

Evolution and Economics: A Two Way Street

Robert Ariel
Department of Economics and Finance
Baruch College

Evolution and Economics: A Two Way Street

I. Introduction

When I was starting my undergraduate career concentrating in the natural sciences, the work of Newton, Darwin, Maxwell and Einstein was held up as the archetype of how science ought to be done. By implication, other methodologies, such as those of the social sciences, were somehow of a lesser order. Later, as I expanded my studies into the social sciences and eventually came to teach economics and finance, I found that the methods of the social sciences although different had their own internal logic. Perhaps more significantly, I became more convinced of the unity of knowledge, which is a prerequisite for the borrowing of results and methods from one academic discipline for use in another.

This chapter is about one such kind of borrowing, namely between economics and evolutionary theory. The theme of this chapter is simply stated: There are at least two notions which are employed in economics and which may profitably be used to explicate some features of evolutionary theory. In turn, evolutionary theory can cast light on some features of economic behavior.

First, both economics and evolution employ the notion of optimization, whether optimization of profit of a business, or optimization of a species for its environmental niche. However, the notion of ‘optimum’ is by no means a simple one, and some distinctions between different kinds of optimization which are employed in economics find parallels in, and can be used to clarify similar notions in evolutionary theory.

Second, economics recognizes two methods of analysis, the Prescriptive, and the Descriptive ¹.

The Prescriptive mode seeks to give advice on how economic agents can best achieve their goals, for example, how a business can maximize its profits. To this end, economics has developed a large technical apparatus which enables it to say that one method of achieving a goal is more effective than another. In an example which will be described below, an experiment is described in which this same technical apparatus from economics is applied to a biological behavior (specifically, the foraging behavior of bumblebees) in order to test whether bumblebee foraging is, in fact, ‘optimal’, that is, whether evolution has performed as predicted to optimize bumblebee foraging behavior. Thus economic methodology is used to test evolutionary predictions.

The Descriptive mode of economic analysis takes actually observed behaviors (e.g. the fact that people use money rather than rely on barter) and seeks to find explanations for these behaviors. Now the notion of ‘explanation’ is notoriously slippery, if for no other reason than (as anyone who has ever been

¹ This distinction is sometimes characterized by other terms such as between normative and positive economics.

around a three-year old for any period of time knows) a series of ‘why?’ questions has no intrinsic end. Therefore, in much of economics, one ‘explains’ economic behavior by showing that the behavior is more convenient, or more efficient (e.g. it is more convenient to use money than to spend time searching for someone to barter with who needs exactly what you have and who has exactly what you want.) Still, ideally, for those who subscribe at least partially to a reductionist methodology of science, a most effective way to explain an economic behavior is to show that it satisfies some biological desideratum. Later in this chapter, a piece of research on the origin of interest rates will be described which seeks to do exactly this, that is, to explain an economic behavior in terms of the evolutionary ends which it serves. Put differently, the research argues that the economic agents behave as they do because evolution has ‘optimized’ them to do so.

These two themes -- the shared notion of optimization as a key feature of both economics and evolution, and the use of the prescriptive and descriptive aspects of economics to provides methods for using one discipline to test or explain features in the other -- are explored more fully in the remainder of this chapter.

II. Optimization in Economics and in Evolution

While many casual students of evolutionary theory take it as claim that natural selection acts to make species ‘better’, on the contrary the following argument if taken literally is false: In a “nature red in tooth and claw”, only the fittest survive. Therefore, the existence of a species is proof that it is fittest (“optimized”) for its environment. Its behavior (e.g. the number of offspring per brood) must be such as to guarantee that the greatest number live to reproduce themselves.

On the contrary, evolutionary biologists caution against the unwarranted assumption that existing species are somehow optimized through evolutionary competition.

But the opposite claim is also false: Species are not here by accident. The mere existence of some species says something about its fitness: No variety of sheep which is genetically predisposed to run towards wolves is likely to be found. Therefore, it seems the truth lies somewhere between the two polar opposite claims: “Species characteristics are optimal.” and “Species characteristics are random.”

Economic theory recognizes several distinct notions which may profitably be applied in explicating the sense in which an evolutionary equilibrium is ‘optimal’, since economics recognizes that market determined outcomes, although perhaps not absolutely optimal, but still may be better than random outcomes.

Consider the following two varieties of optima: the global optimum, and a local optimum. To use a visual analogy, water in a mountain lake is at the lowest level relative to the local surroundings, but compared to the full distance down the mountain, its elevation may be very high, indeed. So the pool may be a local optimum resting point for the water, but only the bottom of the mountain is the global optimum.

In economics this distinction between a local and a global optimum has proven very useful in explaining economic phenomena. For example, if one incrementally improves, e.g. a manufacturing process, the process may become more profitable ... one may move towards a local optimum. However, this local optimum may not be the global optimum i.e. the cheapest manufacturing method may not be achievable by “evolutionary steps” but may require a complete change in method. In like fashion, a species may evolve towards greater evolutionary fitness, but even after it has arrived at an optimum, there is no requirement that the optimum be a global optimum.² Thus while evolutionary theory does not guarantee that that which is, is best (i.e. is optimal), it does strongly suggest that existing species are either at, or moving towards local optima. Evolutionary theory thus positions the actually existent as not optimum, but also not as arbitrary; not necessarily the best, but not random.

III. Prescriptive and Descriptive Economics, and the Lessons For and From Evolution

The social science of economics (in its two branches of prescriptive and descriptive economics) and the biological paradigm of evolutionary theory contain at least two methodological areas of overlap:

Prescriptive Economics: One task of economics is to recommend how an economic agent can obtain the largest quantity of some desired product (e.g. some good or service) given limited resources. For example, the sole brick maker in a town wants to maximize his happiness from his job. The inputs to the process are clay and his labor. Should he seek to make the largest number of bricks possible? Or will he be better off making fewer bricks, selling them for a higher price due to their scarcity, and moreover enjoy the extra leisure? Given the environment within which the business operates, economics can make recommendations on how best to achieve the desired outcome. Sometimes achieving the highest output

² In illustration of how the notion of local optimum can help to explicate both economic and biological phenomena, consider the following examples drawn from these two fields: In manufacturing, an economic local optimum was attained by a Soviet chandelier factory which was compensated based on the total tonnage of lighting fixtures produced. The factory determined that the local optimum was attained by building few but big fixtures which, unfortunately, weighted so much that they pulled down the walls and ceilings to which they were attached. Perhaps an analogous kind of perverse “compensation” is responsible for some of the ‘expensive’ ostentatious male mating displays such as large nests, heavy antlers or excessive plumage. Whatever benefits of ‘health’ may be displayed by a vigorous male in building a large nest, one can not help think that the energy could be put to better use (although, of course, one can never know for sure.) Still the notion of a local optimum provides a ready model for these seemingly odd treadmills of seemingly useless effort that some species seem to be on. It also provides a model for the long term persistence of some highly improbable species; in our terminology these may be a local optima and may persist until another species occupies a ‘lower’ local optimum, much as the Soviet chandelier factory persisted in manufacturing gargantuan light fixtures. Students of economics are already familiar with the seeming uselessness of peacock displays. The analogy between this well-known case, renders more plausible the seemingly useless economic behaviors occasionally engaged in by both individuals and corporations.

will require cooperation among several economic agents (e.g. a brick maker and a house builder) and then economics can describe the process whereby the overall benefits are divided between these two.

Descriptive Economics: Most commonly brick makers will not come to economists before starting their businesses. Instead, economists will encounter ongoing economic organizations, and then seek to explain why the organization has the structure and method of operation it does, by assuming that the organization is attempting to optimize some quantity (e.g. profits from a business.) While this process may be fairly straightforward for some sorts of business organizations, for others (such as consortiums or more complicated structures) the reason for the form of organization may be harder to discern. In addition, a developing branch of economics seeks to explain the structure and operation of organizations which are not traditionally thought of as falling within the purview of economics, such as the economic benefits of family groupings as a form of social organization.

Both the Prescriptive and Descriptive methods of economics offer parallels to, and shed light on, issues in evolutionary theory.

Evolutionary Theory Can Assist Prescriptive Economics: The starting point for prescriptive economic analysis is the assumption that economic agents want something (e.g. both money and leisure in the case of our brick maker.) Presumably these are desired in part because money can be converted into food and shelter, while leisure can be used for rest, for exercise of curiosity, and for socializing with others. Still, this is only a partial answer: Why should the brick maker want to socialize with others? How much income is he willing to sacrifice in order to obtain leisure time? While individuals naturally differ, some broad range of trade-offs appear to characterize most people, and these, presumably, somehow reflect features of our biological makeup, which in turn have been determined by our evolutionary history. Therefore, perhaps evolutionary considerations may shed light on perennial and foundational economic questions.

Descriptive Economics Can Assist Evolutionary Theory:

By the structure of evolutionary theory, the interaction of living creatures and their environment is predicted to maximize the ‘output’ of a product which may broadly be called ‘evolutionary fitness’.

Evolutionary biologist caution against making the leap from the mere fact that a species exists and has some characteristic to the claim that the species has been “optimized” for its environment, and that the characteristic, whatever it is, is somehow ‘the best’ characteristic for that species in its environment. These biologists properly observe that chance plays a significant role in the development of specie characteristics as well.

However, realistically, the mere existence of some species does say something about its fitness: No variety of sheep which is genetically predisposed to run towards wolves is likely to be found. Therefore, it seems the truth lies somewhere between the two claims: “Species characteristics are optimal.” and “Species characteristics are random.”

Indeed, in those few cases when it has proven possible to see how quickly a colony of some species will modify its characteristics in order to increase its survival fitness, the results suggest that adaptation in the direction of improved survivability and reproductive success can occur quite rapidly. The well-known adaptability of micro-organisms to hostile environments (e.g. environments ‘contaminated’ by antibiotics) is the bane of medicine, but even relatively advanced organisms can change characteristics fairly quickly. A particularly dramatic study finds that on a small and isolated island, a newly introduced plant whose seeds are wind-dispersed quickly adapts by reducing or eliminating the wind-dispersal apparatus of the seed: On a small island, wind-dispersed seeds will blow into the ocean and fail to grow. Remarkably, this adaptation in seed structure took less than five generations (less than 9 years) which is far shorter than the geological time frames over which we usually imagine evolutionary change as proceeding³.

In brief, while evolutionary theory does not guarantee that each feature of an organism is ‘optimal’, the existence of some feature bearing strongly on survival can reasonably be construed as prima facie evidence that the feature is at least ‘good enough’. Yet in many such cases, we are unable to understand exactly how this feature contributes to survival.

Economics possesses a large theoretical apparatus for analyzing complex systems and determining the optimal means of achieving some goal. In some cases, this economic apparatus can be turned on a biological system as a way of independently testing if some species characteristic is, in fact, optimal (or at least quite good.)

The following sections present examples of these two kinds of interactions between economics and evolutionary biology: The next section describes an experiment in which the theoretical apparatus of economics is used to determine if a biological behavior is optimal or not.

A later section describes a case in which economics turns to evolutionary biology to resolve a problem in explaining the choices of economic agents.

³ Martin Cody and Jacob Overton, Journal of Ecology 84, 53-64 (1996), reported in Nature, 380 (14march1996) p.103.

IV. Bumblebees: An Empirical Demonstration of Optimization

Processing and subsequently reacting to information about the environment is a nearly ubiquitous characteristic of multi-cellular organisms. (Indeed, in some cases quite sophisticated chemical and light sensing capacities have been observed in single cell organisms.) The initial step in this sequence, namely how organisms obtain input from the environment, continues to be elucidated as the physiology of sense organs or the biochemistry of non-sense organ responses yield to continue investigation.

For some simple responses, the causal chain linking the environmental stimulus and the organism's response are well understood (For example, plant phototropism in some cases seems due to photo-degradation of plant growth factors on the illuminated side of plants. Indeed, even a nematode, an organism with a rudimentary nervous system, has had its reaction to stimuli mapped as what appears to be a simple stimulus-response circuit mediated through a small number of intervening nerve cells.)

However, for even slightly more complicated organisms, the mechanism of neural processing of sensory information is far less well understood, which gap the field of cognitive science seeks to fill. Although parochially we view an understanding of human cognitive processes as the Grail, and although significant information has been gathered on human acquisition of language capabilities and particular skills, both because of its independent importance and also for the light it may shed on human cognitive processes, the study of simpler organisms has become an active area is cognitive research.⁴

Given the importance of an organism 'correctly' responding to its environment, it is natural to suppose that the cognitive process mediating the response is evolutionarily adaptive⁵. The link between mental processes and evolution is suggested in Darwin's writings⁶, and has found theoretical and empirical support⁷. It has been argued that animals display evolutionarily adaptive mental processes uniquely appropriate for their specific environments.

However, once animal investigation moves beyond the simple stimulus-response loop found in the nematode, the cognitive processes or 'rules' which an animal follows becomes harder to elucidate. For complex behaviors (e.g. hunting) we may marvel at the skill and cunning with which the behavior is executed, but it is difficult to show that this particular hunting style is more effective, or confers an evolutionary advantage, relative to other hunting styles. Clearly, in order to show that one hunting style is

⁴ See, e.g. Pearce, J.M. An Introduction to Animal Cognition (Erlbaum, London, 1987)

⁵ For simplicity I will apply the words 'cognitive' and 'mental' to both humans and to other organisms with nervous systems, although some might prefer different language for the two cases, while others might prefer that other words be applied in common to both cases. I do not intend to beg any questions (whether epistemological, theological, or methodological) by this word usage.

⁶ R.J. Richards, Darwin and the Emergence of Evolutionary Theories of Mind and Behavior (University of Chicago Press, Chicago 1987).

⁷ Staddon, J. E. R., Adaptive Behavior and Learning (Cambridge Univ. Press, 1983).

‘the best’, one needs a criterion of optimality. It is in this latter function that economics can make its contribution.

Economics can make prescriptions about what behavior is optimal for an economic agent in order for the agent to achieve its goals. Naturally, goals sometimes differ from agent to agent. However, sometimes the goals seem so inescapable, that an economic prescription to achieve the goal would seem to be an optimal prescription for all agents. For example, economics frequently assumes that all else equal, an individual would prefer a higher income over a lower income, and although this rule certainly has exceptions (e.g. individuals who have ‘enough’ wealth, or perhaps ascetics) an economic theory founded on this assumption would probably be of broad applicability.

However, any claim about goals which is so universal, obviously can not depend crucially on the particular characteristics of the economic agent. Indeed, the presuppositions must be so inescapable, that the theory should be applicable to virtually any goal-seeking agent, including, if applicable, those goal-seeking agents which are not human.

Therefore, if a human and an animal have ‘goals’ which are, in their respective contexts isomorphic, the optimizing rule which economics prescribes to the human agent should likewise be the optimizing rule for the non-human agent.

Moreover, if the animal cognitive processes have been optimized by evolutionary adaptation, the animals actual behavior should conform to the economic optimizing rule.

Finally, because animal behaviors have a virtually unlimited number of features on which scientists could focus, having a pre-existing theory from a separate field such as economics is valuable to narrow the focus on what features should be observed and measured.

Although the above discussion should apply quite generally, it is somewhat problematic to locate non-human agents whose goals can be ascertained with sufficient certainty that their implementation of the optimization rule can be experimentally tested. The difficulty arises from the fact that most organisms may have multiple overlapping goals, such as obtaining food, avoiding predators, being prepared to procreate, acting to fulfill its role in its social structure, and so on, some of which may dictate conflicting behaviors. Therefore, it is necessary to identify a model animal whose goals in a particular situation can be known with sufficient certainty to be able to determine if their behavior is consistent with their acting optimally to obtain that goal (i.e. if their cognitive rules have been evolutionarily adapted to obtain the goal optimally) .

In a very interesting piece of research summarized below, Leslie A. Real⁸ has argued that bumblebees have certain features which make them attractive as a model system for evaluating the effectiveness of their goal-seeking behavior, and has conducted a series of experiments designed to explore the rules which bumblebees follow in foraging. A summary of his argument follows:

First, foraging bumblebees have a very limited goal: The bumblebee's almost exclusive task is gathering honey and pollen as energy sources for the colony. Meanwhile, foraging itself is an energy consuming activity so the strategy bees employ in visiting flowers is a crucial determinant in the net energy contribution to the colony. Colony energy balance, in turn, is tied to colony reproductive success. Since forager bees are sterile they do not need to engage in behaviors to ensure their own reproductive success. Moreover, individual bumblebees do not communicate to other foragers about the locations of nectar-rich flowers. Thus each bee's contribution to the colony's reproductive fitness is directly tied to its own foraging 'strategy', which, since bees do not communicate, can only depend on the bee's own individual learning plus whatever foraging strategy is 'hardwired' in the bee's cognitive architecture. Therefore, a priori one would expect the hardwired foraging strategy to be subject to strong evolutionary pressure, and, within the caveats given earlier, to be 'optimized' for energy gathering.

An Economist Compares the Similarities Between Human Investor's Behavior and Bumblebees' Investment in Foraging

The problem faced by foraging bumblebees is in certain crucial respects isomorphic to the problem faced by a human investor. Much as human investment must make an investment in the hope of obtaining a greater return, so likewise bumblebees must make an 'investment' of energy in order to forage for more energy rich nectar. Just as human investors are not guaranteed success with their investments but rather face a reward versus risk trade-off (i.e. some investments offer safe returns while others offer potentially greater return but also a chance of a small or even negative return), so likewise a foraging bumblebee faces an uncertain reward of nectar. Moreover, just as a human investor's choice of investment will determine what reward-risk trade-off he faces and much as the investor may change investment if the initial choice does not have the desired reward-risk characteristics, so likewise a bumblebee's decision to invest its energy in seeking to gather nectar from some particular type or bunch of flowers also involve a reward-risk trade-off: Some kinds of flowers may provide a small amount of nectar for sure, while other types of flowers may provide variable and uncertain quantities.

⁸ Leslie A. Real, "Animal Choice Behavior and the Evolution of Cognitive Architecture", Science, 253, 981 (30 August 1991).

Given the range of human investors' tastes, it might seem that in making a choice of investments and thereby selecting some particular reward-risk trade-off, any choice could be possible. Indeed, economics is unable to tell investors which is the best investment choice for everyone. However, it is a singular accomplishment of economics that economic theory suffices to constrain the range of rational choices which investors can select. Put differently, every investment involves a reward versus risk tradeoff. Economic theory can not say that one particular tradeoff is the best. It can however, state that one set of choices is rational, while other choices are irrational.

This result, arising from the work of economists Harry Markowitz, James Tobin and Paul Samuelson (all Nobel prize winners) succeeded in converting the vague notion of a 'reward-to-risk trade-off' of investments into a rule based on concrete and measurable quantities.

Initially, the notions of reward and risk must be defined in terms of measurable quantities. For consider: Although investors know that selecting one investment from among others involves a reward versus risk trade-off, this fundamental advice begs the crucial questions: What is reward? What is risk? At first glance, an investor may take reward to be one or more of several quantities. Is reward best measured by: Steady income? High return? Is 'risk' best measured by: Variability of return? Probability of loss? Failure to perform as expected?

The power of the work of Markowitz and others is two fold:

First, it explicitly identifies the quantitative features of investments which constitute risk and reward. Starting from very general principles, they rigorously show that reward and risk can be quantified by quantities well known to all students of statistics: Reward is identified with the expected or mean return of the investment. Risk is identified with the variance of return.

Second, they examine an enormously broad range of different attitudes towards risk (indeed, for such a broad range that we would label as irrational any individual whose taste's fall outside this range, e.g. they would prefer \$1 for sure over a 90% chance of \$1 million and a 10% chance of \$0). They show that over this enormous range of attitudes towards risk, the reward-risk tradeoff can be quantified by a simple linear rule, called by economists "the Mean-Variance Investment Rule"⁹ This Rule quantifies the desirability of an investment by the quantity:

$$E(R) - K * VAR(R)$$

where R is the range of possible returns on the investment, E(R) is the expected or mean return to the investment, and VAR(R) is the variance or 'spread' of returns around the expected quantity. The

⁹. This rule is so important in investment theory it deserves to be quoted since it is likely to have been encountered by many students of economics and finance. However nothing which follows depends specifically on the exact form of the rule. What is important is that it is quantifiable and unambiguous.

constant K is a quantity which varies from investor to investor and which reflects different willingness to bear risk.

Two points bear emphasizing about this investment rule:

First, the fact that ‘risk’ is identified with the variance of returns is not an obvious result. Intuitively, individuals think of risk as the probability of loss. In fact, risk is correctly identified with the variance of returns, that is, as the ‘spread’ of returns squared.

Second, this result is very general and depends on only a limited number of assumptions. Although the presence of a variable parameter K in the equation can accommodate the self-evident fact that different investors have different aversions to risk, for any particular investor, K is just a fixed number. Hence for that investor, the investment which maximizes the simple linear equation quoted above will be the most desirable. Moreover, for every investor, certain choices are ruled out as irrational. For example, lottery tickets, which have a negative expected return, and a very high risk (and thus have a negative worth by this rule) are not rational investments!¹⁰

The Bumblebee as a Model Investor

Earlier, the arguments of Leslie Real were summarized in support of bumblebees as a model animal for investigating foraging behavior. Returning to Real’s work, he explores the optimality of foraging behavior in a series of experiments. Bumblebees when foraging expend energy in visiting and drawing nectar from flowers. Their task as foragers is to gather the highest quantity of energy-rich nectar to carry to the colony, after subtracting their own considerable energy needs to perform the foraging. Naturally, different types of flowers will vary both in terms of the amount of nectar they provide, as well as in the variability of the amount of nectar between flowers of a given type. For example, petunias (e.g.) may contain a uniform amount of one drop of nectar per flower, while honeysuckle may average two drops per flower while individual honeysuckle blossoms may vary from zero to five drops per flower.

The bumblebee thus faces an investment problem: Observations on honeybees show that a single foraging run may involve visits to some 40 flowers. If the bumblebee is foraging in a garden containing both petunias and honeysuckle, should the bumblebee optimally play it safe and visit only petunias with

¹⁰ For example, if a lottery ticket costs \$1, pays off \$1 million to the winner, and has odds of success of 1:5 million, then the expected payoff is negative since you lose \$1 for sure (the ticket cost) but on average will win \$0.20 per ticket for an expected return of -\$0.80. Clearly, this investment with a negative expected return is a poor one. However, now change the example slightly so that the odds of winning are increased to 1:1 million. Now the expected return is no longer negative but rather \$0, i.e. a break-even. However, the Mean-Variance Investment Rule still labels this a bad investment, since although its expected return is break-even, it is still very risky with a virtual certainty of experiencing a loss. In terms of the Mean-variance Rule, the first term, $E(R)$ is zero, while the risk is substantial, so $VAR(R)$ is a positive quantity and hence $E(R) - K*VAR(R)$ will be negative, thereby showing that even buying a break-even lottery ticket is inferior to buying no ticket at all.

their guaranteed one drop of nectar per flower? Or should the bumblebee take the riskier course of visiting honeysuckle with a higher average amount of two drops of nectar per flower, even though many individual blossoms will contain little or no nectar? How should the bee optimally construct the ‘foraging portfolio’ of the forty flowers per run?

The notion that bumblebees will behave ‘rationally’ and follow the same “Mean-Variance Investment Rule” which constrain rational human investment choices may superficially appear to be an unwarranted assumption. However, if a bumblebee is to behave optimally to obtain the greatest value in nectar for its energy investment, it must follow this rule. Moreover, investors -or insect foragers - which fail to obey this rule will systematically lose to other investors -or other foragers - which do follow the rule. One would therefore expect strong evolutionary pressure for this Rule to be “hard wired” into the bumblebee’s foraging strategy in order for the bee successfully to compete in its niche with other insect flower foragers.

If bees follow this rule they will act to maximize the quantity:

$$E(R) - K * VAR(R)$$

where as before $E(R)$ is the average amount of nectar per flower and $VAR(R)$ is the variance of nectar among flowers. In terms of the earlier example, when choosing between petunias or honeysuckle growing side-by-side, bumblebees should choose the type of flower which maximizes this quantity.

The last statement is an empirical claim, and as such it can be tested. However, before describing the experiment, it is significant to note that knowledge of this investment rule tells the scientist what experiment to design, namely one in which the mean and variance of nectar is systematically varied between different sets of flowers. Without some prior information concerning what quantities to focus on, experimenters would have a vast number of quantities to juggle before finding ones which are relevant. (*A priori* which of the following are important? Flower color? Distance to the nearest flower? Flower size? And so on and so on.) Without some theory (in this case, economic theory) to guide construction of the experiment in order to ask some of the right questions, the experiment is less likely to yield informative results.

The experiments to test whether bumblebees follow this investment rule used a model system employing blue and yellow artificial flowers, among which the quantity of nectar could be systematically varied.

Empirical predictions of this model include the consequence that if (e.g.) yellow flowers each contain 2 drops of nectar, while half the blue flowers contain 1 drop and half contain three drops of nectar, then bees will systematically prefer visiting yellow flowers, even though the average amount of nectar is the same in yellow and blue flowers, because blue flowers have a higher variance of nectar

dispersion, and thus would thus be less desirable. Another empirical consequence is that greater variation in nectar content among the blue flowers can be overcome by increasing the average amount of nectar in the blue flowers; at some level of higher average reward, the extra risk in blue flowers will be compensated for and bees will become indifferent between the certain nectar content of yellow flowers and the variable nectar content of blue flowers, and will cease to shun the 'risky' flowers and will visit the two classes indiscriminately. If the variability of nectar in blue flowers is increased still further. they would then be avoided in favor of the safe investment of yellow flowers until the average quantity of nectar in blue flowers is increased yet again. For any level of variability or 'risk' in nectar content of blue flowers, an increase in the average content of nectar in the blue flowers will again make them preferred to yellow flowers

In a series of experiments these predictions were tested, and were found to be well obeyed by actual bumblebee foraging behavior¹¹. In a typical bumblebee foraging run of forty flower visits, at first bumblebees visited blue and yellow flowers indiscriminately, but quickly formed estimates of the average amount of nectar per flower, as well as of the variability of nectar content (with the estimate based on the few most recent flower visited.) During the last twenty flowers of the foraging run, bumblebees concentrated on the class of flowers which theory said was preferable, either the safe yellow flowers with constant nectar content if the reward among the risky blue flowers was insufficient, or else the blue flowers if they provided more than enough compensation for the flower-to-flower variability of the reward they offered. In essence, bumblebees proved to be rational 'investors' of the time and energy resources which they expend in foraging.

Interestingly, in a parallel series of experiments performed by Real and others¹², bumblebees do display a cognitive shortcoming: They base their estimates of the average amount of nectar per flowers based on only a small number of the most recently visited flowers. This short-run, myopic estimate of the values (which bumblebees appear to share with other species, e.g. certain birds) can be compared with the number of observations which mathematical statistics tells us is the true number needed to form reasonable estimates, yielding the conclusion that bees form their estimates much too quickly. Interestingly, substantial experimental evidence shows that humans, also, overweight recent observations in forming estimates of probabilities of different outcomes¹³. It is unknown why bumblebees (or humans) possess this systematic bias. The fact that this same bias exists across such different species is

¹¹ Naturally, in different runs blue and yellow flowers were alternated as being the 'safe' or 'risky' choices. The actual quantity of nectar in flowers are described as 'drops', but the actual quantities employed are closer to a few tenths of a 'drop'. The artificial flowers were scattered randomly within a laboratory 'field'.

¹² Real, op. cit. and references therein.

itself intriguing and suggests searching for a deeper biological basis for this common error in information processing.

Implication of Bumblebees as Model Investors and the Use of Economics to Explain their Behavior

Evolutionary theory has sometimes been tarred with the criticism that the theory lacks empirical content since the phrase ‘survival of the fittest’ begs the question: “Fittest for What?” by inviting the only possible response ‘the fittest to survive.’ Although this criticism has been answered in a number of ways (e.g. by noting that the slogan ‘survival of the fittest’ is a meta-theoretical statement about theory construction, and not a component of the theory itself), this experiment with bumblebees takes on the criticism on its own territory. A separate science, economics, has independently determined a rule for optimum investment behavior. Although one must not confuse ‘fittest’ with ‘optimal’ or ‘best’, still, as argued earlier it is hard to construe the empirical content of evolutionary theory as predicting anything other than a direction of change of species characteristics, namely as predicting at least movement towards an optimum. It is therefore comforting that the actual behavior of bumblebees follow an optimization rule derived independently by the completely different discipline, namely by financial economics. Moreover, quite apart from the experimental result that bumblebees do or do not follow an optimal investment rule, the fact that it is meaningful to conduct this experiment, i.e. meaningful to examine whether the optimal rule was or was not followed, by itself constitutes evidence that evolutionary theory is not immune from experimental test because the theory is fundamentally tautological. These conditions apply here: Economic theory provides a clear rule for optimally investing in risky assets, i.e. for behavior when reward is not certain but rather is risky. Forager bumblebees have as a single task the gathering of nutrients from flowers when the reward from any given flower is not certain but rather is risky. The relevant characteristics of the two problems are sufficiently isomorphic that the solution from the one case carries over to the other case.

In sum: Bumblebees are good investors. And the fact that they can experimentally be shown to be good investors provides some weight to the claim that evolutionary adaptation will tend towards ‘optimal’ behaviors. Moreover, the empirical examination of a prediction of evolutionary theory refutes the claim that the theory is devoid of empirical content.

Finally, as an interesting sidelight, students of finance who first learn Markowitz’s Mean-Variance Investment Rule are quick to realize that many individual investors who buy a few stocks here and there wind up with portfolios which fall far short of the optimal given by Markowitz’s Rule. When I

¹³ Thaler, Richard H. (ed) Advances in Behavioral Finance (Russell Sage Foundation, New York, 1995)

describe to them the Mean-variance efficient foraging behavior of bumblebees, they are amused that the untutored investor displays less “investment” skill than an insect.

V. Economics Borrows from Evolutionary Theory

In the example of the former section, an economic theory was used to test whether the behavior evinced by an animal conforms to criteria which economic theory has determined to be optimal, since optimality (or at least something approaching it) would be expected in creatures subject to strong competitive pressure.

In this section, the an example is presented exploring the opposite order: Instead of using an economic theory to test whether an animal behavior is optimal as above, here, evolutionary theory will be used to make predictions about economic behavior. This second sort of example, is in one sense more orthodox from a reductionist perspective, since economics, being a social science and hence lower on the reductionist hierarchy ought to be able to be reduced to a ‘higher’ science such as biology. However, reductionism’s limited success in crossing the division between the natural and the social sciences by reducing the latter to the former makes any example where this may be successfully accomplished of special interest. Hence the following example in which a social phenomenon is apparently ‘reduced’ to a feature of evolutionary biology is of particular interest.

Many economic theories and their applications implicitly or explicitly make assumptions about the goals of economic agents, but the reasons for these goals and preferences are taken as exogenous to the theory. In some cases, the ground for the assumed preference is sufficiently straightforward as to generate little controversy. For example, an economic theory which is grounded on the assumption that an economic agent will prefer more money over less money is likely to generate few qualms over the theory’s foundations. However, most economic theories must make more complicated assumptions about the preferences of economic agents, and as a rule these assumed preferences are simply taken as exogenous - as an assumption of the theory - and are justified, if at all, by some heuristic argument to a greater extent than is healthy for the discipline.

Therefore, one growth branch of economics are attempts to ground preferences in more fundamental theories. One such strand of research seeks to ground preferences in theories of psychology. For example, if it is a general psychological characteristic that individuals attribute great importance to very rare occurrences with extremely important outcomes (e.g. to the very small probability of an airplane crash), then an economist may use this fact to explain why individuals buy lottery tickets when standard laws of economics clearly mitigate against these as a good investment.

However, even reliance on ‘laws’ of psychology carry an ad hoc feeling with them. To put the point bluntly, if individuals have certain preferences, one gains no explanatory power for the origin of these preferences by calling them a ‘law’ of psychology.

Therefore, a second effort to explain preferences in economics seek to ground these in more fundamental features of biology. Much as in the earlier example of bumblebee foraging, in which a particular pattern of foraging behavior could be shown from first principles to be optimal for the goal of maximizing energy gathering and thus, arguably, to be hard-wired into the biology of the bee, so likewise, if certain preferences could be shown to contribute to the evolutionary fitness of economic agents, then these preferences could likewise be justified as being grounded in the biology of the economic agent.

Example of Grounding an Economic Preference in Evolutionary Theory: The Origin of Interest Rates

Few quantities are more important or more ubiquitous in economic theories than the interest rate to be paid on the loan of money or some other good (e.g. grain in agrarian societies). In standard economic theory, interest must be paid to a lender to induce the lender to postpone the pleasure which the immediate use of the money could provide. In return, the lender receives a larger sum (both principal and interest) in the future, which can be spent to provide larger amount of pleasure in the future. Put differently, to induce individuals to postpone the immediate consumption today of \$1, one must promise them more than \$1 next year; interest is the ‘bribe’ to induce the lender to postpone immediate gratification¹⁴.

However, a substantial remaining problem is why the interest rate has the value it does. Although the daily news is full of talk about interest rate fluctuations, the history of interest rates dating back to ancient Athens show that averaging over day-to-day fluctuations, in stable societies the real interest rate

¹⁴ It should be clarified that the interest rate being referred to here is the ‘real’ interest rate, that is, the actual increase in the amount of goods that can be purchased. The interest rate one sees posted in a bank, for example, both includes this ‘real’ component, but also includes a component to compensate for the dollar’s loss of purchasing power through inflation. No rational individual will lend a dollar bill for one year if during that year inflation is expected to rob that bill of some purchasing power. Put differently, the investor will at least require the repayment of enough to compensate for the loss of purchasing power. The question at issue is why investors require the repayment of more than merely enough to compensate for loss of purchasing power. Put differently, in those eras when prices are stable, why does a positive interest rate still exist, even when there is no risk of non-repayment as in U.S. Treasury bonds? (Of course, if repayment is uncertain, then a still higher interest rate may be charged. Economists recognize three components in an interest rate: Compensation for risk of non-repayment (if applicable), compensation for loss of purchasing power through inflation, and, finally, a third component which is the actual “time value of money”. The origin and method of calculating the first two components are well

has tended to a value of roughly 3%. The fact that this same value may be found in very different societies and very different times suggest that it arises not from chance, but from some deeper cause. A hard-wired, biological reason is hinted at both by the near universality of this 3% figure, and can be reinforced by the following thought experiment:

People currently lend willingly at a 3% real annual rate, that is, will postpone spending for a year in order to be able to spend 3% more in a year. Now for the thought experiment: Suppose that the typical human life span was only one month long. Then what interest rate would you charge to lend money today if the loan will be repaid in one year to your eleventh generation grandchild? Probably you would not make such a loan, or else would charge some astronomical interest rate. Why? Because except in some abstract sense most people feel little personal attachment to their eleventh generation progeny (except in the sense of continuing a line, or a people, or some other abstract sense.) In contrast, people regularly make loans or investments which will benefit as-yet-unborn children and grandchildren.

This is the emotion-based thought experiment. What about the biology which presumably underlie these emotions?

First, a disposition to save for hard times would presumably confer a survival advantage on those who do so. In brief, the argument is as follows: Suppose that in the dim past, there were two groups or 'clans', such that all members of the Clan A are indolent and produce or gather only what is needed for immediate consumption, while the members of Clan B are genetically inclined to produce in excess of immediate consumption and 'save' the excess foodstuff or tools ('capital'). Given variability in resource availability, Clan B, which is able to achieve a higher, smoother level of consumption may have a higher sustainable rate of population growth. Indeed, in the extreme case, if the Clans are subject to a Malthusian test, it is conceivable that only the members of Clan B, who have saved against this 'rainy day' may survive. Even if Clan A is not extinguished, the higher growth rate of Clan B would lead to Clan B members coming to dominate the entire population, and thus the entire population would come to be characterized by a preference to work and save. Hence, the arguments conclude, a tendency to save would confer an evolutionary advantage, and thus would be selected for in man's evolution.

This argument could explain why people save, but not the constant 3% rate of interest people charge. A number of arguments have been advanced to try to quantify this biological insight in order to yield a quantitative interest rate prediction such as 3% per year. Consider one example ¹⁵ of this line of argument:

understood. It is this third component, the time value of money, whose origin is problematic, and which the present argument seeks to explain.

¹⁵ Rogers, A.R., Evolution of Time Preference by Natural Selection. American Economic Review, June 1994 84(3) 460-81.

Under one interpretation of evolutionary theory, survival advantage should be measured not solely by the number of one's own progeny which survive, but rather by the number of whole or partial copies of an individual's genome which are perpetuated. This interpretation has been used to explain seemingly self-sacrificing altruistic behavior by animals not only to their immediate progeny but also to genetically close relatives.

As applied to humans (and ignoring the genes shared with close relatives), an individual's entire genetic investment is in his or her children, and, once the children pass reproductive age, in grandchildren. Thereafter, the fraction of one's own genome present in an individual in succeeding generations rapidly decreases. The argument now assumes that having saved excess goods (e.g. food, tools, clothes and so on in the pre-agrarian societies in which humans evolved) confers a survival advantage both on oneself and on one's progeny if the saved goods are passed on to them or expended for their benefit. Since the genes shared with individuals in distant generations is small, somewhere between an individual consuming everything immediately on the one hand (which confers no survival advantage on those sharing one's partial genome), or saving for ten generations hence (which will confer a survival advantage on individuals with which one has few genes in common) on the other hand, there will exist an optimal method of distributing savings so that progeny (whether children, grandchildren, or beyond) receive the benefits of the saving in direct proportion to the fraction of the individual's genome they carry. (Notice, by the way, that this method of distributing a survival advantage is what natural selection would 'optimize' for if the propensity to save were a biological trait. In the entire argument, there is no suggestion that individuals consciously choose to distribute their savings in this calculated fashion.)

The working out of the argument employs the technical apparatus of economics. In essence, the argument hinges on the trade off of an individual's desire to accelerate or postpone consumption in order to maximize the individual's pleasure, against actions which maximize the group's overall evolutionary fitness.

The greater pleasure referred to arises from the right to choose the most pleasurable best time to consume some commodity, for example food. As mentioned earlier, at its most fundamental level, "interest" is a bribe paid to a lender to induce the lender to postpone immediate consumption in favor of consumption of a greater amount when the loan (plus interest) is repaid. For example, all else equal, if an individual has to choose between enjoying a large meal today, and enjoying the same meal in a year, most individuals will choose the meal today. It is only if the individual will receive more than one meal in a year (that is, receive 'interest' in the form of extra food) that the meal may be postponed. The size of the 'bribe' is determined by the consumer's (or investor's) "preferences".

As argued above, however, the individual's choice whether to consume now, or consume later, will not be neutral for group survival: Groups whose members are willing to postpone immediate consumption will receive a survival advantage. Hence the "preference" for consumption now versus consumption later, while appearing to both the individual and to economic theory to be a free choice, will, in fact, have been under strong evolutionary pressure to take on the value which optimizes evolutionary fitness, that is, which maximizes the genetic contribution the individual confers on later generations.

To derive quantitative results from these insights requires adding a few additional facts.

First, during the long human history during which willingness to postpone consumption evolved, the growth rate per generation of the human population was effectively zero. (That is, the 'population explosion' is at best a phenomenon of the last few centuries, or, of the last few millennia. This is easy to prove since a growth rate of even 1% per year would lead to an increase in population every 1000 years of 21,000% ! This was never observed in pre-history.)

Second, the limiting factor in human reproduction was the number, and lifetime fecundity, of females.

Combining these two facts, one concludes that during human pre-history, each mother on average had one daughter who reached reproductive age, and who produced on average one female offspring. Each daughter carries half the genes of the mother, so from a genetic point of view, a daughter's offspring (a grandchild) is 'worth' half as much to the mother as one of her own offspring. Now consider as the simplest case a mother making an 'investment' which will pay off exactly one generation later and whose payoff may be used for the benefit of the now-fecund daughter. The one unit of consumption which the mother defers today would presumably contribute to the mother's own reproductive success if consumed. Therefore, for the "investment" to be worthwhile for the mother to make from an evolutionary point of view, the one unit of consumption must have grown to two units before it is consumed by the daughter who shares just half of the mother's genes. To put the point differently, assuming reproductive success is tied to access to resources, then one unit of goods (e.g. food) consumed by the mother will convey the identical long term genetic success as two units consumed by a daughter at exactly the same point in their respective reproductive lives, since the daughter has half the genes of the mother.

Given these assumptions it is easy to calculate the minimum growth rate which the "investment" (i.e. the postponed consumption by the mother) must earn before the mother will make the investment. The requirement is that one unit of consumption grow to two units in one generation. The time between generations in a variety of pre-industrial societies is just under 30 years. Therefore, solving for the

interest rate which will cause the investment to double in size in 30 years yields an interest rate of about 2 ½ % per year. (Some additional refinements of the argument yields a slightly larger value.) These figures are consistent with the values actually observed.

The key point, here, is not whether this particular argument is sound or not. Rather, the argument serves to illustrate that importing premises from evolutionary theory into the apparatus of economics permits deriving answers to long-standing economic questions.

Nor is this example unique. Other cases in which economics has borrowed from evolutionary theory include:

-- an evolutionary and economic “derivation” of the Golden Rule: a proof within economic theory that if populations evolve so as to maximize the ‘utility’ (roughly, the economic correlate of happiness) of its members, then behavior according to the Golden Rule (“Do unto others etc. ...”) will be selected for by evolution within stable populations;^{16 17}

-- the very large body of work by Gary Becker , winner of the Nobel Memorial Prize in Economics, which seeks to explain a variety of human, biological behaviors, in terms of their economic consequences (e.g. family size).¹⁸

VI. Conclusion

At the outset, it was suggested that even though economics and evolution have their roots in radically different approaches to the world, cross-influences already exist, and the possibility of additional cross-borrowings should not be discounted, and, indeed, may profitably be sought.

This chapter cited two examples in which cross-influence between economics and evolutionary theory played crucial roles in arguments in biology and economics, respectively. The economic apparatus for determining optimums can sometimes be used as an independent check on whether evolution has actually optimized (or at least rendered efficient) some important feature of animal behavior, and the foraging behavior of bees was cited as an example of this interaction.

In the second example, one version of evolutionary theory (which claims that behaviors which optimize propagation of the genome will be selected for) was employed as a premise in an economic

¹⁶ Hansson, Ingmar and Charles Stuart, “Malthusian Selection of Preferences” American Economic Review (June 1990), 529-544.

¹⁷ Interestingly, for those familiar with the arguments of John Rawles in his Theory of Justice, this argument bears a striking similarity to his argument for a kind of Golden Rule-like behavior. The analogy lies in that genes do not know into which actual individuals they will be born, and thus operate behind the Veil of Ignorance which Rawles postulates for individuals.

¹⁸ For example, Becker, Gary, Human Capital (University of Chicago Press, 1964 and later editions.)

argument. This argument seeks to provide an explanation for a standing puzzle in economics, namely the origin and magnitude of one of the most important economic quantities, namely interest rates.

It is very significant that economics can draw on the results of another science, in this case, on evolutionary theory, to provide input data for economics' technical apparatus. As economics continues to probe more deeply into understanding the origins of specific economic behaviors, it will inevitably repeatedly encounter places where individual choice is a central variable. Ultimately, many of the most important economic choices are probably rooted in biology, and hence economics can be predicted increasingly to be driven to draw on the biological basis of human choice, which, in the final analysis, must be strongly influenced by evolution-driven optimization.

Finally, the fact that this cross-marriage of a natural and a social science can give rise to conclusions which neither could achieve on its own once again underlines the unity of knowledge, and provides a powerful argument against excessive narrowness of focus both in scholarship and in instruction.