innovation over the product life cycle. *American Economic Review* 86: 562–583
- Leijonhufvud A, Friedman D, Howitt P (1998) *Adaptive process economics*. Draft manuscript, University of Trento, Economics
- Marcel A, Sargent T (1989) Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48: 337–368
- McKelvey RD, Palfrey TR (1995) Quantal response equilibria for normal form games. *Games and Economic Behavior* 10: 6–38
- Mas-Colell A, Whinston MD, Green JR (1995) *Microeconomic theory*. Oxford University Press, NY
- Nelson R, Winter S (1982) *An evolutionary theory of economic change*. Harvard University Press, Cambridge, MA
- Rabin M (1998) Psychology and economics. *Journal of Economic Literature* 36: 11–46
- Selten R (1991) Evolution, learning and economic behavior. *Games and Economic Behavior* 3: 3–34
- Stahl DO (1998) Evidence-based rule learning in symmetric normal form games. *International Journal of Game Theory* (forthcoming)
- Weibull J (1995) *Evolutionary game theory*. MIT Press, Cambridge, MA