

# RECENT DEVELOPMENTS IN EVOLUTION AND ECONOMICS

Larry Samuelson

University of Wisconsin

Nemmers Conference, Northwestern University, May 7, 2004

# I. Introduction

Evolution has a long history in economics. For example:

“When ... the demand price is greater than the supply price ... there is at work an active force tending to increase the amount brought forward for sale. On the other hand, when ... the demand price is less than the supply price ... there is an active force tending to diminish the amount brought forward for sale. When the demand price is equal to the supply price ... it is in equilibrium. ... Such an equilibrium is stable ....” [Marshall (1920)]

“...unless the behavior of businessmen in some way or other approximated behavior that consistent with the maximization of returns, it seems unlikely that they would remain in business for for long. ... The process of ‘natural selection’ thus helps to validate the hypothesis ....”  
[Friedman (1935)]

“...there are forces at work in any actual economy that tend to drive an economy toward an equilibrium if it is not in equilibrium already.”  
[Arrow and Hahn (1971)]

More recently, evolution has appeared in the form of evolutionary game theory, organized around two questions:

- Does adaptive behavior provide a foundation for Nash equilibrium?
- Does adaptive behavior provide a means of selecting among multiple equilibria?

References include (among many others) Samuelson (1997), Sandholm (2006), Vega Redondo (1995), Weibull (1994), and Young (1997). This work will not be our focus today. Instead,

“I want to place an admittedly vague dividing line between the two bodies of research. Within the scope of our discussion, I wish to include models in which decision makers make *deliberate* decisions by applying procedures that guide their reasoning about “what” to do, and probably also about “how” to decide. In contrast, evolutionary models treat agents as automata, merely responding to changing environments, *without deliberating* about their decisions.” (Rubinstein, 1998, pp. 1-2)

The focus here will similarly be on deliberative decision making and bounded rationality.

The approach is to view evolution not as a metaphor for a behavioral adjustment process, but as the process that formed the agents whose behavior we study.

This includes our preferences, information processing procedures, decision rules, and behavior patterns in similar terms.

Evolutionary psychology provides antecedents for this work (e.g., Barkow, Cosmides and Tooby (1992)).

The danger in doing this is that it can be too easy to explain behavior by simply building this behavior into the model (cf. Postlewaite (1998)). The experience of psychology with instinctivism is a cautionary tale.

This is presumably why classical models are so popular and useful.

The hope is that a focus on evolutionary foundations can impose some discipline on this process.

The idea is that the easier it is to construct evolutionary foundations for a model of bounded rationality, and the more readily do these foundations have implications that we can evaluate, the more willing should we be to work with such a model.

If this method is to be useful, there must be more content to constructing an evolutionary foundation than, for example, to simply building behavior into the utility function.

Equivalently, there must be more content to the statement “agents have a taste for  $X$  and there is a good ‘evolutionary’ reason for this taste” than simply “agents have a taste for  $X$ ”.

It is not yet clear that this is the case. (And the experience of evolutionary psychology is mixed (Gould and Lewontin (1979)).) We can build evolutionary models of many things. Are there things we cannot justify? Do evolutionary models have implications we can evaluate? These questions are still open.



How do we proceed?

One line of work focusses on how sensible behavior might have evolved in our original environment that has counterproductive effects in our modern environment.

- How does our environment differ from the evolutionary environment?
- How do these differences give rise to “mismatch” stories - such as our taste for sweet and fatty foods.

This approach is popular in evolutionary psychology, but will not be our main concern.

A second line of work revolves around getting a better understanding of Nature's utility maximization problem.

- Why problem does Nature face in the evolutionary environment? What is the corresponding utility maximization problem?
- What “nonclassical” features might the resulting utility maximization problem exhibit? What aspects of our behavior might these features explain?

A third line of work revolves around how agency problems may lead to interesting features of our behavior.

- Computational issues may force Nature to give us simpler utility functions than her objective. These may include limitations on our abilities to do the required calculations in a timely fashion, or on Nature's ability to code us with the appropriate utility function.

- Information asymmetries may require that we be given a different utility function than Nature's objective. Nature may have information we do not (returning to the previous point) and we may have information Nature does not.

A body of work that phrases such as “the evolution of preferences” calls to mind, that we will not discuss, is *the indirect evolutionary approach*.

The primary difficulty with this approach is that it revolves around some version of an assumption that, when two players meet, they are able to observe each others' preferences.

This is crucial. It allows players with different preferences to face different distributions of opponent actions, because these preferences are observed by opponents and induce the opponents to take different actions.

If players with different preferences face the same distribution of opponent actions, then the highest payoff must accrue to the player whose preferences induce actions closest to those induced by material preferences. There is then no scope for anything other than material preferences.

Preferences thus serve as a form of commitment. Results from indirect evolutionary models are invariably findings of the form that it can be valuable to commit.

What's wrong with thinking of observable preferences as a commitment device? In any equilibrium in which the existing preferences do not match material fitnesses, it would be evolutionarily advantageous to mimic the existing preferences, while actually having material preferences.

Nature thus has an incentive to produce such hybrids. To accommodate this, we need a model in which preferences and signals about preferences evolve.

Once the latter possibility arises, all bets are off. (Robson 1990)

The study of such models remains an interesting open problem.

Work has been done on a variety of topics, including:

- Fairness
- Altruism
- Risk assessment
- Risk preferences (Robson (1992,1996))
- Envy (Samuelson (2004))
- Temptation (Samuelson and Swinkels (2004))
- (Over)confidence (Compte and Postlewaite (2005))

We will have time to examine the first three of these (with references given below).

## II. Fairness

The puzzle: Behavior in bargaining games appears to exhibit a concern for fairness.

Some possible explanations:

- People “play fair”
- People have utility functions that attach a disutility to asymmetry. (Bolton and Ockenfels (2000), Fehr and Schmidt (1999))
- There exists a “phantom future.” (Fehr and Henrich (2003)).

We pursue an alternative here, based on Nature’s attempt to conserve on reasoning resources (Samuelson (2002)).

Consider a population of agents facing a stochastic process that forms them into pairs to engage in a strategic interaction. There are three types of interaction:

- ultimatum game
- alternating-offer bargaining (Rubinstein (1982))
- a tournament.

We think of the tournament as a stand in for “all other games.” In particular, the tournament will be designed to ensure that there is constant pressure on reasoning resources. This is consistent with the Machiavellian hypothesis for human intelligence.



A strategy in the tournament is for player  $i$  to name a sequence  $\{1, 2, \dots, n_i\}$ .

Consider a sequence  $\{\rho_0, \rho_1, \rho_2, \dots\}$  with

$$\begin{aligned}\rho_0 &= \frac{1}{2} \\ \rho_{k+1} &> \rho_k \\ \rho_{k+1} - \rho_k &\leq \rho_k - \rho_{k-1}.\end{aligned}$$

Then if  $n_1 - n_2 = k$ , payoff  $\rho_k \Pi$  goes to agent 1 and  $(1 - \rho_k \Pi)$  to agent 2.

Think of this as a “smoothed” war of attrition.

Behavior is implemented by a finite automaton.

The automaton may have up to three initial states, corresponding to the three games, but these need not be distinct.

States are costly.  $C(n, 0)$  is increasing and strictly convex in  $n$ , with

$$\begin{aligned}\lim_{n \rightarrow -\infty} C(n, 0) &= 0 \\ \lim_{n \rightarrow \infty} C(n, 0) &= \infty \\ \lim_{n \rightarrow \infty} C(n+1, 0) - C(n, 0) &= \infty.\end{aligned}$$

$$\begin{aligned}C(n, \alpha) &= C(n - \alpha, 0) \\ \lim_{\alpha \rightarrow \infty} C(n, \alpha) &= 0.\end{aligned}$$

Demands in the bargaining games are subject to trembles. This is essential to studying backward induction.

We assume that there are no trembles in the tournament, but adding them affects only the notation.

We let  $\gamma$  index the probability of a tremble.

We fix  $\alpha$  and  $\gamma$ , and then examine (Nash) equilibrium automaton choices. We then let  $\gamma \rightarrow 0$  and  $\alpha \rightarrow \infty$ .


What sort of behavior does this model suggest we should expect?

We should expect to see failures of backward induction when there exist alternative Nash equilibria that can be pooled with other games.

Is this result useful? It already allows us to distinguish this model from those based on inequality aversion. Consider the ultimatum and best shot games.

Inequality aversion suggests that offers rejected in the ultimatum game should also be rejected in the best-shot game.

The evidence suggests otherwise (Prasnikar and Roth (1992), also Binmore, Ponti, McCarthy, Samuelson and Shaked (2002) and Falk, Fehr and Fischbacher (2003)).



More generally, the idea behind this model is that people use models to reason about the world.

To understand behavior, we need to understand these models.

Psychologists have some suggestions for the building blocks we may have for such models.

Making more precise models of models, and finding and examining their implications, remains an important challenge for future work.

### III. Altruism

The puzzle: people often incur costs (sometimes large costs, in extreme cases their lives) in providing benefits to others.

This behavior is in tension with our simplest economic models.

It also seems counterproductive from an evolutionary standpoint.

There are numerous semantic pitfalls. See Sober and Wilson (1998) for a discussion.

We take our model of altruism to be the prisoners' dilemma:

		2	
		<i>C</i>	<i>D</i>
1	<i>C</i>	$b - c, b - c$	$-c, b$
	<i>D</i>	$b, -c$	$0, 0$



Evolutionary biologists tend to explain altruism in terms of group selection.

Group selection has been controversial in biology. The survival of cooperation involves a race between:

- Individual selection
- Migration
- Mutation
- Group selection.

There is (now) general agreement that these forces might combine to allow cooperation to survive, but many regard this as empirically implausible.

A key requirement in any group selection model of altruism is that there be some assortativity in the matching process, so that  $C$ 's are more likely to meet  $C$ 's.

Concerns about whether group selection pressures will be strong enough to support altruism are typically arguments about whether there is enough assortativity.

Often geographical separation or isolation is taken to be the basis of group formation. We will look at an alternative here, based on relatedness.

Suppose that every family has two children. Mating is monogamous and there is no inbreeding. Transmission of behavior is sexually haploid. Parents are matched nonassortatively.

Each family has two siblings, who then play the prisoners' dilemma given above, once.

Conditional on being a  $C$ , your sibling has a gene from the same source as yours ( $C$ ) with probability  $\frac{1}{2}$  and the other parent's gene ( $C$  with probability  $p$ ) with probability  $\frac{1}{2}$ .

Conditional on being a  $D$ , your sibling has a gene from the same source as yours ( $D$ ) with probability  $\frac{1}{2}$  and the other parent's gene ( $C$  with probability  $p$ ) with probability  $\frac{1}{2}$ .

Hence,

$$p(C|C) - p(C|D) = \frac{1}{2} + \frac{1}{2}p - \frac{1}{2}p = \frac{1}{2}.$$

Expected payoffs in the prisoners' dilemma are:

$$\begin{aligned} C : & \quad p(C|C)(b - c) + (1 - p(C|C))(-c) \\ D : & \quad p(C|D)(b). \end{aligned}$$

The former is larger if

$$\frac{c}{b} \leq p(C|C) - p(C|D) = \frac{1}{2}.$$

This is a special case of Hamilton's rule (Hamilton (1964), Bergstrom (2002)). Hamilton observed that one should be willing to incur a cost  $c$  in order to confer a benefit  $b$  on another if their coefficient of relatedness,  $r$ , satisfies

$$r > \frac{c}{b}.$$

In a monogamous population without inbreeding, coefficients of relatedness are

- Sibling:  $\frac{1}{2}$
- Half Sibling:  $\frac{1}{4}$
- Cousin:  $\frac{1}{8}$ .

The difference

$$p(C|C) - p(C|D)$$

is commonly referred to as the *index of assortativity*. If matching among parents is not assortative, and reproduction is sexual haploid, then the index of assortativity is the degree of relatedness.

What does this have to do with explaining altruism?

Perhaps nothing. It explains altruism between relatives, but the puzzle is to explain altruism between nonrelatives.

We now turn to a potential informational constraint. How do we identify our degree of relatedness?

- One possibility is physiological. Some animals identify family members by smell.
- Another possibility is behavioral. We may assume we are related to those who act like they are related to us.

If this is the case, what behavior would we expect? That altruism is more commonly exhibited in when people find themselves in circumstances typical of family members.

What evidence do we have?

- Sexual attractiveness of seeming siblings
- Bonding and altruism
- Similarity and altruism

At this point, the prediction is vague and the evidence anecdotal. The challenge is to make them precise.



## IV. Risk Assessment

The puzzle: people appear to be inept at comparing and evaluating risks.

This is true at a personal level. Many people feel safer driving than flying. Some occupations must enforce personal risk-control measures.

It is true at a collective level. There appears to be no coherence to our risk management policies.

We explore here a simple model of why Nature may not evaluate all risks equally. (See Bergstrom (2004), Curry (2001), Gillespie (1973), Houston and McNarmara (1999), and Robson (1996).)

For our first model, we consider a finite initial population consisting of different types.

Let a type  $i$  consist of a vector  $(p_1(i), \dots, p_n(i))$  of probabilities, where  $p_i$  is the probability of having  $i$  offspring. (When zero offspring is a possibility, we get analogous results, conditioned on nonextinction.)

The expected number of offspring is given by:

$$g(i) = \sum_{j=1}^n p_j(i)j$$

Let  $z_t(i)$  be the number of agents in the population consisting of type  $i$  at time  $t$ .

Which type will come to dominate the population? This is a branching process. The basic tool is

**Proposition 1** *With probability 1,*

$$\frac{z_t(i)}{g(i)^t} \rightarrow W$$

*for some nonnegative random variable  $W$ . If extinction is impossible (i.e.,  $p_0(i) = 0$ ), then  $W$  is positive with probability 1.*

The proof exploits the law of large numbers. Notice that  $W$  must be random, since the initial outcomes will have a permanent effect on  $z_t/g^t$ .

Now let us compare two types, 1 and 2. Suppose  $g(1) > g(2)$ . With probability one, there are numbers  $w_1$  and  $w_2$  such that, for any  $\epsilon$  there is a time  $T$  such that, for all  $t > T$ , we have

$$\begin{aligned} z_t(1) &> (w_1 - \epsilon)g(1)^t \\ z_t(2) &< (w_2 - \epsilon)g(2)^t. \end{aligned}$$

But then, for large  $t$ ,

$$\frac{z_t(1)}{z_t(2)} > K \left( \frac{g(1)}{g(2)} \right)^t,$$

and hence type 1 dominates the population.

The conclusion: Nature selects the type with the highest expected number of offspring.

Now suppose that, in each periods, each agent must select a lottery over consumption. Let  $q_k$  be the probability of consumption bundle  $x_k$  and let  $p_j(x_k)$  be the probability of  $j$  offspring when realizing bundle  $x_k$ . Then Nature's objective is

$$\max_x \sum_k \sum_j q_k j p_j(x_k) = \sum_k q_k u(x_k)$$

for some “utility” function  $u$ . Risk attitudes over consumption arise out of the maximization of expected offspring.

Now for a variation on this model, suppose that the environment can be one of two states,  $H$  or  $L$ . Each period, a state is drawn independently, with each state having probability  $\frac{1}{2}$ .

There are two types of agents. A type one agent is sensitive to the state, producing 4 offspring in state  $H$  and 0 offspring in state  $L$ . A type two agent produces 3 offspring with probability  $\frac{1}{2}$  and 0 with probability  $\frac{1}{2}$ , but independently of the state.

We thus have  $g(1) > g(2)$ . Our previous analysis applies to agents of type 2. With positive probability, such agents become extinct, but this probability is less than one. If they do not become extinct, their numbers eventually become proportional (with probability 1) to  $g(2)^t$ .

In contrast, type 1 becomes extinct the first time state  $L$  occurs. Hence, either the entire population becomes extinct, or type 2 eventually dominates the population.

The implication is that aggregate uncertainty is quite different than idiosyncratic uncertainty.

This role for aggregate uncertainty does not hinge upon extinction. Again let there be two equally likely states  $H$  and  $L$  and two types of agents.

Type 1 produces 2 offspring with probability  $\frac{1}{2}$  and 1 offspring with probability  $\frac{1}{2}$ , independently of the current state. Then, by our previous results,  $z_t(1)$  is eventually proportional to  $\left(\frac{3}{2}\right)^t$ .

Type 2 produces 2 offspring in state  $H$  and 1 in state  $L$ . Then we have  $z_t(2) = 2^{n(t)}$ , where  $n(t)$  is the (random) number of times state  $H$  has been drawn in the periods  $0, \dots, T - 1$ .

As  $t$  gets large,  $\frac{n(t)}{t}$  gets very close to  $\frac{1}{2}$ . Then we have

$$\frac{1}{t} \ln z_t(2) = \frac{n(t)}{t} \ln 2 \approx \frac{1}{2} \ln 2 = \ln(\sqrt{2}).$$

Because  $\frac{3}{2} > \sqrt{2}$ , type 1 dominates the population.



More generally, let there be a finite number of states of the environment  $s = 1, \dots, S$ , drawn independently in each period with probabilities  $\rho_1, \dots, \rho_S$ .

Let  $g_s(i)$  be the expected number of offspring of type  $i$  in state  $s$ . Then the basic tool is that, with probability one,

$$\frac{z_T(i)}{\prod_{t=0}^{T-1} g_{s(t)}(i)} \rightarrow W,$$

where  $s(t)$  is the realized state in period  $t$  and  $W$  is a (nondegenerate) random variable, equal to zero (with probability 1) if and only if extinction occurs.

We then have, approximately, for large  $T$ ,

$$z_T(i) = \prod_{t=0}^{T-1} g_{s(t)}(i)W,$$

or, neglecting the constant  $W$ ,

$$\frac{1}{T}z_T(i) = \frac{1}{T} \sum_{t=0}^{T-1} \ln g_{s(t)}(i) = \sum_{s=1}^S \frac{n_s(T)}{T} \ln g_s(i),$$

where  $N(s)$  is the number of times state  $s$  has occurred.

From this, and another law-of-large-numbers argument, we get, for sufficiently large  $T$ ,

$$\frac{1}{T} z_T(i) = \sum_{s=1}^S \rho(s) \ln g_s(i) \equiv \rho(i).$$

From this, we conclude

- If there exists a type  $\rho(i) > 1$ , then the population avoids extinction with probability one, and conditional on nonextinction, the type with the largest value of  $\rho(\cdot)$  dominates the population.

The conclusion is the idiosyncratic and aggregate uncertainty are treated differently.

One first calculates, conditional on each state in the aggregate uncertainty, the expected fitness of each type.

One then calculates takes the expectation, over the aggregate uncertainty, of the logs of these fitnesses.

Conditional on nonextinction, the highest such expectation eventually dominates the population.

In the absence of aggregate uncertainty, the criterion is to maximize the expected number of offspring. If there is only aggregate uncertainty, the criterion is to maximize the expectation of the log of the number of offspring.

This makes the agent look more risk averse over aggregate than idiosyncratic uncertainty.

Again let  $p_j(k)$  be the probability that consuming bundle  $x$  leads to  $j$  offspring, and let  $u(x) = \sum_j j p_j(x)$  be the expected number of offspring from bundle  $x$ . This is independent of the state and common across agents.

Now suppose that the state  $s$  is drawn, and then each individual faces a lottery over consumption bundles that is independent across individuals, given the aggregate state.

Let  $q_k(s)$  be the probability of bundle  $k$  in state  $s$ . Then evolutionary success will accrue to agents selecting the lottery maximizing

$$\sum_{s=1}^S \rho(s) \ln \left( \sum_k \sum_j j q_k(s) p_j(x_k) \right) = \sum_{s=1}^S \rho(s) \ln \left( \sum_k q_k(s) u(x_k) \right)$$

This is no longer an expected utility. In particular, the objective is not linear in the probabilities  $q_i(s)$ .

To illustrate the implications, suppose the utility function is  $u(x) = x$  and  $x \in \mathfrak{R}$ , so the agent is supposedly an expected-payoff maximizer.

Let there be two, equally likely, aggregate states. Suppose the agent has a choice between a deterministic payoff of 3 and a lottery that gives 2 in one state and 4 in the other. The former is strictly preferred. The agent is thus risk averse over aggregate uncertainty. (Note the same lottery with idiosyncratic payoffs would make the agent indifferent to the certain outcome.)

Let there be two lotteries, one giving idiosyncratic payoffs of 2 and 4, each with probability  $\frac{1}{2}$ , and one giving 2 in one state and 4 in the other. The former is strictly preferred. We can then perturb this so that the agent strictly prefers a stochastically dominated lottery.

Once again, we have a potential informational constraint.

How do we identify which risks are aggregate, and which are idiosyncratic?



If this model of risk attitudes is relevant, what behavior would we expect?  
That people tend to be more averse to aggregate than to idiosyncratic risk.

What evidence do we have?

- Psychologists have studied assessment of and reactions to risk.
- A consistent finding is that people are especially fearful of risky situations in which they perceive themselves to have no control. Controlled risk is more acceptable. This is often offered to explain the cars vs. planes distinction.
- Could control be a proxy for aggregate risk?

At this point, the prediction is again vague and the evidence anecdotal.  
The challenge is to make them precise.

## IX. Conclusion

We are interested in economic models that may allow more freedom in specifying preferences and describing preferences than classical models.

The difficulty is that, without some discipline, such models allow us to rationalize anything.

Evolutionary models hold the promise of providing such discipline. It remains to be seen whether this promise will be realized.