Innovation and Selection in Evolutionary Models of Technology: Some Definitional Issues

Revised, November 1997

Joel Mokyr
Department of Economics
Northwestern University
2003 Sheridan Rd., Evanston IL 60208
Phone: (847)491-5693; Fax (847)491-7001
E-mail: J-MOKYR@NWU.EDU

This paper is part of my forthcoming *Neither Chance nor Necessity: Evolutionary Models and Economic History* (Princeton University Press, 1998). Prepared for the Conference on Evolutionary Models in Economics, Jan. 9-12, 1997, Oxford, England. Revised preliminary draft, not to be cited.

Introduction. Alfred Marshall's ([1890] 1930, p. xiv) famous exhortation that the Mecca of the economist lies in economic biology rather than in economic dynamics and Frank Hahn's (1991) more recent prediction that economics will return to its affinities to biology has become something of an obsession to those who want to apply evolutionary theory to certain issues in economics.[1] There is a sense that if an economic idea is to be developed in analogy with another discipline, that *each* element in that other discipline must find a counterpart in the economic model developed. Thus for example Ramstad (1994, p. 82) cites with approval similar criticisms made by John Commons and Edith Penrose, and adds with some exasperation Awhat is the 'organism' that is supposedly 'evolving' in this conception of economic evolution... there is no conception of the firm in the Schumpeterian framework either as genotypes or phenotypes. This being the case, resort to the natural selection metaphor ... entails even more serious problems than those

[1]Apart from the rather obvious question what in the world he could have meant by Aeconomic biology@ there remains the fact that he himself did little or no work to carry out this program (Thomas, 1991). While comparative statics remained an inevitable Atemporary auxiliary@ Marshall (1890, p. xv) felt that Athe central idea of economics ... must be that of living force and movement.@

already noted.@[2] In presenting my earlier work, commentators often wondered if the evolutionary theory of technological change I advocated was an analogy, a simile, a metaphor, or a purely intellectual game.

The idea of evolutionary models outside biology seems to suffer from this need to shoehorn each concept into analogous concepts in evolution. Such correspondences may be instructive and entertaining, but the usefulness of an evolutionary theory of economic change does not depend crucially on *every* element in economics being mapped onto a precise correspondence in evolutionary biology. This paper submits that the main and obvious reason why this is so is that the system Darwin was describing is itself only a special case of a much broader set of dynamic theories that are often named Darwinian in his honor, but which basically need not follow all the restrictive postulates that August Weissman and the neo-Darwinian orthodoxy placed on how the process works in living beings.

In what follows, the evolutionary model refers to a system that describes the history of a population composed of Aevolutionary units@ which I will call *manifest entities*. Entities belong to larger *groups* of identical or very similar units, which in turn may be usefully classified into even larger families. In the biological world, an entity would be a specimen, and the group to which it belongs is the species. The manifested entity is to be distinguished from what may be called the *underlying structure*, which constrains the traits of the entity but does not wholly define it. Every evolutionary system consists of those elements. In

[2]Penrose's (1952) complaint was mostly that evolution represents a selection from random mutations. In economics, on the other hand, innovation presents a conscious effort to alter the environment.

biology, the underlying structure is the genotype, while the manifested entity is the phenotype. In evolutionary epistemology the underlying structure is the knowledge basis, whereas the manifested entity is the cultural unit on which selection occurs, such as words, artistic forms, ideas, customs, and similar. They thus closely correspond to Dawkins=s (1976) idea of Amemes.@ For the economic historian, the most interesting Aunit@ to have evolved in the past is a technological one, the Atechnique.@ I will define this unit with some precision below.

Selection occurs because of superfecundity. Any Darwinian system must select, because there are more manifest entities than can be accommodated. The exact way in which this works differs, of course, from question to question. Yet in a standard Darwinian model there must be selection which gives the system direction. For there to be selection, there must be a need to select. Darwin=s famous insight was that each species produces more offspring than can be accommodated, that the reduction in numbers was not random but directed, and that this Adirectedness@ favored certain traits. In other, non-biological systems, it seems less obvious why there need to be selection. I will address this issue below.

Finally, any evolutionary system has certain dynamic properties which connect the present to the past. Inevitably, the past in some way constrains the present. The change of any variable is by and large local, that is, it is unlikely for each variable that takes a given value to change very much from period to period. There are large debates on how much Avery much@ precisely is, and whether the distribution of possible changes contains, albeit at low probability, changes that can alter the traits of the entity in a dramatic way to the point where we can allow for discontinuities, saltations, and Apunctuating@ events. Such debates are conducted both on a theoretical and an empirical level, and conclusions are likely to differ depending on the context and the nature of the entities under investigation. One question is whether evolutionary systems can be defined as Markov Chains, that is, stochastic processes based on transition probabilities that define how likely a particular variable is likely to change into something else. Markov chains have the property that each variable at time t is wholly dependent only on the value of the variable in times t-1 and the transition probabilities and not on what happened in t-2 and before, given the value in t-1. Clearly this is true for all Darwinian systems even if the number of possible states is huge. Evolutionary systems, however, tend to be asymmetrical. In biology the asymmetry is extreme enough for history to be, for all practical purposes, irreversible. When a species evolves into something else, it rarely if ever reverts back. Birds may well have originated from reptiles, but there is no expectation that they could reverse this process. This means that the matrix of transition probabilities is asymmetrical.

To see how a Darwinian system satisfying these properties could differ from a biological system of living beings, here is a rather extreme example. Consider an urn of infinite size that contains an equal number of black and white balls used to fill an urn of finite size. At each time period we draw two balls at random from the first urn, and deposit a ball in a second urn of constant size according to the following rule: if both balls have the same color, deposit a white ball, if they have different colors, deposit a black one. Then, if the second urn is at its full size remove a ball from it according to the following rule: if the original draw's first ball was white, remove a black ball, otherwise remove a white one. Let the variable we are interested in be the proportion of white balls in the second urn. This crude game is Darwinian in that Ainnovation@ is purely random, that it follows a clear-cut selection rule, and that it embodies the concepts of superfecundity and selection. It is also clear that the composition of the second urn is a Markov Chain and we can make certain (rather uninteresting) statements about it as a function of innovation, selection and superfecundity. A somewhat more interesting Darwinian model is Conway's by-now famous game of life, well-described in Dennett (1995).[3] In this game there is random innovation (at least at the beginning of the game), there is a clear-cut selection function, and there is superfecundity in that if the system becomes overcrowded, cells begin to die off. Note, however, what these Darwinian games do *not* have. There is no clear-cut distinction between genotype and phenotype; there is no obvious genetic mechanism at all and fitness in the traditional sense plays no role. There is, however, an Aunderlying structure@ namely the rules by which the games are played, which constrain but do not determine the outcome. We can also introduce post-selection non-randomness and frequency-dependence without changing the model very much. For instance, in the simple urn model we can select the balls from the first urn according to a rule depending on the composition of the second urn: Divide the first urn into two compartments, one with more black balls and the other with more white balls, then choose the balls from the first compartment if the number of white balls in the second urn exceeds n, otherwise choose them from the second compartment. Furthermore, by endogenizing n, we can impart pre-selection direction on the behavior system (which is of course not permissible in classical Weissmanian evolutionary systems) and still retain a measure of randomness requiring post-innovation selection as well.

Clearly, then, the standard Darwinian model of evolutionary biology is a very special case of a large and diverse set of possible models describing dynamic systems of this nature. What I am proposing to do below is to define sets of technological information or knowledge and characterize the way in which they

---

[3]The way the game is played is that cells in a matrix can be either Aalive@ or Adead@. The rules are determined by the number of the eight adjacent cells in a 3x3 subset that are either alive or dead. If a cell has zero or one neighbors alive, it dies of loneliness if it is alive or stays dead. If it has two live neighbors its state is unaltered. If it has three neighbours, it comes to life if it is dead and stay alive it was so; if it has four or more, it dies of congestion.

behave historically. While technological information is only a subset of all knowledge, it is the one of the most interest to the economic historian, and I will therefore use it as an example.

Techniques and Evolution.

As I have argued in past papers, the unit of analysis that makes sense for students of the economic history of technology is the technique.[4] What is a technique? Essentially a technique is a set of instructions on how to do something that involves production.[5] The set of feasible techniques which I shall call λ is based on a much larger set, the set of useful knowledge, which I shall refer to as Ω. The relation between those two sets is somewhat reminiscent of the relation between the genotype and the phenotype: the genotype constrains what the phenotype can be but does not solely determine it. This analogy between Aknowledge@ and the genome is attractive because we do not know exactly *how* knowledge produces techniques.[6] Yet the analogy also illustrates why the biological paradigm is too limiting for an explanation of cultural phenomena such as technology: in nature there is no feedback from the manifest entity to the underlying structure, that is, the genome is not affected by phenotypical change. All that happens is that phenotypes affect the chances of the entity to survive and reproduce, and through that imparts direction on the

[4]This definition is similar to the Nelson and Winter (1982) definition of the Aroutine.@ It is of course possible to construct evolutionary models that use other basic units of analysis such as firms or even societies, but these models would not be in the tradition of evolutionary epistemology. The choice of unit depends, as it always does, on the questions asked.

[5]Household production, too, falls under this definition: cooking, cleaning, child-care and so on consist of instructions leading to outcomes that enhance economic welfare. For an analysis along these lines, see Mokyr, 1996.

[6]Equivalently, we do not know exactly how phenotypes are generated in nature either. As Dawkins points out (1982, p. 22), *that* they are somehow generated and that genes play an important role in the process is incontrovertible. Lewontin (1992, p. 141) adds that the phenotype corresponding to a genotype is never completely specified and that it would be best to say that to each genotype there corresponds a characteristic distribution of phenotypes and that for

distribution of genotypes in the population. Yet why should we impose such a restriction on *all* evolutionary systems?

---

each phenotype there is more than one genotype, even in a given environment.

The mapping from the set of useful knowledge Ω to the set of feasible techniques λ must be one of the central notions in any evolutionary model of technology. It contains the entire relationship between scientific knowledge and its application, but Ω contains a great deal more than science. The mapping is just as real when techniques are based on custom, superstition, or false theoretical concepts, but still lead to practices that are in use. The phenotype itself produces Atraits@ that determine the likelihood that the technique will be selected.[7] The knowledge base of a given technique can be broad or narrow. In the extreme case the only knowledge in Ω supporting a given technique is that it actually works. In that case we can speak of a *singleton* technique.[8] The non-uniqueness of the mapping from Ω to λ is a central feature of evolutionary systems: we are Anot our genes alone@ and techniques are not just the useful knowledge in

[7]Some scholars would like to regard the *artifact* as the phenotype (Farrell, 1993; Basalla, 1988; Kauffman, 1995; Petroski, 1992). This view, however, seems less than helpful unless one was trying to develop a history of artifacts rather than techniques. Thus there is a technique called Aweaving using a flying shuttle.@ The flying shuttle could be thought off as Athe artifact@ yet it was no more the artifact than the loom itself, the Apickers@ which shot the shuttle back and forth, the cords which the weaver had to pull to shoot the shuttle across, or even the cotton cloth which was being produced. Concrete artifacts can be seen to be awkward units of selection once we realize that techniques sometimes involve instructions of Ahow to@ that may be most useful and yet involve few artifacts such as avoiding exposure to bacteria by washing one's hand before eating.

[8]A good example is the use of Cinchona bark, containing quinine, as a cure for malaria, which was adopted in Europe in the late seventeenth century without the slightest knowledge of how and why and worked against a disease which was totally misunderstood.

existence. To be sure, *some* knowledge has to exist to realize a technique, but it has to combine with other elements to lead to observable techniques. These Aother@ ingredients constitute much of the elements of the economic and social history of technological change.

Much as in the ontogeny of a living being, there is a distinction between the technique, which is a set of instructions, and the final realization of the *outcomes* defined in product characteristic space. Just as genes code for certain proteins which then become building blocks for certain traits and organs, the instructions on how to make a pair of shoes or artificial sweeteners differ conceptually from the product or the outcome itself. Selection picks, in the final analysis, on outcomes, not genes, although as long as there is a one-to-one mapping from one to other -- or something close to it -- this should make little difference. In terms of our definition above, the technique is the group to which the unit or specimen belongs. Each time the technique is used may be compared to a specimen that belongs to this species. A species therefore exists as long as the technique is used by someone somewhere. If nobody uses the technique anywhere, it can be thought of as extinct. Yet unlike the biological concept, extinction is not irreversible, unless the knowledge underlying it is lost as well and cannot be rediscovered. Such cases are rare in the history of technology.[9] The prevalence of a technique is thus measured by the number of its occurrences, and each occurrence is equivalent to an Aentity@ living Aone life.@

In some sense the technique is, to use Dawkins's term, a vehicle for the information underlying it. But this is where the analogy ends: in biology genetic information cannot exist without a living vehicle. Once

[9]Thus the knowledge that had the Su Sung Clock built in China in 1086 was literally forgotten in subsequent centuries ; after the emperor fled the Chin Tartars who invaded his capital he could not reconstruct it, perhaps because the best workman had been carried off; the memory of the device lingered on a bit, but the ability to build it was lost. In the twentieth century, on the basis of the original pictures and some assumptions about features that must have been there (but not depicted) Joseph Needham and his associates reconstructed the clock.

the last specimen of an animal species dies, the genes are just as extinct as their carrier. Knowledge, -however, can survive outside the technique and indeed, as we shall see, it can evolve quite separately from the technique as defined. Moreover, whereas the techniques are a Avehicle@ for the information used, they need some vehicle themselves to exist, such as firms or households, an issue I will return to below. After all, a technique is not used in some passive sense just as a living being=s life is Alived.@ When a technique is used, this happens because some other entity like a firm or a household has consciously chosen to use it, selecting it from a much larger set of potential techniques. To complicate matters further, these vehicles are themselves often the subject of selection. This produces a Ahierarchy@ of selection still highly controversial amongst evolutionary biologists.

As noted, one of the most challenging issues in the history of technical knowledge is the relationship between the underlying information and the technique itself. One could formulate this in terms of engineering problems and roads to their solution, as Vincenti (1990) does in his classic analysis of this problem. But as a general statement I am not persuaded by Vincenti's adoption of Michael Polanyi's notion that operational principles have to be understood before design (p. 208).[10] Many if not most techniques worked through history, at least till 1850, without their designers or users having the slightest idea of their operational principle. Trial and error, serendipidity, and even totally false principles usually led to techniques that

[10]A similar notion is expressed by Margaret Jacob (1997). Jacob submits (p. 131) that A people cannot do that which they cannot understand and mechanization required a particular understanding of nature that came out of the sources of scientific knowledge.@ Clearly people can do Athat which they cannot understand@ unless one places a very stringent definition of what Aunderstanding@ means here. The fact remains that most of pre-1780 physics, chemistry, botany, and medicine provided a poor handle on the processes involved. If Watt and Smeaton and Donkin and Hornblower and Trevithick and the less famous engineers that made the British Industrial Revolution succeeded, it was not because they really Aunderstood@ what they did, but because they were good mechanics, meaning they were dexterous, had been trained and taught to experiment and measure carefully, had developed the skills to communicate with others, and had a good intuition about what might work and what not. Little of this involved the actual principles of the natural processes at work.

worked sufficiently well to survive the selection process. One does not have to know *how* fertilizer or crop rotations enhance next year's yields or *why* quinine cures malaria; the knowledge underlying their use is simply that Ait works.@ Vincenti's use of Laudan's framework of the selection and solution of technological problems, akin to Landes's notion of Achallenge and response,@ followed by choice between rival solutions provides one way to look at the connection between knowledge and technique. The knowledge provides the tools to solve the problem, while the technique embodies the solution. Other mechanisms can be imagined. We could distinguish between techniques that need a deep understanding of the underlying knowledge for it to work, such as nuclear engineering or electronics, and singleton techniques that are based on no more knowledge than Aif you follow these instructions, such and such works.@ Even today, much of our λ, from psychiatry to macroeconomic forecasting, is based on little deep understanding of the basic processes at work. Of course, even when the knowledge basis is required this applies largely to the engineers designing and improving it; for the day-to-day operation no such knowledge may be necessary.

The advantage of using the technique as the unit of selection in an evolutionary analysis of the economic history of technology is that it focuses on a meaningful and important historical entity while satisfying the basic characteristics of the evolutionary system. First, new knowledge (changes in Ω) is introduced through a highly stochastic process. In this regard the entire body of literature following Campbells' seminal work is relevant: knowledge changes by stochastic mutation and selective retention. Such mutations may or may not express themselves on the phenotype, that is, mapped onto λ; most do not. If they do not and yet are selected to be retained in Ω, they could be Aactivated@ (that is, expressed) at a later time as part of adaptation to changing environment especially the arrival of complementary knowledge.[11] Technological innovations thus differ from the purely random process of mutations in the living world, as the search for new techniques is clearly in some sense motivated by needs and opportunities. Yet it is a long way from being a deterministic process: we do not get all the innovations we need because we are constrained by the finiteness of Ω, which is a complicated way of saying that we do not know enough. But more is involved: even if the necessary knowledge exists in Ω, it must occur to somebody to look for the solution, and the technological problem has to be well-formulated. The existence of knowledge in Ω thus does not guarantee that the mapping will occur. Why it does or does not remains one of the central questions in the history of technology.

Second, there is superfecundity in the system in that there are far more techniques that we need. There is more than one way to skin a cat, but for each cat we can only use one, and if there are more ways to skin a cat than there are cats, superfecundity requires selection. Moreover, when a new technique is proposed, users have to choose between the old and the new ways of doing things. This produces clear-cut selection in the model. From a shoemaker choosing between different materials for the soles of shoes to a

[11]One of the best examples is George Cayley's celebrated discovery of the principle of fixed-wing aircraft using the air-resistance as a substitute for the flapping of wings which all airplane designers must understand (Vincenti, 1990, p. 208; Bagley, 1990, pp. 617-19). Early plane designers were aware of Cayley's work (Vincenti, 1990, p. 243; Crouch, 1989, p. 161).

committee of engineers choosing between alternatives of machinery deployed in a modern manufacturing plants, there is a process of selection amongst existing techniques. Yet new variation has to be forthcoming if selection is to continue to play a role. In that regard, the Darwinian model and the problems arising in the historical description of the evolution of production technology provide a nice fit. Whether selection occurs due to a conscious and purposeful Aselector@ making the choice, a purely abstract invisible hand, or even complex fixed rules written as nested conditional rules in a game such as our urns example, is secondary.

Selection mechanisms also apply to general knowledge, but it follows principles quite different from those followed by the set of feasible techniques because the notion of superfecundity does not apply in its simple version: in principle it would be possible to cram the minds, libraries and hard disks of the world with an ever-increasing body of useful knowledge and draw from this store when circumstances require it. In practice, however, large amounts of knowledge is rejected, forgotten, or lost. In part, this is because information storage costs are not quite zero or because information is willfully destroyed or destroyed as a by-product of war or other forms of vandalism (e.g. the burning of the library of Alexandria). In part, this is because some knowledge is incompatible with other knowledge.[12] Phlogiston physics, Ptolemaic cosmology, and the humoral theory of disease no longer play important roles outside the history of science because it is logically impossible to adhere to them and to modern theories at the same time. Yet disbelieving knowledge does not mean we do not have access it and could not revive it if we chose to. In

[12]For an illuminating analysis of how knowledge is selected when such choices have to be made see Durlauf (1997).

short, selection in technology involves a double-layered process of selection: the selection within Ω, accepting and rejecting new knowledge, and the selection within λ, in which actual techniques are selected for actual usage.

How do innovations arise? Despite dissimilarities with living systems, which I will elaborate below, it is clear that the process of generating innovations proposed here shares some important features with the way living beings do. For one thing, the vast majority of all human knowledge, like DNA, is non-coding or Ajunk@ in the sense that it does not apply directly to production. Most scientific (let alone other forms of) knowledge has no applications and does not affect production technology right away although it may be Astored@ and in rare cases called into action when there is a change in the environment or when another complementary invention comes along. Thus most additions to Ω, like mutations, are predominantly Aneutral@ and do not affect the selection criterion one way or another, but may become useful when the environment changes and calls for adaptation (Stebbins, 1982, p. 76). The activation of such previously inert knowledge may be the evolutionary equivalent of what economists think of as Ainduced innovation.@ It is arguable that such neutral changes with no phenotypical effect should not be regarded properly as evolution at all (Maynard Smith 1976, p. 331). General knowledge, too, is being created at a rate much faster than technological knowledge, but if it finds no application in production, it is not a part of technological evolution and might be regarded as Aneutral.@ In that sense the expansion of technology is highly accidental: new knowledge that seems to serve no obvious purpose is nonetheless created and retained, and is available when needed. No wonder that the great Japanese geneticist Motoo Kimura, the proponent of genetic neutralism, proposed to replace the notion of Asurvival of the fittest@ by the concept of Asurvival of the luckiest@ (Kimura, 1993).

At times a mutation occurs that is truly favorable, does not draw much upon previous knowledge, and is right-away Aselected for@. Nature, like technologically creative societies, does not abide by the maxim Aif it ain't broke don't fix it@ (Cziko, 1995, p. 22). Examples of mutations of this kind are the discoveries of the smallpox vaccination by Jenner in 1798 or X-rays by Röntgen in 1895. Neither drew upon any obvious prior knowledge, yet they did not remain dormant for extended time as in Stebbins's view of evolutionary progress. From the outset these inventions were viewed as Afit@ mutations, selected over previously existing techniques.[13] In that regard, at least, they were more like Richard Goldschmidt's Ahopeful monsters@ -- sudden large phenotypical changes possessing higher fitness.

Another evolutionary notion of use to historians of technology is the by now well-understood principle of exaptation, first proposed by Gould and Vrba (1982). The basic idea is that a set of selection criteria that chooses a technique for one reason but which then owes its success and survival due to another trait. In other word, a unit=s function and interaction with the environment can explain its survival but not necessarily why it exists in the first place (Maynard Smith and Szatmáry, 1995, p. 8). One might venture that in the history of technology exaptation is probably more common than in natural history. Many of the most dramatic inventions of the modern age were originally selected for quite different purposes than what eventually turned out their most enduring trait -- consider famous example of the gramophone, originally

[13]Röntgen announced his accidental discovery of aAnew kind of rays@ in Nov. 1895; In March 1896 it was

intended by Edison to serve as a dictaphone (Ohlman, 1990, pp. 720-21).

---

already used to find a part of a broken needle in a woman's finger at the Royal Free Hospital in London.

Another area where technological evolution deviates from the strictures of evolution in living beings is the process of recombination. In nature sexual reproduction means essentially that the diploid genome is a linear combination of the haploid gametes of the two parents. There are probably substantial advantages to this kind of reproduction (suspected to be associated with immune responses to parasites) but it is far from clear that they play a major role in accounting for evolutionary change for two reasons. First, mating can occur only between two very similar creatures (members of the same species). Second, the diploid cell is a *weighted mean* rather than the *sum* of the genetic information embedded in its parents's cells. Genetic information is thus non-additive. The adaptive gains from recombination can thus be substantial but inherently local in nature. Human knowledge and its use, on the other hand, are not so constrained. In that sense what is meant by recombination in technological history means something quite different. In an earlier paper (Mokyr, 1996) I pointed out that recombinations in technological history involve the combination of existing knowledge in new forms, whereas mutations involve the emergence of altogether new knowledge.[14] There is a third way in which innovation can occur namely Ahybrids@ which represent combinations of vehicles rather than of the knowledge embodied in them. By applying a small internal combustion engine on board a hot air balloon we create a new vehicle but not necessarily new underlying knowledge. Again, we see that the rules of evolution apply not only to the techniques but to the vehicles themselves -- natural selection and innovation occur in hierarchies rather than exclusively at the smallest level of analysis. Some writers find this difference between cultural evolution and biology lethal to the analogy: Stephen Jay Gould,

[14]Historical examples of such recombinations are of course abundant. One is the fusee, a conical device borrowed from crossbow designs by fifteenth century watchmaker to equalize the changing force of an unwinding mainspring (White, 1978, p. 308). Watchmaking in its turn was an endless source of ideas for toys, musical boxes and precision machinery requiring precision-made cogs, springs, and gears. For other examples, see Mokyr, 1996.

who has never warmed up to the idea of evolutionary theories outside biology has pointed out that Ahuman cultural change need not even follow genealogical lines -- the most basic requirement of a Darwinian evolutionary process -- for even the most distant cultural lineages can borrow from each other with ease@ (Gould, 1997, p. 52). Yet cultural and technological entities have identifiable parenthoods and thus genealogies. The requirement that in living beings their number be two at most and that the parents have highly compatible gametes underscores the very special and restrictive parameters of Darwinian process in living beings.

Another example of a common feature between evolutionary biology and technology is the non-continuity of entities. There is a finite set of traits and each trait changes often in discontinuous fashion. Techniques and species are discontinuous in the rather trivial sense that we have no more a continuum of intermediate forms between dogs and cats than we have a continuous range of varieties between automobiles and motorcycles.[15] This in and of itself is evidence for some kind of selection. The selection mechanism prefers a given set of traits over alternatives in its immediate environment, so that the evolutionary process settles down on Apeaks@ in the evolutionary landscapes. These fitness surfaces, to use

[15]Cf. Eldredge, 1989, pp. 99-102, who points out that the entire concept of species in biology must imply discontinuities in the smooth continuum of phenotypic diversity. Some intermediate forms between motorcycles and automobiles such as the three-wheeled Ascootmobiles@ of the 1960s have been tried but were soon abandoned or survive in small niches.

Kauffman's (1995, p. 169) term, are correlated, in that nearby points have similar heights and intermediate points have lower fitness, so that a creature that is half-way between two existing species is less likely.[16]

To what extent phylogenetic discontinuity actually plays a major role in natural history has remained a matter of dispute. While the saltative speciation events proposed by Goldschmidt and Otto Schindewolf have been rejected by most evolutionary biologists, it is equally clear that this does *not* mean that the process proceeds exclusively by infinitesimally small steps. In technology, too, major and discontinuous breakthroughs or Amacroinventions@ play an essential role in historical episodes of technological change, as I have argued elsewhere (Mokyr, 1991). Moreover, traits change often together in pleiotropic fashion (if for different reasons), producing different Abundles@of traits. When aspirin was introduced it represented a Apackage@ that simultaneously reduced fever, alleviated pain, and as was later discovered, also prevented heart disease. Of course, more frequently positive traits are bundled up with negative ones, which is when economists speak of externalities or Aside-effects.@

[16] Kauffman, for whom the parallelisms between biological and technological evolution are obvious, adds that Aorganisms, artifacts and organizations evolve on correlated but rugged landscapes@ because optima represent compromise solutions that meet the conflicting constraints of different subproblems (p. 179).

It may be helpful to present one more example of an evolutionary system that satisfies our criteria and that can be used to highlight what is essential to *all* Darwinian systems and what is peculiar to the flora and fauna of our planet. Consider the evolution of language.[17] By our definition each word in a language is an entity or species and each time it is used it is a specimen. A larger group of words together constitutes a language, comparable to higher phyla (although single words can belong to more than one language). Clearly, new words are created continuously through spelling errors, typos, and proposed neologisms. These innovations are created through processes that are at least in part purposeful, so that the process that creates them cannot be said to be entirely random. Yet this process has enough randomness in it to produce superfecundity and require selection processes that weed out mistakes. There are selection criteria that decide which words will be used and which will not, though like in other selection processes, these are not always wholly understood. A word, too, contains underlying information much like a technique or a gene: it is constrained by syntax, grammar and language structure. Communication technology (which is what language really is) is perhaps somewhat closer to production techniques than to living beings, yet it follows

[17]See especially Cavalli-Sforza and Feldman (1981), pp. 19-29; Van Parijs (1981), p. 109. Darwin (1871, p. 466) himself felt that his ideas applied to the development of languages. Like organic beings, he noted, they can be classified into groups, and some are Aselected@ and become dominant while others go extinct and quotes with approval a scholar who noted that there was a constant struggle for life among the words and grammatical forms of each language.

its own specific dynamic.

Selection Units, Replicators, Vehicles and Interactors

To repeat, the main reason why non-biological Darwinian systems should not be seen as drawn *in analogy with* biology is that biological systems represent a special case, and a rather limited one at that, of all possible Darwinian systems. For one thing, consider the main function that the entities subject to the evolutionary process are supposed to have. The main function of the entities upon which evolution operates, be they organisms (as the orthodox Darwinian view has it) or genes (in the views of Dawkins and other ultra-Darwinians) is to reproduce themselves. Yet this seems to be a uniquely biological point of view, less than satisfactory in a general model in which entities are chosen consciously according to some external criterion. The most obvious one in any analysis is the contribution of a technique to some objective function that society tries to maximize. This criterion is then transformed into sub-criteria that agents observe such as firms maximizing profit, workers minimizing physical effort, or the like. But at times technological choices are made by a central agency such as the Federal Communications Commission determining which standard to set for digital HDTV. Similarly, words will be selected by their contribution to an objective function relating to effective communication. Here, too, the selector can be as an invisible hand (a neologism Acatches on@) or a conscious and purposeful organization such as a language academy.

A potentially difficult problem is the definition of the appropriate group to which the entity on which selection occurs belongs. Cultural evolution has no general solution to the classification problem, although intuition suggests certain obvious rules.[18] The set of instructions underlying the manufacturing of car engines, for instance, could be regarded as a group. Just as we recognize that some species are distinct but closely related, we could say that the distance between making a Chevrolet and a Toyota car engine is smaller than the distance between the making of a car engine and a two-stroke lawnmower engine. Yet the latter would still share a greater percent of the underlying information than, say, they share with the instructions on how to perform by-pass surgery. Genetic distance, measured by differences in allele frequencies which can be assessed using enzyme electrophoresis techniques, has provided a rigorous tool to evaluate this distance in living beings. No such exact tool exists for the analysis of technology, and perhaps none is needed. Yet some kind of ordinal measure is suggested by our intuition. Humans and Chimpanzees share 99 percent of their genome; we feel justified in saying that they are closer to each other than to cockroaches. This is confirmable by experiment. Similar comparisons could be imagined in cultural evolution although no experimental methods could lend them rigor.

[18]Thus Cavalli-Sforza and Feldman (1981, p. 174) think that Atwo languages that are mutually incomprehensible are analogous to two different species.@ It seems better to think of a language as a Ahigher@ order than a species; the distinctions between Alanguages@ and Adialects@ are arbitrary in any event. Mutual incomprehensibility seems a weak criterion.

We may define a group to which all entities that display certain traits belong such as all engines, computer programs, or headache medications. This classification involves arbitrariness that is seemingly absent in the natural world. After all, the definition of species revolves around reproductive compatibility, a property that has no obvious equivalent in cultural evolution. The definition of a technique seems more arbitrary: we might define all flights on a Boeing 767 from Chicago to London as members of a single species of all flights on a Boeing 767 or of another species of all flights from Chicago to London. From an operational point of view this seems a less than fatal problem. After all, if we want to say that the frequency of flights between Chicago and London went up relative to the frequency of the Aspecies@ of travel by train-boat combination, the exact *units* in which this is expressed matter little. We can classify techniques by their purpose (headache medications) or the physical principles on which they are based (engines) or even morphological characteristics of the process (textiles). Moreover, here, too, the differences between biology and other evolutionary systems are less than what they appear: the cohesion of biological Aspecies@ is ambiguous as well, and the groupings in evolutionary biology other than species display the same arbitrariness.[19]

[19]For one thing, species as a biological concept at times differ from the species as a taxonomic category. There

are morphologically undistinguishable species yet do not interbreed known as *sibling species*. Conversely, some very different-looking animals like the subspecies of the snake *Elapha obsoleta* do interbreed. Some species interbreed to produce sterile hybrids, yet some ducks and some orchids can be made in captivity to hybridize to produce fertile offspring yet do not exchange genes in nature (Futuyma, 1986, pp. 111-12).

More serious is the problem of replication. Techniques, to use Dawkins's term, are replicators: copies are Amade@ of a technique in period t and these are reproduced in the next period. This is precisely what gives evolutionary systems their dynamic properties. The stochastic process is based on the principle that units reproduce themselves almost exactly. Techniques, however, do not replicate themselves the way organisms do: there are no analogies to the natural phenomena of birth, reproduction, and death. Nothing in the world of production technology resembles the way that genetic information is transmitted from parents to offspring during meiosis. Instead, we have to define somewhat arbitrarily what the Alife@ of a non-living entity looks like (Cavalli-Sforza and Feldman, 1981, p. 14). There are cases in which this seems easy. Consider a technique for growing wheat: it would seem that such a technique is like a specimen Aborn@ at the beginning of the growing season and dies at its end. If the next year the farmer grows the wheat using exactly the same technique, should this be regarded as a very similar but different copy of the specimen, or one and the same specimen? And if it is the same, how about the son of the farmer who learns the techniques from his father and thirty years later is still using the same technique? Or consider flying a passenger from Chicago to London: is the flight of each passenger a Alife@ or the flight itself? And is the same flight next day a different specimen (Ason of UAL flight 928")? The concept of a *generation* is thus arbitrary here -- though the same can be said of many forms of life: is the cutting made from a perennial plant a separate entity or the same specimen? None of this, as Cavalli-Sforza and Feldman point out, should stop us from observing the selection and acceptance of a trait over time and the evolution of the distribution of the frequencies with which entities exist.

Moreover, while we can think of techniques as undergoing phylogenetic development, it is hard to see anything that resembles an ontogeny, that is, the development of the phenotype from the genotype through interaction with the environment. In biology the act of transmission of knowledge occurs instantaneously at the beginning of the life of the entity, and at the end of the life of the vehicle organism it is irretrievably and abruptly lost, like a machine that has to be junked at the very end of its life (a depreciation process known as Aone-horse-shay@ among economists). In technology, on the other hand, the vehicle keeps acquiring knowledge throughout its life while the depreciation of knowledge is often continuous: memories fade, skills blunt, artifacts wear out gradually. Yet the principle of depreciation is similar, and for that reason the technological DNA, too, needs to be passed on over time.[20] But how?

In Richard Dawkins=s formulation, replicators need vehicles. In a sense the technique is a vehicle for the useful knowledge underlying it. But the technique, unlike specimens of animals, as we have seen, is not a concrete entity but, yet it requires in turn a carrier or vehicle itself. What kind of vehicles exist in the history of technology? The answer is less sharp and well-defined than in biology, because technological knowledge is more protean than genes, but they must exist for if they did not, the wheel would have had to be reinvented at any instance of time. The main classes of vehicles are delineated below:

[20]The distinction between *vertical* and *horizontal* transmission is vague here. A technique Areproduces@ when it is repeated, regardless of whether the repeater is the same user, his apprentice, his competitor, or another producer across the ocean.

Artifacts. Because artifacts (or Aperforming devices@) exist for long periods and are used over and over again, they embody to a large extent the information in them. Indeed, a considerable literature exists in the theory of economic growth based on so-called vintage models in which the technology is wholly defined by the date of production of the capital goods in which it is embodied. Thus the technique of hitting a nail into a wooden board is to some extent embodied in the shape of the hammer. Yet there are different ways of handling a hammer and the exact knowledge of where to hold the stem, how hard to hit a nail, or to dip the nail in oil to prevent the wood from splitting. While the artifact imposes a considerable constraint on the range of the technique, it is possible for changes in other knowledge (or other artifacts) to alter the way existing artifacts are used. The dogma that once the artifact is formed its Agenotype@ cannot be altered is thus inapplicable unless one defines the genotype as pertaining to one Aperiod@ in which the artifact is used. While an artifact embodies knowledge, in and of itself it is rarely enough to wholly describe the technique. A piano embodies the knowledge how to make sound, so seeing one is enough to produce *some* sound but more knowledge is required to produce the *Hammerklavier* sonata. People can learn from artifacts by means of reverse engineering: once we see a wheelbarrow or a cotton gin we can understand how they are used and reproduce it.[21] Artifacts can operate as alternative storage places for some kinds of

[21]Dawkins, 1982, p. 175 picks an unusually infelicitous metaphor when he compares the central dogma of embryology to technology. He argues that one cannot dissect a cake and recreate its recipe, and similarly one cannot map back from phenotype to genotype. The argument is of course correct for DNA but not for cakes. In some techniques it is possible to reproduce a technique from the phenotype and to reconstruct the essential information on which it is based – indeed, that is precisely what we mean by reverse engineering.

knowledge. In this sense, of course, it is possible for knowledge to survive and be resurrected even when the other vehicle, human memory, no longer carries it. It also greatly facilitates, as we shall see, the process of innovation.

Storage Devices. Technological information can be stored in non-performing devices such as textbooks, manuals, encyclopedias, or oral traditions of how to produce certain goods can exist independently of whether they are (or even have ever been) used. Although it may appear that this is just one other mechanism through which they learn from each other, most storage devices can transmit information across time as well as across space. Cookbooks are perhaps the paradigmatic storage device since a cooking recipe is a paradigmatic example of a technique in that it provides a set of instructions based on an underlying structure of knowledge. Implicit storage devices also take the form of Vincenti's (1990, p. 210) Anormal configuration.@ It is normally assumed that cars have four wheels, a round steering wheel, and that the gas pedal is lower than the break pedal, and almost all cars are made that way.

Firms. One solution to the wear-and-tear of knowledge is to group them together in an infinitely-lived organization dedicated to carry out a set of techniques. One explanation of the existence of the firm is precisely to perpetuate the technological entities through the production processes it performs. Members of the firm jointly carry the technique and each time the firm produces the goods it specializes in, the firm may be regarded as the vehicle of that technique. Before the Industrial Revolution, the family (or household) fulfilled this function, and economic history has documented families that stored such information and passed it on over time while preventing others from access.[22]

Human Memory. Much of the knowledge set is transmitted from period to period simply because the knowledge is embedded in people's brains and they can remember how the technique was used the last time around. Thus the farmer who grows wheat year after year can be said to have access to knowledge and map into his feasible technique set by doing the same thing over and over again.

Direct Communication. The techniques can be transmitted by the training of individuals in which the G's are passed on from one entity to the other. From the point of analysis here, a farmer who Aknows@ how to grow wheat and does so year after year in identical fashion, can also replicate the technique by teaching his son or apprentice how to grow wheat. If the son then emulates the father precisely, the technique can be said to have replicated. Indeed, the only reason why such learning is necessary is because the vehicle is subject to wear and tear -- if people lived forever with constant ability, learning would be unnecessary for the technique to survive. As it is, such transmission is necessary and we could imagine a system in which all knowledge was passed on only vertically and succeeded or failed by classical Darwinian mechanisms of survival and differential reproduction of Avehicles.@ . On-the-job training, emulation by neighbors, and the transmission of Michael Polanyi's famous Atacit knowledge@ clearly belong in this category.

[22] In some industries, particularly in ironmaking, skills were the traditional realm of dynasties in which technological knowledge was kept as much as possible within the family. See Evans and Rydén (1996).

One potential solution to be explored is to think of technique as an Ainteractor,@ that is a unit that first and foremost interacts in some way with its environment.[23] Every organism, when it exists, interacts with its environment. The interaction then is defined by traits that ultimately determine the chances that the unit has for survival and/or reproduction, but its essence is the interaction, not the replication. Indeed, defining any technique as an interactor seems obvious once one thinks of it; the instructions in the technique lead to the production of a god or service, which are meaningful only in the context of interaction. Hull defines an interactor an entity that interacts as a cohesive whole with the environment in such a way that this interaction leads to differential replication (Hull, 1988, pp. 408-09). The point of an interactor in Hull=s definition is survival more than replication, and survival can only be achieved by a certain amount of success measured by a criterion that feeds into whatever does the selection; a technique that does not produce a workable product at an acceptable price will not survive. It remains to be seen to what extent the interactor differs significantly from Dawkins=s Avehicle@ and whether these concepts will be useful in understanding the history of technology.

Selection and Teleology

[23]The idea of an organism as an interactor dates back to a number of classic articles written by David Hull in the early 1980s. For a summary, see Lloyd, 1993.

The outline of an evolutionary theory of technology is now becoming somewhat clearer. To summarize: because the Aentity@ on which evolution occurs, the technique, is a procedure or a routine, the main actors in the history of technology, human beings, their organizations and artifacts, play somewhat different roles. They are the vehicles that carry each entity from Aperiod@ to Aperiod.@ It is possible to debate whether the unit on which selection occurs is the Aentity@ or the Ainformation@ on which it is based. This is a deep and complex debate in evolutionary biology.[24] But in the history of technology it could also be credibly maintained that the actual unit on which selection occurs is the artifact, the person, or firm which in the present analysis are the vehicles carrying the information needed for production. After all, there seems to be no reason why we should exclude Darwinian models of the type outlined above to determine which firms, people, or artifacts survive. It all depends on the question being asked.[25]

Regardless of whether the notion of a hierarchy of natural selection makes sense in biology, it will be

[24]For good expositions from opposing camps, see Dawkins, 1982, and Eldredge, 1995. I am deliberately ignoring at this stage the further complication of group selection, although such complications can rather well be formulated in this framework.

[25]It might, however, be asked usefully in the case of firms or artifacts how the replication mechanism works and whether there is any kind of replication going on altogether (Penrose, 1952, p. 808).

readily recognized that in production there must be more than one selection process going on at the same time. First, the market selects on the outcomes, that is, which products will be produced, how, and by whom. There is superfecundity here because each society is capable of producing a great deal more products in many more fashions than it actually does. Second, the vehicles themselves select, and here the selection is conscious: they pick and choose from different techniques in λ. Third, firms select on the actual alternative ways of producing the goods. Another level of superfecundity is at the vehicle level: there are a lot more firms that can produce a good or a lot more engineers that can learn of some technique than actually do. This is vehicle selection. Finally, as we have seen, there is some measure of selection at the level of Ω.

Before turning to the question of how we describe technological change in this system, there is another basic way in which technological and biological Darwinian systems differ. In the ultra-Darwinian view, the final purpose of existence is existence itself. A successful replicator, in Dawkins's view is one that has an infinitely lived germ-line. In this sense its raison d'être is wholly recursive. Not all evolutionary biologists agree. In a more economistic view, living beings have a purpose which we could define heuristically as Awell-being.@ In this view, reproduction is important but in the final analysis epiphenomenal to survival and well-being (Eldredge, 1995, pp. 40, 187, 211). Regardless of what the Apurpose of life@ is, in any technological system, there is a deeper and unequivocal purpose for the existence of techniques, namely to increase the utility of human agents. Each technique, when it is applied, serves an Aultimate@ purpose which, while obviously intertwined and correlated with its fitness, can be regarded separately. In other words, economic historians can regard a technique in terms of what it was supposed to do -- produce goods and services -- as well as in its success in reproducing itself. Ultimately any selector will have to be judged by its success in satisfying human needs and the survival of each entity is correlated with that criterion. The correlation is less than perfect however: at times techniques are selected that do not satisfy the objective function of human need as efficiently as others.

It must be a central axiom of any evolutionary view of technology that despite this difference in fundamental objective, Darwinian logic can be carried over to systems that do not follow recursive objective functions but serve some distinct purpose. In this regard, the assessment of Asuccess@ becomes less relativistic and terms like progress acquire a meaning they lack in natural history. For instance, the invention of the safety bicycle in 1885 and its rapid success was not only a success in that the number of this species increased relative to penny-farthing cycles, but also in that it satisfied in a superior fashion an ulterior social objective function of people needing to transport themselves safely and efficiently in the context of their time. Fitness, in an economic framework, has thus two dimensions: the first is the ability of the entity to reproduce itself (or the information contained in it); the other is its ability to contribute to the material well-being or the quality of the existence of a representative individual.[26] The extent to which those two coincide is exactly the stuff of which the economic history of technology is made. Clearly the ability of an innovator to market an invention and to persuade others of the qualities of a new technique contributes to the reproduction chances, even if it does not necessarily coincide with its contribution to social welfare.

[26]This distinction corresponds to Amartya Sen's (1993) distinction between the Aquality of the species@ and the Aquality of the lives we lead.@ While Sen's formulation is controversial, it seems quite clear that one can see the difference between the economist's social welfare function and the biologist's concept of fitness even when applied to techniques.

Innovation and Adaptation

In an evolutionary framework, change occurs through blind variation and selective retention in Ω. The appearance of new techniques is constrained by changes in Ω. The point is that while economics regards changes in knowledge as largely driven by incentives, and the outcome of rationally-driven search processes, an evolutionary theorists regards them as a spontaneous, largely autonomous process, in which knowledge begets more knowledge. The main reason why knowledge is constrained in its expansion is that any evolutionary system has inertive forces which resist change and while this kind of resistance is of course an important source of direction, it almost guarantees that change will rarely be precisely what the system needs.

Do technological entities violate the rules of evolutionary biology by transmitting Aacquired characteristics@ to new generations? While in living beings all genotypic change occurs at conception, this is of course not the case in other systems. Thus the genome can Aacquire@ characteristics during the lifetime of the vehicle and pass them on. Hence the long literature on the ALamarckian@ nature of cultural evolution. Yet the present framework suggests that this debate is largely beside the point because the Alifetime of the vehicle@ is arbitrarily defined. In principle we might define an infinitesimally short period of time during which the technique is used as the Alifetime@ so that learning or new ideas that occur can always be defined as occurring at the start of the technique's existence. In other words, the existence of the entity becomes continuous, and the basic discontinuity that drives the difference between ontogeny and phylogeny disappears.

Instead, the main ALamarckian@ characteristic of technological evolution is that there is feedback from λ to Ω. The underlying structure of knowledge is influenced by the techniques in use.[27] The strict Weissmanian constraints on the evolution of living beings preclude this: there can be no feedback from phenotype to genotype and the motion is unidirectionally from Ω to λ. Moreover, in living beings the change in Ω is always purely random, and directionality is imparted only by selection on the members of λ. In all knowledge system, this randomness in changes in Ω is abandoned, as clearly there is conscious search for new knowledge which imparts an a priori directionality on all change. This preselection does not preclude the operation of an ex post selection mechanism on Ω and well as on λ, as long as the preselection is highly imprecise. In both of these respects, then, technological evolution differs from natural selection in living beings.

The framework developed here is capable of distinguishing between innovation and adaptation. A change in environment (say, due to changes in the availability of complements or substitutes) will cause selection to favor those techniques that result from other parts in Ω. This would be pure substitution, as the system simply switches to another part of existing knowledge. Yet it may also induce search process that expand a different part of Ω. For instance, the appearance of a new disease will focus research on trying to isolate the pathogenic organism and uncover the mode of transmission. Such induced knowledge may or may not be forthcoming. In other cases, expansions in Ω often serve no purpose at all and seem to be uncorrelated with any new Aneed@ of the system. In that sense, the continuing growth in Ω is spontaneous and unpredictable, and in Campbell=s terminology, it is *blind*, ideas Aoccurring to someone@ in unpredictable fashion, much the way Campbell and others suggested it. While these mutations are not really

[27] There are many famous examples of technique influencing knowledge, the most famous one being Sadi-Carnot=s celebrated formulation of the Laws of Thermodynamics while observing steam engines.

Acopying errors,@ they are obviously following some kind of stochastic process subject to the usual constraints of path dependence.

Summary: Information and Selection.

Technological information, as emphasized above, is embodied in techniques, coherent bodies of knowledge which instruct Aagents@ how to engage in production. Is the technique the Aunit@ on which selection occurs? Techniques require information that generates them. As noted, this information can be carried in a person, an artifact, or an organization such as a firm. It can also be preserved in storage devices. While the information is thus isomorphic to genes in some ways, it does not in any obvious way replicate itself and it can remain dormant. The difficulty with Dawkins's idea of Amemes@ as another kind of replicator is that it is Aa unit of information residing in a brain@ (1982, p. 109). Yet when knowledge Areplicates@ by being copied onto another brain, it is not clear that it satisfies his criterion of making an *exact* copy of itself, since it is not obvious that another individual's brain will interpret the words the same way. The only observable phenotype is the technique itself. If two blacksmiths make horseshoes the same way, the techniques belong to the same group. If one learned the technique from the other, the technique can be said to have replicated itself and the frequency of this technique has increased. If the technique makes better or cheaper horseshoes, the successful blacksmith is likely to be imitated or have more apprentices, and thus the fitness of this technique is likely to go up.

Note also the difficulties inherent in the concept of a set of useful knowledge. One individual could develop an industrial chemical process from a deep knowledge of chemical engineering. An imitator could then emulate him, and produce the same good using the same technique but without possessing the knowledge. The first one could then be said to map from the knowledge set Ω to the feasable techniques set λ whereas the other's knowledge consists exclusively of the technique. Clearly not everyone has to know what can be known. Furthermore, most of the knowledge in λ consists of techniques that do not work or do so inefficiently (inside the boundary of λ). On the other hand, ignorance means inflexibility: the imitator who only Aknows@ one technique cannot respond to changes in the environment.

Selection in cultural systems occurs simultaneously at all levels, and forced choices between units of selection as posed by Dawkins and other ultra-Darwinians are inappropriate outside the limited confined of biology. There is no Aselfish gene@ -- in economics all there is are objective functions of economic agents which are then aggregated by markets or other aggregators.

To illustrate how selection affects technology, I adapt a variation on what Dennett (1995, ch. 5) has proposed as the difference between the possible and the actual. Consider the universe of techniques. The largest meta-set is that of *imaginable* techniques. This includes all techniques that could conceivably be concocted by the human mind. The dimensions of this set are huge but in the final analysis it is constrained by the limitations of the human mind. Smaller and wholly contained in this set of techniques *ever imagined* until time t. This, too, is an expanding set, since new science fiction novels and starry-eyed engineering students dream up new techniques. Another subset of the imaginable techniques is the set of *possible* techniques, that is, those that are not contradicted by the laws of nature.[28] This set is not quite the same as the set *believed to be possible* at any time t. Today we believe that travel at speeds exceeding the speed of light and animal breeding through the passing of acquired characteristics are not possible. Yet we cannot

[28]There might well be techniques that are possible but not imaginable, but these would be of little interest to us.

be ontologically sure that they are not. The union of the sets that were ever imagined until time t, possible to the best of our knowledge, and believed to be possible at that time is the set of *potential* techniques at time t. Wholly contained in that set is the set of *feasible* techniques which are not only possible and believed to be so, but also within the technological capabilities of the time.[29] The feasible set is essentially the same as our set λ. It is clear, then, that it, too, is Aselected@ from a larger subset. Within the feasible set lies the *realizable* set. Feasible techniques may not be realizable because of political constraints, social taboos, or the physical absence of a crucial ingredient (say, enriched plutonium). A subset of the realizable techniques is the set of *rational* techniques. These techniques dominate other techniques in their ability to satisfy whatever objective function we impose. There would thus be a different set of rational techniques if we imposed the pure Darwinian fitness criterion (maximize occurrences) or an ulterior social welfare function. It makes no sense to select techniques elsewhere. Finally, there is a set of *optimal* techniques which is that segment of the rational set that is best suited to any given environment.

Needless to say, some of these selection mechanisms appear to be rather trivial logical exercises, a bit like Daniel Dennett's ALibrary of Mendel.@ Yet they highlight the fact that understanding what we observe begins with asking why there are things we do not observe. Mark Ridley (1985, p. 56-57) has noted that non-existent forms of life may be absent either because there was negative selection or because the necessary mutations never appeared. The difference, as Ridley points out, is due either to selection or to constraints. Selection means that they appeared but were selected against. But if they did not appear, is that

[29]Roger Bacon, Guy de Vigevano and Francesco di Giorgio Martini all imagined techniques that were beyond the reach of their times, so they were in the potential but not the feasible sets. See for instance Gille (1969, ch.3; 1978). Gille (1969, p. 42) points out that mechanization was in part the result of a Aspeculative type of thinking; Da Vinci was not the least of its representatives.@

because they *could* not have or because the mutation simply never occurred although it could have? In the world of technological selection, non-observed techniques in any given period could be either totally impossible, possible but never occurred to anyone, occurred but were beyond practical reach, and so on. These are *all* in some sense selection mechanisms. What we actually observe is a tiny sliver of what *could have been*. Except in the light of the theory of natural selection, Theodosius Dobzhansky once observed, nothing in nature makes much sense. This, surely, remains true for the history of technology as well.

REFERENCES

============

Basalla, George. 1988. *The Evolution of Technology*. Cambridge: Cambridge University Press.

Bagley, J.A. 1990. AAeronautics,@ in Ian McNeil, ed., *An Encyclopedia of the History of Technology*, London: Routledge, pp. 609-647.

Cavalli-Sforza, Luigi L. and Feldman, W.M. 1981. *Cultural Transmission and Evolution: a Quantitative Approach*. Princeton: Princeton University Press.

Crouch, Tom. 1989. *The Bishop=s Boys: A Life of Wilbur and Orville Wright*. New York: Norton.

Cziko, Gary. 1995. *Without Miracles: Universal Selection Theory and the Second Darwinian Revolution*. Cambridge, MA: MIT Press.

Darwin, Charles. 1871. *The Descent of Man and Selection in Relation to Sex*. New York: Modern Library edition.

Dawkins, Richard. 1976. *The Selfish Gene.* Oxford: Oxford University Press.

Dawkins, Richard. 1982. *The Extended Phenotype*. San Francisco: W.H. Freeman.

Dennett, Daniel C. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life.* New York: Simon and Schuster.

Durlauf, Steven N. 1997. AReflections on How Economic Reasoning can Contribute to the Study of Science.@ Santa Fe Institute Working Paper 97-05-043.

Eldredge, Niles. 1995. *Reinventing Darwin: The Great Evolutionary Debate.* London: Weidenfeld & Nicholson.

Evans, Chris and Rydén, Göran, ARecruitment, Kinship, and the Distribution of Skill: Bar Iron Production in Britain and Sweden, 1500-1860.@ Presented to a Conference on ATechnological Revolutions in Europe, 1760-1860,@ Oslo, May 31 - June 2, 1996.

Farrell, Christopher J. 1993. "A Theory of Technological Progress." *Technological Forecasting and Social Changes* Vol. 44, pp. 161-78.

Futuyma, Douglas J. 1986. *Evolutionary Biology*. Sunderland, Mass.: Sinauer Publishers.

Gould, Stephen Jay and Vrba, Elisabeth S. 1982. "Exaptation -- a Missing Term in the Science of Form." *Paleobiology* Vol. 8, No. 1, pp. 4-15.

Gould, Stephen Jay. 1997. AEvolution: The Pleasures of Pluralism.@ *New York Review of Books*, June 26, 1997.

Hahn, Frank. AThe next Hundred Years,@ *Economic Journal* Vol. 101, pp. 47-50.

Hull, David, 1988. *Science as a Process*. Chicago: University of Chicago Press.

Jacob, Margaret. 1997. *Scientific Culture and the Making of the Industrial West*. New York: Oxford University Press.

Kauffman, Stuart A. 1995. *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. New York: Oxford University Press.

Kimura, Motoo. 1992. ANeutralism,@ in *Keywords in Evolutionary Biology*, edited by Evelyn Fox Keller and Elizabeth Lloyd. Cambridge, MA: Harvard University Press, pp. 225-230.

Lewontin, Richard. 1993. AGenotype and Phenotype.@ in *Keywords in Evolutionary Biology* edited by Evelyn Fox Keller and Elisabeth A. Lloyd. Cambridge, MA: Harvard University Press, pp. 137-44.

Lloyd, Elisabeth A. 1993. AUnit of Selection.@ in *Keywords in Evolutionary Biology*, edited by

Evelyn Fox Keller and Elizabeth Lloyd. Cambridge, MA: Harvard University Press, pp. 334-340.

Marshall, Alfred [1890] 1930. *Principles of Economics*. London: Macmillan.

Maynard Smith. 1976. AWhat Determines the Rate of Evolution?@ *American Naturalist* Vol. 110, pp 331-38.

Maynard Smith, John and Szathmáry, Eörs. 1995. *The Major Transitions in Evolution*. Oxford and New York: W.H. Freeman.

Mokyr, Joel. 1991. "Was There a British Industrial Evolution?" In *The Vital One: Essays Presented to Jonathan R.T. Hughes*, edited by Joel Mokyr. Greenwich, CT: JAI Press, pp. 253-86.

Mokyr, Joel. 1996. "Evolution and Technological Change: a new Metaphor for Economic History?" in Robert Fox, ed., *Technological Change.* London: Harwood publishers.

Nelson, Richard R. and Winter, Sidney. 1982. *An Evolutionary Theory of Economic Change*. Cambridge: The Belknap Press.

Ohlman, Herbert. 1990. AInformation: Timekeeping, Telecommunications and Audiovidual Technologies,@ in Ian McNeil, ed., *An Encyclopedia of the History of Technology*, London: Routledge, pp. 686-758.

Van Parijs, Philippe. 1981. *Evolutionary Explanation in the Social Sciences: An Emerging Paradigm*. Totowa, J.J.: Rowman and Littlefield.

Penrose, Edith T. 1952. ABiological Analogies in the Theory of the Firm,@ *American Economic Review*, Vol. XLII, No. 5 (Dec.), pp. 804-19.

Petroski, Henry. 1993. *The Evolution of Useful Things*. New York: A. Alfred Knopf.

Ramstad, Yngve. 1994. AOn the Nature of Economic Evolution.@ In Lars Magnusson, ed., *Evolutionary and Neo-Schumpeterian Approaches to Economics*. Boston: Kluwer.

Ridley, Mark. 1985. *The Problem of Evolution*. Oxford: Oxford University Press.

Sen, Amartya. 1993. "On the Darwinian View of Progress." *Population and Development Review* Vol. 19, No. 1 (March), pp. 123-37.

Stebbins, G. Ledyard. 1982. *Darwin to DNA, Molecules to Humanity*. San Francisco: W.H. Freeman.

Thomas, Brinley. 1991. AAlfred Marshall on Economic Biology,@ *Review of Political Economy,* Vol. 3, No. 1 (Jan.), pp. 1-14.

Vincenti, Walter. 1990. *What Engineers Know and How They Know It*. Baltimore: Johns Hopkins Press.

White, Lynn. 1978. *Medieval Religion and Technology*. Berkeley, University of California Press.